Essays in Financial Economics

by

Ali Kakhbod

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Economics
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2021

©Ali Kakhbod, 2021. All right reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Author ....
Department of Economics
June, 2021

Certified by ....
Daron Acemoglu
Institute Professor
Thesis Supervisor

Certified by ....
Ricardo Caballero
Ford International Professor of Economics
Thesis Supervisor

Certified by ....
Leonid Kogan
Nippon Telephone and Telegraph Professor of Management
Thesis Supervisor

Accepted by ....
Amy Finkelstein
John & Jennie S. MacDonald Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Essays in Financial Economics

by

Ali Kakhbod

Submitted to the Department of Economics
on June, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis contains four chapters on liquidity, financial crisis, dynamic pricing and optimal contracting with externalities.

The first chapter studies how transparency (information disclosure), along with long term incentive of informed dealers, affect price informativeness (efficiency), market liquidity and welfare in dynamic over the counter (OTC) markets. We show more transparency, via the public disclosure of additional information about past trades, paradoxically, makes the markets more opaque, by reducing market price informativeness. However, this market opacity creates liquidity and improves welfare. Our policy implications are threefold: (i) forward-looking incentive of informed dealers reduces price efficiency but improves liquidity; (ii) the post-trade public disclosure of prices may have no impact on price efficiency and liquidity; (iii) however, the public disclosure of past transaction orders or volumes reduces price efficiency but improves liquidity. We also derive several testable implications about price efficiency and market liquidity and demonstrate the robustness of our findings in face of a general class of payoff functions, stochastic trading positions, divisible and indivisible orders, finite and infinite trading calendars, and fixed or time-varying fundamentals.

In the second chapter we propose an amplification mechanism of financial crises based on the information choice of investors. Information acquisition always makes investors more likely to act against what is suggested by the prior. Deteriorating public news under an initially strong (weak) prior increases (reduces) the value of private information and induces more (less) information acquisition. Deteriorating public news always increases the probability of a crisis, since the initially strong (weak) prior suggests do-not-attack (attack). This effect is amplified when information choices are endogenous. To enhance financial stability, a policymaker can use taxes and subsidies to affect information acquisition. We also derive testable implications for the magnitude of amplification. This chapter is published in the *Review of Financial Studies (RFS), vol 30.*

In the third chapter we study the problem of optimal dynamic pricing for a monopolist selling a product to consumers in a social network. In the proposed model, the only means of spread of information about the product is via Word of Mouth communication; consumers' knowledge of the product is only through friends who already know about the product's existence. Both buyers and non-buyers contribute to information diffusion while buyers are more likely to get engaged. By analyzing the structure of the underlying endogenous process, we show that the optimal dynamic pricing policy for durable products with zero or negligible marginal cost, drops the price to zero infinitely often. By attracting low-valuation agents with free-offers and getting them more engaged in the spread, the firm can reach out to potential high-valuation consumers in parts of the network that would otherwise remain untouched without the price drops. We provide evidence for this behavior from smartphone app market, where price histories indicate frequent free-offerings. Moreover, we show that despite infinitely often drops of the price to zero, the optimal price trajectory does not get trapped near zero. We demonstrate the validity of our results in face of strategic forward-looking agents, homophily-based engagement in word of mouth, network externalities, and consumer inattention to price changes. We further unravel the key role of the product type in the drops by showing that the price fluctuations disappear after a finite time for a nondurable product. This chapter is published in the *Management Science (MS), vol 64*.

Finally, in the last chapter we study optimal contracting between a firm selling a divisible good that exhibits positive externality and a group of agents in a social network. The extent of externality that each agent receives from the consumption of neighboring agents is privately held and is unknown to the firm. By explicitly characterizing the optimal multilateral contract, we demonstrate how inefficiency in an agent's trade propagates through the network and creates unequal and network-dependent downward distortion in other agents' trades. Furthermore, we describe bilateral contracts (non-linear pricing schemes) and characterize their explicit dependence on the network structure. We show that the firm will benefit from uncertainty in an agent's valuation of other agents' externality. We describe the profit gap between multilateral and bilateral contracts and analyze the consequences of the explicit dependence of the contracts on network structure. When the network is balanced in terms of homogeneity of agents' influence, network structure has no impact on the firm's profit for bilateral contracts. On the other hand, when the influences are heterogeneous with high dispersion (as in core-periphery networks) the restriction to bilateral contracts can result in profit losses that grow unbounded with the size of networks. This chapter is published in the *Journal of Economic Theory (JET), vol 183*.

JEL Classification: G1, G2, G3

Thesis Supervisor: Daron Acemoglu
Title: Institute Professor

Thesis Supervisor: Ricardo Caballero
Title: Ford International Professor of Economics

Thesis Supervisor: Leonid Kogan
Title: Nippon Telephone and Telegraph Professor of Management

Thesis committee members

Professor Daron Acemoglu
Institute Professor
Department of Economics, MIT

Professor Ricardo Caballero
Ford International Professor of Economics
Department of Economics, MIT

Professor Leonid Kogan
Nippon Telephone and Telegraph Professor of Management
Sloan School of Management, MIT

Professor Andrey Malenko
Associate Professor of Finance
Ross School of Business, University of Michigan

Acknowledgments

I am indebted to many people who supported me throughout my PhD.

I am incredibly fortunate to have had Daron Acemoglu, Ricardo Caballero, Leonid Kogan, and Andrey Malenko as my advisors.

I could not have asked for more inspiring and devoted advisor than Daron. I can only admire the depth and clarity of his thought, that have greatly improved all my work. I will strive to emulate his curiosity, his persistence, and his generosity and empathy towards students. I cannot overstate how influential and supportive Daron has been. I am tremendously indebted to him for my personal and professional development.

I am deeply grateful to Ricardo for his profound insights, optimism, endless support, steady encouragement, and his amazing vision on financial models and their connection to reality. His devotion to his students and his generosity with his time will always be an inspiration for me.

I owe a great debt of gratitude to Leonid. Since the spring of 2017 that I took his asset pricing class he has supported me unconditionally. Leonid has been a constant source of inspiration and motivation for me to think about interesting problems. I am especially grateful to him for believing in me and supporting me all throughout these years.

Andrey's research and style have had a tremendous influence on me. I am greatly indebted to him for the long hours he spent discussing, reading, and improving my ideas through sharp suggestions. He supported me unconditionally and provided me with economic ideas and great comments. I have been so fortunate to benefit from Andrey's mentorship and advice throughout my PhD.

I have been fortunate to benefit from conversations with many students and faculty at MIT and elsewhere. Especially I am grateful to Ali Jadbabaie and Asu Ozdaglar who have always been there for me as sources of support, encouragement, and inspiration. I would also like to thank Demos Teneketzis for stimulating my interest in economics.

Throughout the years at the MIT economics department, I have also benefited enourmousely from various discussions with Marios Angeletos, Glenn Ellison, Maryam Farboodi, Drew Fudenberg, Robert Gibbons, Bengt Holmström, Stephen Morris, Parag Pathak, Steve Ross, Alp Simsek, Juuso Toikka, Rob Townsend, Ivan Werning, Alex Wolitzky, Muhamet Yildiz and Haoxiang Zhu. I am also grateful to Amir Kermani and Nadya Malenko for numerous conversations, great comments and suggestions throughout my PhD.

I am lucky to be surrounded by exceptionally brilliant and humble colleagues and friends. Especially I am grateful to Toni Ahnert, Amir Ajorlou, Josh Bosshardt, Oğuzhan Çelebi, Mohammad Fallahi, Aria Ghasemian Sahebi, Ömer Karaduman, Mazi Kazemi, Giacomo Lanzani, Pooya Molavi, Kevin Li, Uliana Loginova and Fei Song. I would also like to thank Gary King for his very kind assistance since my first day at MIT. I also thank Emily Gallagher, Julia Martyn-Shah and Shannon May for their great help during my job market year. For the generous financial support, I am thankful to Noosheen Hashemi, Farzad Nazem and the HAND foundation scholarship program.

Last but not least, I am thankful to my parents Parvin and Ghavam, my brother Payam and my sister Parnaz for their unconditional love and support throughout my life. Countless others deserve mention yet are not named. To all of them, I am grateful.

Contents

List of Figures

List of Tables

Chapter 1

Liquidity, Welfare and Transparency in Dynamic OTC markets

1.1 Introduction

A large portion of trade occurs in dealer-based over-the-counter (OTC) markets for almost all financial assets (e.g., municipal bonds).[1] A dealer-based OTC market does not use a centralized trading mechanism, such as an auction or limit-order book, to aggregate bids and offers and to allocate trades. An OTC trade negotiation is initiated when a trader (retailer, customer) contacts the dealer (broker, market maker) and asks for terms of trade. Next, the dealer typically provides a price to a trader. He can then choose to accept it or reject it. For this reason, OTC markets are relatively opaque and traders are somewhat in the dark about the most attractive available terms and about whom to contact for those

This chapter is coauthored with Fei Song. I am extremely grateful to Daron Acemoglu, Ricardo Caballero, Leonid Kogan and Andrey Malenko for their invaluable guidance and support. We also greatly thank Marios Angeletos, Ehsan Ebrahimy, Glenn Ellison, Maryam Farboodi, Tom George, Amir Kermani, Praveen Kumar, Nadya Malenko, Stephen Morris, Borghan Narajabad, Parag Pathak, Juuso Toikka, Muhamet Yildiz, Iván Werning and Haoxiang Zhu for their great comments and suggestions.

[1]For example fixed income securities are mainly traded over-the-counter, such as swaps, bonds and repos. For instance, Nagel (2016) finds that 95% of electronic swaps trades are over-the-counter. Moreover, Tuttle (2014) shows 16.99% of total dollar volume (18.75% of share volume) of National Market System (NMS) stocks is executed OTC without the involvement of an alternative trading systems. Common justifications for OTC trading include: regulatory barriers, nonstandardization and asset complexity.

terms.[2]

A dealer-based OTC market is also complex because dealers have private information about asset values that may dynamically change over time with the arrival of good and bad news (economic states).[3] This happens because the dealer frequently observes the order flow, negotiates trade terms and gathers information, while the trader (her customer) usually has more limited opportunities to trade and thus relatively less information about recent economic states. In this case, such disparity in market access is relatively common knowledge and tends to convey a bargaining power to the dealer (this particularly holds in municipal bonds).[4]

A common concern about OTC markets is their opaqueness. Given the important role that OTC markets played in the global financial crisis,[5] many regulators have attempted to shed some light on those so-called dark markets. Perhaps the most notable reform aiming to increase transparency was the U.S. Dodd-Frank Act, implemented after the 2008 financial crisis.

The transparency requirements of the U.S. Dodd-Frank Act (through Trade Reporting and Compliance Engine (TRACE))[6] aim to promote market stability (e.g., by improving market liquidity and efficiency) through two types of regulations:

(i) the public disclosure of some aggregate information on trading volumes;

(ii) the public disclosure of previous transaction prices.

How does information disclosure (through the above two types) affect transparency (price

[2]For empirical evidences see for example Green, Hollifield and Schürhoff (2007*a*), Ashcraft and Duffie (2007) and Massa and Simonov (2003). There are also several empirical works that show dealers quote prices have smaller spreads to customers who are likely to be uninformed (Linnainmaa and Saar (2012)), and OTC trades are less informative compared to trades on exchanges (Bessembinder and Venkataraman (2004)). For more institutional backgrounds see Section 1.2.

[3]Good news tends to have a positive effect on markets and one can see asset value's drift rising while its volatility falling soon after the news come out in the open. This story is quite opposite after a bad news. In fact, bad news tends to have a negative effect on markets so that asset value's drift falling while its volatility rising soon after the news come out in the open. See for example Kothari and Warner (1997), Fama (1998), Daniel, Hirshleifer and Subrahmanyam (1998) and Hong, Lim and Stein (2000) that all have excellent synopses of the literature on stock price reactions to various events.

[4]For more see Section 1.2. There are several empirical evidences supporting imperfect competition in a dealer-based OTC market where a monopolistic dealer offers quote prices to (unsophisticated) customers. For example, see Green, Hollifield and Schürhoff (2007*a*) that document dramatic variation across investors in the prices paid for the same municipal bond. See also Massa and Simonov (2003) that report dispersion in the prices at which different dealers trade the same Italian regulator bonds.

[5]For a reference, see Duffie (2012) and Duffie (2017).

[6]See Title VII, Wall Street Reform and Consumer Protection Act.

informativeness), liquidity, and welfare? Can these mandates paradoxically lead to more market opacity (i.e., less market price informativeness)? Do post-price and post-order disclosures act the same? If not, what are the differences?

In this paper we provide answers to these questions.

In Section 3.2, we offer a general dynamic model of tradings. In our model asset values (e.g., dividends or PV of investments) change over times in a regime-switching framework.[7] There are two sources of uncertainty that directly contribute to the asset value diffusions: (1) volatility shocks, and (2) drift shocks. To incorporate these uncertainties, we model the diffusion of the asset values by a linear dynamic process and embed an ergodic Markov chain into it to model the economic states diffusion.[8] The current economic state is only known by the dealer. The dealer has the market power to quote prices and liquidity traders (customers) are the potential sellers or buyers of the asset in each period.[9] At each time, given the dealer's quote price a trader decides to accept or reject the proposed price offer. The dealer is strategic, forward-looking and risk-neutral, and traders are strategic, myopic and risk-averse. Therefore, a trader is motivated to trade by his risk sharing motive. We focus on analyzing perfect Bayesian equilibrium in this dynamic trading game.

Section 1.4 briefly considers the static (one-period) variation of the model (i.e., dealer only cares about maximizing her within-period payoffs). This static variation specifies equilibria that are useful to analyze the full dynamic model. In fact, we show depending on how desperate a trader is, there are only two types of equilibria in the static model. When a trader is sufficiently conservative and dislikes the volatility in asset values, meaning that trader's risk-aversion is sufficiently high so that he is so desperate to hedge against uncertainty shocks, there exists an opaque (pooling) equilibrium in which dealer fully **conceals** her private information. Of course, hiding information, whenever possible, is best the dealer (ex-ante, i.e., before observing her private signal) can do, as by doing so she gains the maximal information rent from her private information. Moreover, market liquidity (trade activity) and risk sharing in this opaque equilibrium increase, because trades occur in both good times and bad times. What if a trader's motive to hedge against uncertainty shocks is moderate? This introduces a condition to support the second equilibrium. In this case, it is too costly for the dealer to conceal her private information. As a result, it is impossible to

[7]Change in the asset value is due to underlying economic state that may change over time. In subsequent sections we allow the state to be fixed.

[8]Therefore, our model also includes business cycles and macro economic shocks (boom and recession).

[9]In subsequent sections we allow traders' trading positions to stochastically change from a seller to a buyer and vice-versa from period to period.

persuade the trader to accept the uniform (uninformative) price. Consequently, she should **reveal** her private information about the economic state, leading to a revealing (separating) equilibrium. Importantly, in contrast to the previous opaque equilibrium, now trades only occur in trader's favorable times (that is, in good times if traders are sellers and in bad times if traders are buyers). As a result, the amount of market liquidity (trade activity) falls, so do the amount of risk sharing and the dealer's ex-ante static profit.

We note that in this static case, since market participants act as if this is a one-shot game (trade), information disclosure (i.e., whether TRACE is in place or not) has no bite on the structure (and construction) of the pooling equilibrium. In the next two sections we consider the full dynamic model where dealer is forward looking (i.e., the dealer also cares about future cash flows by discounting her continuation payoffs).

Section 1.5 analyzes the case where the TRACE is in place (i.e., the history of past trades is fully observable) and dealer is forward looking. We show in this case dealer conceals her information more than the the static case (implying that dealer's myopia improves market efficiency as it increases price informativeness). This is because when the informed dealer is forward-looking, and TRACE is in place, the above static opaque (pooling) equilibrium becomes easier to sustain. To see this, note that a deviation to decline a transaction in trader's unfavorable times has a future cost for the dealer. Even though she can avoid the loss in the current period, since TRACE is in place, the following traders observe her deviation and will expect to play the revealing (separating) equilibrium, in all future periods. Traders will adjust their expectations and the dealer can only collect at most her revealing equilibrium payoff. Such payoff is achievable when she reveals her private information in all future periods and is lower than her on-path continuation payoff. Such loss of continuation payoffs discourages dealer's deviation and makes the opaque pricing scheme easier to sustain under the dynamic trading setting. Therefore, in compare to the static case, price opacity increases when dealer is forward-looking and TRACE is in place.[10] Moreover, market liquidity and dealer's ex-ante profit will both rise.[11]

[10]Hence, this result suggests that the long-term incentive harms financial stability by reducing the market price informativeness, which is in contrast with the Dodd-Frank section 956. That section strongly encourages long-term incentive based compensation schemes for managers (e.g., dealers), inducing them to become more forward-looking.

[11]We also show that there exist an equilibrium in which the informed forward-looking dealer fully reveals her private information in all the trades. For such equilibria to exist, traders' risk-aversion should be high enough to induce them to trade with the informed dealer in trader's unfavorable time as well. As a result, in contrast to the previous informative equilibrium, trade occurs in both times. However, whenever this equilibrium exists, so does a dynamic opaque one. As the dynamic opaque equilibrium results in the maximum dealer's ex-ante static payoff, this dynamic fully revealing equilibrium is never chosen.

How does TRACE affect the structure of the pooling equilibrium when dealer is forward looking? Section 1.6 addresses this question. There we add an important component of the model that the past history of the **volumes and prices** may not be fully available (i.e., **private history**), and dealer is forward-looking. Particularly, in line with the U.S. Dodd-Frank Transparency Act of 2010, for corporate and municipal bonds and swaps (which follow TRACE's definition of transparency and require the public dissemination of post-trade transaction information regarding price and volume),[12] we consider three cases:

1. Past prices are not observable (lack of post-trade price transparency).[13]

2. Both past prices and past transaction orders (volumes or trades) are not observable.

3. Past prices are not observable, but traders observe signals about past transaction orders (some aggregate information on past trading volumes).

We show that the consequence of public information dissemination of **past transaction prices** can be different from that of **past trading volumes**. In the first case we show that the lack of knowledge about past prices has no impact on the structure of the above dynamic opaque equilibrium (specified in Section 1.5). Therefore, post-trade price transparency does not necessarily improve price efficiency. This result holds because the knowledge of past transaction orders are **sufficient statistics** for a previous deviation. One can implement the dynamic opaque equilibrium with no change when only order information is available. As a consequence, with post trade price transparency, trade activity may not increase (consistent with the empirical evidence in Asquith, Covert and Pathak (2013)).

More surprisingly, in the second and the third cases, we show that the public information disclosure about previous transaction orders (volumes), paradoxically, makes the market more opaque by reducing price informativeness. The intuition, which we call the *reputation building (or commitment device) mechanism*, is as follows. The informed dealer has an incentive to achieve a *reputation* of *no-revelation-history* so that in future periods she can extract the information rent and take advantage of traders' hedging motive. Therefore, the availability of past trade details enables such a reputation building and provides the dealer an incentive to hide her private information and maintain the *no-revelation-history*.

[12] Similar reforms have been proposed for public transactions reporting in swap execution facilities (SEFs). Japan and Europe (in a more ambitious framework known as MiFID II and MiFIR) have followed a similar course as the United States (Duffie (2017)).

[13] Post-trade price transparency for (almost) all U.S. corporate bonds and some other fixed-income instruments have actually been mandated by the SEC since 2002, via the Transaction Reporting and Compliance Engine (TRACE).

In addition, since in pooling equilibrium trades occur in both good and bad times, post-trade public disclosure regarding *volumes* via TRACE can improve the liquidity (trading volume), (consistent with the empirical evidence in Bessembinder and Maxwell (2008))

In Section 1.7 we consider social welfare and dealer's profit. We show that more information disclosure about past trades increases expected welfare and expected surplus of the dealer.[14] This is mainly because, as shown above, with more information disclosure about past trades, the opaque (pooling) trading equilibrium is easier to sustain. Therefore, this result shows there exists a trade-off between price informativeness, liquidity, and social welfare in the OTC markets (e.g., municipal and corporate bond markets). Thence, welfare can be decreased if regulators also care about price informativeness, implying that implementing TRACE may have mixed effects on welfare. In this regard, to enhance financial transparency via price informativeness and simultaneously boost market liquidity and achieve the highest feasible social welfare we propose a policy to *randomly audit dealers*. We show that increasing auditing intensity can force the dealer to reveal her private information about economic states more frequently and, as a result, this leads to more price informativeness.

How robust are the above insights? In Section 1.8 we demonstrate the robustness of our model in several important extensions. First, we show that our main conclusions do not depend on deliberate functional forms and can go beyond mean-variance utility functions. Second, we extend the analysis to the case where traded orders are divisible. We show divisible trades have no impact on the opaque equilibrium structure. However, in contrast to the discrete order case, in the revealing equilibrium trade also occurs in trader's unfavorable times (but its amount is still strictly less than that in the other times). But other than this nothing changes structurally. Third, we allow trader's demand shock to stochastically change over times between a seller and a buyer.[15] We show that our results are robust to such a change. Fourth, we present a sufficient and necessary condition for the existence of a semi-opaque trading equilibrium in static game. In the semi-opaque trading equilibrium, the dealer always trades with traders in their favorable times and mixes between trading and not trading in the other times. The expected ex-post social welfare of this class of equilibrium lies between that of opaque equilibrium and that of informative equilibrium, and increases in the probability of trading. This result reinforces our previous one that there is a trade-off

[14] This result may appear surprising, as it goes against the general lesson of contract theory that less disclosure gives more information rent to the party with private information. However, it holds because with more information disclosure about past trades, opaque trading equilibrium is easier to sustain, impairing the price efficiency and improving dealer's ex-ante profit.

[15] That is, traders' demand shock is modeled as a binary stochastic process changing over time between seller and buyer.

between price informativeness and social welfare. Finally, we extend the analysis to the case where the trading calendar is finite, the fundamental is fixed, and traders have heterogenous risk aversions, which is their private information. Given all these important changes, we still manage to show that the main insights of the model continue to hold. Thence, these extensions together show that our main takeaways are robust to different changes in model specification.

Finally, Section 1.9 concludes.

1.1.1 Related Literature

This paper is part of the growing literature on dealer-based OTC markets. Prior literature on OTC markets focuses on the dealers' ability to contract with customers (Grossman and Miller (1988)), discrimination based on order size (Seppi (1990)), welfare effects of decentralized trading (Malamud and Rostek (2017), Babus and Parlatore (2017)), price movements in OTC markets when block orders are large (Grossman (1992)), searching for good price in OTC markets with multiple dealers (Zhu (2012)), random search and matching in large markets among a continuum of traders (Duffie, Gârleanu and Pedersen (2005), Lagos and Rocheteau (2009)), private information transmission in a single pairwise exchange via a mechanism design approach (Shimer and Werning (2019)), "small" informed traders trading with uninformed dealers learning from aggregate trade (Lester et al. (2018)), dynamic price discovery and (potentially) wrong learning with sequential trading (Kakhbod and Song (2020*a*)), speed competition and segmentation between multiple dealers in illiquid markets (Çelebi, Kakhbod and Li (2020)), adverse selection with search frictions and discrete trading opportunities (Guerrieri and Shimer (2014)) and less competition in equity trades that are sent to dark pools (Zhu (2014)).[16] In contrast to these works, our focus here is on the role of public information disclosure in price efficiency (informativeness), market liquidity and social welfare with informed, forward looking dealers.[17]

[16]See Duffie (2012) for an excellent review of the literature on dealer based OTC markets.

[17]Our work also contributes to the literature of information disclosure in financial markets. See Verrecchia (2001), Goldstein and Sapra (2013) and Goldstein and Yang (2017) for excellent reviews on this topic. For example, Back, Liu and Teguia (2020) show disclosure of the terms of a transaction reduces the dealer's rents when she disposes of the inventory in a second transaction. A few recent papers present models in which disclosure can harm price informativeness, although through different channels than ours. For example Goldstein and Bond (2015) and Goldstein and Yang (2018) construct noisy REE models and show that the type of information being disclosed is important in determining whether disclosure is desirable. In particular, when disclosure is about a variable that the real decision maker cares to learn, disclosure can harms price informativeness. In the same realm of REE models, Banerjee, Davis and Gondhi (2018) show that lowering the cost of information concerning the fundamentals may induce traders to learn more about noise. As

Moreover, our market liquidity results contribute to the literature of market liquidity with financial frictions. In particular, other studies have also argued that recent liquidity in bond markets has been improved by crisis-induced regulations (see e.g. Section 619 part of the Dodd-Frank Act) (Adrian et al. (2013), Duffie (2017)) and capital and leverage ratio requirements for banks (Caballero (2010)). In this regard alternative theories of market liquidity have been proposed. See Brunnermeier and Pedersen (2005) for predatory trading, Gorton and Pennacchi (1990) for adverse selection in secondary debt markets, Shleifer and Vishny (1992) and Kiyotaki and Moore (1997) for financial stress of the natural buyer, Caballero and Krishnamurthy (2008) for Knightian uncertainty, Guerrieri and Lorenzoni (2009) for precautionary behavior on business cycle fluctuations, Bebchuk and Goldstein (2011) and Goldstein (2012) for panics, self-fulfilling crisis and credit freezes, and Ahnert and Kakhbod (2018) for costly private-information acquisition. In contrast to these theories, our results on improving market liquidity is due to a new channel based on the extent of available information about past transactions, dealer's long-term incentive, and traders' hedging motives.[18]

Finally, this work contributes to the literature on market microstructure. There is an earlier literature that considers the liquidity demand side to be more informed (e.g., Glosten and Milgrom (1985), Easley and O'Hara (1987), Admati and Pfleiderer (1988)).[19] Importantly, however, in sharp contrast to these important works, here dealer is informed, strategic and forward looking, further, we isolate the distinct impacts of disclosing past prices and orders on price informativeness, liquidity and welfare in dynamic dark markets.

a result, information efficiency may decrease. Amador and Weill (2010) construct a monetary model and show that releasing public information about productivity shocks can reduce welfare through reducing price informativeness.

[18]The comparison between public offers and private offers in our paper is related to the repeated game literature with imperfect monitoring (Abreu, Pearce and Stacchetti (1990), Abreu, Milgrom and Pearce (1991), Fudenberg, Levine and Maskin (1994), Fudenberg and Levine (1994)). However, rather than focusing on the characterization of all possible perfect public equilibrium payoffs, we are more interested in when opaque (pooling) equilibria can hold. We show that hidden past orders can encourage dealer to share her private information. For other relevant liquidity papers see Bosshardt and Kakhbod (2020*a*), Bosshardt and Kakhbod (2020*b*).

[19]This paper is also related to the literature of the market efficiency in strategic informed trading, dated back to Kyle (1985*a*,*b*)'s seminal articles. Wang (1993, 1994) consider an infinite-horizon model where competitive insiders receive information on a firm's dividends over time in steady-state. They show that risk-neutral competitive insiders will reveal their private information instantly whereas risk-aversion can reduce their trading aggressiveness, leading to a slower information revelation. Back and Pedersen (1998) consider a finite-horizon model with a monopolistic informed insider and show that the insider reveals her information gradually. Chau and Vayanos (2008) consider the market efficiency in an infinite-horizon model with a monopolistic insider trading with competitive dealers and noisy traders as well. They discover that the insider chooses to reveal her information quickly, as the market approaches continuous trading.

1.2 Institutional Background

The main features of the model, we present in the next section, are particularly consistent with the municipal bond (and corporate bonds with high search cost) traded on OTC markets. In this section we provide some institutional backgrounds and facts about OTC markets particularly municipal bond markets and properties of market participants in these markets. These facts are the main components of the model we present in the next section.

OTC market is huge, many trades occur in these markets (e.g., corporate bonds, asset backed securities (ABS), money markets) and municipal bonds (see Figure 1-1).

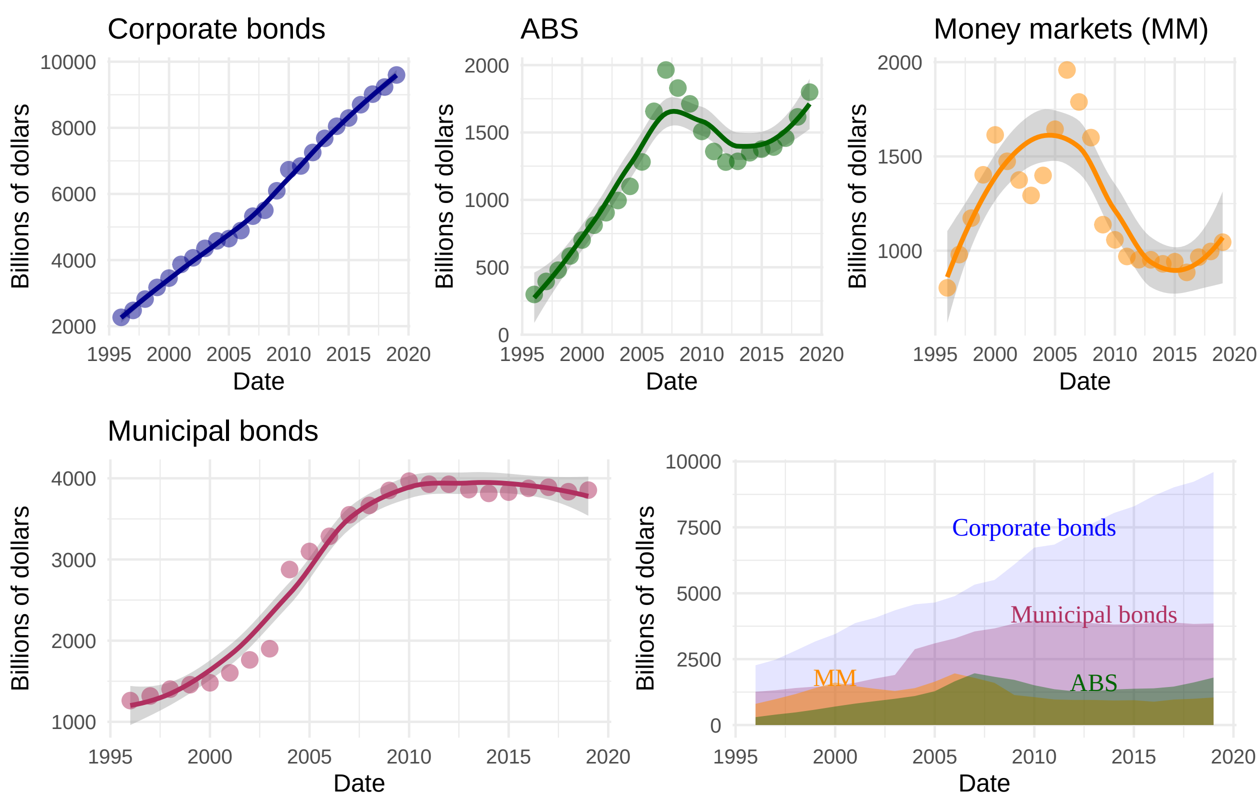

Figure 1-1: This panel shows the evolution of corporate bonds, ABS, money markets and municipal bonds in Billions of dollars outstanding, trading in OTC markets.

Municipal bonds (background). Municipal bonds (munis) are bonds issued by local governments or government sponsored entities to finance such capital projects as construction of highways, stadiums, hospitals and schools. The municipal securities market is a vast market with enormous diversity. According to Bloomberg, approximately $3.8 trillion ($3800 billion) of municipal securities are outstanding (see Figure 1-1). Investors' income from municipal bond coupons are exempt from federal tax, and in most cases, state tax (if the holder lives in the issuing state). 80% of municipal bonds are AAA rated or insured, which makes default very unlikely, compared with corporate bonds. Even in the rare occasional default, the recovery rate is usually very high. The no tax, low risk features make municipal bonds especially attractive to investors in high tax brackets. The long term nature of the bonds also makes municipal bonds an ideal instrument for planning retirement.

Customers in municipal bonds. Retail investors hold most of the municipal bonds. Individuals directly hold 38% of all municipal bonds in the first quarter of 2006, according to Bond Market Association. Including open-end and closed-end mutual funds, individuals hold more than 56% of the market. The large retail clientele leads to a unique features of the municipal bond market that is retail investors (customers) are less sophisticated than their institutional counterparts.

Feature 1: **Customers are less informed/sophisticated.** Retail investors' lack of sophistication is a direct consequence of their lack of information and technologies. Retail investors (customers) are not well informed about the activities of municipalities, and they also lack the ability to accurately value bonds given their call or put options or other complex features. Retail investors usually do not have direct access to information centers such as Bloomberg, where searching, pricing, calculating tax and viewing current offerings can be done easily, nor do they have a large contact list to solicit bids. These limit their ability to bargain with dealers for better prices.

Dealers/brokers' role. Municipal bonds are traded in OTC markets. Buyers and sellers do not trade directly. All transactions are conducted through middlemen — dealers/brokers/market makers (Nasdaq is also an OTC market, however, buyers and sellers can directly transact through limit orders). The existence of an intermediary reduces the inefficiency of bilateral search and enhances liquidity. However, dealers have large market power in these markets.

Feature 2: **Dealers' local monopoly power and private information.** Indeed, there are several factors that hinder the municipal bond market from being competitive. First, municipal market is highly fragmented. With over 1.3 million bonds outstanding, the municipal bond market easily dwarfs the equity market, corporate bond and treasury bond markets, in terms of number of instruments. Liquidity at any moment in time, is thinly spread over the vast number of bonds. In fact, most bonds hardly trade after issuance. Thus, direct competition among dealers is naturally low. Second, the municipal bond market has been notoriously opaque (see Green, Hollifield and Schürhoff (2007*a*,*b*)), more on opacity in OTC markets comes in the next paragraph.

In addition, in these markets dealers are more informed than the customers. Particularly, Green, Hollifield and Schürhoff (2007*a*,*b*) used a mixture model to uncover the hidden variable "informedness" that is not observable by econometricians, but known to the dealers, and were able to calculate dealer's profit against uninformed buyers/customers. They provide first-hand evidence of dealer's discriminative pricing based on perceived sophistication.

Feature 3: **Opacity and regulations.** Finally, OTC markets in general are opaque. And, their opaqueness is a common concern. Many regulators have attempted to improve transparency in these markets. The most notable reform aiming to increase transparency was the U.S. Dodd-Frank Act, implemented after the 2008 financial crisis.

The transparency requirements of the U.S. Dodd-Frank Act (through TRACE) aim to improve market transparency through two types of regulations: (i) the public disclosure of some aggregate information on trading volumes; (ii) the public disclosure of previous transaction prices of standardized derivatives. For some OTC markets, such as those for U.S. corporate and municipal bonds, regulators have mandated post-trade transparency (price and volume) through publicly announcing an almost complete record of transactions shortly after they occur. Empirical analyses about implications of such mandates in bond markets, through TRACE, are documented in Edwards, Harris and Piwowar (2007), Green, Hollifield and Schürhoff (2007*a*), Green, Hollifield and Schürhoff (2007*b*), Bessembinder and Maxwell (2008), Goldstein, Hotchkiss and Sirri (2007), Green, Li and Schürhoff (2010) and Asquith, Covert and Pathak (2013). See Duffie (2012) for an excellent review of transparency requirements of the U.S. Dodd-Frank Act.

The above *Features 1-3* hold in the model we present in the next section.

1.3 Model

We consider an infinite-horizon dynamic trading game between an informed, risk-neutral and forward-looking **dealer** (broker, market maker) and a series of uninformed, risk-averse and myopic **traders** (retailer).[20,21] Time is discrete $t \in \{1, 2, 3, \cdots\}$. At each period, a trader comes to the dealer, either possessing an asset or desiring an asset to hedge his other investments. The market asset value a_t changes with the arrival of good and bad news over time.[22] Suppose θ_t represents the underlying economic **state**, which is known to the dealer but not traders. The dealer has the bargaining power and in each period offers a take-it-or-leave-it price p_{t,θ_t} to the trader. The trader then makes a decision o_t of whether to accept the dealer's offer or not. The size of the asset traded in each period is normalized to 1, and we extend the analysis to divisible orders in Section 1.8.2.

The main components of the model are discussed in details below.

1.3.1 Good Times and Bad Times: Asset Value Dynamics

The asset value at time t, denoted by a_t, is given by the following dynamic:

$$a_t = \varphi a_{t-1} + \text{innov}_t,$$

where $\varphi \in [0, 1]$ is the persistence coefficient, and innov_t denotes the stochastic **innovation** in the asset value at time t:

$$\text{innov}_t = J_{\theta_t} + \sigma_{\theta_t} z_t. \tag{3.1}$$

z_t (idiosyncratic shock) is an independent and normally distributed process with mean zero and variance one, i.e. $z_t \sim \mathcal{N}(0, 1)$. θ_t is an ergodic Markov chain that takes values from $\Theta \equiv \{\text{g}, \text{b}\}$, where g stands for good times and b for bad times.[23] Consistent with the empirical literature, in good times the asset value has a higher mean and a lower variance, while in bad times the asset value has a lower mean and a higher variance. J_{θ_t} denotes the state-dependent **drift** and σ_{θ_t} denotes the state-dependent **volatility** of the innovation

[20] In this framework the dealer acts as a broker/market maker.

[21] A finite horizon version of the model is analyzed in the extension section.

[22] Time "t" asset can be different from time t' asset. That is, trading assets between the broker and retailers may change over time.

[23] In section 1.8.5 we consider the case where the fundamental is fixed.

shock at time t when the economic state is θ_t. In other words, we assume

$$J_{\mathrm{g}} > J_{\mathrm{b}} \quad \text{and} \quad \sigma_{\mathrm{g}} \leq \sigma_{\mathrm{b}}. \tag{3.2}$$

We assume that the state transition follows a Markov process (illustrated in Figure 1-2):

$$\begin{aligned} &\mathrm{Prob}\{\theta_t = g|\theta_{t-1} = g\} = \alpha_g, && \mathrm{Prob}\{\theta_t = b|\theta_{t-1} = g\} = 1 - \alpha_g, \\ &\mathrm{Prob}\{\theta_t = g|\theta_{t-1} = b\} = \alpha_b, && \mathrm{Prob}\{\theta_t = b|\theta_{t-1} = b\} = 1 - \alpha_b, \end{aligned} \tag{3.3}$$

where θ_t is the economic state in period t and $\alpha_g, \alpha_b \in (0, 1)$. In addition, z_t and θ_t are independent, i.e. $z_t \perp \theta_t$, for all t. We assume that $J_g, J_b, \sigma_g, \sigma_b, \alpha_g, \alpha_b$ and φ are common knowledge to both parties. Moreover, at the end of period t, asset value a_t and the economic state θ_t become publicly observable after the trade.[24]

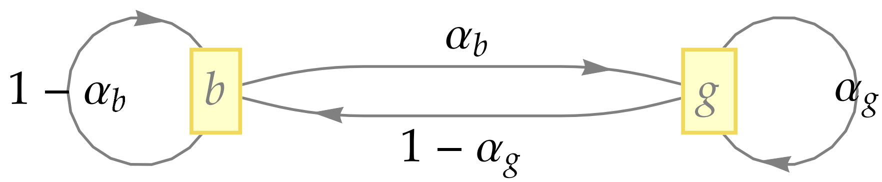

Figure 1-2: Markov chain dynamics (transition probabilities) between good times g and bad times b.

1.3.2 Demand Shocks

At the beginning of period t, the trader's demand shock χ_t is independently drawn from $\{-1, 1\}$ such that

$$\mathrm{Prob}\{\chi_t = 1\} = \beta \in [0, 1], \quad \mathrm{Prob}\{\chi_t = -1\} = 1 - \beta,$$

where $\chi_t = 1$ represents that the trader coming at period t possesses an extra unit of risky asset to sell, while $\chi_t = -1$ represents that he is in need of a unit of risky asset,[25] and without purchasing one, he has to pay the realized value of the asset to someone else. The trading position χ_t is also observable to the dealer.

[24] A variation of the model in which the fundamental is not revealed to future traders is analyzed in the extension section.

[25] In other words, $\chi_t = -1$ ($\chi_t = 1$) means that the retailer is willing to short (long) the asset.

Hence, at period t, a trader has three order decisions $o_t \in \{-1, 0, 1\}$: $o_t = -1$ means that the trader accepts the bid offer and sells the asset to the dealer;[26] $o_t = 1$ means that the trader accepts the ask offer and buys the asset from the dealer; $o_t = 0$ represents that the trader declines the dealer's offer. Given the trader's trading position, a trader who is in need of an asset will never accept a bid offer from the dealer, and a trader who possesses an extra asset will not like to accept an ask offer. This puts restriction on the feasibility of $\{o_t\}_{t=1}^{\infty}$:

Case (1): When $\chi_t = 1$, the trader is a potential seller. Hence, $o_t \in \{0, -1\}$, i.e., $o_t = 0$ means keeping the asset and $o_t = -1$ means selling the asset.

Case (2): When $\chi_t = -1$, the trader is a potential buyer. Hence, $o_t \in \{0, 1\}$, where $o_t = 0$ means rejecting dealer's price offer and $o_t = 1$ means buying the asset from the dealer.

1.3.3 Information Disclosure: Public vs. Private history

In order to analyze the effect of information disclosure, we compare two variants about the observability of the past history of trades.

In our **public history** model, the past history of trades (including prices and transaction orders) is publicly observable. In this case TRACE is in place. In other words, we assume that past transactions are public and future traders can view prices offered to previous traders and whether they are accepted or not.

The assumption that previous trade history is publicly observable is relaxed in several directions in Section 1.6, where we present the **private history** variant of our model. In Section 1.6.1 we discuss the case where although past prices are unobservable, future traders can still perfectly learn whether a previous trade happens, that is, only the bilateral transaction price is kept private from the informed dealer and the trader who received it. In Section 1.6.2 we further relax that assumption and examine a case where future traders can only observe imperfect signals about whether trade happens before. As we will show in Section 1.6, learning previous prices has no effect on the equilibrium behaviors. It is the information of orders (volumes) that matters.

[26] Here we use negative value of o_t to represent that by selling the asset, the trader actually loses one unit of the asset.

1.3.4 Beliefs

Let h^{t-1} denote the past history of trades available at the beginning of period t. Each trader is rational. Strictly speaking, at the beginning of each period t, he Bayesian updates his prior belief about the economic state $\theta_t = (J_{\theta_t}, \sigma_{\theta_t})$ from the past history, h^{t-1}, and the new price offer, p_{t,θ_t}, from the dealer,

$$\xi(p_{t,\theta_t}; h^{t-1}) \equiv \text{Prob}\{\theta_t = g|h^{t-1}, p_{t,\theta_t}\} = \frac{\text{Prob}(\theta_t = g|h^{t-1})\text{Prob}(p_{t,\theta_t}|\theta_t = g, h^{t-1})}{\sum_{\theta_t} \text{Prob}(\theta_t|h^{t-1})\text{Prob}(p_{t,\theta_t}|\theta_t, h^{t-1})}.$$

1.3.5 Payoffs

In our main model we assume that dealer is risk-neutral and trader has mean-variance preferences. In Section 1.8.1 we will show that our main conclusion does not depend on specific functional forms.

Traders' payoffs. Each trader's end-of-period wealth is given by

$$w_t^{\chi_t} = \chi_t a_t + o_t \cdot (a_t - p_{t,\theta_t}), \tag{3.4}$$

The myopic, rational trader is **risk-averse**. His ex-ante utility at the beginning of period t, given the available information about past trades, and the price offer, p_{t,θ_t}, from the dealer, is given by

$$\mathbb{E}\left[w_t^{\chi_t}\middle|h^{t-1}, p_{t,\theta_t}\right] - \frac{\rho}{2}\text{Var}\left[w_t^{\chi_t}\middle|h^{t-1}, p_{t,\theta_t}\right], \tag{3.5}$$

where ρ denotes his risk-aversion coefficient. Therefore, as trader dislikes the volatility of his end-of-period wealth, he has an incentive to trade with the dealer due to this risk sharing motive.

Dealer's payoff. The end-of-period wealth of the **risk-neutral** dealer (market maker) is also given by

$$u_t \equiv (p_{t,\theta_t} - a_t) \cdot o_t.$$

The dealer, instead, is **forward-looking** and her utility, given the history h^{t-1} (that depends on the underlying information disclosure protocol) and her private information about the economic state θ_t, is given by:

$$\begin{aligned} U_t &= (1-\delta)\mathbb{E}\left[\sum_{s\geq t}\delta^{s-t}(p_{s,\theta_s}-a_s)\cdot o_s \middle| h^{t-1},\theta_t\right] \\ &= (1-\delta)\mathbb{E}[u_t|h^{t-1},\theta_t] + \delta\mathbb{E}[U_{t+1}|h^{t-1},\theta_t], \end{aligned} \tag{3.6}$$

where $\delta \in [0,1)$ is dealer's discount factor. $\delta = 0$ represents the case where the dealer becomes fully myopic like traders (retailers).

1.3.6 Equilibrium Concept and Refinement Criterion

Throughout this paper, the solution concept considered is perfect Bayesian equilibrium, defined below.

Definition 1.3.1. *A perfect Bayesian equilibrium (PBE)* $\{p^*(\cdot), o^*(\cdot), \xi^*(\cdot)\}$ *consists of the dealer's price offer,* $p^*(\cdot)$*, the trader's order decision,* $o^*(\cdot)$*, and the trader's posterior belief about the economic state* θ_t*,* $\xi^*(\cdot)$*, such that the following properties hold:*

- *The dealer chooses her optimal price offer* $p_{t,\theta_t} \equiv p^*(\theta_t; h^{t-1})$ *to maximize her expected utility (see Eq.* (3.6)*) given the transaction history* h^{t-1} *(that depends on the underlying protocol we study), her private information about* θ_t*, and trader's order decision* $o^*(\cdot)$*.*

- *Given any public history* h^{t-1} *and the price offer* p_t *from the dealer, the trader updates his posterior about the economic state* θ_t *according to the Bayes's rule, whenever it applies.*

- *The trader chooses his optimal order decision* $o^*(p_t, \xi^*; h^{t-1}) \in \{1, 0, -1\}$ *to maximize his expected utility (see Eq.* (3.5)*), given the price offer* p_t*, the public transaction history* h^{t-1}*, and his posterior belief* $\xi^*(\cdot)$ *about the underlying economic state* θ_t*.*

In the case of multiple equilibria, since the informed dealer moves first, intuitively, we apply a refinement specifying the maximal ex-ante utility for the dealer. We call such equilibria *maximal PBE.* This refinement follows the same spirit as the MOP criterion proposed by Mailath, Okuno-Fujiwara and Postlewaite (1993).[27] Moreover, the maximal PBE not only

[27] In fact, as shown in the proof, in the static case, the MOP criterion in Mailath, Okuno-Fujiwara and

are optimal for the informed party who moves first, they also generate the highest ex-ante social welfare. A formal version of the selection criterion is presented below.

Definition 1.3.2 (MOP Criterion and Maximal PBE)**.** *We say that a pure strategy PBE* $(p(\cdot), o(\cdot), \xi(\cdot))$ ***defeats*** *another pure PBE* $(p'(\cdot), o'(\cdot), \xi'(\cdot))$ *if and only if*

- *in a one-time signaling game, for any trader's initial prior belief* ξ_0*, we have*

$$U(p(\cdot), o(\cdot), \xi(\cdot)) \geq U(p'(\cdot), o'(\cdot), \xi'(\cdot));$$

- *in a repeated signaling game with perfect monitoring, along the equilibrium path, for any* t*, we have the following relationship between the continuation payoffs,*

$$U_t(p(\cdot), o(\cdot), \xi(\cdot)) \geq U_t(p'(\cdot), o'(\cdot), \xi'(\cdot)),$$

  and the inequality is strict for some t*.*

We call a pure PBE a ***maximal PBE*** *if it is defeated by no other pure PBE. A PBE outcome is called a maximal PBE outcome if it is obtained through a maximal PBE.*

With the above setups, we are ready to show how the public disclosure of past trades (via TRACE), dealer's long-term incentive and trader's hedging motives can drastically change the nature of market price informativeness, market liquidity and welfare in dynamic OTC markets.

Comment 1.3.1. For the ease of explanation, we first study a case where $\beta = 1$, that is, traders are always potential sellers who possess risky assets (see Section 1.3.2). Then in Section 1.8.3, we show that this case is exactly mirrored by the one where sellers are buyers and need risky assets to hedge, i.e., $\beta = 0$. We also analyze the situation where traders' trading positions stochastically change from period to period, i.e., $\beta \in (0, 1)$.

1.3.7 Plan

Moving forward, the plan to analyze the model (in terms of characterizing the pooling equilibrium) is as follows. We first consider the static case in which the dealer is myopic

Postlewaite (1993) picks the same set of maximal PBE. Several signaling models use this concept (or stronger versions of it), e.g., Hartman-Glaser (2017) and Carapella and Williamson (2015).

(that is $\delta = 0$), in this case pooling equilibrium is not affected by the information disclosure. Next, we consider the dynamic cases, where $\delta > 0$: we first consider the case where the past information about trades is fully observable—we call it the public history case—this is equivalent to the scenario where the TRACE program is in place. Analyzing this case serves as a benchmark to consider the next case which is our main environment. In that case, following the different types information disclosure we mentioned before, observing past trades information for future traders is partial—and we call it the private history case.

The comparison between the public and private history cases (in terms of the pooling equilibrium structures) identifies the role of TRACE (or information disclosure) on the liquidity, price informativeness and welfare. Finally, we show that when the dealer is forward looking and *no* information about past trades is available, we get back to the static case in terms of the structure of the pooling equilibrium.

All together, with analyzing these cases we identify the role of information disclosure (or TRACE) along with the long-term incentive of informed dealers on price informativeness, liquidity and welfare (see Figure 1-3).

1.4 Static Trading

In this section we briefly consider the static version of the model, i.e., $\delta = 0$ meaning that dealer only cares about maximizing her within-period payoffs. This static case will introduce two equilibria that are useful to study the main dynamic model we study in the next sections.

Before the trade in each period, the dealer learns the underlying economic state θ_t, which affects the asset value a_t.[28]

1.4.1 Opaque Static Trading Equilibrium

As formally shown later in Section 1.7.3, any pricing strategy that increases traders' uncertainty and hides her private information allows the dealer to extract more information rent from them. We call such a *pooling* equilibrium where the dealer can hide her private information about current underlying economic state an *opaque static trading equilibrium* (OSTE).

[28] The case where in each period t the dealer is (like traders) fully uninformed about the realization of the economic state θ_t is considered in Appendix 1.10.

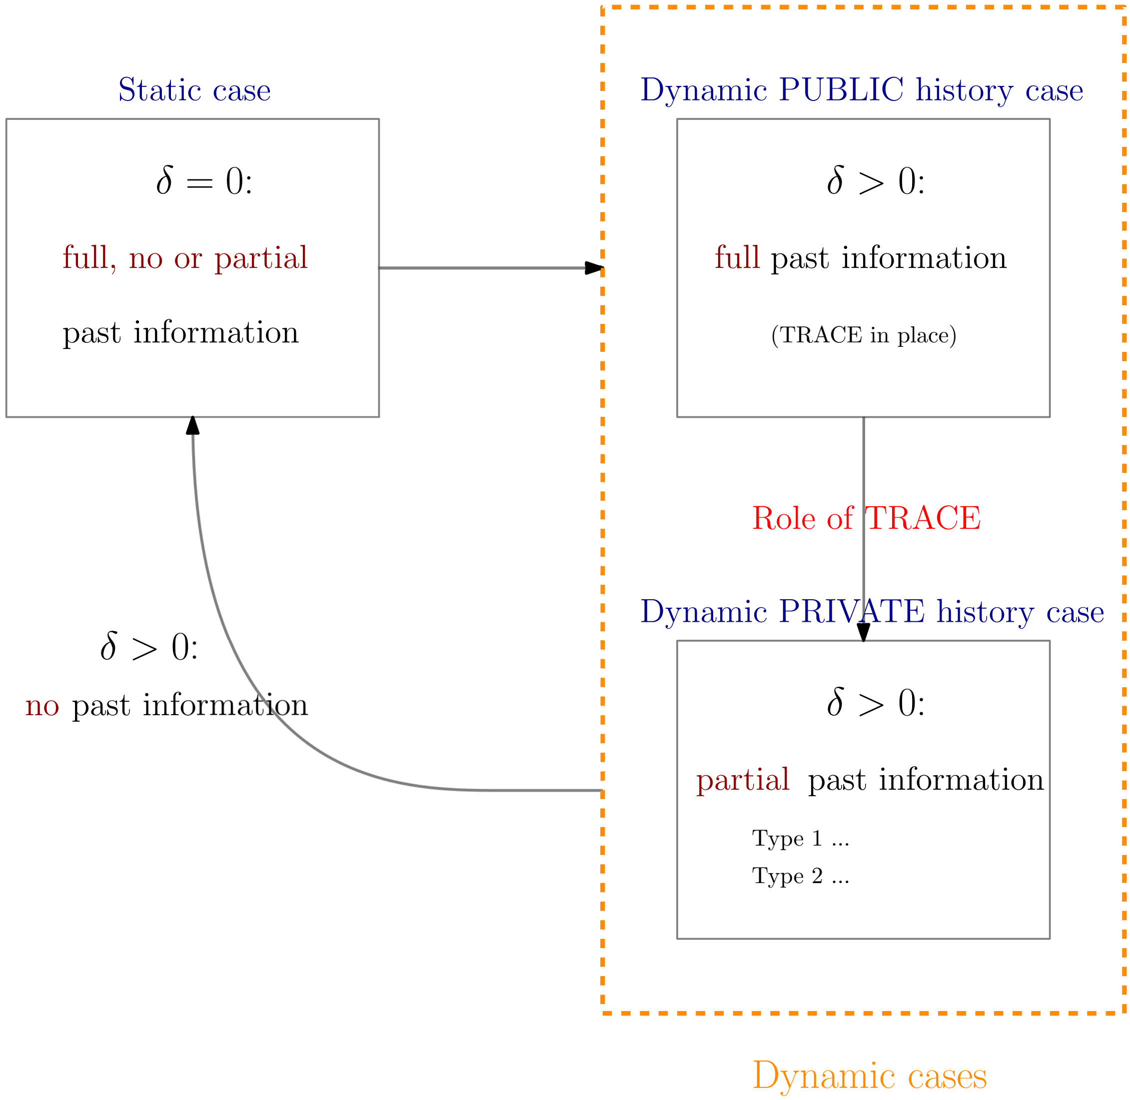

Figure 1-3: This panel shows the plan to analyze the model.

Definition 1.4.1. *If the dealer and traders are both myopic, then an equilibrium is an opaque static trading equilibrium (*OSTE*) if and only if on the equilibrium path, the dealer's pricing strategy* $\{p_t(h^{t-1})\}_{t=0}^{\infty}$ *is independent of her private knowledge about the drift and the volatility in asset values, that is, independent of the current underlying economic state* θ_t*.*

In other words, dealer offers opaque prices that are same in good times and in bad times. The following proposition specifies the necessary and sufficient condition for the existence of an OSTE.

Proposition 1.4.1. *The opaque static trading equilibrium* (OSTE) *exists if and only if*

$$\rho \geq \rho^{\text{OSTE}} = 2\frac{\alpha_{\max}(J_g - J_b)}{\alpha_{\max}(1-\alpha_{\max})(J_g - J_b)^2 + \alpha_{\max}\sigma_g^2 + (1-\alpha_{\max})\sigma_b^2} \tag{4.1}$$

where $\alpha_{\max} = \max\{\alpha_g, \alpha_b\}$.[29]

Proof. See Appendix 1.12. ∎

Proposition 1.4.1 shows that opaque pricing strategy is feasible if and only if traders are conservative enough to hedge their risky assets via trading with the dealer. In other words, opaque pricing equilibrium can be sustained if traders' risk-aversion coefficient is sufficiently large (i.e., $\rho \geq \rho^{\text{OSTE}}$) and traders, who sufficiently dislike volatilities in their asset values, will accept relatively low offers from the dealer. In bad times, the dealer (buyer) may have an incentive to deter the trade if it is not profitable. Therefore, when traders are risk-averse enough, the possibility to buy the asset at a low enough price then makes the transaction in bad times more likely to be profitable and prevent the rejection from the dealer.

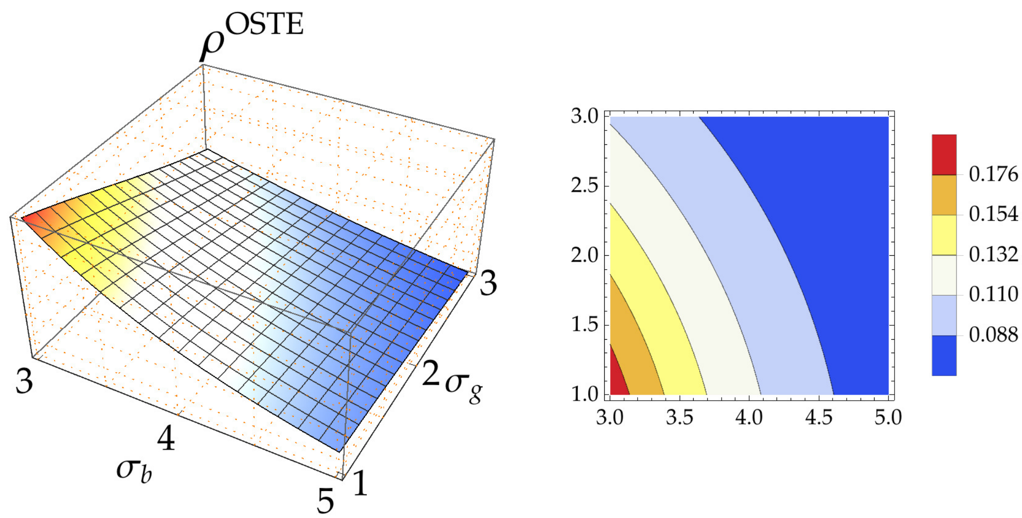

Figure 1-4: This chart plots ρ^{OSTE} for $\alpha_g = \alpha_b = \frac{1}{2}$ and $J_g - J_b = 1$. The area above the surface represents the region where an OSTE exists. From the graph, ρ^{OSTE} decreases in the asset innovation volatilities in both good and bad times (i.e., σ_g and σ_b).

Proposition 1.4.1 also implies some intutive comparative statics, summarized in Corollary 1.4.1 (see also Figure 1-4). For example, as the asset innovation becomes more volatile,

[29] The reason that $\alpha_{\max}$ appears in the ρ^{OSTE} is because the trader's individual rationality constraint needs to be satisfied when he believes the current economic is good both with probability α_g and with probability α_b. Here the larger one is binding. See the proof for more details.

either in good or in bad times (i.e., σ_g or σ_b increases) traders are more likely to accept a certainty price offer and an OSTE exists even if traders are less risk-averse. In other words, expectedly, the pooling equilibrium becomes easier to sustain and the cutoff risk-aversion coefficient ρ^{OSTE} decreases.[30]

Corollary 1.4.1. *The threshold ρ^{OSTE} always decreases in the asset innovation volatilities in both good and bad times (i.e., σ_g and σ_b). It also monotonically increases in $\alpha_{\max}$.*

Finally, it is worth noting that among all the OSTE, there exists one providing the dealer with the highest ex-ante expected payoff. By this opaque pricing strategy, dealer not only obtains the benefit of insurance due to the residual risk $\sigma_{\theta_t} z_t$ (which is equal to $\frac{\rho}{2}\text{Var}[\sigma_{\theta_t} z_t | h^{t-1}]$) but also receives her information rent for shocks in the drift of asset innovations J_{θ_t} (which is equal to $\frac{\rho}{2}\text{Var}[J_{\theta_t}|h^{t-1}]$). Hence, overall her ex-ante surplus in this Pareto dominant OSTE becomes

$$\begin{aligned} U_t^{\mathsf{OSTE}} &= \frac{\rho}{2}\left[\text{Var}[\sigma_{\theta_t} z_t | h^{t-1}] + \text{Var}[J_{\theta_t}|h^{t-1}]\right] \\ &= \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2]. \end{aligned}$$

From now on we denote this payoff as dealer's OSTE payoff.

1.4.2 Informative Static Trading Equilibrium

What happens if traders are not so risk-averse to hedge against their uncertainty shocks? One scenario is that they may reject the opaque offers that are sufficiently low. However, since dealer's outside option is normalized to zero, she will always be better off to trade as long as it provides non-negative interim gains. This drives our attention to another kind of equilibrium where dealer sacrifices her information rent to induce trade in good times. We call such a *separating* or *revealing* equilibrium an *informative static trading equilibrium* (ISTE).

Two important observations are worth noting.

- Our first observation is that $p_t(g; h^{t-1}) \equiv p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. On the one hand, this

[30] In addition, the equilibrium is harder to sustain with a larger $\alpha_{\max}$. To see the intuition, consider an extreme case where $\alpha_g = \alpha_b$. If the trader believes the current period is more likely to be in good times, then his uninformative expectation about the asset is closer to the higher one in good times, hence the trader is more likely to refuse the trade in bad times. Therefore, the dealer's individual rationality constraint in bad times is harder to fulfill and the equilibrium is easier to break down.

is trader's (seller's) highest evaluation of the asset, given any belief. That is, trader values his asset most when he is most optimistic and believes now is in good times for sure. That is, he will accept any offer weakly above this, and it's suboptimal for the dealer to offer a price strictly above it. Therefore, in any ISTE, $p_t^g \leq \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. On the other hand, in any ISTE, dealer offers different prices at different times, and a Bayesian trader, after observing p_t^g, will believe that the current period is in good times. Therefore, his individual rationality constraint implies that $p_t^g \geq \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. Together we have $p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ in any informative equilibrium where dealer chooses to price discriminatingly.

- Our second observation is that in an ISTE, trade only occurs in good times. Because otherwise in good times the dealer can offer a lower price p_t^b and persuade the trader that the current economic state is bad and he should sell his asset at a lower price. By doing so she can purchase her asset at a lower price. To discourage such a deviation, trade cannot happen in bad times.

Therefore, we just show all ISTE are payoff equivalent and it is without loss of generality to restrict on-path prices as follows.

Definition 1.4.2. *If the dealer and the trader are both myopic, then an equilibrium is called informative static trading equilibrium (*ISTE*) if and only if the dealer chooses a pricing strategy* $\{p_t(\theta_t, h^{t-1})\}$ *such that* $p_t(g, h^{t-1}) = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ *and* $p_t(b, h^{t-1}) < \varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2$.

Given the above observations, the next proposition characterizes the necessary and sufficient condition for the existence of ISTE.

Proposition 1.4.2. *Suppose the dealer is myopic (i.e.* $\delta = 0$*) and is informed about the current economic state (i.e.* θ_t*). If and only if*

$$\rho < \rho^{\mathsf{ISTE}} = 2\frac{J_\mathrm{g} - J_\mathrm{b}}{\sigma_\mathrm{g}^2}, \tag{4.2}$$

there exists an ISTE*.*

Proof. See Appendix 1.12. ∎

The intuition behind (4.2) is as follows. Similar to the argument that the dealer in good times should not be encouraged to mimic the dealer in bad times, by symmetry dealer in bad times also should not have the incentive to offer p_t^g to induce trade. This provides a lower bound on the price offered in good times, p_t^g, which characterizes the upper bound of trader's risk-aversion coefficient ρ for the existence of an ISTE.

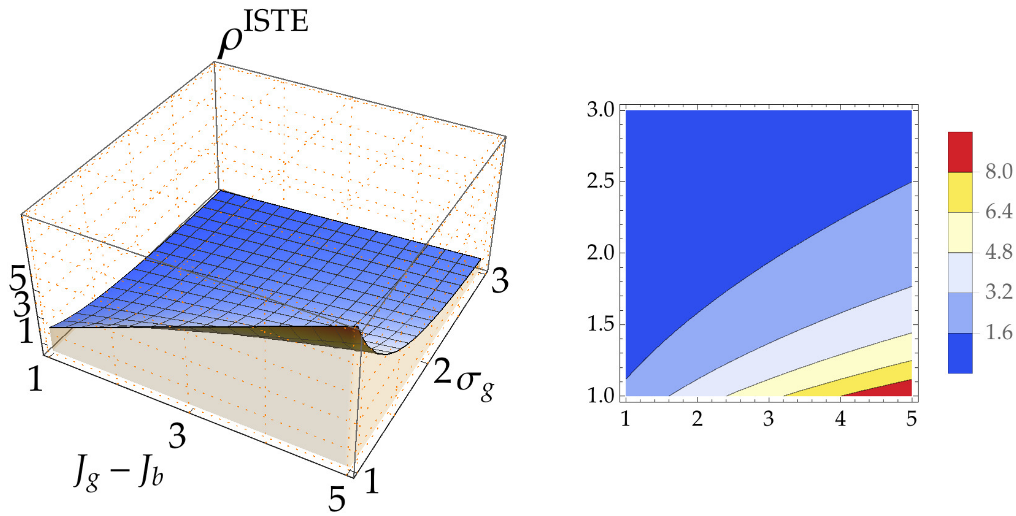

Figure 1-5: This chart plots ρ^{ISTE} for $\alpha_g = \alpha_b = \frac{1}{2}$. The gray area below the surface represents the region where ISTE exists. It also shows ρ^{ISTE} increases in the spread $J_g - J_b$ and decreases in the asset volatility in good times σ_g. Unlike ρ^{OSTE}, however, ρ^{ISTE} is independent of σ_b.

Proposition 1.4.2 also implies some comparative statics, summarized in Corollary 1.4.2 (See also Figure 1-5). For example, it shows that the threshold ρ^{ISTE} always decreases in the asset innovation volatility in good times, i.e., σ_g, and increases in the jump spread $J_g - J_b$. Hence, decreasing the volatility σ_g and increasing the spread $J_g - J_b$ both expand the domain of trader's risk-aversion coefficient (i.e., its hedging motive) for which an ISTE exists. Most importantly, since trade only occurs in good times, this threshold is independent of asset volatility in bad times, i.e., σ_b.

Corollary 1.4.2. *The threshold ρ^{ISTE} decreases in the asset innovation volatility in good times, i.e. σ_g. It does* not *depend on the asset innovation volatility in bad times, i.e., σ_b. Finally, it increases in the spread $J_g - J_b$.*

It is also worth noting that, since in any ISTE trade only occurs in good times, the dealer's ex-ante surplus purely comes from her insurance due to the residual risk in good times, that is, $\alpha_{\theta_{t-1}} \frac{\rho}{2} \sigma_g^2$. Since in ISTE the dealer reveals her private information about economic states, her ex-ante surplus is lower than that in the opaque equilibrium. Moreover, the price offer in good times increases in the drift of good times shock J_g and decreases in the risk-aversion coefficient ρ, as well as the volatility in good times σ_g.

Finally, the next corollary compares the thresholds ρ^{OSTE} and ρ^{ISTE}.

Corollary 1.4.3. *When the informed dealer is myopic, the threshold of* ISTE *is strictly larger than that of* OSTE*. In other words*

$$\rho^{\mathsf{ISTE}} > \rho^{\mathsf{OSTE}}. \tag{4.3}$$

In addition, for any $\rho \geq \rho^{\mathsf{OSTE}}$, there is a unique maximal PBE outcome, achieved only via an OSTE*. For any $\rho \in [\rho^{\mathsf{OSTE}}, \rho^{\mathsf{ISTE}}]$, such an* OSTE *defeats the* ISTE*.*

Proof. See Appendix 1.12. ∎

Corollary 1.4.3 implies that $\rho^{\mathsf{ISTE}} > \rho^{\mathsf{OSTE}}$. Hence, for all ρ between ρ^{OSTE} and ρ^{ISTE}, both types of equilibria (i.e., OSTE and ISTE) exist in the static trading game. Nevertheless, according to the MOP criterion, there always exists an OSTE that defeats the ISTE, making the OSTE the unique maximal PBE. In addition, as shown later, the former also generates higher social surplus and for any risk-aversion coefficient ρ, there exists an OSTE that Pareto dominates ISTE (see figure 1-6).

Comment 1.4.1. The results in this section hold no matter whether the history of past trades are public or private. In other words, if the dealer is myopic, only static trading equilibria discussed in this section, as well as the semi-opaque ones in Section 1.8.6, will be played.

1.5 Dynamic Trading: Public History

As discussed before when dealer is myopic information disclosure about past trades does not affect the (threshold) structure of the pooling equilibrium. What happens when dealer is forward looking and TRACE is in place (i.e., the past history of trades is observable)?

So far we have shown that when traders are sufficiently conservative to hedge against uncertainty shocks, an informed but myopic dealer can conceal her private information about economic states (OSTE). Next, motivated by Dodd-Frank (section 956), we consider how long-term incentives of dealers affect the sustainability of opaque pricing strategies and the transparency in OTC markets. In particular, we ask: does the long-term incentive of informed dealers improve market price informativeness? To answer this question, we extend the benchmark case and allow the informed dealer to be *forward-looking*, i.e., she also cares about future cash flows and has a positive discount factor $\delta > 0$.

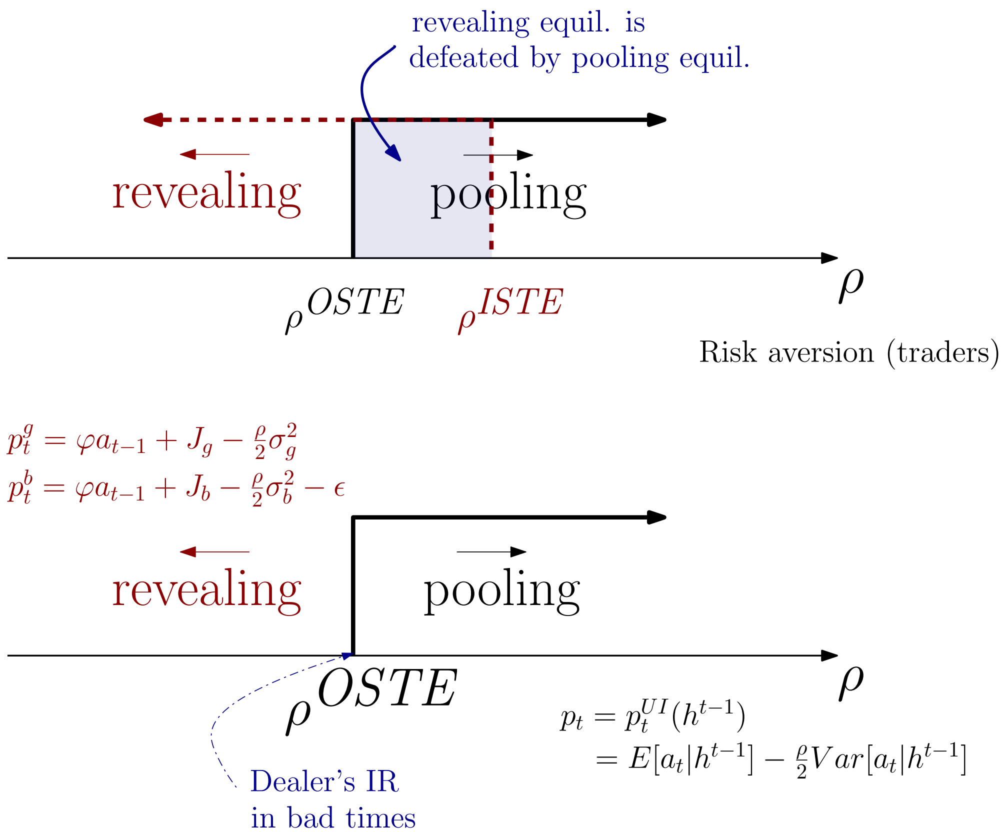

Figure 1-6: This chart plots the equilibria when the informed dealer is myopic. Depending on the extent of trader's hedging motive, there are two equilibria. When his risk-aversion coefficient is low (i.e., low hedging demand or $\rho < \rho^{\mathsf{ISTE}}$), the dealer reveals her information but trade only occurs in good times, reducing liquidity, risk sharing, and the dealer's ex-ante surplus. When trader's risk-aversion coefficient is high (i.e., high hedging demand or $\rho > \rho^{\mathsf{OSTE}}$), the dealer will conceal her private information and trades happen in both times, leading to the dealer's highest ex-ante surplus. These thresholds (i.e., ρ^{ISTE} and ρ^{OSTE}) are pinned down (explicitly) by drifts and volatilities of both times, as well as transition probabilities of the regime switch. When both equilibria are available (i.e., $\rho^{\mathsf{OSTE}} \leq \rho \leq \rho^{\mathsf{ISTE}}$) then the revealing equilibrium (ISTE) is defeated by the pooling equilibrium (OSTE).

1.5.1 Opaque Dynamic Trading Equilibrium

This section shows that the dealer's forward-looking incentive enables the threat of the loss of future profits and provides another device to deter her deviation, making opaque pricing equilibria easier to sustain and reducing the price efficiency. We also characterize the sufficient and necessary condition for the existence of such equilibria.

Section 1.4 shows that OSTE holds in the one-shot game when traders are risk-averse enough ($\rho \geq \rho^{\mathsf{OSTE}}$). Thus, it also consists of a subgame perfect equilibrium in the dynamic trading game. Corollary 1.4.3 implies that $\rho^{\mathsf{ISTE}} > \rho^{\mathsf{OSTE}}$, so opaque pricing is a static Nash equilibrium for $\rho > \rho^{\mathsf{ISTE}}$. We now focus our attention on a non-trivial case $\rho \leq \rho^{\mathsf{ISTE}}$ and examine whether or not the opaque equilibrium can be sustained under the dynamic setting. That is, throughout this section, we make the following assumption.

Assumption 1.5.1. *Assume $\rho \leq \rho^{\mathsf{ISTE}}$ and an informative static trading equilibrium (*ISTE*) always exists.*

Next, we consider a class of equilibrium with forward-looking dealer and along the equilibrium path, the dealer can still conceal her private information about the current economic state and provide an opaque price at any history. We call such an equilibrium an **opaque dynamic trading equilibrium** (ODTE).

Definition 1.5.1. *If the informed dealer is forward-looking and the traders are uninformed and myopic, then a PBE is called an opaque dynamic trading equilibrium (*ODTE*) if and only if on the equilibrium path, in each period, the dealer offers a price $p_t(h^{t-1})$ that does not depend on her private information about the economic state θ_t.*

Since in an ODTE price offers reveal *no* information about the economic states, dealer's offers can only depend on the public history. To sustain such an equilibrium to the greatest extent, that is, to deter the decline of the trade by the dealer (buyer) in bad times, one needs to implement the harshest punishment for such deviations.

First, the harshest punishment should involve grim trigger strategies. That is, once the dealer deviates, she will be punished in all future periods. Second, the deviator's expected one-shot payoff in the punishment stage should be as low as possible. In order to have robust results against any off-path construction, we want to use the minmax profile. But we also need to make sure that there is no incentive to deviate during the punishment stage. Fortunately, in this dynamic trading game, as shown in the following lemma, this minmax profile turns out to be the ISTE, which is a Nash equilibrium of the one-shot game and has the subgame perfect property in the dynamic setting.

Lemma 1.5.1. *For any conditional belief $\xi(p_t;\cdot)$, the dealer's ex-post* ISTE *payoffs consist of the min-max static payoffs given that the trader best responses to her strategy. Specifically,* *we have*

$$\min_{\xi(p_t;h^{t-1})} \max_{p_g,p_b} \min_{o_t \in \mathsf{BR}_{\xi(p_t;h^{t-1})}} (\varphi a_{t-1} + J_g - p_g) \cdot \Big(- o_t(p_g) \Big) = \frac{\rho}{2}\sigma_g^2$$

$$\min_{\xi(p_t;h^{t-1})} \max_{p_g,p_b} \min_{o_t \in \mathsf{BR}_{\xi(p_t;h^{t-1})}} (\varphi a_{t-1} + J_b - p_b) \cdot \Big(- o_t(p_b) \Big) = 0.$$

Proof. See Appendix 1.12. ∎

Then when the dealer is forward-looking, declining a transaction in bad times has future costs for her and makes it more costly to deviate. In fact, after a low price and no trade in bad times, although the dealer can avoid the loss in current period, in all future periods, the following traders observe her deviation and will expect to play ISTE equilibrium. The potential traders will adjust their expectations and thus the dealer can collect at most the ISTE payoff, only achievable when she reveals her private information about the economic states in all future periods. As a result, in all future periods, she is unable to hide her private information anymore and hence her expected continuation value is lowered after the deviation. Such concern discourages dealer's deviation and makes the opaque equilibrium easier to sustain under this dynamic trading setting, reducing market price informativeness.

The following proposition captures this idea and characterizes the necessary and sufficient condition for the existence of an opaque dynamic trading equilibrium (ODTE). We leave the specific characterization of the threshold, the construction of such an equilibrium and the necessity proof in the Appendix.

Proposition 1.5.1. *There exists a $\rho^{\mathsf{ODTE}} \leq \rho^{\mathsf{OSTE}}$ such that if and only if $\rho \geq \rho^{\mathsf{ODTE}}$, there exists an opaque dynamic trading equilibrium (*ODTE*) in which the dealer always conceals her private information about current economic state and trades occur in both* good *times and* bad *times. Moreover, for any $\rho \geq \rho^{\mathsf{ODTE}}$, there is a unique maximal PBE outcome, achieved only through an* ODTE*. Specifically, for any $\rho \in [\rho^{\mathsf{ODTE}}, \rho^{\mathsf{ISTE}}]$, there exists an* ODTE *that defeats and Pareto dominates the PBE where* ISTE *is played in every period.*

Proof. See Appendix 1.12. ∎

Finally, the comparison between ρ^{ODTE} and ρ^{OSTE} implies that ODTE is easier to sustain than OSTE. Therefore, long-term incentives of dealers may reduce financial transparency via the decrease of market price efficiency (see also Figure 1-7).

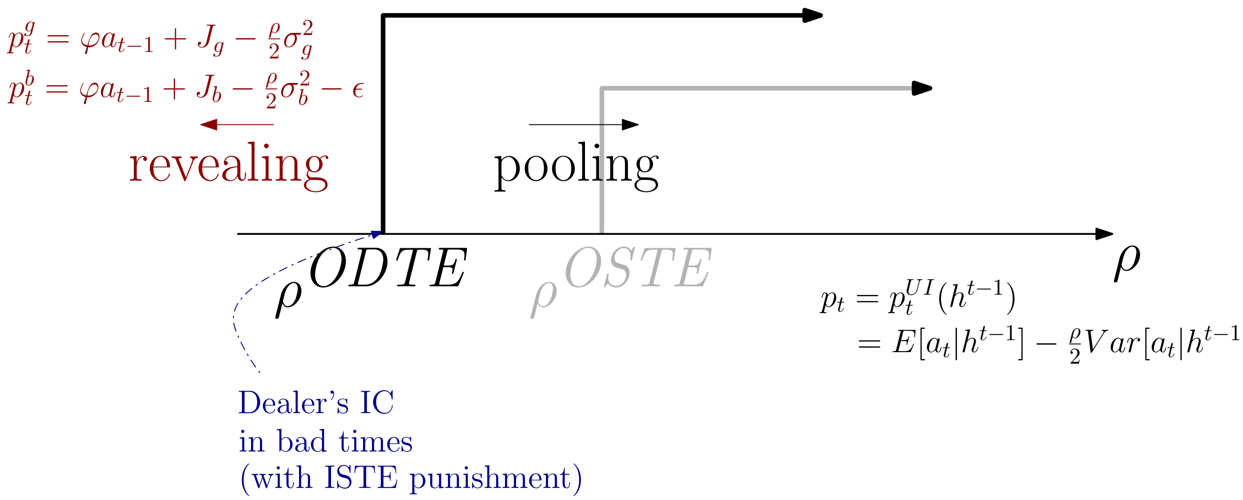

Figure 1-7: This panel shows that $\rho^{\mathsf{ODTE}} \leq \rho^{\mathsf{OSTE}}$. Hence, when the dealer is forward-looking (i.e. $\delta > 0$), an opaque pricing equilibrium is easier to sustain, reducing price efficiency.

Comment 1.5.1. For completeness, we also characterize another kind of equilibrium where traders always distinctly price the assets in good times and in bad times. We call it an informative dynamic trading equilibrium (IDTE) and leave the details of the construction and the discussion in Appendix 1.11. To sum up, under some technical assumption, whenever both types of equilibria exist, there is always an ODTE that defeats and Pareto dominates any IDTE.

1.6 Information Disclosure via TRACE: Private vs. Public History

To analyze the effect of information disclosure on OTC markets, we now consider an important extension of our model where the informed dealer is still forward-looking but previous trade history is **not** publicly available. Particularly, in line with the Dodd-Frank transparency Act of 2010 for municipal bonds, corporate bonds and swaps (implemented by TRACE and requiring the public dissemination of post-trade transaction information regarding price and volumes), we assume that previous economic states are still publicly observable and discuss two kinds of imperfect monitoring about previous transactions:

- The first one is that although future traders cannot observe exact price offers provided by the dealer, they can perfectly observe whether there is a transaction between the dealer and the trader in a previous period.

- The second kind of imperfect monitoring involves the situation where future traders can observe neither previous price offers nor whether or not there is a transaction in a previous period. Nevertheless, they can obtain a noisy signal y_t about o_t, whether a trade occurs in period t or not. Suppose y_t is independently drawn from a distribution $F_{o_t}(\cdot)$.

In particular, in this section, we ask:

> Does TRACE improve financial transparency by increasing price efficiency (informativeness)? Moreover, how does TRACE affect market liquidity?

1.6.1 Private Prices and Public Volumes History

The next proposition shows that in the first case, the ODTE constructed for Proposition 1.5.1 can still hold. In fact, from the analysis of ODTE, one can learn that offering a price weakly higher than $\varphi a_{t-1} + J_g - \frac{\varrho}{2}\sigma_g^2$ will not discourage traders from accepting the offers, but the dealer derives lower profits and trigger future punishments. Therefore, such strategies are always sub-optimal. Then the only kind of possible profitable deviation is to offer a price strictly lower than $\varphi a_{t-1} + J_g - \frac{\varrho}{2}\sigma_g^2$ and other than on-path price p_t^*. But in such a deviation, traders will always decline the offer and there is no transaction. Therefore, if the information about o_t is publicly observable, then players can perfectly monitor whether there are previous deviations and use the same punishment device as that in the ODTE. In other words, o_t is a sufficient statistic for previous deviations.

Proposition 1.6.1. *If future traders cannot observe previous price offers, but can perfectly observe whether or not there is a trade in any previous period, then there exists an opaque dynamic trading equilibrium if and only if $\rho \geq \rho^{\mathsf{ODTE}}$.*

Proof. See Appendix 1.12. ∎

Therefore, past orders are sufficient statistics for future traders to determine whether a dealer deviates or not. Proposition 1.6.1 implies that post trade price transparency does not affect the existence of opaque trading equilibrium. Releasing these details to the public does not make it easier or harder to elicit dealers' inside knowledge about the economic states. Moreover, since in any opaque trading equilibrium, trades occur in both good times and bad times, the post trade price transparency via TRACE need not decrease market liquidity.

1.6.2 Private Prices and Private Volumes History

We now turn our focus on a more interesting case, where future traders cannot observe previous price offers, but can have noisy signals about whether a transaction happens or not in any previous period. We show that the opaque equilibrium has similar structure as before and compare its threshold with the ones in previous sections.

We model this case as in the imperfect monitoring literature,[31] where all pure strategy Nash equilibria are payoff equivalent to the set of perfect public equilibria. Thus, it is without loss of generality to focus on the equilibria where players use public strategies only. Formally speaking, we only consider cases where both the dealer and the trader choose their strategies and beliefs based on past public signals $y^t \equiv \{y_1, ..., y_t\}$ and past economic states $\{\theta_1, ..., \theta_{t-1}\}$.

We are interested in how private history will affect the existence of opaque dynamic trading equilibrium. In opaque dynamic trading game, on the equilibrium path the dealer will conceal her private information about the current economic state and trades always occur. To deter dealer's deviations to the greatest extent, we need to implement the harshest punishment for her off the equilibrium path. According to Lemma 1.5.1, the informative static trading equilibrium (ISTE) will be played forever after the deviation.

Denote Y_{-1t} as the set of signals that upon observing, the dealer still offers an opaque price and traders will form corresponding beliefs. Denote Y_{0t} as the set of the rest signals which upon observing, the players will play an ISTE equilibrium. Denote $f_i^t \equiv \text{Prob}(y_t \in Y_{-1t} | o_t = i), i = 0, -1$. Ideally if we can observe o_t directly, then simply setting $y_t = o_t$ can give us $f_{-1}^t = 1$ and $f_0^t = 0$. Another extreme case is that the signal is completely uninformative, that is, $F_{o_t=0} = F_{o_t=-1}$, then we will have $f_0^t = f_{-1}^t$ no matter how we partition the signal space Y.

To simplify the analysis and make the model tractable, we restrict our attention to the equilibria where the division of signals is fixed from period to period. So we can drop the subscript and instead write Y_{-1}, Y_0, f_{-1} and f_0. We, specifically, define the opaque pricing equilibrium associated with this partition as a **private history equilibrium**.

Definition 1.6.1. *Consider an environment where future traders can view neither the details of previous offers nor whether or not there is a transaction at period t, but just a noisy signal $y_t \in Y$ about o_t. A private history equilibrium consists of a set of strategies $p(y^t)$ and beliefs*

[31] Here rather than solve the whole game for all possible public equilibrium, as in Fudenberg and Levine (1994), we, instead, focus on when opaque ones can be sustained.

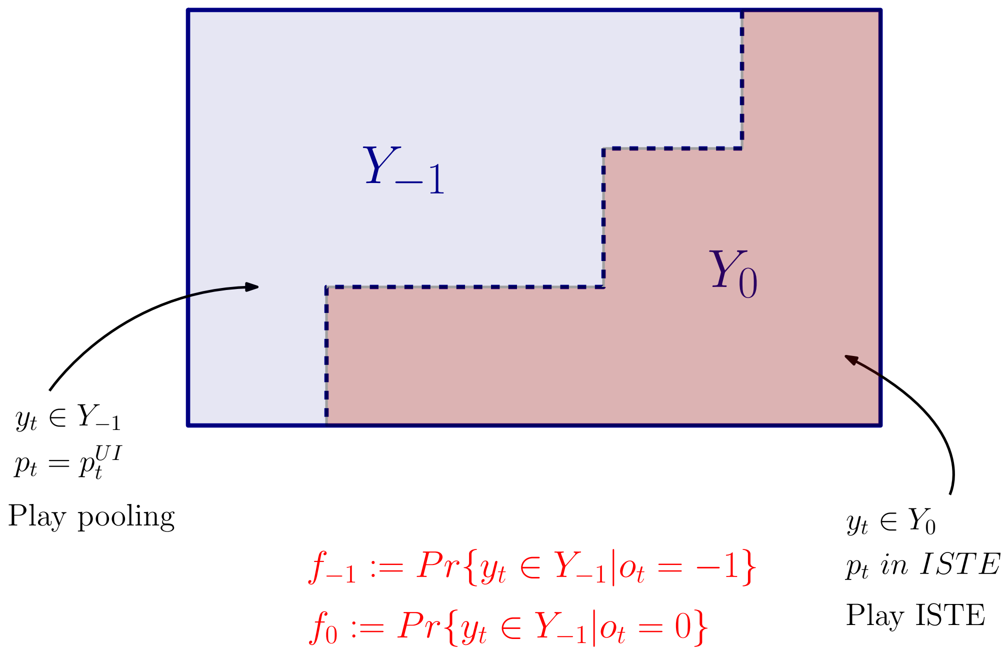

Figure 1-8: Signal partitions and the private history equilibrium structure.

$\xi(\cdot|y^t)$ *such that for some non-empty subset* $Y_{-1} \subset Y$*, if* $y^s \in Y_{-1}, \forall s \leq t-1$*, then in period* t*, the dealer chooses to offer a price* $p_t^*(h^{t-1})$ *that is independent of her private information about* θ_t*.*

Then, we ask if opaque pricing strategies can still form an equilibrium under this private history setup. The answer is that the private history equilibrium defined above exists when trader's risk-aversion coefficient is above a certain threshold.

Proposition 1.6.2. *Consider a case where future traders can only observe whether a trade occurs or not at period* t *via a noisy signal* y_t*, which is independently drawn from the distribution* $F_{o_t}(\cdot)$*. There exists a* $\rho^{\mathsf{private}} \in [\rho^{\mathsf{ODTE}}, \rho^{\mathsf{OSTE}}]$ *such that the following statement is true: if and only if* $\rho \geq \rho^{\mathsf{private}}$*, then there exists a private history equilibrium defined above. Moreover, for any* $\rho \in [\rho^{\mathsf{private}}, \rho^{\mathsf{ISTE}}]$*, there exists a private history equilibrium that defeats and Pareto dominates the equilibrium where* ISTE *is played in every period.*

Proof. See Appendix 1.12. ∎

Proposition 1.6.2 has important policy implications for whether or not the regulation institution should require dealers and traders to report information about private transactions to the public. On the one hand, in private history equilibrium, occasionally unlucky

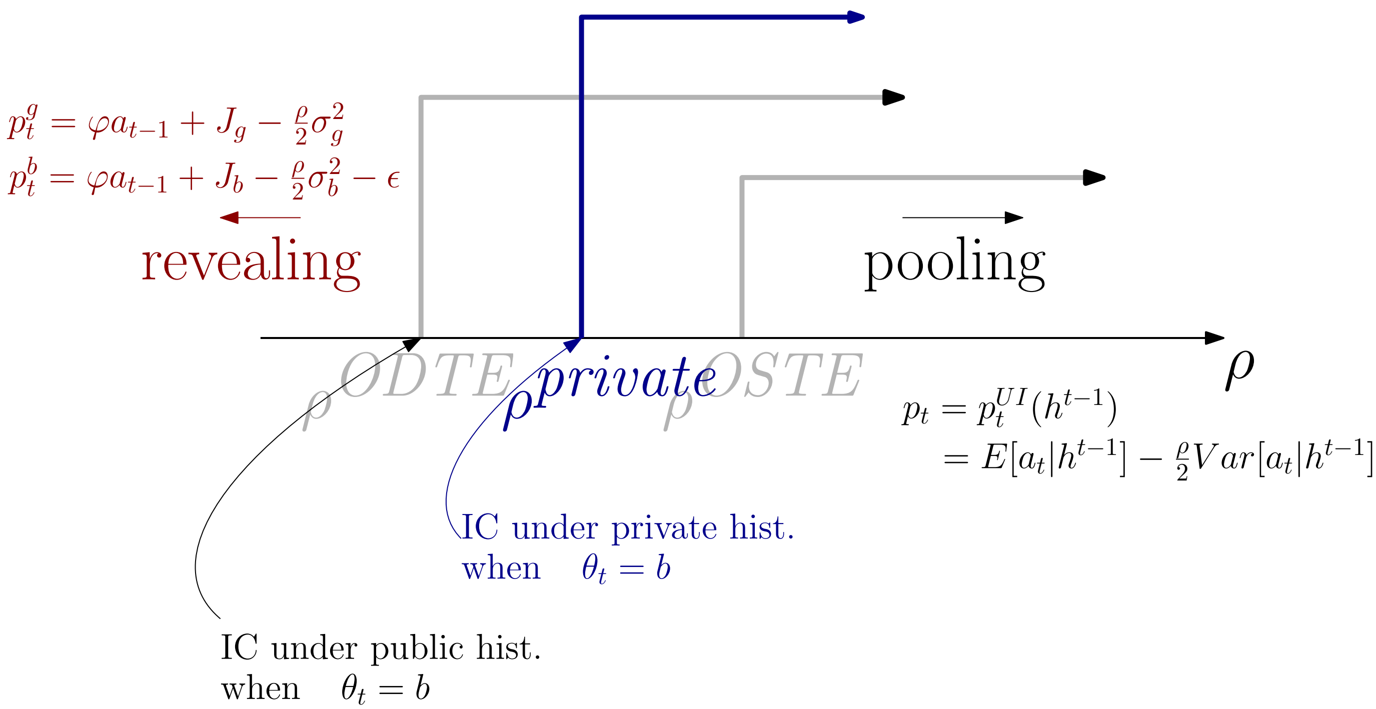

Figure 1-9: This chart plots three kinds of opaque pricing equilibria, OSTE (when the dealer is myopic), ODTE (when the dealer is forward-looking and previous orders are publicly observable) and private history equilibrium (when the dealer is forward-looking but previous orders are coarsely observed). We always have that ρ^{private} lies between ρ^{ODTE} and ρ^{OSTE}.

signals are produced and the game moves into the punishment phase, where ISTE is played. Therefore, even when traders are sufficiently risk-averse ($\rho > \rho^{\mathsf{private}} > \rho^{\mathsf{ODTE}}$) such that both ODTE and private history equilibrium can be sustained, the private history equilibrium still leads to more price transparency due to its occasional derailment to the punishment stage. In conclusion, when traders are intermediately risk-averse, requiring dealers and traders to accurately report past transaction orders will alter the equilibrium outcome drastically from ISTE to ODTE. When traders are sufficiently risk-averse and opaque pricing strategies can be sustained before and after this regulation, it will reduce the occurrence of informative pricing outcomes. From both points of view, asking the dealer to fully reveal her previous transaction orders (volumes) paradoxically impairs transparency and informativeness of the market prices.

On the other hand, from private history to public history, with more information about past transaction orders, the domain shrinks of traders' risk-aversion coefficient where opaque trading equilibria can be supported. In fact, the more informative signals of transaction orders become, the more credible the following claim from the dealer:

> "I will not deviate in the current period because if I did so, in following periods the future traders would figure this out and punish me for this deviation."

The intuition, which we call the reputation building (or commitment device) mechanism, is as follows. The informed dealer has an incentive to achieve a reputation of no-revelation-history so that in future periods she can extract the information rent and take advantage of traders' hedging motive. Therefore, the availability of past trade details enables such a reputation building and provides the dealer an incentive to hide her private information and maintain the no-revelation-history. Consequently, opaque pricing strategies along the equilibrium path can be supported. This is why for $\rho \in (\rho^{\mathsf{ODTE}}, \rho^{\mathsf{private}})$, public history of past orders enables opaque pricing and impairs the post price transparency in the market.

In conclusion, our analysis above of post-trade disclosure regarding prices and volumes shows that post-trade transparency via TRACE, paradoxically, makes markets more opaque, by reducing market price efficiency.

Finally, as in Section 1.6.1, we obtain a similar conclusion about market liquidity (trade activity and volume). This result is immediate because in any opaque equilibrium, trades occur in both good and bad times, thus post-trade disclosure regarding volumes via TRACE need not reduce market liquidity, and may even increase it.

1.6.3 Signal Structure and Price Informativeness

Sometimes it is unrealistic to completely ban or completely release the information about past transactions. The next proposition discusses the effect of one type of change in signal structure on the existence of opaque pricing equilibrium, as well as price transparency.

Proposition 1.6.3. *For a fixed f_0, the cutoff of the existence of a private history equilibrium ρ^{private} characterized in Proposition 1.6.2 decreases in f_{-1}. If $f_{-1} - f_0 = 1$, then ρ^{private} coincides with ρ^{ODTE}. If $f_{-1} - f_0 = 0$, then ρ^{private} coincides with ρ^{OSTE}.*

Proof. See Appendix 1.12. ∎

Hence, making the signal more distinctive by purely increasing f_{-1} will relax the requirement on traders' risk-aversion coefficient, and the opaque pricing equilibrium can be sustained with less risk-averse traders. An increase of f_{-1} also leads to more lenient strategies, lower frequency of triggering the punishment stage and a more transparent pricing scheme. In other words, the effect of this signal structure change is twofold. First it makes opaque pricing equilibrium easier to hold. Second, even with sufficiently risk-averse traders, an increase in f_{-1} makes it less likely to have a signal in Y_0, so less likely to have price transparency.

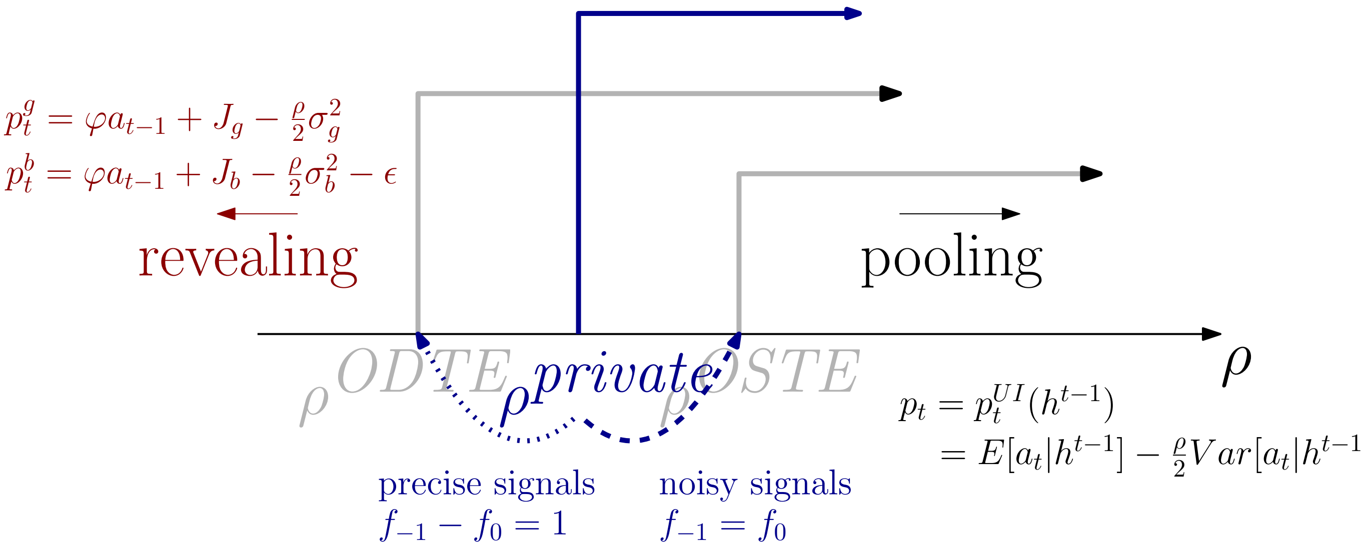

Figure 1-10: When past signals are very noisy, in the Blackwell sense, (i.e., $f_{-1} = f_0$) then ρ^{private} increases to ρ^{OSTE}, in contrast, when past signals are informative, (i.e., $f_{-1} - f_0 = 1$) then ρ^{private} decreases to ρ^{ODTE}.

Finally, it is worth noting that with certain conditions, our result in this section reinforces the message in Kandori (1992). He shows that as signals of past actions are improved in Blackwell's sense, the set of PBE payoffs expands. If f_0 is fixed and f_{-1} decreases, past history signals become less informative, and ρ^{private} will increase. Since there always exists a private history equilibrium Pareto dominating the ISTE, for any ρ, the equilibrium payoff set weakly decreases (for intermediate ρ such decrease is strict). In other words, less informative signals shrink the possible equilibrium payoff set.

1.7 Policy

To explore potential policy interventions, in this section we consider a policymaker who is concerned about transparency in OTC markets. We discuss policy implications of our results and provide suggestions such as *stochastically auditing* dealers to improve price informativeness in these markets. We also compare social welfare of all equilibria discussed above and show that there is a trade-off between price informativeness and social welfare in OTC markets.

1.7.1 Information Releasing

Proposition 1.6.1 and Proposition 1.6.2 together imply that if the policymaker (regulator) wants to increase price transparency in OTC markets, one way is to control what kind of information future traders can get access to. Simply restricting the disclosure of details of past transactions is not enough. To delink the dynamic incentive of dealers, it is necessary to hide whether or not there are transactions, or information of orders, as well.

One way to implement the above strategy is to ban the discussion of past transactions with future traders. The regulator can also forbid the use of past deals to advertise and attract new traders. Moreover, the regulator can encrypt the identity or past order history of dealers when they are approached by traders.

1.7.2 Stochastic Auditing

If restricting the disclosure of information about past orders is hard to implement in some circumstances, the regulator can also improve price informativeness of the OTC market via supervision and punishment. As shown in this subsection, increasing the auditing intensity will make the opaque equilibrium harder to sustain.

In this section we model the case where the regulator steps in to combat the information hiding of dealers. At the beginning of each period, the regulator randomly audits a fraction of dealers and checks if they provided opaque, uninformative orders in last period. If caught, then a dealer is prohibited from participating in financial activities and all her future payoffs are gone.

Specifically, at the beginning of period $t+1$, the regulator randomly picks up q fraction of dealers to check. As past information about economic states, asset values and price offers are all available to the regulator, it can then check if in period t, the dealer offers a price of $p_t^g \equiv \varphi a_{t-1} + J_g - \frac{\varrho}{2}\sigma_g^2$ or something strictly below $p_t^b \equiv \varphi a_{t-1} + J_b - \frac{\varrho}{2}\sigma_b^2$. If not, then the regulator can ban the participation of this dealer into future trades. Otherwise the dealer remains in the market. We call this q the auditing intensity.

We show in the following proposition that as auditing intensity increases, the opaque equilibrium becomes harder to sustain.

Proposition 1.7.1. *Let q represent the regulator's auditing intensity, that is, with probability q, at the end of a period t, a dealer is audited. If she provides a price weakly above $\varphi a_{t-1} + J_b - \frac{\varrho}{2}\sigma_b^2$ but not $\varphi a_{t-1} + J_g - \frac{\varrho}{2}\sigma_g^2$, then all her future payoffs are forfeited. Denote P_q as* *the set of trader's risk-aversion coefficients when* ODTE *can be sustained. If* $q_1 < q_2$*, then* $P_{q_2} \subseteq P_{q_1}$*. Specifically,* $P_0 = [\rho^{\mathsf{ODTE}}, \infty]$ *and* $P_1 \subsetneqq [\rho^{\mathsf{OSTE}}, \infty]$*.*

Proof. See Appendix 1.12. ∎

In other words, stochastic auditing can also force the dealer to reveal her private information about economic states more frequently and as a result, lead to more price transparency in the market. This coincides with the general intuition of censorship literature that the higher the detection effort, the less likely individuals will misbehave. The regulator can then choose the optimal auditing intensity given her desire of price informativeness and the limited auditing resources.

Comment 1.7.1 (Maximizing Price Efficiency and Liquidity). In fact, if the informative dynamic trading equilibrium (IDTE) exists (see Remark 1.5.1 and Appendix 1.11), then there is a policy that can achieve *maximal* price efficiency as well as market liquidity at the same time. The regulator can punish the dealer who does not follow the IDTE. Specifically, if the regulator wants to implement an IDTE where in bad times the dealer collect a payoff of π^b (π^b satisfies the condition in Appendix-Proposition 1.11.1), then the regulator can ban all future participation of the dealer if she is observed not to provide on-path prices $p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ in good times or $p_t^b = \varphi a_{t-1} + J_b - \pi^b$ in bad times. In such an equilibrium, the dealer is pricing discriminatingly and the offered price reflects her private information. Moreover, in IDTE the trade also occurs in bad times, the market liquidity is the same as when there is no auditing. Such policy also gives a higher level of social surplus than the one we constructed above. In fact, even though this social surplus is still below the highest possible level, it achieves the maximum the regulator can get if he wants full price efficiency. The intuition is that post-trade price transparency discourages traders to hedge and lowers the benefit of trade, impairing social surplus and implying that there is a trade-off between price efficiency and social welfare.

1.7.3 Welfare

In this section we compare social welfare and the dealer's profits of different equilibria specified above. We first point out that there is a trade-off between social welfare and price informativeness in OTC markets. We then show that with more transparency and public disclosure of trade volumes, the privately informed dealer can gather a higher expected profit.

The latter result appears surprising, as it goes against the general lesson of contract theory that less disclosure gives more information rent to the party with private information.

Our results hold for the following reason. In our model more public disclosure about past trades enables the dealer to hide her private information in the current period. Consequently, opaque trading equilibrium is easier to sustain, impairing price efficiency and improving total social welfare. The dealer's expected profits depend on the division of this total surplus between her and the trader. But for each equilibrium under certain level of transparency of past orders, with more public information about past trade volumes, there always exists an equilibrium that Pareto dominates the old one, and the dealer's expected profit increases.

We now analyze the social welfare in two kinds of equilibrium. In ISTE, the dealer fully reveals her private information and the trade only occurs in good times. As a result, the ex-ante welfare reduces to its minimum level. In opaque equilibria (either static or dynamic), however, the ex-ante welfare achieves its maximum. The difference between these extreme scenarios, in regions where both types are available, is what we call the **welfare gap**. The next proposition shows the exact amount of this welfare gap.

To make the social surplus comparable, we here define the social surplus as the sum of the dealer's and trader's per-period expected payoffs given their equilibrium behaviors. Formally speaking, W^{ODTE} represents the social surplus in a certain period given that both parties follow their on-path ODTE strategies. We are going to compare this with the social surplus in another equilibrium in the same period.

Proposition 1.7.2. *In period t, the expected social welfare for each equilibrium is explicitly given by:*

$$W^{\mathsf{OTE}} \equiv W^{\mathsf{OSTE}} = W^{\mathsf{ODTE}} = \frac{\rho}{2}\left[\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2\right],$$
$$W^{\mathsf{ISTE}} = \frac{\rho}{2}\left[\alpha_{\theta_{t-1}}\sigma_g^2\right].$$

In the region where ISTE *and either* OSTE *or* ODTE *exist (i.e., $\rho \in [\rho^{\mathsf{ODTE}},\ \rho^{\mathsf{ISTE}}]$), the welfare gap is*

$$W^{gap} \equiv W^{\mathsf{OTE}} - W^{\mathsf{ISTE}} = \frac{\rho}{2}\left[(1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2\right],$$

which is independent of σ_g and is convex and increasing in σ_b as well as the spread $J_g - J_b$.

Proof. See Appendix 1.12. ∎

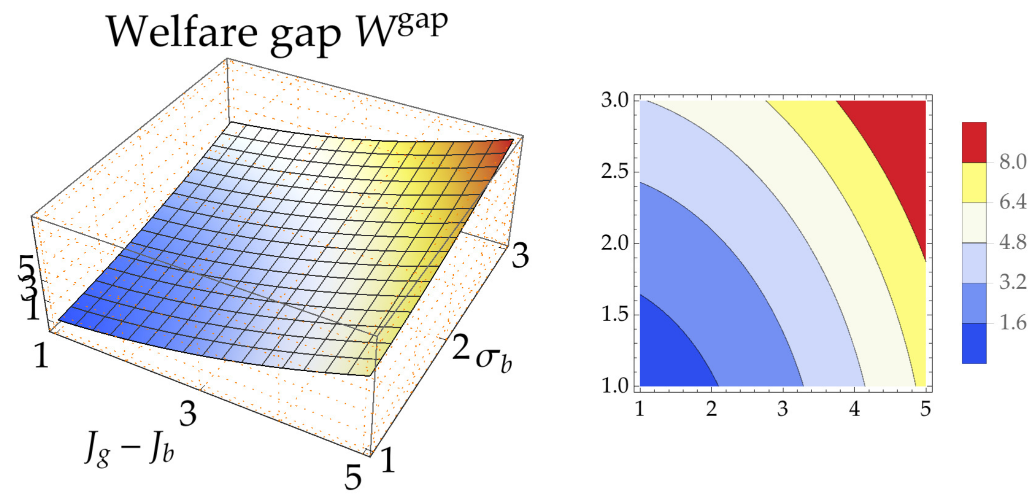

Figure 1-11: This chart plots the welfare gap W^{gap}. It shows that the welfare gap is increasing in the spread $J_g - J_b$ and the asset volatility in bad times σ_b. Note that W^{gap} is independent of the asset volatility in good times σ_g.

The ex-ante social welfare of ISTE is smaller than that of OSTE or ODTE for two reasons. First, there is no trade in bad times in ISTE. Although the dealer still has a higher evaluation than the trader, who possesses the asset, the incentive compatibility constraint prevents both parties from exchanging the asset and leads to a loss of the social welfare, captured by $\frac{\rho}{2}(1 - \alpha_{\theta_{t-1}})\sigma_b^2$. The residual term in the welfare gap, $\frac{\rho}{2}\alpha_{\theta_{t-1}}(1 - \alpha_{\theta_{t-1}})(J_g - J_b)^2$ comes from the information rent. In fact, in informative equilibrium traders learn about economic state θ_t before the transaction (after observing dealer's bid price offer). Therefore, the only uncertainty in his evaluation comes from the residual risk, and this reduce of risk decreases benefits of trade as well as his desire to hedge. As a result, social surplus coming from the trade decreases and we lose the social insurance derived from the uncertainty of drift of innovation.

It then follows immediately that the welfare gap is independent of the asset volatility in good times (i.e., σ_g) but monotonically increases in the asset volatility in bad times σ_b and the spread $J_g - J_b$ (see also Figure 1-11). This is because the dealer gains the insurance from residual risks (i.e., due to $\sigma_{\theta_t} z_t$) in good times in both informative equilibrium and opaque equilibrium. Hence, it cancels out in the welfare gap.

We next discuss the dealer's expected profits. As shown in previous sections, there exist multiple opaque equilibria, and each differs from the others on specific shares of the total surplus the dealer and the trader can get. Therefore, the dealer's profits with different levels of public disclosure of past trades are not comparable. However, we can fix this technical difficulty by comparing the equilibrium with the same division ratio between the dealer and the trader. By doing so, for example, in opaque equilibrium the total surplus is higher than that in informative equilibrium, and as a result, dealer can gain higher profits in an opaque equilibrium than in an informative one. Next, recall that Proposition 1.6.2 shows that if information about past orders is available, the dealer can use such signals as a commitment device and opaque equilibrium is easier to hold. Therefore, disclosing the dealer's information about past order volumes can actually increase her expected payoffs. Surprisingly, such result is apart from the general lesson of contract theory, where private information always provides information rent for the insider.

1.8 Robustness and Extensions

In this section we extend our model in several other directions and analyze the robustness of our main findings. In Section 1.8.1, we show that our main conclusions do not depend on deliberate functional forms and can go beyond the mean-variance preference. Section 1.8.2 extends the analysis to the case where traded orders are divisible. In Section 1.8.3 we allow trader's demand shock to stochastically change over times, meaning that traders change from a seller to a buyer or vice-versa from period to period. In Section 1.8.5 we consider the model when the fundamental is fixed and the risk-aversion coefficients change over time. Finally, Section 1.8.6 characterizes a semi-pooling equilibrium in the one-shot trading game.

1.8.1 More General Forms of Utilities

In this section we show that our main conclusions are not the result of particular functional forms and can be applied to more general utility functions. But with more general forms of utilities, we need a standard to compare the risk-aversion of different agents. In this paper, we introduce the following definition and argue that all the results in this paper still hold if we relax the mean-variance utility form and use the following standard to compare the risk preference of traders.

Definition 1.8.1. *Suppose $\tilde{x}_g$ and $\tilde{x}_b$ are lotteries with normally distributed outcomes, $\tilde{x}$ is a lottery of lotteries, such that with probability α, the outcome is drawn according to the lottery $\tilde{x}_g$ and with probability $1-\alpha$, the outcome is drawn according to the lottery $\tilde{x}_b$. An* *utility function V* ***exhibits more risk-aversion*** *than an utility function U if and only if*

$$0 \leq CE_U(\tilde{x}_g) - CE_V(\tilde{x}_g) \leq CE_U(\tilde{x}) - CE_V(\tilde{x}),$$

where $CE_f(\cdot)$ is the certainty equivalence function for $f(\cdot)$, that is, $CE_f(\tilde{y}) = f^{-1}(\mathbb{E}(f(\tilde{y})))$ for any random variable $\tilde{y}$ and increasing function f.

The intuition behind the above definition is as follows. $\tilde{x}$ is a more volatile lottery than $\tilde{x}_g$, and if the agent is more risk-averse, then he would further dislike risks and value $\tilde{x}_g$ over $\tilde{x}$ more than a less risk-averse agent. In other words, more risk-averse trader requires a higher risk premium for the risky asset than less risk-averse traders. The difference is more salient for asset with a higher risk level. For example, in our baseline model, the utility function is mean-variance utility and

$$CE_U(\tilde{x}) = \mathbb{E}(\tilde{x}) - \frac{\rho^U}{2}\mathrm{Var}(\tilde{x}).$$

Then a mean-variance utility function V exhibits more risk-aversion than a mean-variance utility function U if and only if $\rho^V \geq \rho^U$, which coincides with the traditional comparison of risk-aversion for mean-variance utilities.

With the above definition, we can derive the following propositions and argue that the general message of our analysis is not altered by this generalization.

Proposition 1.8.1. *(a) If* OSTE *exists with trader's utility function U, then there still exists an* OSTE *when trader's utility function is V and V exhibits more risk-aversion than U.*

(b) If ISTE *exists with trader's utility function V, then there still exists an* ISTE *when trader's utility function is U and V exhibits more risk-aversion than U.*

(c) If ODTE *exists with trader's utility function U, then there still exists an* ODTE *when trader's utility function is V and V exhibits more risk-aversion than U.*

(d) If a private history equilibrium exists with trader's utility function U, then there still exists a private history equilibrium when trader's utility function is V and V exhibits more risk-aversion than U.

Proof. See Appendix 1.12. ∎

Therefore, the structure of the equilibria described in previous sections (OSTE, ODTE, ISTE and private history equilibrium) all remain the same. If traders are extremely risk-averse, then opaque equilibrium can be sustained, no matter whether it is under static setting, dynamic setting with public history, or dynamic setting where the order history is imperfectly observed. The informative pricing strategies can be sustained when traders' risk-aversion coefficients are sufficiently low.

Proposition 1.8.2. *Fix a trader's utility function U,*

(a) if both an OSTE *and an* ISTE *exist, then there exists a private history equilibrium;*

(b) if both a private history equilibrium and an ISTE *exist, then there exists an* ODTE.

Proof. See Appendix 1.12. ∎

Proposition 1.8.2 implies that the relative relation between OSTE, ODTE, ISTE and private history equilibrium also remain the same. Specifically speaking, dynamic trading environment provides a commitment device for the dealer and makes the opaque pricing strategies easier to sustain, while the imperfect observation of past volumes impairs the supervision by future traders on the dealer's action, reduces the cost of her deviation and increases the potential of the breakdown of an opaque equilibrium.

1.8.2 Divisible Orders

In this section we extend our analysis to the case where traded orders can be divisible, i.e., $-o_t \in [0, 1]$.[32] At the beginning of period t, the informed dealer offers a take-it-or-leave-it bundle (o_t, p_t) to the trader, where o_t and p_t are, respectively, the order (volume) and and the per-unit (bid) price. As we next show, nothing fundamental changes in terms of the structure of opaque equilibria. For the informative static trading equilibrium (ISTE), however, things work differently and the trade now can also occur in bad times, in contrast to the discrete order case presented in Proposition 1.4.2, where the trade only occurs in good times.

If selling o_t unit of his risky asset to the dealer, a trader with the prior belief ξ_t can collect

[32] Recall that previously when a trader is a seller, $o_t \in \{0, -1\}$, so in the corresponding divisible case $-o_t \in [0, 1]$.

an ex-ante payoff of

$$p_t\Big(-o_t\Big)+\Big(1-\Big(-o_t\Big)\Big)\left[\varphi a_{t-1}+\xi_t J_g+(1-\xi_t)J_b\right]$$
$$-\frac{\rho}{2}\Big(1-\Big(-o_t\Big)\Big)^2\left[\xi_t(1-\xi_t)(J_g-J_b)^2+\xi_t\sigma_g^2+(1-\xi_t)\sigma_b^2\right].$$

OSTE. We start with the analysis of the opaque static trading equilibrium. We claim that the one constructed in Section 1.4 remains an equilibrium. Since in an OSTE the dealer's offer contains no information about the current economic state, after observing it the trader still holds his prior beliefs, $\xi_t=\alpha_{\theta_{t-1}}$.

In this case the dealer solves the following mechanism design problem:

$$\max_{(-o_t)\in[0,1],p_t}\Big(-o_t\Big)(\varphi a_{t-1}+\alpha_{\theta_{t-1}}J_g+(1-\alpha_{\theta_{t-1}})J_b-p_t)$$
$$s.t.\quad p_t\Big(-o_t\Big)+\Big(1-\Big(-o_t\Big)\Big)\left[\varphi a_{t-1}+\alpha_{\theta_{t-1}}J_g+(1-\alpha_{\theta_{t-1}})J_b\right]$$
$$-\frac{\rho}{2}\Big(1-\Big(-o_t\Big)\Big)^2\left[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2+\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2\right]\geq 0$$

The trader's individual rationality (IR) constraint provides a lower bound for p_t. Plugging this into the objective function, the problem is equivalent to

$$\max_{(-o_t)\in[0,1]}\quad -\frac{\rho}{2}\Big(1-\Big(-o_t\Big)\Big)^2\left[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2+\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2\right]+$$
$$\Big(\Big(-o_t\Big)+1-\Big(-o_t\Big)\Big)(\varphi a_{t-1}+\alpha_{\theta_{t-1}}J_g+(1-\alpha_{\theta_{t-1}})J_b)$$
$$\Leftrightarrow\max_{(-o_t)\in[0,1]}\quad -\frac{\rho}{2}\Big(1-\Big(-o_t\Big)\Big)^2\left[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2+\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2\right]$$

Therefore, the optimal order to provide is still $o_t=-1$, which coincides with the equilibrium offers in the discrete case. The cutoff for the existence of such an equilibrium remains the same.

ISTE. Next we look at the ISTE. We first restrict that trader's risk-aversion coefficient in the range where ISTE exists for the discrete order case, i.e., $\rho\leq\rho^{\mathsf{ISTE}}=2\frac{J_g-J_b}{\sigma_g^2}$. Then denote $(o_{t,\theta_t},p_{t,\theta_t}),\theta_t\in\{g,b\}$ as an ISTE. The dealer maximizes

$$\Big(-o_{t,g}\Big)(\varphi a_{t-1}+J_g-p_{t,g})+\Big(-o_{t,b}\Big)(\varphi a_{t-1}+J_b-p_{t,b})$$

subject to the following constraints

$$
\begin{aligned}
&\varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2 \leq (-o_{t,b})p_{t,b} + (1-(-o_{t,b}))[\varphi a_{t-1} + J_b] - \frac{\rho}{2}(1-(-o_{t,b}))^2\sigma_b^2 && \textbf{trader's IR at } \theta_t = b \\
&\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2 \leq (-o_{t,g})p_{t,g} + (1-(-o_{t,g}))[\varphi a_{t-1} + J_g] - \frac{\rho}{2}(1-(-o_{t,g}))^2\sigma_g^2 && \textbf{trader's IR at } \theta_t = g \\
&(-o_{t,b})(\varphi a_{t-1} + J_b - p_{t,b}) \geq 0 && \textbf{dealer's IR at } \theta_t = b \\
&(-o_{t,g})(\varphi a_{t-1} + J_g - p_{t,g}) \geq 0 && \textbf{dealer's IR at } \theta_t = g \\
&(-o_{t,b})(\varphi a_{t-1} + J_b - p_{t,b}) \geq (-o_{t,g})(\varphi a_{t-1} + J_b - p_{t,g}) && \textbf{dealer's IC at } \theta_t = b \\
&(-o_{t,g})(\varphi a_{t-1} + J_g - p_{t,g}) \geq (-o_{t,b})(\varphi a_{t-1} + J_g - p_{t,b}) && \textbf{dealer's IC at } \theta_t = g
\end{aligned}
$$

We claim that with divisible orders, the ISTE constructed in previous sections fail to exist. Moreover, there exists another informative equilibrium that Pareto dominates it. We prove this by contradiction. Suppose $o_{t,g} = -1$ and $p_{t,g} = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. The dealer's problem then becomes

$$
\begin{aligned}
\max_{-o_{t,b}, p_{t,b}} \quad & (-o_{t,b})(\varphi a_{t-1} + J_b - p_{t,b}) \\
\text{s.t.} \quad & (-o_{t,b})[p_{t,b} - \varphi a_{t-1} - J_b + \frac{\rho}{2}\sigma_b^2(2-(-o_{t,b}))] \geq 0 && (8.1) \\
& (-o_{t,b})[\varphi a_{t-1} + J_b - p_{t,b}] \geq 0 && (8.2) \\
& (-o_{t,b})(\varphi a_{t-1} + J_g - p_{t,b}) \leq \frac{\rho}{2}\sigma_g^2 && (8.3) \\
& (-o_{t,b})(\varphi a_{t-1} + J_b - p_{t,b}) \geq \frac{\rho}{2}\sigma_g^2 + J_b - J_g && (8.4)
\end{aligned}
$$

In the above program, inequality (8.1) is from trader's IR constraint at $\theta_t = b$, inequality (8.2) is dealer's IR constraint at $\theta_t = b$, inequality (8.3) is from dealer's IC constraint at $\theta_t = b$, and inequality (8.4) is from dealer's IC constraint at $\theta_t = g$.

We show that $o_{t,b} = 0$ fails to be a solution to dealer's optimization problem. To see this, consider an $-o_{t,b} \in (0, 1)$, and a price $p_{t,b} < \varphi a_{t-1} + J_b$. If we can check that this pair satisfies conditions (8.1)-(8.4), then it gives the dealer positive payoff in bad times and dominates the no-trade bundle where $o_{t,b} = 0$.

For this bundle to work we need

$$
\begin{aligned}
p_{t,b} &\geq \varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2(2-(-o_{t,b})) && (8.5)\\
p_{t,b} &< \varphi a_{t-1} + J_b && (8.6)\\
p_{t,b} &\geq \varphi a_{t-1} + J_g - \frac{\rho}{2(-o_{t,b})}\sigma_g^2 && (8.7)\\
p_{t,b} &\leq \varphi a_{t-1} + J_b - \frac{\frac{\rho}{2}\sigma_g^2 + J_b - J_g}{(-o_{t,b})} && (8.8)
\end{aligned}
$$

Since $\rho \leq \rho^{\mathsf{ISTE}}$, $-\frac{\rho}{2}\sigma_g^2 - J_b + J_g \geq 0$. Condition (8.8) is not binding. We also notice that for $-o_{t,b}$ small enough, $\varphi a_{t-1} + J_g - \frac{\rho\sigma_g^2}{2(-o_{t,b})} < \varphi a_{t-1} + J_b$. Therefore, one can pick a $p_{t,b}$ in the interval

$$
\left[\max\{\varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2(2-(-o_{t,b})), \varphi a_{t-1} + J_g - \frac{\rho}{2(-o_{t,b})}\sigma_g^2\}, \varphi a_{t-1} + J_b\right],
$$

then conditions (8.1)-(8.4) are all satisfied, and the dealer can still collect positive payoffs in bad times. This equilibrium gives the dealer a higher payoff than that in the discrete case.

ODTE and private history equilibrium. We claim the ODTE and the private history equilibrium discussed in this paper still exist. From the analysis of OSTE, we learn that there are no immediate benefits to offer divisible orders $o_t \in (0,1)$. Furthermore, deviating to divisible orders can trigger future punishment, although this punishment is less severe when continuous quantity is allowed. Hence, the dealer does not want to deviate and offer something with $-o_t \in (0,1)$ in the repeated setting.

The discussion in the previous sections then directly follows. ODTE and private history equilibrium still exist for some $\rho \geq \rho^{\mathsf{ODTE}'}$ and $\rho \geq \rho^{\mathsf{private}'}$, respectively. Moreover, we still have the following relationship

$$
\rho^{\mathsf{ODTE}'} \leq \rho^{\mathsf{private}'} \leq \rho^{\mathsf{OSTE}}.
$$

1.8.3 Different Trading Positions

In this section we allow traders' positions to change stochastically between sellers and buyers. Specifically, we assume demand shocks follow an i.i.d. process such that

$$
\text{Prob}\{\chi_t = 1\} = \beta \in (0,1), \quad \text{Prob}\{\chi_t = -1\} = 1-\beta.
$$

The trading position χ_t is public information for both the dealer as well as the traders since period t.

OSTE. In a static model, if $\chi_t = 1$, i.e., the trader is a potential seller, this is exactly what happens in our baseline model and results stay the same. Proposition 1.8.3 shows what will happen when $\chi_t = -1$, that is, when the trader is a potential buyer. The opaque equilibrium has similar structure as in previous sections.

Proposition 1.8.3. *When the trader's trading position is $\chi_t = -1$, that is, he is in demand of a unit of asset and needs to buy it from the dealer, there exists an opaque static trading equilibrium* (OSTE) *if and only if*

$$\rho \geq \rho^{\mathsf{OSTE}}_{\chi_t=-1} \equiv 2\frac{(1-\alpha_{\min})(J_g - J_b)}{\alpha_{\min}\sigma_g^2 + (1-\alpha_{\min})\sigma_b^2 + \alpha_{\min}(1-\alpha_{\min})(J_g - J_b)^2}, \quad \alpha_{\min} = \min\{\alpha_g, \alpha_b\}$$

Proof. See Appendix 1.12. ∎

ISTE. As in OSTE, the analysis of $\chi_t = 1$ case for the ISTE is exactly same as in our baseline model. Proposition 1.8.4 shows that the informative equilibrium also shares the same structure as the case when traders are potential sellers.

Proposition 1.8.4. *When the trader's trading position is $\chi_t = -1$, that is, he is in demand of a unit of asset and needs to buy it from the dealer, there exists an informative static trading equilibrium* (ISTE) *if and only if*

$$\rho \leq \rho^{\mathsf{ISTE}}_{\chi_t=-1} \equiv 2\frac{J_g - J_b}{\sigma_b^2}$$

Proof. See Appendix 1.12. ∎

As shown by the following Corollary, the comparison relationship between $\rho^{\mathsf{ISTE}}_{\chi_t=-1}$ and $\rho^{\mathsf{OSTE}}_{\chi_t=-1}$ remains unchanged.

Corollary 1.8.1. *When the dealer is myopic and is a potential buyer, we have*

$$\rho^{\mathsf{ISTE}}_{\chi_t=-1} > \rho^{\mathsf{OSTE}}_{\chi_t=-1}$$

Proof. See Appendix 1.12. ∎

In this informative equilibrium, trade only occurs in bad times and the dealer can collect an expected payoff of $\frac{\rho}{2}\sigma_b^2$ during bad times and 0 otherwise.

ODTE. We now consider the dynamic game, i.e., $\delta > 0$. Similarly as in our baseline model, we construct an equilibrium as below.

- If the dealer has not deviated in previous periods, then

  1. the dealer offers a price $p_t(\theta_t, \chi_t, h^{t-1}) = p_t^*(h^{t-1}, \chi_t)$; if $\chi_t = 1$, then this is a bid offer; if $\chi_t = -1$, then this is an ask offer;
  2. the trader holds his prior belief after observing $p_t^*(h^{t-1}; \chi_t)$ and for all other prices, believes for sure that the economic state is unfavorable to the dealer, i.e.,

$$\xi(p_t; h^{t-1}) = \begin{cases} \alpha_{\theta_{t-1}} & \text{if } p_t = p_t^*(h^{t-1}; \chi_t) \\ 1 & \text{if } p_t \neq p_t^*(h^{t-1}; \chi_t), \chi_t = 1 \\ 0 & \text{if } p_t \neq p_t^*(h^{t-1}; \chi_t), \chi_t = -1 \end{cases} , \forall h^{t-1};$$

  3. if $\chi_t = 1$, then the trader will accept any bid price that is weakly above $p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ or exactly at $p_t^*(h^{t-1}, \chi_t = 1)$;
  4. if $\chi_t = -1$, then the trader will accept any ask price that is weakly below $p_t^b = \varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$ or exactly at $p_t^*(h^{t-1}, \chi_t = -1)$.

- If the dealer deviates previously, then both parties observe this and play an informative static trading equilibrium (ISTE).

  1. If $\chi_t = 1$, then in this ISTE, the trade only occurs in good times and the dealer can collect an expected payoff of $\frac{\rho}{2}\sigma_g^2$ during good times and 0 otherwise.
  2. If $\chi_t = -1$, then in this ISTE, the trade only occurs in bad times and the dealer can collect an expected payoff of $\frac{\rho}{2}\sigma_b^2$ during bad times and 0 otherwise.

Following the same technical steps as in the analysis of ODTE, we can show the following result.

Proposition 1.8.5. *Given that $\sigma_g^2 - \sigma_b^2 + (1 - \alpha_{\max})(J_g - J_b)^2 \geq 0$. Suppose the trader has a trading position χ_t in period t, then there exists a $\rho_{\chi_t}^{\mathsf{ODTE}} \leq \rho_{\chi_t}^{\mathsf{OSTE}}$ such that if and only if $\rho \geq \rho_{\chi_t}^{\mathsf{ODTE}}$, there exists an* ODTE*.*

Proof. See Appendix 1.12. ∎

Private history equilibrium. The construction of the private history equilibrium now becomes

- If $y^s \in Y_1$ for all $s \leq t-1$, then
  1. the dealer offers a price $p_t(\theta_t, \chi_t, h^{t-1}) = p_t^*(h^{t-1}, \chi_t)$; if $\chi_t = 1$, then this is a bid offer; if $\chi_t = -1$, then this is an ask offer;
  2. the trader holds his prior belief after observing $p_t^*(h^{t-1}; \chi_t)$ and for all other prices, believes for sure that the economic state is unfavorable to the dealer, i.e.,

$$\xi(p_t; h^{t-1}) = \begin{cases} \alpha_{\theta_{t-1}} & \text{if } p_t = p_t^*(h^{t-1}; \chi_t) \\ 1 & \text{if } p_t \neq p_t^*(h^{t-1}; \chi_t), \chi_t = 1 \\ 0 & \text{if } p_t \neq p_t^*(h^{t-1}; \chi_t), \chi_t = -1 \end{cases}, \forall h^{t-1};$$

  3. if $\chi_t = 1$, then the trader will accept any bid price that is weakly above $p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ or exactly at $p_t^*(h^{t-1}, \chi_t = 1)$;
  4. if $\chi_t = -1$, then the trader will accept any ask price that is weakly below $p_t^b = \varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$ or exactly at $p_t^*(h^{t-1}, \chi_t = -1)$.
- If $y^s \in Y_0$ for some $s \leq t-1$, then both parties observe this and play an informative static trading equilibrium (ISTE).
  1. If $\chi_t = 1$, then in this ISTE, the trade only occurs in good times and the dealer can collect an expected payoff of $\frac{\rho}{2}\sigma_g^2$ during good times and 0 otherwise.
  2. If $\chi_t = -1$, then in this ISTE, the trade only occurs in bad times and the dealer can collect an expected payoff of $\frac{\rho}{2}\sigma_b^2$ during bad times and 0 otherwise.

The analysis of the private history equilibrium when traders have independent trading positions across periods follows the same spirit of that of ODTE. In fact, we can show the following result.

Proposition 1.8.6. *Given that $\sigma_g^2 - \sigma_b^2 + (1-\alpha_{\max})(J_g - J_b)^2 \geq 0$. Suppose that the trader has a trading position χ_t in period t, then there exists a $\rho_{\chi_t}^{\mathsf{private}} \in [\rho_{\chi_t}^{\mathsf{ODTE}}, \rho_{\chi_t}^{\mathsf{OSTE}}]$ such that if and only if $\rho \geq \rho_{\chi_t}^{\mathsf{private}}$, there exists a private history equilibrium.*

Proof. See Appendix 1.12. ∎

The structure of ODTE and the private history equilibrium, as well as the comparative relations between ρ^{ODTE}, ρ^{private} and ρ^{OSTE} remains unchanged. Therefore, different trading positions do not change the message our main conclusions send.

1.8.4 Bid-Ask Spread

In this section we consider the bid-ask spread in good times and in bad times for both equilibria. First, in an opaque equilibrium, the dealer offers the same price $p_t^{*\mathsf{UI}}(\theta_{t-1}, \chi_t)$ in both times. Then the bid price p^{bid} is the equilibrium price when $\chi_t = -1$, or when the trader is a potential seller:

$$\begin{aligned} p^{bid} &= p_t^{*\mathsf{UI}}(\theta_{t-1}, \chi_t = -1) \\ &= \varphi a_{t-1} + \alpha_{\theta_{t-1}} J_g + (1-\alpha_{\theta_{t-1}}) J_b - \frac{\rho}{2}(\alpha_{\theta_{t-1}} \sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2). \end{aligned}$$

Similarly the ask price p^{ask} is the one when $\chi_t = 1$ or when the trader is a potential buyer:

$$\begin{aligned} p^{ask} &= p_t^{*\mathsf{UI}}(\theta_{t-1}, \chi_t = 1) \\ &= \varphi a_{t-1} + \alpha_{\theta_{t-1}} J_g + (1-\alpha_{\theta_{t-1}}) J_b + \frac{\rho}{2}(\alpha_{\theta_{t-1}} \sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2). \end{aligned}$$

Therefore, the bid-ask spread in an opaque equilibrium becomes:

$$p^{ask} - p^{bid} = \rho(\alpha_{\theta_{t-1}} \sigma_g^2 + (1 - \alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1 - \alpha_{\theta_{t-1}})(J_g - J_b)^2).$$

In an informative equilibrium, however, the trade never occurs in bad times when the trader possesses a risky asset ($\chi_t = -1$) or in good times when the trader is in demand of a risky asset ($\chi_t = 1$). If the dealer can learn about the trader's trading position before making an offer, then such knowledge actually decreases liquidity in OTC markets.

If the dealer is uninformed about whether she is facing a seller or a buyer, then she must offer a bid price and an ask price that makes the trader indifferent between accepting and rejecting. The bid-ask spread now becomes, in good times:

$$p_t^g(\theta_{t-1}, \chi_t = 1) - p_t^g(\theta_{t-1}, \chi_t = -1) = \rho\sigma_g^2;$$

whereas in bad times:

$$p_t^b(\theta_{t-1}, \chi_t = 1) - p_t^b(\theta_{t-1}, \chi_t = -1) = \rho\sigma_b^2.$$

As mentioned above, if the dealer learns about the trading position of the other party, then liquidity increases and the bid-ask spread will be larger than the above ones and the market becomes less efficient. This is because the dealer can deliberately offer non-acceptable prices to deter trades that are not profitable.

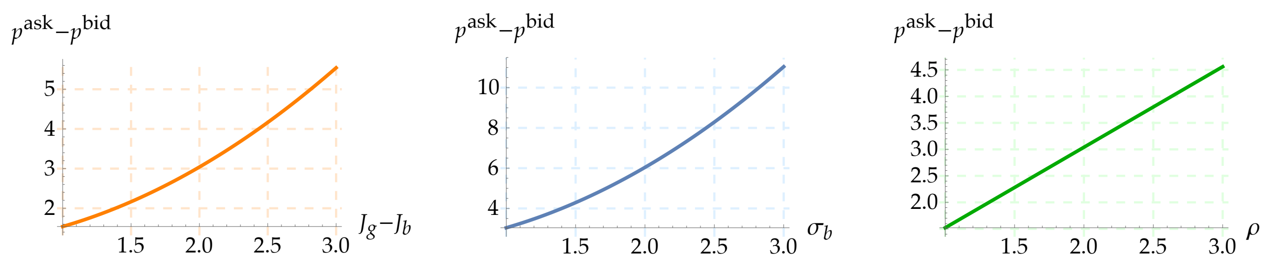

Figure 1-12: Bid-ask spread in opaque equilibria increases in the risk-aversion ρ, the asset volatility σ_b and σ_g, and the jump spread $J_g - J_b$. In the plots $\alpha_g = \alpha_b = \frac{1}{2}$.

1.8.5 Fixed fundamental and variable risk aversion

In this section, in contrast to the benchmark model, we assume the trading calendar (horizon) is finite, i.e., $N < \infty$ periods. Moreover, the state $\theta \in \{g, b\}$ is *fixed*. So the underlying asset value follows the following process: $a_t = J_\theta + \sigma_\theta z_t \equiv a(\theta, z_t)$, where the drift and the volatility of the stochastic shock depend on the underlying fixed economy θ. We also assume the persistence coefficient, for simplicity and without loss of generality, is zero, i.e., $\varphi = 0$.

At the beginning of the game, Nature chooses $\theta = g$ with probability $\alpha^* \in (0, 1)$, which consists of the common prior of all traders (sellers). Since θ is fixed, to induce traders to trade with the dealer, we, however, let traders' risk-aversion coefficients change over time. Precisely, we assume that there are two types of sellers (traders): those who are highly risk-averse (high-type sellers) and those who are not so risk-averse (low-type sellers), with corresponding risk-aversion coefficients ρ_h and ρ_ℓ ($\rho_h > \rho_\ell \geq 0$). At the beginning of the trading, each trader's type is independently drawn by Nature. With probability γ, he becomes a high-type. A seller's risk-averse attitude (type) is privately observed only by himself and the delaer, and is not learned by future sellers.

Consistent with the histories considered in the benchmark model, we analyze three types of histories: public history, no-history and order history. Public history corresponds to the regulations requiring the disclosure of past prices and past orders. In the no-history both dealer's price offers and trader' decisions are not observable in future periods. This corresponds to the situation where no market transparency regulations are implemented. The third case, order history (which is a form of private history), is where the order information is publicly announced but dealer's price offer is kept private between parties in the transaction. In other words, future sellers (traders) can only observe whether there is a trade or not but not on what prices the dealer and a previous seller agree.

In contrast to the benchmark model, since here θ is fixed and is never revealed, we further need to define trader t's belief about θ, given the history of trades available to him. Recall that at the beginning of the game, Nature chooses $\theta = g$ with probability $\alpha^* \in (0, 1)$. In the following periods, we denote $\alpha_t \equiv \text{Prob}(\theta = g|h^{t-1})$ and $\xi_t \equiv \text{Prob}(\theta = g|h^{t-1}, p_t)$ as seller's corresponding posterior belief before and after observing dealer's current offer. That is, before observing any past trade details and the current price offer, each seller believes that $\theta = g$ with probability $\xi_0 = \alpha^*$. Then, α_t is seller's posterior belief about θ given only h^{t-1}, whereas ξ_t is seller's belief given both h^{t-1} and p_t.

For the next results to hold, we need to impose the following mild assumptions.

Assumption 1.8.1 (monotonicity). *Denote $\pi(\rho, \alpha)$ as the static surplus from the trade when $\theta = b$ if seller's (trader's) risk-aversion coefficient is ρ and posterior belief is α. Specifically,*

$$\pi(\rho, \alpha) = \frac{\rho}{2}[\alpha\sigma_g^2 + (1-\alpha)\sigma_b^2 + \alpha(1-\alpha)(J_g - J_b)^2] - \alpha(J_g - J_b).$$

Then $\pi(\rho, \alpha)$ monotonically decreases in α for any $\rho > 0$.

One can verify that the monotonicity assumption holds if $J_g - J_b$ is small enough or $\sigma_g^2 - \sigma_b^2$ is large enough.

Assumption 1.8.2 (consistency). *If the dealer ever reveals her private information before (offering a discriminating price that depends on θ under the public history, or declining the trade under the order history), then she will propose the same offer p_t for all future high-type sellers, no matter whether the revelation happens on-path or off-path.*

Assumption 1.8.2 is a natural assumption, because it says that if the dealer ever reveals her private information to the market (i.e., uninformed sellers (traders) fully observes the true value of the fundamental θ) then it becomes very costly for the dealer to manipulate prices anymore. Particularly, it can be due to high fines imposed by regulators, for example, SEC (securities and exchange commission) imposes high fines for price manipulation, and securities laws and related SEC rules broadly prohibit fraud in the purchase and sale of securities, and the Securities Exchange Act of 1934, Section 9, specifically makes it unlawful to manipulate security prices.

Do the main results of the benchmark model hold in this framework? We next show that the main insights of the benchmark model continue to hold when θ is fixed and the trading calendar (horizon) is finite.

No-history

Consider the no-history environment in which future traders learn nothing about trade details in a previous period. Then, the following result holds.

Proposition 1.8.7. *There exists $\rho^{pool.} < \rho^{sep.}$ such that:*

(i) ***if and only if*** *the seller's (trader's) risk-aversion coefficient $\rho \geq \rho^{pool.}$, there exists a fully pooling equilibrium where the dealer can hide her private information and offer a uniform price that is independent of θ;*

(ii) ***if and only if*** *the seller's (trader's) risk-aversion coefficient $\rho \leq \rho^{sep.}$, there exists a fully separating equilibrium where the dealer can offer a discriminating price that depends on her private information about θ. Moreover, in the separating equilibrium, when $\theta = b$ the dealer will offer a price low enough such that no trade occurs.*

Proof. See Appendix 1.12. ∎

The intuition is as follows. For the first part of Proposition 1.8.7, opaque pricing strategy requires a strong hedging motive from sellers. That is, sellers need to sufficiently dislike their volatility in their asset values, and have a low enough reservation price. Otherwise it becomes too costly for the dealer to buy the asset from them, especially when $\theta = b$. As a result, a high reservation price provides an incentive for the dealer to decline the trade and avoid the loss. If dealer's evaluation of the asset when $\theta = b$ is lower than that of the seller, the fully pooling equilibrium breaks. For the second part, in a fully separating equilibrium, no trade will occur when $\theta = b$ because otherwise when $\theta = g$ the delaer has an incentive to pretend it is in bad times, and dealer's incentive compatibility (IC) constraint fails. For this reason, when the dealer starts to price discriminatingly, she always offers a low enough price when $\theta = b$ so that no trade occurs. This result is consistent with the insights of the benchmark model.

Throughout the rest of the section, we focus on a non-trivial case where $\rho_\ell < \rho^{pool.} \leq \rho_h$.

Public history

In this section we show that compared to the private history, under the public history the opaque pricing strategy is easier to sustain. Hence, the availability of past trade details actually decreases market transparency. Without loss of generality, we fix ρ_h as a parameter, and characterize such conditions in terms of ρ_ℓ.

To formalize the conditions we need to introduce "k-opaque" equilibrium. In a "k-opaque" equilibrium, the dealer prices opaquely in the first k period, after which she starts pricing discriminatingly if a low-type trader (seller) comes. The formal definition is below.

Definition 1.8.2. *Under the public history or order history, a PBE is called a "k-opaque" equilibrium if and only if:*

- *(hide before period k) the dealer's on-path price offer is independent of θ for $t = 1, ..., k$, no matter whether the seller (trader) is high- or low-type;*
- *(hide for high-type) the dealer's on-path price offer is independent of θ for periods $t = k + 1, ..., N$ if the seller (trader) is high-type;*
- *(reveal for low-type after period k) the dealer's on-path price offer is different when $\theta = g$ and $\theta = b$ for periods $t = k + 1, ..., N$ if the seller (trader) is low-type.*

Basically, "k-opaque" equilibrium is a more general concept that nests both a fully pooling and a fully separating equilibrium. When $k = N$, the dealer hides her private information in all periods $t = 1, .., N$, and this is actually a fully pooling equilibrium. When $k = 0$, the dealer reveals her information whenever a low-type seller arrives since the first period, coinciding with the fully separating equilibrium.

Proposition 1.8.8 characterizes the necessary and sufficient condition for the existence of a "k-opaque" equilibrium. From the proposition we learn that the sustainability of a "k-opaque" equilibrium under the public history is similar to that of a fully pooling equilibrium under the no-history. Both exist if and only if low-type traders (sellers) are risk-averse enough. Moreover, the more risk-averse sellers are, the longer the dealer can hide her private information about the underlying economy and offers a uniform price, consistent with the insights of the benchmark model.

Proposition 1.8.8. *Under the public history, that is, when future sellers (traders) can observe past orders as well as prices, there exists a sequence* $\{\rho_k\}_{k=1}^N$ *such that*

- $\rho_N = \rho^{pool.}$
- $\{\rho_k\}_{k=1}^N$ *increases in* k*;*
- *for* $k = 1, ..., N$*, there exists a "k-opaque" equilibrium if and only if* $\rho_\ell \geq \rho_k$*.*

Proof. See Appendix 1.12. ∎

The above result shows that, consistent with insights from the benchmark model, with the public disclosure of past trade details, the opaque pricing strategy becomes easier to sustain. As previously discussed, the intuition is that the forward-looking dealer now also cares about her future profits. If $\theta = g$, then the dealer never wants to deviate. However, when $\theta = b$, if she defects and refuses to offer the opaque price specified in the equilibrium, then such deviation is observed by all future sellers and the buyer loses her reputation of "no-revelation-history". All future sellers can now adjust their beliefs to punish her for this deviation. For example, if they adjust their belief to be the most optimistic one,[33], then the highest feasible profit the buyer can extract from future sellers becomes $\max\{\pi(\rho, 1), 0\}$ (see Assumption 1.8.1 for the definition of $\pi(\cdot, \cdot)$), which is less than her on-path payoff $\pi(\rho, \alpha^*)$. Hence, the access to past trade details enables the dealer to build up a reputation of "no-revelation-history" and decreases dealer's incentive to reveal her private information.[34]

Order history

In this section we discuss what happens under the order history, where future traders can only observe in a previous period whether there is a trade or not but not what transaction

[33]Specifically, future sellers hold a belief that the economy state $\theta = g$ if an off-path action is observed.

[34]We note that the analysis for the public history case can be applied to a more realistic case. In such a case, if the trader declines the offer and no trade happens, then trade details of that period are not publicly announced for future traders. That is, results of this subsection can still hold even if future traders do not observe the offered price in an earlier, no-trade, period. This is because observing no information of a period implies that no trade happened in that period. Such order information is a sufficient static for the informed dealer's deviation and future uninformed traders (sellers) do not need the price information to distinguish whether or not the buyer follows the on-path equilibrium strategy. Therefore, the reputation building mechanism is not weakened and the constructed equilibria in this section can still hold.

prices are. We still focus on "k-opaque" equilibrium. The following proposition summarizes the result.

Proposition 1.8.9. *Under the order history, that is, when future sellers (traders) can only observe whether there is a trade or not in a previous period but not on which prices both parties agree, there exists a sequence* $\{\hat{\rho}_k\}_{k=1}^N$ *such that*

- $\{\hat{\rho}_k\}_{k=1}^N$ *increases in* k *and* $\hat{\rho}_N = \rho^{pool.}$*;*
- *for* $k = 1, ..., N$*, there exists a "k-opaque" equilibrium if and only if* $\rho_\ell \geq \hat{\rho}_k$*.*
- $\hat{\rho}_N = \rho_N = \rho^{pool.}$*,* $\hat{\rho}_{N-1} = \rho_{N-1}$*, and for any* $k = 1, ..., N-2$*,* $\hat{\rho}_k > \rho_k$*.*

Proof. See Appendix 1.12. ∎

From Proposition 1.8.9, the sustainability of a "k-opaque" equilibrium under the order history is similar to that under the public history. Both require traders risk-averse enough for the opaque pricing strategy to sustain. Moreover, Proposition 1.8.9 implies that restricting the information to which future traders can get access increases the chance of a discriminating and informative price from the dealer. In other words, at a certain period, releasing past orders can also make the opaque pricing strategy easier to sustain and leads to less transparency in the current market price.

The intuition behind this major distinction between public and order history is the seller's belief updating process. In the "k-opaque" equilibrium under the order history, after period $k+1$, if a decline of the trade has not been observed yet, then simply observing whether there is a transaction or not now becomes informative for sellers (traders). In fact, the decline of the trade happens only when $\theta = b$ and the trader is low-type. Because future traders can not infer previous traders' types, they can not distinguish whether the transaction happens is due to $\theta = g$, or due to the fact that a high-type seller comes when $\theta = b$. As the two scenarios happen with different probabilities, Bayesian traders start updating their beliefs after period k. As time goes by, if the decline of trade is still not observed, traders tend to be more optimistic about the underlying economy. The monotonicity assumption then implies that dealer's on-path payoff is smaller, making the opaque pricing harder to sustain under the order history than under the public history, consistent with the main insight of the benchmark model.

Dimnishing bargaining power

In this section we extend our model to relax the full bargaining power assumption.[35]

One concern of the reputation building mechanism (the comparison between public and private history) is that in reality the disclosure of past trade details may decrease dealer's bargaining power. If so, then dealer's future periods' payoffs decrease, which impairs the benefits of reputation building and makes the opaque pricing (pooling) strategy harder to sustain. As a result, such bargaining power concern potentially diminish the negative effect of the post-trade transparency on price informativeness in the OTC market. In other words, one may worry that the major mechanism we discussed above, the reputation building mechanism, is overturned by such bargaining power concern. As a result, the disclosure of past trade information may be beneficial to the price informativeness in the current period.

In this subsection, nevertheless, we mechanically model the fact that past trade details may affect the bargaining process between two parties and show that the reputation building mechanism still dominates. Specifically, we assume that the dealer can collect ψ_k fraction of the total surplus if there are k periods of past trade details available to both parties. For example, under the private history, the dealer collects ψ_0 fraction of the total payoff in every period. With more past trade details available, the seller (trader) may have a stronger bargaining power while trading with the dealer. Therefore, we make the following assumption.

Assumption 1.8.3. $\psi_0 \geq \psi_1 ... \geq \psi_{N-1}$.

We redo the analysis of equilibrium sustainability under the private history, the public history and the order history. We start with the private history. As past trade history is still unavailable, the buyer still lacks the tool to build a "no-revelation-history" reputation and extract the information rent in the future period. Thence, the cutoff $\rho^{pool.}$ in Proposition 1.8.7 remains unchanged.

We now look at the dynamic trading game. Although future periods' payoffs are decreased by a lower bargaining power, the on-path ones are still higher than the off-path ones. Thus, the benefit of reputation building still exists, and the post-trade transparency can still promote opaque pricing, compared to the static or the private history case. Proposition 1.8.10 formalizes this observation.

[35] For empirical evidences supporting imperfect (monopolisitic) competition in securities (OTC) market, see e.g., Green, Hollifield and Schürhoff (2007*a*), Ashcraft and Duffie (2007) and Massa and Simonov (2003).

Proposition 1.8.10. *Given Assumption 1.8.3, the following statements are true.*

1. *Under the private history, if and only if $\rho \geq \rho^{pool.}$, there exists a fully pooling equilibrium where the dealer can hide her private information about θ and offer a uniform price in each period.*

2. *Under the public history, for any $k = 1, ..., N$, there exists $\rho_k^{Barg.} \in [\rho_k, \rho^{pool.}]$ such that if and only if $\rho_\ell \geq \rho_k^{Barg.}$, there is a "$k-$opaque" equilibrium;*

3. *Under the order history, for any $k = 1, ..., N$, there exists $\hat{\rho}_k^{Barg.} \in [\min\{\rho_k^{Barg.}, \hat{\rho}_k\}, \rho^{pool.}]$ such that if and only if $\rho_\ell \geq \hat{\rho}_k^{Barg.}$, there is a "$k-$ opaque" equilibrium.*

Proposition 1.8.10 implies that the disclosure of past trade details have two-fold effect. On the one hand, it enables the dealer to build a "no-revelation-history" and extract future information rent and therefore facilitates the opaque equilibrium. This effect reduces market transparency, price informativeness and price discovery in the OTC market. On the other hand, the availability of past trade details decreases dealer's bargaining power and therefore, lowering future payoffs and impairing her benefit from reputation building. As a result, it partially cancels out the former effect. However, Proposition 1.8.10 shows that the former dominates the latter, and in general the disclosure of past trade details is harmful for the current price transparency. Finally, we also want to point out that the introduction of bargaining power does not alter our main results, including the structure of "$k-$opaque" equilibria and the comparison of the sustainability of such equilibria under different histories.

1.8.6 Semi-opaque Equilibrium

Finally, in this section we consider the existence of semi-opaque equilibrium. In static trading game it is possible to have a semi-opaque equilibrium. For completeness, in this section we present the sufficient and necessary condition for the existence of such kind of semi-pooling equilibrium.

Definition 1.8.3 (semi-opaque static trading equilibrium)**.** *Assume the dealer is informed and myopic, then $(p(\cdot), o(\cdot), \xi(\cdot))$ forms a semi-opaque static trading equilibrium if and only if*

1. *When the current economic state is good, the dealer offers a price that induces transaction,*

$$p_t(\theta_t = g; h^{t-1}) = \varphi a_{t-1} + J_b.$$

2. *When the current economic state is bad, the dealer mixes between hiding her information (offering the same price as in good times) and declining the trade (offering a low enough price),*

$$p_t(\theta_t = b; h^{t-1}) = \begin{cases} \varphi a_{t-1} + J_b & \text{with probability } q \\ \varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2 - \epsilon & \text{with probability } 1-q \end{cases}.$$

3. *After observing $\varphi a_{t-1} + J_b$ or $\varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2 - \epsilon$, trader Bayesian updates his belief. After observing all other prices, trader believes that the current economic state is in good times.*

$$\xi_t(p_t; h^{t-1}) = \begin{cases} \frac{\alpha_{\theta_{t-1}}}{\alpha_{\theta_{t-1}} + (1-\alpha_{\theta_{t-1}})q} & \text{if } p_t = \varphi a_{t-1} + J_b; \\ 0 & \text{if } p_t = \varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2 - \epsilon; \\ 1 & \text{otherwise.} \end{cases}$$

4. *Trader accepts any price offer weakly above $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ or exactly at $\varphi a_{t-1} + J_b$. That is, $o_t(p_t) = -1$ if and only if $p_t \geq \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ or $p_t = \varphi a_{t-1} + J_b$.*

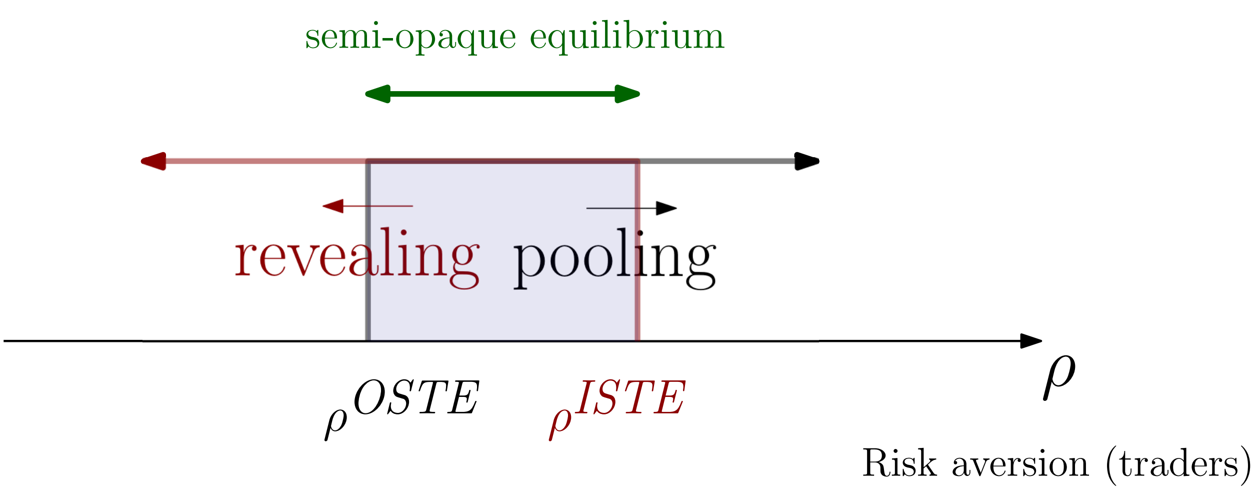

Figure 1-13: This Figure depicts OSTE, ISTE as well as the semi-opaque static trading equilibrium. When traders are sufficiently risk-averse ($\rho > \rho^{\mathsf{OSTE}}$), only OSTE can be sustained. When traders are not very risk-averse ($\rho < \rho^{\mathsf{ISTE}}$), only ISTE exists. When traders have intermediate risk-aversion coefficients ($\rho^{\mathsf{ISTE}} \leq \rho \leq \rho^{\mathsf{OSTE}}$), all three kinds of equilibria exist.

In other words, in such semi-opaque equilibrium, with probability q the dealer can hide her private information and offer an uninformative price, but with probability $(1-q)$ she will decline the trade in bad times. The indifferent condition implies that in bad times, both declining the trade and offering the opaque price should give her the same ex-post payoff (which is 0 in this case). It then follows that this opaque price has to be her evaluation of the asset in bad times. Finally trader's IR constraint implies that

$$\varphi a_{t-1} + \hat{\alpha} J_g + (1-\hat{\alpha}) J_b - \frac{\rho}{2}[\hat{\alpha}(1-\hat{\alpha})(J_g - J_b)^2 + \hat{\alpha}\sigma_g^2 + (1-\hat{\alpha})\sigma_b^2] \leq \varphi a_{t-1} + J_b,$$

where $\hat{\alpha} = \frac{\alpha_{\theta_{t-1}}}{\alpha_{\theta_{t-1}} + (1-\alpha_{\theta_{t-1}})q}$. This characterizes the condition for the existence of semi-opaque static trading equilibrium, as shown in the following Proposition.

Proposition 1.8.11. *If and only if trader's risk-aversion coefficient ρ lies in the interval $(\rho^{\mathsf{OSTE}}, \rho^{\mathsf{ISTE}}]$, there exists a semi-opaque static trading equilibrium. Moreover, there exists a semi-opaque static trading equilibrium that Pareto dominates the informative static trading equilibrium (ISTE), and is Pareto dominated by an opaque static trading equilibrium (OSTE). In conclusion,*

$$W^{\mathsf{ODTE}} = W^{\mathsf{OSTE}} > W^{\mathsf{semi}} > W^{\mathsf{ISTE}}.$$

Finally, the expected social surplus at time t in a semi-opaque static trading equilibrium increases in the trading probability in bad times, q.

Proof. See Appendix 1.12. ∎

Hence, for $\rho \in (\rho^{\mathsf{OSTE}}, \rho^{\mathsf{ISTE}}]$, both opaque equilibrium and informative equilibrium exists, and Proposition 1.8.11 shows that there exists equilibria in between. Since the dealer now does not fully reveal or fully hide her private information about the current economic state, the social surplus of such semi-opaque equilibrium lies between that of the opaque one and that of the informative one.

1.9 Conclusion

After the 2008 financial crisis, new regulations requiring mandatory transparency (via TRACE) have been implemented in many financial markets. In this paper, we study dynamic OTC markets with an informed, forward-looking, strategic and risk-neutral dealer (broker, market maker), and uninformed, myopic, strategic and risk-averse traders (retailers, customers). We develop a tractable, yet rich, model to study the role of dealer's long-term incentive (i.e., dealer's myopia) and TRACE (i.e., the public information disclosure of trade history (price and trading volumes)) on price informativeness (efficiency), market liquidity and social welfare.

We show that forward-looking incentive of informed dealers reduces market price informativeness but improves market liquidity. More importantly, we show more transparency via the public disclosure of additional information about past trades (particularly, orders or volumes, not prices), paradoxically, makes the markets more opaque, by reducing market price informativeness. However, this opacity (induced by opaque prices) is beneficial, because it increases market liquidity and welfare. Therefore, there exists a trade-off between price informativeness, liquidity, and welfare in the OTC markets. To enhance financial transparency via price informativeness and simultaneously boost market liquidity and achieve highest feasible social welfare, an effective way is to randomly audit the dealers.

The above insights are robust and continue to hold in several extensions of the model. Particularly, we demonstrate the robustness of our findings in face of a general class of payoff functions, stochastic trading positions, divisible and indivisible orders, finite and infinite trading calendars, and fixed or time-varying fundamentals.

Thanks to our explicit characterizations, we also derive several testable implications about how asset volatilities, extent of jump spread in good times and bad times, and the risk sharing motives (trader's risk-aversion) affect price informativeness and market liquidity in the OTC markets. Finally, the qualitative predictions of our model are consistent with much of the empirical evidence observed in corporate bond (with high search costs) and municipal bond markets.

Appendix

1.10 Uninformed Dealer

In this section, we briefly consider the set up in Section 1.4 in a simple case where in each period t the dealer is (like traders) fully **uninformed** about the realization of the economic state θ_t.

Clearly, when the dealer is uninformed about the economic state θ_t in period t, her price offer reveals *no* extra information about the drift and volatility (i.e. J_{θ_t} and σ_{θ_t}, respectively) of the innovation part (i.e. innov_t) in the asset value a_t. Thus, in this fully risk sharing environment, the dealer obtains all the ex-ante insurance due to residual risks $\sigma_{\theta_t} z_t$ as well as the drift innovation shocks J_{θ_t}. Therefore, in each period, the uninformed dealer, denoted by UI, offers a bid price that makes the trader indifferent between selling and keeping the asset:

$$\begin{aligned} p_t^{*\mathsf{UI}}(h^{t-1}) &= \mathbb{E}[a_t|h^{t-1}] - \frac{\rho}{2}\mathrm{Var}[a_t|h^{t-1}] \\ &= \varphi a_{t-1} + \mathbb{E}[J_{\theta_t}|h^{t-1}] - \frac{\rho}{2}\left[\mathrm{Var}[\sigma_{\theta_t} z_t|h^{t-1}] + \mathrm{Var}[J_{\theta_t}|h^{t-1}]\right] \\ &= \varphi a_{t-1} + J_b + \alpha_{\theta_{t-1}}(J_g - J_b) - \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2]. \end{aligned}$$

Any bid price above $p_t^{*\mathsf{UI}}$ is ex-ante individually rational for traders and will be accepted. Given this optimal pricing offer $p_t^{*\mathsf{UI}}$, the dealer's ex-ante payoff becomes

$$\begin{aligned} U_t^{\mathsf{UI}} &= \frac{\rho}{2}\left[\mathrm{Var}[\sigma_{\theta_t} z_t|h^{t-1}] + \mathrm{Var}[J_{\theta_t}|h^{t-1}]\right] \\ &= \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2]. \end{aligned} \tag{10.1}$$

The above analysis is independent of dealer's discount factor ρ. That is, no matter whether the dealer is myopic or forward-looking, she will choose to offer a price $p_t^{*\mathsf{UI}}$ and obtains all the (ex-ante) insurance.

It then immediately follows that the dealer's ex-ante surplus monotonically increases with trader's hedging motive (captured by his risk-aversion coefficient ρ), the variance of residual risks (captured by σ_g and σ_b) and the variance of drift shocks (captured by $J_g - J_b$). The following lemma summarizes these comparative statics.

Lemma 1.10.1. *If the dealer (either myopic or forward-looking) is uninformed about the current economic state θ_t, then she can obtain all the (ex-ante) insurance due to residual risks as well as innovation shocks. In addition, her ex-ante surplus monotonically increases with trader's hedging motive ρ, volatility of good times and bad times (i.e. $\sigma_\mathrm{g}, \sigma_\mathrm{b}$ respectively) and the additional drift of good times compared to that of bad times $J_\mathrm{g} - J_\mathrm{b}$.*

Proof of Lemma 1.10.1. When the dealer is uninformed about the economic state θ_t in period t, she cannot depend her price offer on it. In other words, dealer's price offer reveals *no* information about the drift and volatility (i.e. J_{θ_t} and σ_{θ_t}, respectively) in the asset value a_t. Therefore, we focus our attention on an equilibrium such that along the equilibrium path, the trader will not update his posterior belief of θ_t from the prior, $\xi(\theta = g|\theta_{t-1})$, and will accept any price offers that is no lower than the evaluation of a_t based on this posterior. Thus, in each period t, the uninformed dealer, denoted by UI, in order to maximize her ex-ante payoff, offers a minimum price that makes the trader indifferent between selling and keeping the asset:

$$
\begin{aligned}
p_t^{*\mathsf{UI}}(\theta_{t-1}) &= \mathbb{E}[a_t|h^{t-1}] - \frac{\rho}{2}\mathrm{Var}[a_t|h^{t-1}] \\
&= \varphi a_{t-1} + \mathbb{E}[J_{\theta_t}|h^{t-1}] - \frac{\rho}{2}\left[\mathrm{Var}[\sigma_{\theta_t} z_t|h^{t-1}] + \mathrm{Var}[J_{\theta_t}|h^{t-1}]\right] \\
&\overset{(a)}{=} \varphi a_{t-1} + \alpha_{\theta_{t-1}} J_g + (1-\alpha_{\theta_{t-1}})J_b - \frac{\rho}{2}(\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2)
\end{aligned}
\tag{10.2}
$$

where (a) fallows because

$$
\begin{aligned}
\mathbb{E}[J_{\theta_t}|h^{t-1}] &= J_b + \alpha_{\theta_{t-1}}(J_g - J_b), \\
\mathrm{Var}[\sigma_{\theta_t} z_t|h^{t-1}] &= \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2, \\
\mathrm{Var}[J_{\theta_t}|h^{t-1}] &= \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2.
\end{aligned}
\tag{10.3}
$$

In this equilibrium, the dealer will offer $p_t^{*\mathsf{UI}}$. Then we can calculate her ex-ante surplus as:

$$
\begin{aligned}
U_t &= (1-\delta)\mathbb{E}[u_t|h^{t-1}, \theta_t] + \delta\mathbb{E}[U_{t+1}|h^{t-1}] \\
&= (1-\delta)\mathbb{E}\left[\sum_{s\geq 0}\delta^s\left((a_{t+s} - p^{*\mathsf{UI}}_{t+s,\theta_{t+s}})\ \mathbf{1}(o_{t+s}(p^{*\mathsf{UI}}_{t+s,\theta_{t+s}}) = \mathrm{Sell})\right)\middle| h^{t-1}\right] \\
&= (1-\delta)\sum_{s\geq 0}\delta^s\mathbb{E}\left[(a_{t+s} - p^{*\mathsf{UI}}_{t+s,\theta_{t+s}})|h^{t-1}\right] \\
&= \frac{\rho}{2}(1-\delta)\sum_{s\geq 0}\delta^s\left[\mathrm{Var}[\sigma_{\theta_{t+s}} z_{t+s}|h^{t-1}] + \mathrm{Var}[J_{\theta_{t+s}}|h^{t-1}]\right] \\
&= \frac{\rho}{2}(1-\delta)\sum_{t=0}^{\infty}\delta^t\left[\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2\right]
\end{aligned}
$$

The comparative statics results then immediately follow. ∎

1.11 Informative Dynamic Trading Equilibrium

In this section, we also characterize the informative dynamic trading equilibrium where the dealer always distinctly prices the assets in good times and in bad times. This equilibrium (whenever is available) can be applied as a policy to improve financial transparency through the price informativeness and at the same time can attain the maximal market liquidity and social welfare. We next specify when this equilibrium exists. Then we argue that it may not be played without external force by regulators. Finally, following the discussions of Section 1.7.2, we explain how regulators can implement it. But we need to point it out that whenever both exist, there is an ODTE that defeats and Pareto dominates any IDTE.

We already know that as OSTE, ISTE is also a Nash equilibrium of the one-shot game, which is subgame perfect and remains as the PBE in the dynamic trading setting. We now look at whether this enrichment of the environment can provide us some other kinds of equilibria. For example, in ISTE, transactions only occur in good times. We now ask whether in the dynamic trading setting, can trade also occur in bad times while the dealer reveals her private information of economic states and prices the asset differently in good times and bad times?

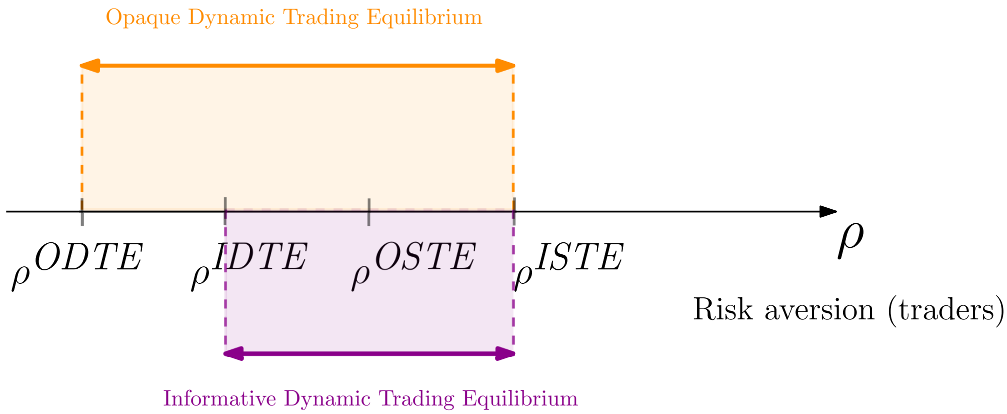

Figure 1-14: This chart considers dynamic equilibria when insider is forward-looking and $\alpha_g \geq \alpha_b$. Importantly, $\rho^{\mathsf{ODTE}} < \rho^{\mathsf{IDTE}} < \rho^{\mathsf{OSTE}}$, and the region in which ODTE exists covers the region for which IDTE exists.

Suppose that in a PBE, at period t, on the equilibrium path the dealer offers distinct prices $p_t^g(h^t)$ and $p_t^b(h^t)$ under good times and bad times, respectively. Then after observing an offer of $p_t^g(h^t)$, the trader will believe that the underlying economic state is good with certainty. His evaluation about the asset will be $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$, and he will reject any offer below this value. But also remember that this value is the highest possible evaluation from traders. Therefore, all price offers above $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ are dominated by $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$, and to induce transaction and gather positive profits in each period, the dealer will not price below this value. As a result, we learn that $p_t^g(h^t) = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ for any price history and in good times the dealer always claims a profit of $\frac{\rho}{2}\sigma_g^2$.

If in such an equilibrium, trade only occurs in good times, then in bad times dealer claims zero profit, which is the same as the ISTE payoff. Therefore, we cannot punish dealer's deviation by using payoffs in future periods. As a result, the problem is equivalent to that under one-shot game and this kind of equilibrium is equivalent to ISTE.

Then the next question to ask is, with the dynamic trading environment, can trade also occur in bad times? The following proposition summarizes the necessary and sufficient condition for the existence of such equilibria. Then the following corollary shows that whenever such an IDTE is played, as long as $\alpha_g \geq \alpha_b$, there exists an ODTE. In other words, if economic states are persistent, then IDTE is always Pareto suboptimal.

AppendixProposition 1.11.1. *If there exists a* $\pi^b \in [0, \frac{\rho}{2}\sigma_b^2]$ *such that*

$$(1 - \frac{\delta(1-\alpha_g)}{(1-\delta)(1-\delta(\alpha_g - \alpha_b))})\pi^b \leq \frac{\rho\sigma_g^2}{2} - J_g + J_b \leq [1 + \delta\frac{1 - \delta\alpha_g - (1-\delta)\alpha_b}{(1-\delta)(1-\delta(\alpha_g-\alpha_b))})]\pi^b,$$

*then there exists an informative dynamic trading equilibrium (*IDTE*) where the dealer fully reveals her private information about current economic state by providing distinct price offers at different times, and trade occurs in both* good *times and* bad *times.*

Proof. See the following Appendix for the proof. ∎

Appendix-Corollary 1.11.1. *If $\alpha_g \geq \alpha_b$, then whenever an* IDTE *exists, there also exists an* ODTE*.*

Proof. See the following Appendix for the proof. ∎

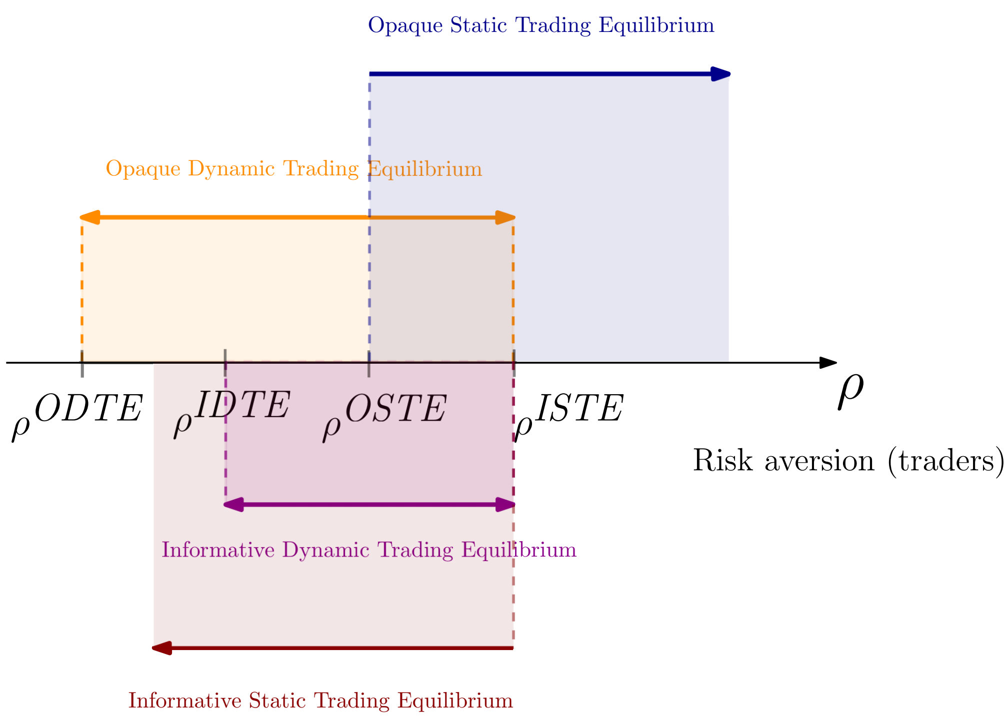

Figure 1-15: This chart considers all the equilibria (when previous orders are publicly observable). In particular, this plot shows how the dealer's forward-looking incentive can lead to market inefficiency by offering more opaque trade prices. When the economic state is persistent (i.e., $\alpha_g \geq \alpha_b$), we always have $\rho^{\mathsf{ODTE}} < \rho^{\mathsf{IDTE}} < \rho^{\mathsf{ISTE}}$.[36] When traders' risk-aversion is sufficiently low (i.e., $\rho < \rho^{\mathsf{ODTE}}$), the only equilibrium available is ISTE, where the dealer reveals her private information in each period.

We now consider the social welfare in IDTE. As trade also occurs in bad times, it gives a higher level of social welfare than that of ISTE. However, as the dealer's offer still reveals her private information, the trader has lower incentive to hedge his asset and the benefit of the trade is impaired. It then follows that the social welfare under an IDTE is still below the maximum level as in OSTE and ODTE. The next Proposition formally displays this result.

[36] Figure 1-15 depicts the circumstance where $\rho^{\mathsf{IDTE}} < \rho^{\mathsf{OSTE}}$, but it is also possible that $\rho^{\mathsf{IDTE}} \geq \rho^{\mathsf{OSTE}}$

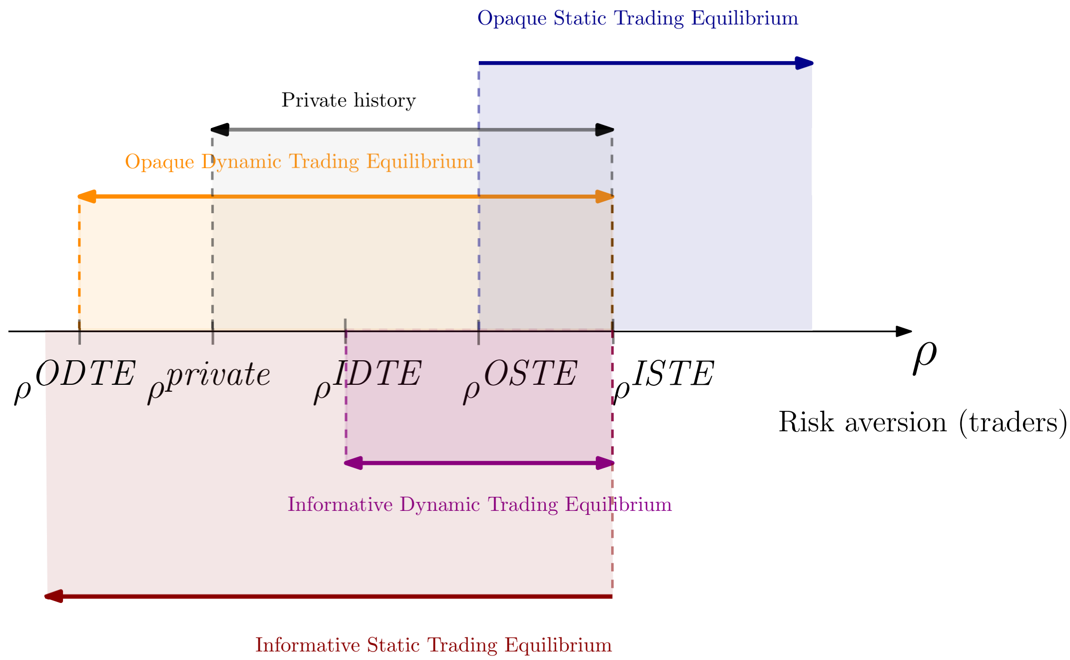

Figure 1-16: All equilibria (pooling, separating and private history) discussed in the benchmark model.

AppendixProposition 1.11.2. *If an* IDTE *exists, then the expected social welfare in period t is*

$$W^{\mathsf{IDTE}} = \frac{\rho}{2}[\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2].$$

Moreover, if the dealer is pricing discriminatingly, then this is the most social welfare one can get.

Proof. See the following Appendix for the proof. ■

Finally, as mentioned in Section 1.7.2, whenever it exists, the regulator can implement an IDTE, which provides the maximum price efficiency and market liquidity, and the highest social welfare a market with post-price transparency can get. Specifically, in each period t, the regulator can randomly check a dealer with probability q. The regulator can get access to all past economics states and price offers. If the dealer has been caught not pricing p_t^g in good times or not pricing p_t^b in bad times, then she is banned to participate in all future trade. Similarly, we show that the market transparency increases with the auditing intensity q.

AppendixProposition 1.11.3. *If for certain parameters and the auditing intensity q, there exists an* IDTE*, then for any $q' \in [q, 1]$, there also exists an* IDTE*.*

Proof. See the following Appendix for the proof. ■

1.12 Appendix [Proofs]

Proof of Proposition 1.4.1. We present our proof in two steps. First, we construct an OSTE to show the "if" part of the statement. As we will discuss later, this equilibrium also happens to be the OSTE with the highest dealer's ex-ante expected payoff. Second, we characterize dealer's incentive compatibility (IC) constraints and trader's individual rationality (IR) constraints and prove the "only if" part.

(if part:)

We prove by construction. We show that the following strategies and beliefs form a PBE:

1. In each period, the dealer offers a price $p_t(\theta_t; h^{t-1}) = p_t^{*\mathsf{UI}}(\theta_{t-1})$.
2. The trader believes that the current underlying economic state is good with probability $\alpha_{\theta_{t-1}}$. That is, $\xi_t(p_t; h^{t-1}) = \alpha_{\theta_{t-1}}, \forall p_t, h^{t-1}$.
3. The trader will accept a price offer as long as it is weakly higher than $p_t^{*\mathsf{UI}}(\theta_{t-1})$.

To show this is a PBE, we first show it's sequential rational. That is, $p_t^{*\mathsf{UI}}(\theta_{t-1})$ is a best response to the trader's belief $\xi_t(p_t; h^{t-1})$ and his order decision. In fact, if the dealer wants to trade the asset with the trader, the lowest price she can ask should make him indifferent between keeping or selling the asset. This cutoff price is $p_t = p_t^{*\mathsf{UI}}(\theta_{t-1})$. Condition (4.1) implies that $p_t^{*\mathsf{UI}}(\theta_{t-1})$ does not exceed $\varphi a_{t-1} + J_b$, and hence is weakly smaller than $\varphi a_{t-1} + J_g$. Therefore, trading is weakly preferred by the dealer than not trading, in both good times and bad times. Whence, the dealer has no other pricing strategies that are more optimal.

To check for the consistency, actually we are going to show a stronger version of consistency. We will show that constructed beliefs will satisfy the consistency requirement under PBE. We show this by constructing the following sequence of totally-mixed strategies: $\{p^n(\theta_{t-1})\}$ such that $p^n(\theta_{t-1})$ puts $(1 - \frac{1}{n})$ probability on $p_t^{*\mathsf{UI}}(\theta_{t-1})$ and uniformly takes values from all other possible strategies. The sequence of beliefs are $\xi^n = \xi_t$. Again, since $\{p^n(\theta_{t-1})\}$ does not depend on θ_t, observing the on-path price reveals zero information about θ_t. Consequently, the Bayesian updating posterior after observing any price remains at the prior level ξ_t. In other words, the constructed $(p_t^{*\mathsf{UI}}(\theta_{t-1}), \xi_t(p_t; h^{t-1}))$ is consistent.

Therefore, we just show the constructed strategies and beliefs form a PBE.

(only if part:)

Suppose $\{p_t(h^{t-1})\}$ and $\xi_t(p_t; h^{t-1})$ forms an equilibrium and trade always occurs on the equilibrium path. After observing the equilibrium price $p_t(h^{t-1})$, the trader will update his belief according to the Bayes' rule. That is,

$$\xi_t(p_t(h^{t-1}); h^{t-1}) = \frac{\text{Prob}(\theta_t = g)\text{Prob}(p_t = p_t(h^{t-1})|\theta_t = g)}{\sum\limits_{\theta_t = g,b} \text{Prob}(\theta_t)\text{Prob}(p_t = p_t(h^{t-1})|\theta_t)} = \frac{\alpha_{\theta_{t-1}}}{\alpha_{\theta_{t-1}} + 1 - \alpha_{\theta_{t-1}}} = \alpha_{\theta_{t-1}}.$$

In fact, dealer's offer does not depend on the underlying economic state θ_t, hence contains no information about it. After observing this offer, trader will not update his prior.

The dealer should offer the trader a high enough price such that he is willing to sell his asset. In other words, the price offered by the dealer should be weakly larger than the trader's evaluation of the asset, given his on-path belief $\xi_t(p_t(h^{t-1}); h^{t-1}) = \alpha_{\theta_{t-1}}$. Thus the trader's individual rationality constraint implies that

$$p_t(h^{t-1}) \geq p_t^{*\mathsf{UI}}(\theta_{t-1}). \tag{12.1}$$

On the other hand, the dealer should not have an incentive to offer a low price and prevent the trade, both in good times and in bad times. Hence the individual rationality constraints are:

$$p_t(h^{t-1}) \leq \mathbb{E}(a_t|\theta_{t-1} = g) = \varphi a_{t-1} + J_g \tag{12.2}$$

$$p_t(h^{t-1}) \leq \mathbb{E}(a_t|\theta_{t-1} = b) = \varphi a_{t-1} + J_b \tag{12.3}$$

Combine equation 12.1 and 12.3 together we get

$$\begin{aligned}
& p_t^{*\mathsf{UI}}(\theta_{t-1}) \leq p_t(h^{t-1}) \leq \varphi a_{t-1} + J_b \\
\Leftrightarrow\ & p_t^{*\mathsf{UI}} - \varphi a_{t-1} - J_b \leq 0 \\
\Leftrightarrow\ & \alpha_{\theta_{t-1}}(J_g - J_b) - \frac{\rho}{2}\left[\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2\right] \leq 0, \quad \forall \theta_{t-1} \\
\Leftrightarrow\ & \rho \geq \rho^{\mathsf{OSTE}} \equiv 2\frac{\alpha_{max}(J_g - J_b)}{\alpha_{\max}(1-\alpha_{\max})(J_g - J_b)^2 + \alpha_{\max}\sigma_g^2 + (1-\alpha_{\max})\sigma_b^2}.
\end{aligned}$$

■

Proof of Proposition 1.4.2. We prove by constructing the following strategies and beliefs:

1. The dealer offers

$$p_t(\theta_t, h^{t-1}) = \begin{cases} p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2 & \text{if } \theta_t = g; \\ p_t^b = \varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2 - \epsilon & \text{if } \theta_t = b. \end{cases}$$

2. At observing a price offer p_t, the trader (seller) will believe that the current underlying economic state is good unless $p_t < p_t^g$, in which case he will believe that $\theta_t = b$.

3. At period t, the trader accepts the offer if and only if the price is weakly above p_t^g.

First, we show that trader's strategies are sequential rational given the pricing strategy of the dealer and his own belief of trader. If he is offered p_t^b, then he will believe it is in bad times and his evaluation of the asset will be strictly greater than p_t^b. He will reject an offer p_t^b. For other price offers, his evaluation of the asset will be p_t^g. Therefore, he will accept any price above p_t^g, and reject any price below p_t^g.

Second, we show dealer's strategies constructed above are sequential rational. In fact, in good times, given the trader's belief and strategy:

- if the dealer offers a price at p_t^g, then she can collect an expected payoff of $\frac{\rho}{2}\sigma_g^2 > 0$;

- if the dealer offers a price higher than p_t^g, then the trader will still accept it. But now the dealer is worse off since she pays more for the asset;

- if the dealer offers a price lower than p_t^g. If it is p_t^b, the trader will reject it since it is smaller than the trader's evaluation given his belief that the current period is in bad times. If it is not p_t^b, then it is still lower than p_t^g, which is the trader's evaluation given his belief that it is in good times. Hence, the dealer will get 0, which is strictly lower than $\frac{\rho}{2}\sigma_g^2$, what she can get from offering p_t^g.

Similarly, in bad times:

- if the dealer offers a price at p_t^b, then the trader will reject it and the dealer gets 0;

- if the dealer offers a price other than p_t^b and the trader still rejects it, then the dealer still gets 0;

- if the dealer offers a price other than p_t^b but the trader accepts it, then since the trader believes it is in good times, this price should be weakly higher than p_t^g. But now the dealer's expected payoff has become

$$\varphi a_{t-1} + J_b - p_t^g < 0$$

  due to the inequality condition 4.2.

Finally, to show consistency, now construct totally mixed strategies $\{p^n(\theta_t)\}$: $p^n(g)$ puts $(1-\frac{1}{n})$ probability on p_t^g, and takes vales from all other possible strategies uniformly; whereas $p^n(b)$ puts $(1-\frac{1}{n^2})$ probability on p_t^g, and takes vales from all other possible strategies uniformly. We now show that the limit of the Bayesian posteriors given these strategies will be the one we constructed above. In fact, if observing a price other than p_t^g or p_t^b, then the posterior becomes

$$\begin{aligned}
\lim_{n\to\infty} \text{Prob}(\theta_t = g|p_t, h^{t-1}) &= \lim_{n\to\infty} \frac{\alpha_{\theta_{t-1}}\text{Prob}(p_t|\theta_t = g)}{\alpha_{\theta_{t-1}}\text{Prob}(p_t|\theta_t = g) + (1-\alpha_{\theta_{t-1}})\text{Prob}(p_t|\theta_t = b)} \\
&= \lim_{n\to\infty} \frac{\alpha_{\theta_{t-1}}\frac{1}{n}}{\alpha_{\theta_{t-1}}\frac{1}{n} + (1-\alpha_{\theta_{t-1}})\frac{1}{n^2}} \\
&= \lim_{n\to\infty} \frac{\alpha_{\theta_{t-1}}}{\alpha_{\theta_{t-1}} + (1-\alpha_{\theta_{t-1}})\frac{1}{n}} = 1.
\end{aligned}$$

If observing a price offer at p_t^g, then the Bayesian posterior becomes

$$\begin{aligned}\lim_{n\to\infty}\text{Prob}(\theta_t = g|p_t^g, h^{t-1}) &= \lim_{n\to\infty}\frac{\alpha_{\theta_{t-1}}\text{Prob}(p_t^g|\theta_t = g)}{\alpha_{\theta_{t-1}}\text{Prob}(p_t^g|\theta_t = g) + (1-\alpha_{\theta_{t-1}})\text{Prob}(p_t^g|\theta_t = b)}\\ &\geq \lim_{n\to\infty}\frac{\alpha_{\theta_{t-1}}(1-\frac{1}{n})}{\alpha_{\theta_{t-1}}(1-\frac{1}{n}) + (1-\alpha_{\theta_{t-1}})\frac{1}{n^2}}\\ &= \lim_{n\to\infty}\frac{\alpha_{\theta_{t-1}}(n^2-n)}{\alpha_{\theta_{t-1}}(n^2-n) + (1-\alpha_{\theta_{t-1}})} = 1.\end{aligned}$$

If observing a price offer at p_t^b, then the Bayesian posterior becomes

$$\begin{aligned}\lim_{n\to\infty}\text{Prob}(\theta_t = g|p_t^b, h^{t-1}) &= \lim_{n\to\infty}\frac{\alpha_{\theta_{t-1}}\text{Prob}(p_t^b|\theta_t = g)}{\alpha_{\theta_{t-1}}\text{Prob}(p_t^b|\theta_t = g) + (1-\alpha_{\theta_{t-1}})\text{Prob}(p_t^b|\theta_t = b)}\\ &\leq \lim_{n\to\infty}\frac{\alpha_{\theta_{t-1}}\frac{1}{n}}{\alpha_{\theta_{t-1}}\frac{1}{n} + (1-\alpha_{\theta_{t-1}})(1-\frac{1}{n^2})}\\ &= \lim_{n\to\infty}\frac{\alpha_{\theta_{t-1}}}{\alpha_{\theta_{t-1}} + (1-\alpha_{\theta_{t-1}})(n-\frac{1}{n})} = 0.\end{aligned}$$

Therefore, Bayesian posteriors calculated from $\{p^n(\theta_{t-1})\}$ approaches to the trader's beliefs, and we have proved that the construction at the beginning of this proof forms a PBE. ∎

Proof of Corollary 1.4.3. As shown, respectively, in Propositions 1.4.1 and 1.4.2

$$\rho^{\mathsf{OSTE}} = 2\frac{(J_g - J_b)}{(1-\alpha_{\max})(J_g - J_b)^2 + \sigma_g^2 + \frac{1-\alpha_{\max}}{\alpha_{\max}}\sigma_b^2},$$

and

$$\rho^{\mathsf{ISTE}} = 2\frac{J_g - J_b}{\sigma_g^2}$$

By a little algebra it is immediate that

$$\frac{(J_g - J_b)}{(1-\alpha_{\max})(J_g - J_b)^2 + \sigma_g^2 + \frac{1-\alpha_{\max}}{\alpha_{\max}}\sigma_b^2} < \frac{J_g - J_b}{\sigma_g^2},$$

completing the proof of the first part.

We now show that the constructed OSTE is a maximal PBE. The first observation is that in static game, there are only two types of equilibria, fully separating ones and fully pooling ones. For $\theta_t = g$, as shown in the analysis of ISTE the equilibrium price is $p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ Given $\rho \geq \rho^{\mathsf{OSTE}}$, it is not hard to check that $p_t^{*\mathsf{UI}} < p_t^g$. For fully pooling equilibrium, $p_t^{*\mathsf{UI}}$ is the lowest price offer the trader can accept. Therefore, the constructed equilibrium generates the maximal payoff for the dealer with private information $\theta = g$. Moreover, conditional on $\rho \geq \rho^{\mathsf{OSTE}}$, when $\theta = b$, the dealer also collects non-negative profits in the constructed equilibrium, which is higher than 0, what she can get in fully equilibrium. Hence, the constructed equilibrium also maximizes dealer's payoff when $\theta = b$.

In fact, we can actually show that when $\rho \in (\rho^{\mathsf{OSTE}}, \rho^{\mathsf{ISTE}}]$, the MOP criterion in Mailath, Okuno-Fujiwara and Postlewaite (1993) selects the same set of maximal PBE. Mailath, Okuno-Fujiwara and Postlewaite (1993) defines the MOP criterion in the following way.

Definition 1.12.1 (maximal PBE)**.** *A pure PBE* $(p(\cdot), o(\cdot), \xi(\cdot))$ *defeats another pure PBE* $(p'(\cdot), o'(\cdot), \xi'(\cdot))$ *if and only if there exists a price* $\hat{p}$ *such that*

1. $\forall \theta : p'(\theta) \neq \hat{p}$ *and* $K \equiv \{\theta | p(\theta) = \hat{p}\} \neq \emptyset$*;*

2. $\forall \theta \in K : U(p(\cdot)|\theta) \geq U(p'(\cdot)|\theta)$*, and* $\exists \theta \in K : U(p(\cdot)|\theta) > U(p'(\cdot)|\theta)$*;*

3. $\exists \theta \in K : \xi'(\theta|\hat{p}) \neq \frac{\alpha_\theta \pi(\theta)}{\sum_{\theta=g,b} \alpha_\theta \pi(\theta)}$ *for any* $\pi : \{g, b\} \to [0, 1]$ *satisfying*

   - $\theta' \in K$ *and* $U(p(\cdot)|\theta') > U(p'(\cdot)|\theta') \Rightarrow \pi(\theta') = 1$*;*
   - $\theta' \notin K \Rightarrow \pi(\theta') = 0$*.*

We say a pure PBE is undefeated if there is no other pure PBE that defeats it.

To see that an OSTE defeats an ISTE in the sense of Definition 1.12.1, we fix a price $\hat{p} = p_t^{*\mathsf{UI}}$. Then $K = \{g, b\}$. For $\forall \theta \in \{g, b\}$, $U_t^{\mathsf{ISTE}}(\theta = g) = \varphi a_{t-1} + J_g - p_t^g = \frac{\rho}{2}\sigma_g^2$, $U_t^{\mathsf{ISTE}}(\theta = b) = 0$; and

$$\begin{aligned} & U_t^{\mathsf{OSTE}}(\theta) \\ = \; & \varphi a_{t-1} + J_\theta - p_t^{*\mathsf{UI}} \\ = \; & \begin{cases} (1-\alpha_{\theta_{t-1}})(J_g - J_b) + \frac{\rho}{2}[\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2] & \text{if } \theta = g; \\ -\alpha_{\theta_{t-1}}(J_g - J_b) + \frac{\rho}{2}(\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2) & \text{if } \theta = b. \end{cases} \end{aligned}$$

If $\rho > \rho^{\mathsf{OSTE}}$, then one can easily check that $U_t^{\mathsf{OSTE}}(\theta = b) > U_t^{\mathsf{ISTE}} = 0(\theta = b)$. When $\theta = g$,

$$\begin{aligned} & U_t^{\mathsf{OSTE}}(\theta = g) - U_t^{\mathsf{ISTE}}(\theta = g) \\ = \; & (1 - \alpha_{\theta_{t-1}})[J_g - J_b + \frac{\rho}{2}(\sigma_b^2 - \sigma_g^2 + \alpha_{\theta_{t-1}}(J_g - J_b)^2)] \\ > \; & 0. \end{aligned}$$

Second, as $U_t^{\mathsf{OSTE}}(\theta) > U_t^{\mathsf{ISTE}}(\theta)$ for both $\theta = g$ and $\theta = b$, the only feasible $\pi : \{g, b\} \to [0, 1]$ is that $\pi(\theta) = 1$, for all $\theta \in \{g, b\}$. Then

$$\frac{\alpha_{\theta_{t-1}}\pi(g)}{\alpha_{\theta_{t-1}}\pi(g) + (1-\alpha_{\theta_{t-1}})\pi(b)} = \alpha_{\theta_{t-1}} \neq \xi^{\mathsf{ISTE}}(\theta = g|p_t^{*\mathsf{UI}}) = 1;$$
$$\frac{(1-\alpha_{\theta_{t-1}})\pi(b)}{\alpha_{\theta_{t-1}}\pi(g) + (1-\alpha_{\theta_{t-1}})\pi(b)} = 1 - \alpha_{\theta_{t-1}} \neq \xi^{\mathsf{ISTE}}(\theta = b|p_t^{*\mathsf{UI}}) = 0.$$

Finally, in static case there are only two types of pure strategy equilibria, OSTE and ISTE, and the former defeats the latter, it is obviously that OSTE are the only equilibria that survive the refinement in Mailath, Okuno-Fujiwara and Postlewaite (1993). ∎

Proof of Proposition 1.5.1. We prove the proposition in three steps. First, we show the "if" part by construction. Second, we prove the "only if" part by showing the off path payoff in our construction above is the harshest feasible punishment one can put for deviation. Third, we show that he cutoff for the existence of ODTE is weakly below that of OSTE. Finally, we show the constructed ODTE is a maximal PBE when there is equilibrium multiplicity.

(I) if part

We will first construct a set of strategies and beliefs and characterize the sufficient condition for this set to become a PBE. We then show that more risk-averse traders will make these conditions easier to sustain and this sufficient condition provides a lower bound for ρ, trader's risk-aversion coefficients.

We now show the following construction forms a PBE.

- If the dealer has not deviated in previous periods, then
    1. the dealer offers a price $p_t(\theta_t, h^{t-1}) = p_t^{*\mathsf{UI}}(\theta_{t-1})$;
    2. the trader believes that the current underlying economic state is good with probability $\alpha_{\theta_{t-1}}$ i.e., $\xi(p_t; h^{t-1}) = \alpha_{\theta_{t-1}}, \forall p_t, h^{t-1}$;
    3. the trader will accept any price offers that is weakly above $p_t^{*\mathsf{UI}}(\theta_{t-1})$.
- If the dealer deviates previously, then both parties observe this and play an informative static trading equilibrium (ISTE).

(i) consistency

It is easy to check that the belief construction follows the Bayes rule whenever it applies.

(ii) sequential rationality

First, there is no incentives for trader to deviate or dealer to deviate off the equilibrium path. It follows immediately that given his belief, the trader has no incentive to deviate on the equilibrium path. Since ISTE is a one-shot Nash equilibrium, both parties have no incentives to deviate if previous deviation has been observed.

Second, for dealer along the equilibrium path, at period t, offering a price strictly higher than $p_t^{*\mathsf{UI}}(\theta_{t-1})$ will not change trader's response, but will lower her payoff for the current period and trigger punishment in all future periods. As a result, she has no incentive to do that.

Third, the only possible deviation left is that the dealer may want to offer a price lower than $p_t^{*\mathsf{UI}}(\theta_{t-1})$ and refuse the trade. Again in good times, this lead to a loss of positive profit in the current period, and trigger punishment in all future periods. Hence it is not profitable to deviate in good times.

Finally, in bad times, dealer's incentive compatibility constraint becomes:

$$U_t^{\mathsf{ODTE}}(b,\theta_{t-1}) \geq \underbrace{(1-\delta)\times 0}_{\substack{\text{dealer's static payoff}\\ \text{after rejecting the trade}}} + \underbrace{\delta[\alpha_b U_t^{\mathsf{off}}(g) + (1-\alpha_b)U_t^{\mathsf{off}}(b)]}_{\substack{\text{dealer's continuous}\\ \text{off-path payoff}}}, \forall\theta_{t-1} \tag{12.4}$$

where $\forall\theta_t, \theta_{t-1}$,

$$\begin{aligned} U_t^{\mathsf{ODTE}}(\theta_t,\theta_{t-1}) &= (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p_t^{*\mathsf{UI}}(\theta_{t-1})) \\ &\quad +\delta[\alpha_{\theta_t} U_{t+1}^{\mathsf{ODTE}}(g,\theta_t) + (1-\alpha_{\theta_t})U_{t+1}^{\mathsf{ODTE}}(b,\theta_t)] \\ U_t^{\mathsf{off}}(\theta_t) &= (1-\delta)\mathbf{1}(\theta_t = g)\cdot\frac{\rho}{2}\sigma_g^2 + \delta[\alpha_{\theta_t} U_{t+1}^{\mathsf{off}}(g) + (1-\alpha_{\theta_t})U_{t+1}^{\mathsf{off}}(b)]. \end{aligned}$$

are dealer's on-path and off-path continuous payoffs respectively, given the current economic state is θ_t and the last period economic state is θ_{t-1}.

(iii) lower bound of ρ

We now show that condition (12.9) provides a lower bound for ρ, trader's risk-aversion coefficient.

First, observe that both

$$\varphi a_{t-1}+J_g-p_t^{*\mathsf{UI}}(\theta_{t-1})=(1-\alpha_{\theta_{t-1}})(J_g-J_b)+\frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2+\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2]$$

and

$$\varphi a_{t-1}+J_b-p_t^{*\mathsf{UI}}(\theta_{t-1}) = \alpha_{\theta_{t-1}}(J_b-J_g)+\frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2+\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2]$$

are independent of the calender time t. Since both dealer's static payoffs and her on-path equilibrium strategies solely depend on the current economic state θ_t and last period economic state θ_{t-1}, so do her on-path continuous payoffs in this equilibrium . Let's denote $\Lambda^{\theta_t,\theta_{t-1}} \equiv \varphi a_{t-1}+J_{\theta_t}-p_t^{*\mathsf{UI}}(\theta_{t-1})$ and $U^{\mathsf{ODTE}}(\theta_t,\theta_{t-1}) \equiv U_t^{\mathsf{ODTE}}(\theta_t,\theta_{t-1})$. One can then solve the on-path payoffs recursively from the equation system:

$$U^{\mathsf{ODTE}}(\theta_t,\theta_{t-1}) = (1-\delta)\Lambda^{\theta_t,\theta_{t-1}} + \delta[\alpha_{\theta_t}U^{\mathsf{ODTE}}(g,\theta_t) + (1-\alpha_{\theta_t})U^{\mathsf{ODTE}}(b,\theta_t)], \forall\theta_t,\theta_{t-1}.$$

Or in the matrix form,

$$\begin{bmatrix} 1-\delta\alpha_g & 0 & -\delta(1-\alpha_g) & 0 \\ -\delta\alpha_g & 1 & -\delta(1-\alpha_g) & 0 \\ 0 & -\delta\alpha_b & 1 & -\delta(1-\alpha_b) \\ 0 & -\delta\alpha_b & 0 & 1-\delta(1-\alpha_b) \end{bmatrix} \cdot \begin{bmatrix} U^{\mathsf{ODTE}}(g,g) \\ U^{\mathsf{ODTE}}(g,b) \\ U^{\mathsf{ODTE}}(b,g) \\ U^{\mathsf{ODTE}}(b,b) \end{bmatrix} = (1-\delta)\cdot\begin{bmatrix} \Lambda^{g,g} \\ \Lambda^{g,b} \\ \Lambda^{b,g} \\ \Lambda^{b,b} \end{bmatrix}.$$

Similarly $U_t^{\mathsf{off}}(\theta_t)$ is stationary and can be solved from

$$\begin{bmatrix} 1-\delta\alpha_g & -\delta(1-\alpha_g) \\ -\delta\alpha_b & 1-\delta(1-\alpha_b) \end{bmatrix} \cdot \begin{bmatrix} U^{\mathsf{off}}(g) \\ U^{\mathsf{off}}(b) \end{bmatrix} = (1-\delta) \begin{bmatrix} \frac{\rho}{2}\sigma_g^2 \\ 0 \end{bmatrix}$$

The following lemma shows that an increase in ρ will relax inequality (12.9), the condition for the existence of an ODTE. In fact, the exact cutoff ρ^{ODTE} is obtained when equality holds.

Lemma 1.12.3.1. If for trader's risk-aversion coefficient ρ_1, the PBE constructed above forms to be an ODTE, then for any $\rho_2 > \rho_1$, it still remains as an ODTE if trader's risk-aversion coefficient is ρ_2.

Proof. Given the construction and the analysis right after, we only need to check that dealer's IC constraint in bad times still holds for ρ_2. Fix any realization of economic states $\{\theta_\tau\}_{\tau=t}^{\infty}$ starting from the current period t, then at any future period $\tau > t$, the loss due to the deviation punishment becomes $\Lambda^{\theta_\tau,\theta_{\tau-1}} - \mathbf{1}(\theta_\tau = g) \cdot \frac{\rho_2}{2}\sigma_g^2$. When $\theta_\tau = g$, this difference

$$\begin{aligned} \Lambda^{g,\theta_{\tau-1}} - \frac{\rho_2}{2}\sigma_g^2 &= (1-\alpha_{\theta_{\tau-1}})(J_g - J_b) + \frac{\rho_2(1-\alpha_{\theta_{\tau-1}})}{2}[\alpha_{\theta_{\tau-1}}(J_g - J_b)^2 + \sigma_b^2 - \sigma_g^2] \\ &\geq (1-\alpha_{\theta_{\tau-1}})(J_g - J_b) + \frac{\rho_1(1-\alpha_{\theta_{\tau-1}})}{2}[\alpha_{\theta_{\tau-1}}(J_g - J_b)^2 + \sigma_b^2 - \sigma_g^2] \end{aligned}$$

is greater than the deviation punishment with less risk-averse traders. Similarly when $\theta_\tau = b$, then the difference

$$\begin{aligned} \Lambda^{b,\theta_{\tau-1}} - 0 &= \alpha_{\theta_{\tau-1}}(J_b - J_g) + \frac{\rho_2}{2}[\alpha_{\theta_{\tau-1}}(1-\alpha_{\theta_{\tau-1}})(J_g - J_b)^2 + (1-\alpha_{\theta_{\tau-1}})\sigma_b^2 + \alpha_{\theta_{\tau-1}}\sigma_g^2] \\ &\geq \alpha_{\theta_{\tau-1}}(J_b - J_g) + \frac{\rho_1}{2}[\alpha_{\theta_{\tau-1}}(1-\alpha_{\theta_{\tau-1}})(J_g - J_b)^2 + (1-\alpha_{\theta_{\tau-1}})\sigma_b^2 + \alpha_{\theta_{\tau-1}}\sigma_g^2] \end{aligned}$$

is also greater than that under the case of ρ_1.

Furthermore, at period t, the deviation gain becomes $0 - \Lambda^{b,\theta_{t-1}}$ and decreases in ρ, implying the deviation gain is smaller with more risk-averse trader.

Therefore, with more risk-averse trader, or as ρ increases, dealer's current gain from declining to trade is smaller, but she now faces harsher punishment in all future periods in all realization of economic state shocks. Hence, condition (12.4) is easier to sustain with more risk-averse trader. In other words, if condition (12.4) holds for ρ_1, then it also holds for ρ_2 and with ρ_2 there exists an ODTE as constructed. ∎

(II) only if part

First, we observe that ISTE is the harshest punishment for dealer's deviation. Then, we show that the ODTE constructed in the "if" part is actually the one easiest to sustain and provides a lower bound for the existence of an ODTE (even without the highest dealer's payoff restriction).

Before we move on to prove that, the first observation is that .

Proof of Lemma 1.5.1. In bad times, first observe that the dealer can always offer a low enough price to decline the trade and collect an ex-post payoff of 0.

Next, we show that for the following belief, she can do no better than that. Suppose the trader believes that the current period is in good times unless $\varphi a_{t-1} + J_b$ is offered, at which he believes it is in bad times, then he will only accept any offer above $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ or exactly at $\varphi a_{t-1} + J_b$. Offering a price weakly higher than $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ leads to non-positive payoff, while offering $\varphi a_{t-1} + J_b$ leads to zero payoff. Hence, with this trader belief and trader best responding to it, the best case dealer's ex-post payoff is 0. In other words,

$$\max_{p_b} \min_{o_t \in \mathsf{BR}_{\xi(p_t;\cdot)}} (\varphi a_{t-1} + J_b - p_b) \cdot \Big(- o_t(p_b)\Big) \geq 0, \quad \forall \xi(p_t;\cdot).$$

In good times, first observe that the dealer's ex-post payoff will be at least $\frac{\rho}{2}\sigma_g^2$. In fact, for any $\xi(p_t;\cdot)$,

$$\begin{aligned}
\mathbb{E}_{\xi(\cdot|p)}(a_t|h^{t-1}) &= \varphi a_{t-1} + J_b + \xi(p_t;\cdot)(J_g - J_b) - \frac{\rho}{2}\xi(p_t;\cdot)(1-\xi(p_t;\cdot))(J_g - J_b)^2 \\
&\quad - \frac{\rho}{2}[\xi(p_t;\cdot)\sigma_g^2 + (1-\xi(p_t;\cdot))\sigma_b^2] \\
&\leq \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2.
\end{aligned}$$

Therefore, the dealer can always offer a price $p_g' = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ to confirm the order from the trader and guarantee an ex-post payoff of $\frac{\rho}{2}\sigma_g^2$ for sure.

Next, this is the best she can get if the trade believes that the current period is always in good times. As only prices weakly above p_g' will be accepted by the trader. So we have,

$$\max_{p_g} \min_{o_t \in \mathsf{BR}_{\xi(p_t;\cdot)}} (\varphi a_{t-1} + J_g - p_g) \cdot \Big(- o_t(p_g)\Big) \geq \frac{\rho}{2}\sigma_g^2, \quad \forall \xi(p_t;\cdot).$$

Finally, these ex-post payoffs can be obtained in ISTE. In conclusion, for any arbitrary beliefs trader may hold, dealer's ex-post ISTE payoffs ($\frac{\rho}{2}\sigma_g^2$ in good times and 0 in bad times) are no larger than the min-max static payoffs given that the trader best responses to her strategy. ∎

We now go back to prove the "only if" part. Suppose in an ODTE the dealer offers p_t^* at time t, and trade always occurs. Therefore, the trader learns no information via observing the on-path price history, and for any consistent belief system, her posteriors should stay the same as the prior, that $\xi(p_t^*; h^{t-1}) = \alpha_{\theta_{t-1}}$. Therefore, for the trader to accept this offer, we need $p_t^* \geq p_t^{*\mathsf{UI}}(\theta_{t-1})$.

From the dealer's perspective, the continuous payoff $\hat{U}_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1})$, given the current and the last economic states are θ_t and θ_{t-1}, will be no more than $U^{\mathsf{ODTE}}(\theta_t, \theta_{t-1})$, because in each period starting from t, the dealer is offering a price greater than the on-path price in ODTE constructed in the if part above. Therefore, we have

$$
\begin{aligned}
& U^{\mathsf{ODTE}}(b,\theta_{t-1}) \\
\geq\ & \hat{U}_t^{\mathsf{ODTE}}(b,\theta_{t-1}) \\
\overset{(1)}{\geq}\ & (1-\delta)\times 0+\sum_{k=1}^{\infty}\delta^k \\
& \sum_{h^{t+k}|h^t}\Pr(h^{t+k}|h^t)\max_{p_g,p_b}\min_{o_{t+k}\in\mathsf{BR}_{\xi_{t+k}(p_t;\cdot)}}(\varphi a_{t+k-1}+J_{\theta t+k}-p_{t+k}(\theta_{t+k},h^{t+k-1}))\cdot \\
& \qquad\qquad\Big(-o_{t+k}(p_{t+k}(\theta_{t+k},h^{t+k-1}))\Big) \\
\overset{(2)}{\geq}\ & (1-\delta)\times 0+\delta[\alpha_b U^{\mathsf{off}}(g)+(1-\alpha_b)U^{\mathsf{off}}(b)]
\end{aligned}
$$

Inequality (1) follows from the incentive compatibility of the dealer in bad times. inequality (2) comes from Lemma 1.5.1, that the lowest min-max payoffs in each following history are weakly larger than ISTE payoffs, which are the off-path payoffs. Therefore, whenever an ODTE exists, condition (12.9) holds, and solving the above inequality we get $\rho\geq\rho^{\mathsf{ODTE}}$.

(III) $\rho^{\mathsf{OSTE}}\geq\rho^{\mathsf{ODTE}}$:

we prove the following claim. That is, whenever an OSTE can sustain, so can an ODTE.

Lemma 1.12.4. *Fix all parameter values, whenever there is an* ODTE *equilibrium, there also exists an* OSTE *equilibrium.*

Proof of Lemma 1.12.4. Suppose σ^* is such an OSTE, then dealer's IC constraint in bad times implies that $\varphi a_{t-1}+J_b-p_t(h^t)\geq 0$. Trader's IR constraint implies that $p_t(h^{t-1})\geq p_t^{*\mathsf{UI}}(\theta_{t-1})$. Therefore,

$$\varphi a_{t-1}+J_b-p^{\mathsf{UI}}(\theta_{t-1})\geq 0$$

Consider the ODTE constructed in the "if" part of this proof. From above we know this construction consists of an equilibrium in our dynamic trading game if the following inequality holds, $\forall\theta_{t-1}$,

$$
\begin{aligned}
& (1-\delta)(\varphi a_{t-1}+J_b-p^{\mathsf{UI}}(\theta_{t-1}))+\delta[\alpha_b U^{\mathsf{ODTE}}(g,b)+(1-\alpha_b)U^{\mathsf{ODTE}}(b,b)] \\
\geq\ & (1-\delta)\times 0+\delta[\alpha_b U^{\mathsf{off}}(g)+(1-\alpha_b)U^{\mathsf{off}}(b)]. \qquad (12.5)
\end{aligned}
$$

But from Lemma 1.5.1, $U^{\mathsf{off}}(\theta_t)$ is the lowest possible dealer's ex-post payoff given traders best response to dealer's offers. Therefore, in all future periods and at all future histories, dealer derives weakly lower payoff on the punishment stage than on the equilibrium path.

$$U^{\mathsf{ODTE}}(\theta_t,b)\geq U^{\mathsf{off}}(\theta_t),\forall\theta_t$$

and inequality (12.5) is satisfied. ∎

(IV) check for the MOP criterion

First, in the ODTE constructed at the beginning of the proof, σ^{ODTE}, along the equilibrium path, opaque pricing is always provided by the dealer, thus the maximal ex-ante social surplus is achieved in every period from our analysis in Section 1.7.3. The next observation is that in σ^{ODTE}, the dealer collects all the social surplus. Thus, along the equilibrium path, she achieves the maximal payoff she can ever get in a PBE. Therefore, this PBE consists of an maximal PBE, which defeats and Pareto dominates the equilibrium where ISTE is played in every period.

More importantly, we show that this maximal PBE outcome is unique and only achieved via an ODTE. We prove by contradiction. Suppose not and there is another PBE $\sigma' = (p'(\cdot), o'(\cdot), \xi'(\cdot))$ that is a maximal PBE. In this PBE, in some periods, along the equilibrium path, the dealer offers discriminatingly. Suppose t is the earliest one of such a period. Then again follow the same analysis in Section 1.7.3, as in period t the dealer reveals her private information about θ_t, she will collect strictly lower ex-ante payoff than in σ^{ODTE}. The continuation payoff of σ^{ODTE}, as it already achieves the highest feasible level, is also weakly higher than that of σ'. Thus, we have $U_t(p'(\cdot), o'(\cdot), \xi'(\cdot)) < U_t(p^{\mathsf{ODTE}}(\cdot), o^{\mathsf{ODTE}}(\cdot), \xi^{\mathsf{ODTE}}(\cdot))$, which contradicts with our definition of maximal PBE and completes the proof.

∎

Proof of Proposition 1.6.1. Proof follows from the text in the main body. ∎

Proof of Proposition 1.6.2. As in the proof of Proposition 1.5.1, we first present the "if" part, followed by the proof of "only if" part, and then compare the threshold of trader's risk-aversion coefficient in this private order history setup with those of ODTE and OSTE. Finally, we check for maximal PBE.

(I) if part

We prove the "if" part by construction in three steps. First let's fix a partition (Y_{-1} and Y_0) of the signal space Y, and denote the corresponding probabilities as $f_i = \text{Prob}(y_s \in Y_{-1} | o_t = i), i = -1, 0, \forall s$. We then construct specific type of private equilibrium displayed below. Next, we check for consistency and the sequential rationality, and characterize $\rho^{\mathsf{private}}(f_{-1}, f_0)$. Finally, we search over all partitions of the signal space Y to find the lower bound of trader's risk-aversion coefficient ρ^{private} for the existence of an opaque equilibrium in private order history setup.

The private equilibrium we construct is as follows:

1. If $y^s \in Y_{-1}, \forall s \leq t-1$, then in period t
    - the dealer chooses to offer a price $p_t^*(h^{t-1})$;
    - the trader holds his prior belief after observing $p_t^*(h^{t-1})$ and believes for sure that the economic state is good for all other prices, i.e.,

$$\xi(p_t; h^{t-1}) = \begin{cases} \alpha_{\theta_{t-1}} & \text{if } p_t = p_t^*(h^{t-1}) \\ 1 & \text{if } p_t \neq p_t^*(h^{t-1}) \end{cases}, \forall h^{t-1};$$

- traders will accept any prices weakly above p_t^g or exactly at $p_t^*(h^{t-1})$, and decline all other prices.

2. If $y^s \in Y_0$ for some $s \le t-1$, then at period t both parties play an ISTE equilibrium.[37]

Next, to see sequential rationality, we characterize dealer's incentive compatibility constraint in bad times and check other individual rationality and incentive compatibility constraints for both players at all times are fulfilled. One can check that no profitable deviations off the path. On the equilibrium path, the trader is best responding and the dealer has no incentive to offer a higher price at all times or to offer a lower price in good times. Therefore, the only condition left to check becomes the incentive compatibility constraint of the dealer in bad times: for $\forall \theta_{t-1}$,

$$\begin{aligned} U_t^{\mathsf{private}}(b, \theta_{t-1}; f_{-1}, f_0) \quad \ge \quad & \underbrace{(1-\delta)\times 0}_{\text{dealer's static payoff after rejecting the trade}} \\ & + \underbrace{\delta \cdot f_0 \cdot [\alpha_b U_{t+1}^{\mathsf{private}}(g, b; f_{-1}, f_0) + (1-\alpha_b) U_{t+1}^{\mathsf{private}}(b, b; f_{-1}, f_0)]}_{\text{dealer's continuous if } y_t \in Y_{-1}} \\ & + \underbrace{\delta \cdot (1-f_0) \cdot [\alpha_b U_{t+1}^{\mathsf{off}}(g) + (1-\alpha_b) U_{t+1}^{\mathsf{off}}(b)]}_{\text{dealer's continuous payoff if } y_t \in Y_0} \end{aligned} \tag{12.6}$$

where $\forall \theta_t, \theta_{t-1}$,

$$\begin{aligned} U_t^{\mathsf{private}}(\theta_t, \theta_{t-1}; f_{-1}, f_0) \quad = \quad & (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p_t^{*\mathsf{UI}}(\theta_{t-1})) \\ & + \delta f_{-1} \cdot [\alpha_{\theta_t} U_{t+1}^{\mathsf{private}}(g, \theta_t; f_{-1}, f_0) + (1-\alpha_{\theta_t}) U_{t+1}^{\mathsf{private}}(b, \theta_t; f_{-1}, f_0)] \\ & + \delta(1-f_{-1}) \cdot [\alpha_{\theta_t} U_{t+1}^{\mathsf{off}}(g) + (1-\alpha_{\theta_t}) U_{t+1}^{\mathsf{off}}(b)]. \end{aligned}$$

and $U_t^{\mathsf{off}}(\theta_t)$ is the continuous payoff on the punishment stage or in the ISTE equilibrium, which is no different from that in public history environment.

Similarly we have

$$\Lambda^{\theta_t, \theta_{t-1}} = \begin{cases} (1-\alpha_{\theta_{t-1}})(J_g - J_b) - \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2] & \text{if } \theta_t = g, \\ \alpha_{\theta_{t-1}}(J_g - J_b) - \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2] & \text{if } \theta_t = b, \end{cases}$$

so by symmetry we have $U^{\mathsf{private}}(\theta_t, \theta_{t-1}; f_{-1}, f_0) \equiv U_t^{\mathsf{private}}(\theta_t, \theta_{t-1}; f_{-1}, f_0)$.

Rearranging we get condition 12.6 is equivalent to

$$\begin{aligned} & (1-\delta)\Lambda^{b, \theta_{t-1}} + \delta(f_{-1} - f_0)[\alpha_b U^{\mathsf{private}}(g, b; f_{-1}, f_0) + (1-\alpha_b) U^{\mathsf{private}}(b, b; f_{-1}, f_0)] \\ \ge \quad & \delta(f_{-1} - f_0) \cdot [\alpha_b U_{t+1}^{\mathsf{off}}(g) + (1-\alpha_b) U_{t+1}^{\mathsf{off}}(b)]. \end{aligned} \tag{12.7}$$

$\rho^{\mathsf{private}}(f_{-1}, f_0)$ is constructed as the smallest ρ such that condition 12.7 is satisfied, i.e., the

[37]Here we assume once the game moves into the punishment stage, it never comes back. Future works can study more complicated structure like the one in Ellison (1994), where there is low probability that the game moves back to the first stage from the punishment stage.

ρ that makes the equality holds in condition 12.7. Following the similar steps as in Lemma 1.12.3.1, we can claim that for any $\rho \geq \rho^{\mathsf{private}}(f_{-1}, f_0)$, the dealer's IC constraints in bad times are satisfied. The reason is that an increase in ρ reduces the deviation gain in current period as well as enlarges the punishment loss in all future periods at all realizations.

The check of consistency is the same as in the proof of Proposition 1.5.1. Therefore, the strategies and beliefs constructed at the beginning of this proof form a perfect Bayesian equilibrium.

Finally, we find the optimal partition of the signal space Y that derives the lowest possible $\rho^{\mathsf{private}}(f_{-1}, f_0)$:

$$\rho^{\mathsf{private}} = \min_{\substack{Y_{-1}, Y_0 \ s.t. \\ Y_{-1} \cap Y_0 = \emptyset, Y_{-1} \cup Y_0 = Y, Y_{-1} \neq \emptyset}} \rho^{\mathsf{private}}(f_{-1}, f_0).$$

(II) only if part

To check that $\rho \geq \rho^{\mathsf{private}}$ is necessary for the existence of an opaque dynamic trading equilibrium, fix any ODTE σ in this model with imperfect observation of previous orders. Suppose in this equilibrium the dealer offers p_t on the equilibrium path, and the probability of not triggering a punishment (staying on-path) is f'_{-1} and f'_0, given that the dealer not deviates and deviates at period t correspondingly. We adopt the same idea as in Proposition 1.5.1. Because playing ISTE forever is the harshest punishment that can be put on the dealer, and $p_t^{*\mathsf{UI}}(\theta_{t-1})$ is the lowest price the dealer can charge to still keep the trader. We have

$$\begin{aligned}
& U_t^{\mathsf{private}}(b, \theta_{t-1}; f'_{-1}, f'_0) \\
= \ & (1-\delta)(\varphi a_{t-1} + J_b - p_t) + \delta f'_{-1}[\alpha_b U_{t+1}^{\mathsf{private}}(g, b; f'_{-1}, f'_0) + (1-\alpha_b) U_{t+1}^{\mathsf{private}}(b, b; f'_{-1}, f'_0)] \\
& +\delta(1-f'_{-1}) \cdot [\min_{\xi(p;\cdot)} \max_{p_g, p_b} \min_{o_t \in \mathsf{BR}_{\xi(p;\cdot)}} \alpha_b(\varphi a_{t-1} + J_g - p_g) \cdot \Big(- o_t(p_g)\Big) \\
& \qquad\qquad +(1-\alpha_b)(\varphi a_{t-1} + J_b - p_b) \cdot \Big(- o_t(p_b)\Big)] \\
\overset{(1)}{\geq} \ & (1-\delta) \times 0 + \delta f'_0[\alpha_b U_{t+1}^{\mathsf{private}}(g, b; f'_{-1}, f'_0) + (1-\alpha_b) U_{t+1}^{\mathsf{private}}(b, b; f'_{-1}, f'_0)] + \\
& \delta(1-f'_0) \cdot [\min_{\xi(p;\cdot)} \max_{p_g, p_b} \min_{o_t \in \mathsf{BR}_{\xi(p;\cdot)}} \alpha_b(\varphi a_{t-1} + J_g - p_g) \cdot \Big(- o_t(p_g)\Big) \\
& \qquad\qquad +(1-\alpha_b)(\varphi a_{t-1} + J_b - p_b) \cdot \Big(- o_t(p_b)\Big)],
\end{aligned}$$

$$\begin{aligned}
\overset{(2)}{\Leftrightarrow} \ & (1-\delta)(\varphi a_{t-1} + J_b - p_t) + \delta f'_{-1}[\alpha_b U_{t+1}^{\mathsf{private}}(g, b; f'_{-1}, f'_0) + (1-\alpha_b) U_{t+1}^{\mathsf{private}}(b, b; f'_{-1}, f'_0)] \\
& +\delta(1-f'_{-1}) \cdot \frac{\rho}{2} \alpha_b \sigma_g^2 \\
\geq \ & (1-\delta) \times 0 + \delta f'_0[\alpha_b U_{t+1}^{\mathsf{private}}(g, b; f'_{-1}, f'_0) + (1-\alpha_b) U_{t+1}^{\mathsf{private}}(b, b; f'_{-1}, f'_0)] \\
& +\delta(1-f'_0) \cdot \frac{\rho}{2} \alpha_b \sigma_g^2
\end{aligned}$$

Hence

$$
\begin{aligned}
&(1-\delta)\Lambda^{b,\theta_{t-1}}+\delta(f_{-1}-f_0)[\alpha_b U^{\mathsf{private}}(g,\theta_t;f_{-1},f_0)+(1-\alpha_b)U^{\mathsf{private}}(b,\theta_t;f_{-1},f_0)-\frac{\rho}{2}\alpha_b\sigma_g^2] \\
\overset{(3)}{\geq}\ & (1-\delta)(\varphi a_{t-1}+J_b-p_t)+\delta(f_{-1}-f_0)\cdot \\
&\qquad [\alpha_b U^{\mathsf{private}}_{t+1}(g,b;f_{-1},f_0)+(1-\alpha_b)U^{\mathsf{private}}_{t+1}(b,b;f_{-1},f_0)-\frac{\rho}{2}\alpha_b\sigma_g^2] \\
\overset{(4)}{\geq}\ & (1-\delta)(\varphi a_{t-1}+J_b-p_t)+\delta(f'_{-1}-f'_0)\cdot \\
&\qquad [\alpha_b U^{\mathsf{private}}_{t+1}(g,b;f'_{-1},f'_0)+(1-\alpha_b)U^{\mathsf{private}}_{t+1}(b,b;f'_{-1},f'_0)-\frac{\rho}{2}\alpha_b\sigma_g^2] \\
\geq\ & 0.
\end{aligned}
$$

Here inequality (1) comes from dealer's IC constraint in bad times, equivalence (2) is from Lemma 1.5.1, inequality (3) comes from the fact that $p_t \geq p_t^{*\mathsf{UI}}$, and inequality (4) is due to the construction of f_{-1} and f_0, and the fact that $\frac{\rho}{2}\alpha_b\sigma_g^2$ is the min-max payoff.

Therefore, we derive a necessary condition for the existence of opaque equilibrium in private order history setup and provide a lower bound ρ^{private} for ρ.

(III) $\rho^{\mathsf{ODTE}} \leq \rho^{\mathsf{private}} \leq \rho^{\mathsf{OSTE}}$:

First, to show that $\rho^{\mathsf{private}} \leq \rho^{\mathsf{OSTE}}$, notice that for any ρ such that there exists an opaque static trading equilibrium, then there exists a private history equilibrium where the dealer is forward-looking and imperfectly observing about whether the trade occurs. In fact, given an OSTE equilibrium, on the equilibrium path, both players choose the OSTE strategies and beliefs, off the equilibrium path, both players choose the ISTE strategies and beliefs. As shown previously in the proof, consistency follows with similar construction. There is no incentives for the trader to deviate, nor for both players if the signal to play off-path is revealed. It also follows immediately that the dealer has no incentives to enhance the price offer, or lower it in good times. Therefore, the only condition left to check is that dealer's IC constraint in bad times, i.e., condition (12.7). However, given the on-path strategies is an OSTE, we have $\Lambda^{b,\theta_{t-1}} \geq 0$. Offering a lower price to deter trade in bad times doesn't even provide a static gain, and enhances the probability of punishment in future periods. Hence, no dealer will do this in bad times and the condition is satisfied. We just show that whenever $\rho \geq \rho^{\mathsf{OSTE}}$, a private history equilibrium exists, implying that $\rho^{\mathsf{private}} \leq \rho^{\mathsf{OSTE}}$.

Second, to check that $\rho^{\mathsf{private}} \geq \rho^{\mathsf{ODTE}}$, we show that for any $\rho \geq \rho^{\mathsf{private}}$, there exists an ODTE. In fact, consider the equilibrium we constructed at the beginning of this proof. To see that this forms an ODTE, we only need to check dealer's IC constraint in bad times. That is,

$$(1-\delta)\Lambda^{b,\theta_{t-1}}+\delta[\alpha_{\theta_t}U^{\mathsf{ODTE}}(g,\theta_t)+(1-\alpha_{\theta_t})U^{\mathsf{ODTE}}(b,\theta_t)] \geq \delta\frac{\rho}{2}\alpha_b\sigma_g^2. \tag{12.8}$$

Now notice with the same strategies and beliefs, in ODTE, if the dealer follows the on-path pricing rule, the probability that a punishment is triggered is higher than that in private history equilibrium. Therefore, we have $U^{\mathsf{ODTE}}(\theta_t,\theta_{t-1}) \geq U^{\mathsf{private}}(\theta_t,\theta_{t-1})$. Dealer's IC constraint in bad times in private history equilibrium implies that

$$
\begin{aligned}
& (1-\delta)\Lambda^{b,\theta_{t-1}} + \delta f_{-1}[\alpha_b U^{\mathsf{private}}(g,b) + (1-\alpha_b)U^{\mathsf{private}}(b,b)] + \delta(1-f_{-1})\cdot\frac{\rho}{2}\alpha_b\sigma_g^2 \\
\geq\ & (1-\delta)\times 0 + \delta f_0[\alpha_b U^{\mathsf{private}}(g,b) + (1-\alpha_b)U^{\mathsf{private}}(b,b)] + \delta(1-f_0)\cdot\frac{\rho}{2}\alpha_b\sigma_g^2
\end{aligned}
$$

Thus, we can show that

$$
\begin{aligned}
& (1-\delta)\Lambda^{b,\theta_{t-1}} + \delta[\alpha_{\theta_t}U^{\mathsf{ODTE}}(g,\theta_t) + (1-\alpha_{\theta_t})U^{\mathsf{ODTE}}(b,\theta_t) - \frac{\rho}{2}\alpha_b\sigma_g^2] \\
\geq\ & (1-\delta)\Lambda^{b,\theta_{t-1}} + \delta[\alpha_{\theta_t}U^{\mathsf{private}}(g,\theta_t) + (1-\alpha_{\theta_t})U^{\mathsf{private}}(b,\theta_t) - \frac{\rho}{2}\alpha_b\sigma_g^2] \\
\geq\ & (1-\delta)\Lambda^{b,\theta_{t-1}} + \delta(f_{-1}-f_0)[\alpha_{\theta_t}U^{\mathsf{private}}(g,\theta_t) + (1-\alpha_{\theta_t})U^{\mathsf{private}}(b,\theta_t) - \frac{\rho}{2}\alpha_b\sigma_g^2] \\
\geq\ & 0,
\end{aligned}
$$

which is exactly dealer's bad times IC constraint for the existence of ODTE. Therefore, whenever $\rho \geq \rho^{\mathsf{private}}$, the equilibrium constructed at the beginning of this proof also consists of an ODTE, implying that $\rho^{\mathsf{ODTE}} \leq \rho^{\mathsf{private}}$.

(IV) comparison with the always-ISTE equilibrium

Similarly, as shown in the analysis of Section 1.7.3, the private history equilibrium generates strictly higher social surplus than ISTE. Moreover, the equilibrium we constructed in this proof allows the dealer to collect all the social surplus from the trade. Thus, such a PBE defeats (in the sense of higher ex-ante dealer's payoff in every period in equilibrium) and Pareto dominates the equilibrium where ISTE is always played.

∎

Proof of Proposition 1.6.3. We first show the second half of the proposition. If $f_{-1} - f_0 = 0$, then condition 12.7 is equivalent to $\Lambda^{b,\theta_{t-1}} \geq 0$, which is the sufficient and necessary condition for the existence of OSTE, and hence ρ^{private} coincides with ρ^{OSTE}. If $f_{-1} - f_0 = 1$, then condition 12.7 is equivalent to condition 12.4, which is the necessary and sufficient condition for ODTE to hold. Whence ρ^{private} coincides with ρ^{ODTE}.

We now prove the first half of the proposition. Fix f_0, an increase in f_{-1} implies that in all future periods, the punishment stage is triggered less frequently. As dealer's static on-path payoff is strictly greater than that off the path, an increase in f_{-1} thus increases her continuation payoff. Condition (12.7) is equivalent to

$$
\begin{aligned}
& (1-\delta)\Lambda^{b,\theta_{t-1}} + \\
& \delta(f_{-1}-f_0)\Big[\alpha_b(U^{\mathsf{private}}(g,b,f_{-1},f_0)-U^{\mathsf{off}}(g)) + (1-\alpha_b)(U^{\mathsf{private}}(b,b,f_{-1},f_0)-U^{\mathsf{off}}(b))\Big] \\
\geq\ & 0
\end{aligned}
$$

As f_{-1} increases and f_0 remains the same, $f_{-1} - f_0$ increases, and the deviation punishment $U^{\mathsf{private}}(\theta_t, b, f_{-1}, f_0) - U^{\mathsf{off}}(\theta_t)$ increases. Therefore, the LHS increases and the constraint becomes more slacked. The opaque equilibrium is therefore easier to sustain. ∎

Proof of Proposition 1.7.1. Denote the on-path and off-path continuation payoffs are $U^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}; q)$ and $U^{\mathsf{off}}(\theta_t, \theta_{t-1}; q)$ given the current and the last economic states are θ_t and θ_{t-1}, and the auditing intensity is q. Then along the equilibrium path, the probability that the dealer still stays in the game for the next period is $(1-q)$, and the discount factor is δ. It's as if along the path her discount factor is $\delta(1-q)$ rather than δ. Off the equilibrium path, she will pass the censorship even if being audited, and everything remains the same. That is, $\forall \theta_t, \theta_{t-1}$,

$$\begin{aligned} U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}; q) &= (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p_t^{*\mathsf{UI}}(\theta_{t-1})) \\ &\quad + \delta(1-q)[\alpha_{\theta_t} U_{t+1}^{\mathsf{ODTE}}(g, \theta_t) + (1-\alpha_{\theta_t}) U_{t+1}^{\mathsf{ODTE}}(b, \theta_t)], \\ U_t^{\mathsf{off}}(\theta_t) &= (1-\delta)\mathbf{1}(\theta_t = g) \cdot \frac{\rho}{2}\sigma_g^2 + \delta[\alpha_{\theta_t} U_{t+1}^{\mathsf{off}}(g) + (1-\alpha_{\theta_t}) U_{t+1}^{\mathsf{off}}(b)]. \end{aligned}$$

Similarly as the proof of Proposition 1.5.1, the sufficient and necessary condition for the existence of ODTE is dealer's IC constraint:

$$U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}; q) \geq (1-\delta) \times 0 + \delta[\alpha_{\theta_t} U_t^{\mathsf{off}}(g) + (1-\alpha_{\theta_t}) U_t^{\mathsf{off}}(b)], \forall \theta_t, \theta_{t-1}. \tag{12.9}$$

We first show the first part of the statement. We show that whenever an ODTE exists for a pair of parameters (ρ, q_2), then it also exists for a pair of parameters (ρ, q_1) if $q_1 < q_2$. To show this, we show that $U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}; q)$ decreases in q. In fact, fix a period t and economic states θ_t, in all future period $\tau \geq t+1$, the expected dealer's payoff from that period, conditional on economic state in period $\tau - 1$ is $\theta_{\tau-1}$ and the dealer is still in the game, is that

$$\begin{aligned} &\alpha_{\theta_{\tau-1}} \Lambda^{g, \theta_{\tau-1}} + (1-\alpha_{\theta_{\tau-1}}) \Lambda^{b, \theta_{\tau-1}} \\ &= \alpha_{\theta_{\tau-1}}(\varphi a_{\tau-1} + J_g - p^{*\mathsf{UI}}(\theta_{\tau-1})) + (1-\alpha_{\theta_{\tau-1}})(\varphi a_{\tau-1} + J_b - p^{*\mathsf{UI}}(\theta_{\tau-1})) \\ &= \frac{\rho}{2}[\alpha_{\theta_{\tau-1}}(1-\alpha_{\theta_{\tau-1}})(J_g - J_b)^2 + \alpha_{\theta_{\tau-1}}\sigma_g^2 + (1-\alpha_{\theta_{\tau-1}})\sigma_b^2] \\ &> 0, \quad \forall \tau \geq t+1, \forall \theta_{\tau-1}. \end{aligned}$$

Therefore, all future periods expected payoff is constant over q and is positive. The weights of all future periods, or in other words, the probability to stay in the game for a future period and collect this payoff, decrease in q. Therefore, the continuation payoff $U_t(\theta_t, \theta_{t-1}; q)$, as a weighted sum of all future expected payoffs, also decrease in q.

So a decrease in q enlarges the left hand side of (12.9) and relaxes dealer's IC constraints, making ODTE easier to sustain.

We next show that $P_0 = [\rho^{\mathsf{ODTE}}, \infty]$ and $P_1 \subsetneqq [\rho^{\mathsf{OSTE}}, \infty]$. The former one is obvious as when $q = 0$, the definition of $U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}; 0)$ coincides with that of $U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1})$ and the rest just follows the proof of Proposition 1.5.1. To see the latter one, notice that $U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}; 1) = (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p^{*\mathsf{UI}}(\theta_{t-1}))$.

From the proof of Proposition 1.5.1,

$$\begin{bmatrix} U^{\mathsf{off}}(g) \\ U^{\mathsf{off}}(b) \end{bmatrix} = (1-\delta)\begin{bmatrix} 1-\delta\alpha_g & -\delta(1-\alpha_g) \\ -\delta\alpha_b & 1-\delta(1-\alpha_b) \end{bmatrix}^{-1} \cdot \begin{bmatrix} \frac{\rho}{2}\sigma_g^2 \\ 0 \end{bmatrix}.$$

Hence,

$$\begin{aligned} \alpha_{\theta_t} U_{t+1}^{\mathsf{off}}(g) + (1-\alpha_{\theta_t}) U_{t+1}^{\mathsf{off}}(b) &= \begin{bmatrix} \alpha_{\theta_t} & 1-\alpha_{\theta_t} \end{bmatrix} \cdot \begin{bmatrix} U^{\mathsf{off}}(g) \\ U^{\mathsf{off}}(b) \end{bmatrix} \\ &= (1-\delta)\begin{bmatrix} \alpha_{\theta_t} & 1-\alpha_{\theta_t} \end{bmatrix} \cdot \begin{bmatrix} 1-\delta\alpha_g & -\delta(1-\alpha_g) \\ -\delta\alpha_b & 1-\delta(1-\alpha_b) \end{bmatrix}^{-1} \cdot \begin{bmatrix} \frac{\rho}{2}\sigma_g^2 \\ 0 \end{bmatrix} \\ &= \frac{((1-\delta)\alpha_{\theta_t} + \delta\alpha_{\theta_b})}{1-\delta(\alpha_g-\alpha_b)} \frac{\rho}{2}\sigma_g^2. \end{aligned}$$

Dealer's IC constraint in bad times becomes

$$(1-\delta)(\varphi a_{t-1} + J_b - p_t^{*\mathsf{UI}}(\theta_{t-1})) \geq \delta \frac{\alpha_b}{1-\delta(\alpha_g-\alpha_b)} \frac{\rho}{2}\sigma_g^2 > 0.$$

Therefore, whenever the above IC holds, the IC constraints for OSTE also holds and OSTE exists. Moreover, when $\rho = \rho^{\mathsf{ODTE}}$, the left hand side becomes 0 and dealer's IC in bad times fail to sustain. In other words, we have $P_1 \subsetneqq [\rho^{\mathsf{OSTE}}, \infty]$. ■

Proof of Proposition 1.7.2. In opaque equilibrium, along the equilibrium path in period t, traders learn no information from the price offers, therefore, his ex-post belief about the economic state is still his prior at the beginning of period t: $\xi_t = \alpha_{\theta_{t-1}}$. His ex-post utility after the trade becomes

$$p_t - \varphi a_{t-1} - \alpha_{\theta_{t-1}} J_g - (1-\alpha_{\theta_{t-1}}) J_b + \frac{\rho}{2}(\alpha_{\theta_{t-1}}(\alpha_{\theta_{t-1}}(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2)).$$

Dealer's ex-post utility becomes

$$\varphi a_{t-1} + J_{\theta_t} - p_t.$$

At the beginning of period t, the expectation of this ex-post utility becomes

$$\varphi a_{t-1} + \alpha_{\theta_{t-1}} J_g + (1-\alpha_{\theta_{t-1}}) J_b - p_t.$$

Therefore, in period t the expected ex-post surplus in opaque equilibrium, becomes

$$\begin{aligned} W^{\mathsf{OTE}} &= p_t - \varphi a_{t-1} - \alpha_{\theta_{t-1}} J_g - (1-\alpha_{\theta_{t-1}}) J_b + \frac{\rho}{2}(\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \\ &\quad \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2)) + \varphi a_{t-1} + \alpha_{\theta_{t-1}} J_g + (1-\alpha_{\theta_{t-1}}) J_b - p_t \\ &= \frac{\rho}{2}(\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2)). \end{aligned}$$

In informative equilibrium, along the equilibrium path in period t, if it is in bad times, then no trade happens and the ex-post surplus is 0. If it is in good times, then trader learns about θ_t after observing the price offer of good times. His ex-post utility becomes

$$p_t^g - \varphi a_{t-1} - J_g + \frac{\rho}{2}\sigma_g^2.$$

The dealer's ex-post utility becomes

$$\varphi a_{t-1} + J_g - p_t^g.$$

Thence, the social welfare in good times becomes $\frac{\rho}{2}\sigma_g^2$. At the beginning of time t, the expected ex-post utility becomes $\frac{\rho}{2}\alpha_{\theta_{t-1}}\sigma_g^2$.

∎

Proof of Proposition 1.8.1. (a) In our construction of OSTE, dealer will offer a price that makes an uninformed trader indifferent in both good times and bad times. From the proof of Proposition 1.4.1, the sufficient and necessary condition for the existence of OSTE is that

$$CE(a_t|\xi_t = \alpha_{\theta_{t-1}}) \le \varphi a_{t-1} + J_b.$$

If OSTE exists for utility function U, then we have

$$CE_U(a_t|\xi_t = \alpha_{\theta_{t-1}}) \le \varphi a_{t-1} + J_b.$$

Therefore,

$$CE_V(a_t|\xi_t = \alpha_{\theta_{t-1}}) \le CE_U(a_t|\xi_t = \alpha_{\theta_{t-1}}) \le \varphi a_{t-1} + J_b,$$

and OSTE also exists for utility function V.

(b) In our construction of ISTE, dealer will offer a price $CE(a_t|\xi_t = 1)$ in good times and offer a price low enough to decline trade in bad times. From the proof of Proposition 1.4.2, the sufficient and necessary condition for the existence of such kind of equilibrium becomes

$$\varphi a_{t-1} + J_b \le CE(a_t|\xi_t = 1).$$

If a utility function V exhibits more risk-aversion than U, then

$$CE_U(a_t|\xi_t = 1) \ge CE_V(a_t|\xi_t = 1).$$

If there exists an ISTE with trader's utility function V, then

$$CE_V(a_t|\xi_t = 1) \ge \varphi a_{t-1} + J_b.$$

Therefore,

$$CE_U(a_t|\xi_t = 1) \ge CE_V(a_t|\xi_t = 1) \ge \varphi a_{t-1} + J_b,$$

and ISTE exists if traders have utility function U.

(c) Similarly as the proof of Lemma 1.12.3.1, we only need to check that dealer's IC constraint in bad times is relaxed if the trader's utility function exhibits more risk-aversion. That is, one only needs to check that the deviation gain from the current period decreases and the punishment losses from all future periods increase. That is,

$$-[\varphi a_{t-1} + J_b - CE_V(a_t|\xi_t = \alpha_{\theta_{t-1}})] \leq -[\varphi a_{t-1} + J_b - CE_U(a_t|\xi_t = \alpha_{\theta_{t-1}})]$$

and for $\forall \tau > t, h^\tau$,

$$\begin{aligned} & [\varphi a_{\tau-1} + J_{\theta_\tau} - CE_V(a_\tau|\xi_\tau = \alpha_{\theta_{\tau-1}})] - [\varphi a_{\tau-1} + J_{\theta_\tau} - CE_V(a_\tau|\xi_\tau = 1)] \\ \geq \; & [\varphi a_{\tau-1} + J_{\theta_\tau} - CE_U(a_\tau|\xi_\tau = \alpha_{\theta_{\tau-1}})] - [\varphi a_{\tau-1} + J_{\theta_\tau} - CE_U(a_\tau|\xi_\tau = 1)] \end{aligned}$$

Equivalently,

$$CE_V(a_t|\xi_t = \alpha_{\theta_{t-1}}) \leq CE_U(a_t|\xi_t = \alpha_{\theta_{t-1}})$$

and for $\forall \tau > t$,

$$CE_V(a_\tau|\xi_\tau = 1) - CE_V(a_\tau|\xi_\tau = \alpha_{\theta_{\tau-1}}) \geq CE_U(a_\tau|\xi_\tau = 1) - CE_U(a_\tau|\xi_\tau = \alpha_{\theta_{\tau-1}})$$

These directly come from Definition 1.8.1.

(d) The proof follows the similar steps as in (c).

■

Proof of Proposition 1.8.2. In the proof of Proposition 1.6.2, replace $\Lambda^{\theta_t,\theta_{t-1}}$ with $\varphi a_{t-1} + J_{\theta_t} - CE_U(a_t|\xi_t = \alpha_{\theta_{t-1}})$, replace $\frac{\rho}{2}\sigma_g^2$ with $\varphi a_{t-1} + J_g - CE_U(a_t|\xi_t = 1)$, and the rest of the proof remains unchanged. ■

Proof of Proposition 1.8.3. In any opaque equilibrium, dealer's price offer reveals no extra information about the underlying economic state. Similar to the argument in OSTE when traders are potential sellers, an Bayesian trader will hold the same prior belief after observing the equilibrium price offer. Hence, along the equilibrium path, if trader does not accept the dealer's proposed price offer and $o_t = 0$, using (3.4), then his payoff at the end of period t becomes

$$\begin{aligned} & \mathbb{E}[-a_t|h^{t-1}, p_t] - \frac{\rho}{2}\text{Var}[-a_t|h^{t-1}, p_t] \\ = \; & -(\varphi a_{t-1} + \mathbb{E}[J_{\theta_t}|h^{t-1}]) - \frac{\rho}{2}\left[\text{Var}[\sigma_{\theta_t} z_t|h^{t-1}] + \text{Var}[J_{\theta_t}|h^{t-1}]\right] \\ \overset{(a)}{=} \; & -(\varphi a_{t-1} + J_b + \alpha_{\theta_{t-1}}(J_g - J_b)) \\ & -\frac{\rho}{2}\left[\alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2 + \alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2\right], \end{aligned}$$

where (a) follows from the fact that trader's on-path posterior coincides with the prior. Hence, trader's indifference condition implies he will accept the offer if and only if $-p_t$, his end-of-period wealth for accepting the offer, is larger than $\mathbb{E}[-a_t|h^{t-1},p_t]-\frac{\rho}{2}\text{Var}[-a_t|h^{t-1},p_t]$. That is, the highest price for the trader to accept is

$$\begin{aligned}p_t^{*\mathsf{UI}}(\theta_{t-1},\chi_t\!=\!-1) \;\equiv\; & \varphi a_{t-1}+J_b+\alpha_{\theta_{t-1}}(J_g-J_b)+\\ & \frac{\rho}{2}\left[\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2+\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2\right].\end{aligned}$$

Since trades occur in both good times and bad times, the dealer's IR constraints impose that $p_t^{*\mathsf{UI}}(\theta_{t-1},\chi_t=-1)-(\varphi a_{t-1}+J_g)\geq 0$ and $p_t^{*\mathsf{UI}}-(\varphi a_{t-1}+J_b)\geq 0$. Hence, dealer's IR in good times is binding and gives us

$$\rho\geq 2\frac{(1-\alpha_{\theta_{t-1}})(J_g-J_b)}{\alpha_{\theta_{t-1}}\sigma_g^2+(1-\alpha_{\theta_{t-1}})\sigma_b^2+\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g-J_b)^2},\forall\theta_{t-1}\in\{g,b\}$$

Since the right hand side is decreasing in $\alpha_{\theta_{t-1}}$, we have

$$\rho\geq\rho^{\mathsf{OSTE}}_{\chi_t=-1}\equiv 2\frac{(1-\alpha_{\min})(J_g-J_b)}{\alpha_{\min}\sigma_g^2+(1-\alpha_{\min})\sigma_b^2+\alpha_{\min}(1-\alpha_{\min})(J_g-J_b)^2},\quad \alpha_{\min}=\min\{\alpha_g,\alpha_b\}$$

∎

Proof of Proposition 1.8.4. At any ISTE, along the equilibrium path, dealer's price offer reveals all the information about the economic state. To avoid dealer to offer a good time ask offer in bad times, transaction can only occur in bad times and in good times traders always decline the offer.

We can still construct trader's off-path belief such that he believes the current period is always in bad times. Then in good times, we need to make sure if dealer deviates and offers a lower price, then either it is below $\varphi a_{t-1}+J_g$ so the dealer is experiencing a loss, or the trader will reject this price offer. Given his off path belief, then the highest price he can accept becomes

$$\begin{aligned}-p_t^b &= \mathbb{E}[-a_t|h^{t-1},p_t^b]-\frac{\rho}{2}\text{Var}[-a_t|h^{t-1},p_t^b]\\ p_t^b &= \varphi a_{t-1}+J_b+\frac{\rho}{2}\sigma_b^2.\end{aligned}$$

In other words, no price is above $\varphi a_{t-1}+J_g$ and below $\varphi a_{t-1}+J_b+\frac{\rho}{2}\sigma_b^2$ because otherwise the dealer in good times would rather offer a price in between. Equivalently, this requires

$$\varphi a_{t-1}+J_g\geq\varphi a_{t-1}+J_b+\frac{\rho}{2}\sigma_b^2$$

$$\rho\leq\rho^{\mathsf{ISTE}}_{\chi_t=-1}\equiv 2\frac{J_g-J_b}{\sigma_b^2}$$

One can check that $p_t^b = \min\{\varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2, \varphi a_{t-1} + J_g\}$ and $p_t^g > \varphi a_{t-1} + J_g + \frac{\rho}{2}\sigma_g^2$ satisfy dealer's IC constraint in bad times as well as dealer's and trader's IR constraints at all times. ■

Proof of Corollary 1.8.1. As shown, respectively, in Propositions 1.8.3 and 1.8.4

$$\rho_{\chi_t=-1}^{\mathsf{OSTE}} = 2\frac{(J_g - J_b)}{\alpha_{\min}(J_g - J_b)^2 + \frac{\alpha_{\min}}{1-\alpha_{\min}}\sigma_g^2 + \sigma_b^2},$$

and

$$\rho_{\chi_t=-1}^{\mathsf{ISTE}} = 2\frac{J_g - J_b}{\sigma_b^2}$$

By a little algebra it is immediate that

$$\frac{(J_g - J_b)}{(\alpha_{\min}(J_g - J_b)^2 + \frac{\alpha_{\min}}{1-\alpha_{\min}}\sigma_g^2 + \sigma_b^2} < \frac{J_g - J_b}{\sigma_b^2},$$

completing the proof. ■

Proof of Proposition 1.8.5. The proof closely maps the three steps of the proof of Proposition 1.5.1. That is, we first show the "if" part by showing that whenever there exists an ODTE for traders' risk preference coefficient ρ_1, then for any $\rho_2 > \rho_1$, there exists an ODTE for for traders' risk preference coefficient ρ_2. We then show the "only if" part by showing that for any condition belief the trader may have, when trader is a potential buyer, than the dealer's ex-post ISTE payoffs are no larger than the min-max static payoffs given that the trader best responses to her strategy. Finally, we show that $\rho_{\chi_t}^{\mathsf{ODTE}} \leq \rho_{\chi_t}^{\mathsf{OSTE}}$ by showing that whenever an OSTE exists, then there must exist an ODTE.

Let's take a close look of each of the three steps. First, denote $U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}, \chi_t)$ as the on-path continuation payoff when the current and last period economic states are θ_t and θ_{t-1} respectively and trader of the current period has a trading position of χ_t. Then these payoffs can be recursively defined as follows. For $\forall \theta_t, \theta_{t-1}, \chi_t$,

$$\begin{aligned} U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}, \chi_t) \quad = \quad & (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p_t^{*\mathsf{UI}}(\theta_{t-1}, \chi_t))\chi_t + \\ & \delta\beta\alpha_{\theta_t} U_{t+1}^{\mathsf{ODTE}}(g, \theta_t, 1) + \delta\beta(1-\alpha_{\theta_t})U_{t+1}^{\mathsf{ODTE}}(b, \theta_t, 1) + \\ & \delta(1-\beta)\alpha_{\theta_t} U_{t+1}^{\mathsf{ODTE}}(g, \theta_t, -1) + \delta(1-\beta)(1-\alpha_{\theta_t})U_{t+1}^{\mathsf{ODTE}}(b, \theta_t, -1). \end{aligned}$$

Denote $U_t^{\mathsf{off}}(\theta_t, \theta_{t-1}, \chi_t)$ as the off-path continuation payoff when the current economic state is θ_t and trader of the current period has a trading position of χ_t. Off the equilibrium path, For $\forall \theta_t, \theta_{t-1}, \chi_t$,

$$\begin{aligned} U_t^{\mathsf{off}}(\theta_t, \chi_t) \quad = \quad & (1-\delta)\mathbf{1}(\theta_t = g)\mathbf{1}(\chi_t = 1)\frac{\rho}{2}\sigma_g^2 + \mathbf{1}(\theta_t = b)\mathbf{1}(\chi_t = -1)\frac{\rho}{2}\sigma_b^2 \\ & \delta\beta\alpha_{\theta_t} U_{t+1}^{\mathsf{Off}}(g, 1) + \delta\beta(1-\alpha_{\theta_t})U_{t+1}^{\mathsf{off}}(b, 1) + \\ & \delta(1-\beta)\alpha_{\theta_t} U_{t+1}^{\mathsf{off}}(g, -1) + \delta(1-\beta)(1-\alpha_{\theta_t})U_{t+1}^{\mathsf{off}}(b, -1). \end{aligned}$$

Then if $\chi_t = 1$, a myopic dealer may want to deviate and decline trade in bad times if OSTE fails to exist. Then the binding constraint in this scenario is the dealer's IC in bad times:

$$\begin{aligned} U_t^{\mathsf{ODTE}}(b, \theta_{t-1}, 1) \quad \geq \quad & (1-\delta)\cdot 0 + \delta\beta\alpha_b U_{t+1}^{\mathsf{off}}(g, 1) + \delta\beta(1-\alpha_b)U_{t+1}^{\mathsf{off}}(b, 1) + \\ & \delta(1-\beta)\alpha_b U_{t+1}^{\mathsf{off}}(g, -1) + \delta(1-\beta)(1-\alpha_b)U_{t+1}^{\mathsf{off}}(b, -1). \end{aligned}$$

Then if $\chi_t = -1$, a myopic dealer may want to deviate and decline trade in good times if OSTE fails to exist. Then the binding constraint in this scenario is

$$\begin{aligned} U_t^{\mathsf{ODTE}}(g, \theta_{t-1}, -1) \quad \geq \quad & (1-\delta)\cdot 0 + \delta\beta\alpha_g U_{t+1}^{\mathsf{off}}(g, 1) + \delta\beta(1-\alpha_g)U_{t+1}^{\mathsf{off}}(b, 1) + \\ & \delta(1-\beta)\alpha_g U_{t+1}^{\mathsf{off}}(g, -1) + \delta(1-\beta)(1-\alpha_g)U_{t+1}^{\mathsf{off}}(b, -1). \end{aligned}$$

We now show that whenever the above two conditions hold for ρ_1, then for any $\rho_2 > \rho_1$, they still hold. A sufficient condition is that for any realization of $\{\theta_t\}$ and $\{\chi_t\}$, the current deviation gain decreases in ρ, whereas the future punishment loss in each period increases in ρ. Specifically, the current deviation gain is

$$-(\varphi a_{t-1} + J_{\theta_t} - p_t^{*\mathsf{UI}}(\theta_{t-1}, \chi_t))\chi_t \propto -\frac{\rho}{2}[\alpha_{\theta_t}(1-\alpha_{\theta_t})(J_g - J_b)^2 + \alpha_{\theta_t}\sigma_g^2 + (1-\alpha_{\theta_t})\sigma_b^2],$$

and decreases in ρ. For any realization of future period's economic state θ_τ and trading position χ_τ, the deviation punishment becomes

$$\varphi a_{\tau-1} + J_{\theta_\tau} - p_t^{*\mathsf{UI}}(\theta_{\tau-1}, 1) - \frac{\rho}{2}\sigma_g^2 \propto \frac{\rho(1-\alpha_{\theta_{\tau-1}})}{2}[\alpha_{\theta_{\tau-1}}(J_g - J_b)^2 + \sigma_b^2 - \sigma_g^2]$$

if $\chi_t = 1$ or

$$-\varphi a_{\tau-1} - J_{\theta_\tau} + p_t^{*\mathsf{UI}}(\theta_{\tau-1}, -1) - \frac{\rho}{2}\sigma_b^2 \propto \frac{\rho\alpha_{\theta_{\tau-1}}}{2}[(1-\alpha_{\theta_{\tau-1}})(J_g - J_b)^2 - \sigma_b^2 + \sigma_g^2]$$

if $\chi_t = -1$. The first one increases in ρ always; the second one increases in ρ given our parameter assumption.

Therefore, as ρ increases, the current economic gain from deviation decreases, whereas the future punishment caused by deviation increases, making deviation less profitable and the equilibrium easier to sustain.

Second, we prove the "only if" part by showing that when $\chi_t = -1$, the ISTE payoffs remain as minmax payoffs and consist of the harshest feasible punishment one can put on deviation. Formally speaking, similar to the case where $\chi_t = 1$, we will show that when

$\chi_t = -1$,[38]

$$\begin{aligned}
\min_{\xi(p;\cdot)} \max_{p^g, p^b} \min_{o_t \in \mathsf{BR}_{\xi(p;\cdot)}} (p^g - \varphi a_{t-1} - J_g) \cdot o_t(p^g) &= 0 \\
\min_{\xi(p;\cdot)} \max_{p^g, p^b} \min_{o_t \in \mathsf{BR}_{\xi(p;\cdot)}} (p^b - \varphi a_{t-1} - J_b) \cdot o_t(p^b) &= \frac{\rho}{2}\sigma_b^2
\end{aligned}$$

In fact, in good times, no matter what belief trader may have, dealer can always offer a high enough price to decline the trade and collect a zero payoff. If trader has a belief that the economic state is in bad times unless observing an offer of $\varphi a_{t-1} + J_g$, then he will only accept any offer below $\varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$ or exactly at $\varphi a_{t-1} + J_g$. Offering a price below $\varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$ leads to a negative payoff for the dealer; for price above $\varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$, either trader will decline it or trader will accept it but the dealer's payoff is exactly 0. In conclusion, in good times, given the trader's belief and his optimization order decision over that belief, the best the dealer can collect is 0.

In bad times, the dealer can always offer $\varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$, the trader will always accept it, and the dealer can collect a payoff of $\frac{\rho}{2}\sigma_b^2$. Now let the trader has the same belief that the current period is always in good times, then he will only accept offers weakly below $\varphi a_{t-1} + J_b + \frac{\rho}{2}\sigma_b^2$. Consequently, $\frac{\rho}{2}\sigma_b^2$ is dealer's best case ex-post payoff given this belief and trader best responding to it.

The rest of the "if only" proof then follows the logic of that in Proposition 1.5.1. Since the equilibrium we constructed has already employ the harshest punishment, it gives the lower bound for the traders' risk-aversion coefficient, ρ, for the existence of an opaque equilibrium.

Third, we show that $\rho_{\chi_t}^{\mathsf{ODTE}} \leq \rho_{\chi_t}^{\mathsf{OSTE}}$. Since along the equilibrium path, dealer always offer uninformative prices and off the equilibrium path, dealer has to offer differently in bad times and in good times and hence release her information about the economic states, the on-path continuation payoff $U_t^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}, \chi_t)$ is no less than the off-path continuation payoff $U_t^{\mathsf{off}}(\theta_t, \chi_t)$ due to this information rent. Consequently, if given the parameter values, an OSTE exists, then in the current period the deviation leads to a loss, in all future periods, it triggers the punishment and leads to lower continuation payoffs. As a result, the dealer will never have an incentive to deviate in the dynamic game and therefore, the constructed ODTE is valid. ∎

Proof of Proposition 1.8.6. Since the signal of past trade is independent of trader's trading positions. The analysis of the private history equilibrium follows exactly the same steps as those in the scenario where trader is a potential seller. The exact mapping between the case where $\chi_t = 1$ and the case $\chi_t = -1$ is illustrated in the proof of Proposition 1.8.5. ∎

Proof of Proposition 1.8.7. We first show the necessary condition of the existence for a fully pooling equilibrium and characterize $\rho^{pool.}$. We then prove that $\rho \geq \rho^{pool.}$ is sufficient for the existence of a fully pooling equilibrium by construction. Next, we analyze seller's strategy

[38]Recall that when $\chi_t = -1$ the trader is a buyer so the corresponding order o_t belongs to $\{0, 1\}$.

under a fully separating equilibrium. Finally, we prove the comparison relation between $\rho^{pool.}$ and $\rho^{sep.}$.

(i) **Necessity of $\rho \geq \rho^{pool.}$**

If the buyer offers an opaque price, then after observing it, the seller will still hold his prior belief that $\theta = g$ with probability α^*. To sustain such an opaque pricing strategy and guarantee that trade occurs, dealer's evaluation about the asset should be weakly higher than that of the seller (trader), in both good and bad times. Specifically, it is the $\theta = b$ case that binds and such a fully pooling equilibrium exists if

$$\begin{aligned} & \mathbb{E}(a(b, z_t)) \geq \mathbb{E}(a(\theta, z_t)|\alpha^*) - \frac{\rho_t}{2}\mathrm{Var}(a(\theta, z_t)|\alpha^*), \\ \Leftrightarrow \quad & \rho_t \geq \rho^{pool.} \equiv 2\frac{\mathbb{E}(a(\theta, z_t)|\alpha^*) - \mathbb{E}(a(b, z_t))}{\mathrm{Var}(a(\theta, z_t)|\alpha^*)}, \end{aligned}$$

which characterizes $\rho^{pool.}$.

(ii) **Sufficiency of $\rho \geq \rho^{pool.}$**

We construct the following strategies and beliefs and verify that they consist of a PBE given $\rho \geq \rho^{pool.}$. Specifically, in period t,

- the dealer offers a price $p_t^{opaque} \equiv \mathbb{E}(a(\theta, z_t)|\alpha^*) - \frac{\rho_t}{2}\mathrm{Var}(a(\theta, z_t)|\alpha^*)$;
- Conditional on observing p_t^{opaque}, the seller holds his prior belief that $\Pr(\theta = g) = \alpha^*$. Otherwise he believes that the economy is good for sure;
- the seller accepts a price offer if and only if it is weakly above his evaluation of the risky asset given the above belief.

It is easy to check that the seller's belief updating follows Bayes' rule whenever it applies and the seller is maximizing his expected utility given his belief. To see that the dealer has no incentive to deviate, let us consider her options. She can

- offer the equilibrium price p_t^{opaque} and earns non-negative payoff, guaranteed by condition $\rho \geq \rho^{pool.}$;
- offer a price lower than p_t^{opaque}; however doing so leads to the rejection by the seller, as he adjusts his belief to $\Pr(\theta = g) = 1$ and requires at least $\mathbb{E}(a(g, z_t)) - \frac{\rho_t}{2}\mathrm{Var}(a(g, z_t)) > p_t^{opaque}$;
- offer a price higher than p_t^{opaque}; but no matter how seller responds to it, such a deviation leads to lower payoff of the dealer.

Hence it is also suboptimal for the dealer to deviate.

(iii) **Dealer's Strategy under a Fully Separating Equilibrium**

First, in a fully separating equilibrium, trade can only occur when $\theta = g$, because otherwise in good times the dealer can offer a bad times price offer and convince the seller to sell his asset at a lower price. Such profitable deviation leads to the break of the equilibrium. Next, after observing the good times price offer, the seller will update his posterior to believe that $\theta = g$ for sure. Notice that this is the most optimistic belief he can ever hold and the seller is willing to accept any price offer weakly above

$$p_t^g = \mathbb{E}(a(g, z_t)) - \frac{\rho_t}{2}\text{Var}(a(g, z_t)).$$

Therefore, there is no reason for a rational dealer to offer something strictly higher. In conclusion, in a fully separating equilibrium, in good times the dealer will offer exactly p_t^g and in bad times offer something low enough to deter trade. Eventually, the only thing left to check is that dealer in bad times should not have an incentive to offer a good times price and induce a profitable trade. Thence, her incentive-compatibility (IC) constraint when $\theta = g$ implies that

$$\mathbb{E}(a(b, z_t)) \leq p_t^g = \mathbb{E}(a(g, z_t)) - \frac{\rho_t}{2}\text{Var}(a(g, z_t));$$
$$\Leftrightarrow \quad \rho_t \leq \rho^{sep.} \equiv 2\frac{\mathbb{E}(a(g, z_t)) - \mathbb{E}(a(b, z_t))}{\text{Var}(a(g, z_t))},$$

which characterizes $\rho^{sep.}$.

To see the sufficiency, construct seller's belief such that they believe $\theta = g$ unless they observe a price lower than their bad times evaluations. It is easy to check that dealer's pricing strategy above, along with this belief, forms a PBE.

(iv) $\rho^{pool.} \leq \rho^{sep.}$

To see that $\rho^{pool.} \leq \rho^{sep.}$, we observe that $\mathbb{E}(a(g, z_t)) > \mathbb{E}(a(\theta, z_t)|\alpha^*)$ and $\text{Var}(a(\theta, z_t)|\alpha^*) \geq \text{Var}(a(g, z_t))$. The inequality then follows immediately.

∎

Proof of Proposition 1.8.8. We show the proof in three steps. First, we construct a set of strategies and beliefs and show sufficient conditions for them to consist of a PBE. We characterize ρ_k from the binding constraint. Second, the necessity of $\rho_\ell \geq \rho_k$ comes from the fact that in the construction we have already used the harshest possible punishment. Finally, we show comparison relationships between different bounds.

(I) Characterize ρ_k

First, let us consider the following set of strategies and beliefs:

- if there is no deviation in previous periods, then
  - for $t = 1, ..., k$, the dealer offers $p_t = \mathbb{E}(a(\theta, z_t)|\alpha^*) - \frac{\rho_t}{2}\text{Var}(a(\theta, z_t)|\alpha^*)$;

– for $t = k+1, ..., N$, the dealer offers

$$p_t \begin{cases} = \mathbb{E}(a(\theta, z_t)|\alpha_t) - \frac{\rho_h}{2}\text{Var}(a(\theta, z_t)|\alpha_t) & \text{if } \rho_{t'} = \rho_h, t' = 1, ..., t \\ = \mathbb{E}(a(g, z_t)) - \frac{\rho_h}{2}\text{Var}(a(g, z_t)) & \text{if } \rho_t = \rho_h \text{ and } \rho_{t'} = \rho_\ell \\ & \text{for some } t' \in [k+1, t-1] \ ; \\ = \mathbb{E}(a(g, z_t)) - \frac{\rho_\ell}{2}\text{Var}(a(g, z_t)) & \text{if } \rho_t = \rho_\ell \text{ and } \theta = g \\ < \mathbb{E}(a(b, z_t)) - \frac{\rho_\ell}{2}\text{Var}(a(b, z_t)) & \text{if } \rho_t = \rho_\ell \text{ and } \theta = b \end{cases}$$

– after observing an on-path price offer, the seller updates his posterior based on past trade history as well as the current period price offer;

– sellers accept any offer weakly above their evaluation about the asset, given their posterior belief;

- if there is deviation observed in previous periods, then sellers will hold the belief that $\theta = g$ and the buyer offers

$$p_t \begin{cases} = \mathbb{E}(a(g, z_t)) - \frac{\rho_t}{2}\text{Var}(a(g, z_t)) & \text{if } \rho_t = \rho_h \text{ or } \theta = g \\ < \mathbb{E}(a(b, z_t)) - \frac{\rho_t}{2}\text{Var}(a(b, z_t)) & \text{if } \rho_t = \rho_\ell \text{ and } \theta = b \end{cases}; \tag{12.10}$$

  sellers update their beliefs whenever the Bayes' rule applies, or believe that $\theta = g$ otherwise; sellers accept any offer weakly above their evaluation about the asset, given the above constructed beliefs.

To check the validity of "k-pooling" equilibrium, notice that sellers' beliefs already follow the Bayes' rule and sellers' actions are sequential rational. The only condition to check is that dealer's behavior is sequential rational. We look at all kinds of histories.

Let us first focus on off-path histories. If the buyer is ever observed to deviate before, then by offering the price in equation (12.10), she can earn non-negative profits if the seller is high-type or $\theta = g$, and zero if the seller is low-type and $\theta = b$. Hence, when $\theta = g$ or there comes a high-type, the buyer has no incentive to offer $\mathbb{E}(a(b, z_t)) - \frac{\rho_t}{2}\text{Var}(a(b, z_t))$. Similarly as in the proof of Proposition 1.8.7, offering something higher will lower her current period payoff. If she offers something strictly lower than $\mathbb{E}(a(b, z_t)) - \frac{\rho_t}{2}\text{Var}(a(b, z_t))$, then Bayes' rule does not apply and the seller will believe that the economy is good. Hence, he will reject the trade and also lead to a loss of buyer's current period payoff. Hence, deviating will cause the loss of current payoff and no change in future payoffs. A rational buyer will not deviate off-path.

We then consider possible deviations during period $t = 1, ..., k$. Again, we rule out possible deviations such as offering a higher price. But given seller's beliefs, offering a lower price will lead to his rejection. In good times or if the seller is high-type, such deviations will result in a loss of current profit and trigger future punishment. Hence, only when facing a low-type seller in bad times, can the dealer have an incentive to decline trade and avoid current period loss. In other words, these cases are the only incentive compatibility (IC) constraints we need to check. Denote V_n as buyer's continuation payoff when $\theta = b$ if she keeps offering opaque prices in previous periods and there are n sellers coming. In these periods V_n follows the dynamics:

$$V_n = \gamma\pi(\rho_h, \alpha^*) + (1-\gamma)\pi(\rho_\ell, \alpha^*) + \delta V_{n-1},\ n = N-k+1, ..., N. \tag{12.11}$$

buyer's IC constraints are

$$\pi(\rho_\ell, \alpha^*) + \delta V_{n-1} \geq 0 + \sum_{i=1}^{n-1} \delta^i \gamma\pi(\rho_h, 1), \quad n = N-k+1, ..., N. \tag{12.12}$$

We now look at on-path history h^{t-1} where $t \geq k+1$. Let us look at this case by case.

(i) If no low-type seller comes between period $k+1$ and t (inclusive), then in period t, the dealer offers the uninformative price. Similarly as in the proof of Proposition 1.8.7, as $\rho_h > \rho^{pool.}$, deviation will cause a current loss and will trigger future punishment.

(ii) The second case is that at least one low-type seller comes between period $k+1$ and period $t-1$ (inclusive), but the current period comes a high-type seller. According to the construction, the buyer is supposed to offer $p_t = \mathbb{E}(a(g,z_t)) - \frac{\rho_h}{2}\text{Var}(a(g,z_t))$ and will make positive profit in the current period. Offering higher price will lower the current profit. Offering lower price will let the seller adjust his belief to $\text{Prob}(\theta = g) = 1$ and reject the offer. So all possible deviation leads to the loss of current profit and trigger future punishment.

(iii) The third case is when in the current period comes a low-type seller. Similarly, then given sellers' belief construction, the buyer can not buy the asset at a price lower than $\mathbb{E}(a(g,z_t)) - \frac{\rho_\ell}{2}\text{Var}(a(g,z_t))$ and the equilibrium strategy is already optimal.

Hence, condition (12.12) is a sufficient condition for the validity of the constructed PBE. Now let us show that $n = N-k+1$ is a binding one. We will prove this by induction. let us first finalize the dynamics of V_n: for $n = 1, ..., N-k$, the dynamics of V_n follows

$$V_1 = \gamma\pi(\rho_h, \alpha^*) \tag{12.13}$$

$$V_n = \gamma\pi(\rho_h, \alpha^*) + \gamma\delta V_{n-1} + (1-\gamma)\sum_{i=1}^{n-1} \delta^i \gamma\pi(\rho_h, 1). \tag{12.14}$$

We then introduce an auxiliary sequence $\{\underline{V}_n\}_{n=1}^N$, show that $V_n \geq \underline{V}_n, \forall n = 1, ..., N$ and explicitly characterize the value of this auxiliary sequence.

$$\begin{aligned}
\underline{V}_n &= V_n, \quad n = 1, ..., N-k \\
\underline{V}_n &= \gamma\pi(\rho_h, \alpha^*) + \gamma\delta\underline{V}_{n-1} + (1-\gamma)\sum_{i=1}^{n-1} \delta^i \gamma\pi(\rho_h, 1), \quad n = N-k+1, ..., N
\end{aligned}$$

Then, one can easily check by induction that for $n = 1, ..., N$,

$$\underline{V}_n = [1 + (\delta\gamma) + ... + (\delta\gamma)^{n-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] + (1 + \delta + ... + \delta^{n-1})\gamma\pi(\rho_h, 1).$$

Claim 1.12.1. If $\pi(\rho_\ell, \alpha^*) + \delta V_{N-k} = \pi(\rho_\ell, \alpha^*) + \delta \underline{V}_{N-k} \geq \sum_{i=1}^{N-k} \delta^i \gamma\pi(\rho_h, 1)$, then following claims are true.

- Claim $A_n : V_n \geq \underline{V}_n, \quad n = 1, ..., N;$

- Claim $B_n : \pi(\rho_\ell, \alpha^*) + \delta V_n \geq \pi(\rho_\ell, \alpha^*) + \delta \underline{V}_n \geq \sum_{i=1}^{n} \delta^i \gamma\pi(\rho_h, 1), \quad n = N - k, ..., N$

Proof of the Claim. We prove by induction. First, the statement B_{N-k} and statements $A_1, .., A_{N-k}$ are true. We then prove by induction the rest of claims are true.

First, suppose A_{n-1} and B_{n-1} are true, where $n > N - k$, then

$$\begin{aligned} V_n &= \gamma\pi(\rho_h, \alpha^*) + (1 - \gamma)\pi(\rho_\ell, \alpha^*) + \delta V_{n-1} \\ &\overset{(i)}{\geq} \gamma\pi(\rho_h, \alpha^*) + \gamma\delta V_{n-1} + (1 - \gamma) \sum_{i=1}^{n-1} \delta^i \gamma\pi(\rho_h, 1) \\ &\overset{(ii)}{\geq} \gamma\pi(\rho_h, \alpha^*) + \gamma\delta \underline{V}_{n-1} + (1 - \gamma) \sum_{i=1}^{n-1} \delta^i \gamma\pi(\rho_h, 1) \\ &\overset{(iii)}{=} \underline{V}_n, \end{aligned}$$

where inequality (i) is due to statement B_{n-1}, inequality (ii) comes from statement A_{n-1} and equality (iii) is the definition of $\underline{V}_n$. Hence, statement A_n is true.

Next suppose that statement A_n and B_{n-1} are true, $n > N - k$, then

$$\begin{aligned}
& \pi(\rho_\ell, \alpha^*) + \delta V_n - \sum_{i=1}^{n} \delta^i \gamma \pi(\rho_h, 1) \\
\overset{(iv)}{\geq} \quad & \pi(\rho_\ell, \alpha^*) + \delta \underline{V}_n - \sum_{i=1}^{n} \delta^i \gamma \pi(\rho_h, 1) \\
\overset{(v)}{=} \quad & \pi(\rho_\ell, \alpha^*) + \delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{n-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
& \quad +(\delta + ... + \delta^n)\gamma\pi(\rho_h, 1) - \sum_{i=1}^{n} \delta^i \gamma \pi(\rho_h, 1) \\
= \quad & \pi(\rho_\ell, \alpha^*) + \delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{n-2}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
& \quad +\delta(\delta\gamma)^{n-1} \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
\overset{(vi)}{=} \quad & \pi(\rho_\ell, \alpha^*) + \delta \underline{V}_{n-1} - \sum_{i=1}^{n-1} \delta^i \gamma \pi(\rho_h, 1) + \delta(\delta\gamma)^{n-1} \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
\overset{(vii)}{\geq} \quad & \delta(\delta\gamma)^{n-1} \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
\overset{(viii)}{\geq} \quad & 0,
\end{aligned}$$

where inequality (iv) is from statement A_n, equalities (v) and (vi) use the explicit form of $\underline{V}_n$, inequality (vii) is because of statement B_{n-1} and finally inequality (viii) comes from the monotonicity of $\pi(\cdot, \alpha)$. Thence, statement B_n is true.

Therefore, by induction, statements $A_1, ..., A_N$ and $B_{N-k}, ..., B_N$ are all true. ■

From the above claim, we can then conclude that the binding constraint is

$$\pi(\rho_\ell, \alpha^*) + \delta V_{N-k} \geq \sum_{i=1}^{N-k} \delta^i \gamma \pi(\rho_h, 1),$$

which gives a lower bound of ρ_ℓ, denoted as ρ_k:

$$\pi(\rho_k, \alpha^*) + \delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{N-k-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] = 0. \tag{12.15}$$

(II) The Necessity of $\rho_\ell \geq \rho_k$

To show the necessity of $\rho_\ell \geq \rho_k$, we only need to show that, in the construction of equilibrium, the harshest deviation punishment has already been imposed. In fact, we can show that the off-path payoffs consist of the min-max payoff for the dealer.

Claim 1.12.2. For any conditional belief $\xi(p_t; \cdot)$, the dealer's off-path payoffs consist of the min-max static payoffs given that the seller best responses to her strategy. Specifically, for any h^{t-1},

$$\min_{\xi(p_t;h^{t-1})} \max_{p_g,p_b} \min_{o_t\in \mathsf{BR}_{\xi(p_t;h^{t-1})}} (\mathbb{E}(a(g,z_t))-p_g)\cdot\Big(-o_t(p_g,\xi(\cdot),h^{t-1})\Big) = \frac{\rho_t}{2}\mathrm{Var}(a(g,z_t));$$

and

$$\min_{\xi(p_t;h^{t-1})} \max_{p_g,p_b} \min_{o_t\in \mathsf{BR}_{\xi(p_t;h^{t-1})}} (\mathbb{E}(a(b,z_t))-p_b)\cdot\Big(-o_t(p_b,\xi(\cdot),h^{t-1})\Big) = \max\{\mathbb{E}(a(b,z_t))-\mathbb{E}(a(g,z_t))+\frac{\rho_t}{2}\mathrm{Var}(a(g,z_t)),0\}.$$

Proof. We prove the claim case by case.

- Case I: when $\theta = b$ and the seller is low-type.

  First observe that the buyer can always offer a low enough price to decline the trade and collect an ex-post payoff of 0.

  Next, we show that for the following belief, she can do no better than that. Suppose the seller believes that the current period is in good times unless $\mathbb{E}(a(b,z_t))$ is offered, at which he believes it is in bad times, then he will only accept any offer above $\mathbb{E}(a(g,z_t)) - \frac{\rho_\ell}{2}\mathrm{Var}(a(g,z_t))$ or exactly at $\mathbb{E}(a(b,z_t))$. Offering the former leads to non-positive payoff, while offering the latter leads to zero payoff. Hence, conditional on this seller's belief and seller's best response function, dealer's highest ex-post payoff is 0. Combining with the first observation,

$$\min_{\xi(p_t;h^{t-1})} \max_{p_b} \min_{o_t\in \mathsf{BR}_{\xi(p_t;\cdot)}} (\mathbb{E}(a(b,z_t))-p_b)\cdot\Big(-o_t(p_b,\xi(\cdot),h^{t-1})\Big) = 0.$$

- Case II: when $\theta = g$.

  First observe that the dealer's ex-post payoff will be at least $\frac{\rho_t}{2}\mathrm{Var}(a(g,z_t))$. In fact, for any $\xi(p_t;\cdot)$,

$$\begin{aligned}\mathbb{E}(a(\theta,z_t)|\xi(\cdot|p)) - \frac{\rho_t}{2}\mathrm{Var}(a(\theta,z_t)|\xi(\cdot|p)) &= \mathbb{E}(a(b,z_t)) - \pi(\rho_t,\xi(\cdot|p))\\ &\le \mathbb{E}(a(b,z_t)) - \pi(\rho_t,1)\\ &= \mathbb{E}(a(g,z_t)) - \frac{\rho_t}{2}\mathrm{Var}(a(g,z_t)).\end{aligned}$$

  Therefore, the buyer can always offer a price $p_t^g = \mathbb{E}(a(g,z_t)) - \frac{\rho_t}{2}\mathrm{Var}(a(g,z_t))$ to confirm the order from the seller and guarantee an ex-post payoff of $\frac{\rho_t}{2}\mathrm{Var}(a(g,z_t))$ for sure. Next, this is the best she can get if the seller believes that the current period is always in good times. As only prices weakly above p_t^g will be accepted by the seller. So together we have,

$$\min_{\xi(p_t;h^{t-1})} \max_{p_g,p_b} \min_{o_t\in \mathsf{BR}_{\xi(p_t;h^{t-1})}} (\mathbb{E}(a(g,z_t))-p_g)\cdot\Big(-o_t(p_g,\xi(\cdot),h^{t-1})\Big) = \frac{\rho_t}{2}\mathrm{Var}(a(g,z_t)).$$

- Case III: when $\theta = b$ and the seller is high-type.

  First observe that the buyer can guarantee a non-negative payoff

$$\mathbb{E}(a(b, z_t)) - \mathbb{E}(a(g, z_t)) + \frac{\rho_t}{2}\mathrm{Var}(a(g, z_t))$$

  by offering a price at the level of seller's highest feasible evaluation p_t^g (defined above).

  Next, consider the following belief where the seller always believe $\theta = g$ and rejects any offer lower than p_t^g, then this is the lowest price the buyer can successfully acquire the asset. Thence, similarly we have

$$\begin{aligned}\min_{\xi(p_t;h^{t-1})} \max_{p_g, p_b} \min_{o_t \in \mathsf{BR}_{\xi(p_t;h^{t-1})}} & (\mathbb{E}(a(g, z_t)) - p_g) \cdot \Big(- o_t(p_g, \xi(\cdot), h^{t-1}) \Big) \\ &= \mathbb{E}(a(b, z_t)) - \mathbb{E}(a(g, z_t)) + \frac{\rho_t}{2}\mathrm{Var}(a(g, z_t)).\end{aligned}$$

■

With the availability of Claim 2, let us finish the proof of necessity. In any "k-opaque" equilibrium with consistency assumption, consider a following history h^{k-1}, when the dealer has always offered uniform price in period 1,2,..., $k-1$. We will show that buyer's IC constraint in period k when $\theta = b$ leads to $\rho_\ell \geq \rho_k$. If the dealer ever prices discriminatingly, then in any future period, she can get a profit of $\hat{\pi}_H$ from the high-type and $\hat{\pi}_L$ from the low-type seller. Then,

$$\hat{\pi}_H \geq \pi(\rho_h, 1), \hat{\pi}_L \geq 0.$$

As otherwise, the buyer can deviate to offer $p_t^g = \mathbb{E}(a(g, z_t)) - \frac{\rho_h}{2}\mathrm{Var}(a(g, z_t))$ for the high-type and decline the trade for low-type after revealing her private information. Then the dynamics of V_n in equilibrium σ is as follows.

$$V_n = \gamma\pi(\rho_h, \alpha^*) + \gamma\delta V_{n-1} + (1-\gamma)\sum_{i=1}^{n-1}\delta^i[\gamma\hat{\pi}_H + (1-\gamma)\hat{\pi}_L], \ n = 1, ..., N-k.$$

One can verify by induction that

$$\begin{aligned} V_n &= [1 + (\delta\gamma) + ... + (\delta\gamma)^{n-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\hat{\pi}_H - (1-\gamma)\hat{\pi}_L] \\ &\quad + (1 + \delta + ... + \delta^{n-1})[\gamma\hat{\pi}_H + (1-\gamma)\hat{\pi}_L], \quad n \leq N-k. \end{aligned}$$

Then dealer's IC constraint in period k when $\theta = b$ becomes

$$\begin{aligned} & \pi(\rho_\ell, \alpha^*) + \delta V_{N-k} \geq \sum_{i=1}^{N-k}\delta^i[\gamma\hat{\pi}_H + (1-\gamma)\hat{\pi}_L]; \\ \Rightarrow\ & \pi(\rho_\ell, \alpha^*) + \delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{N-k-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\hat{\pi}_H - (1-\gamma)\hat{\pi}_L] \geq 0; \\ \Rightarrow\ & \pi(\rho_\ell, \alpha^*) + \delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{N-k-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \geq 0, \end{aligned}$$

which exactly results in $\rho_\ell \geq \rho_k$.

(III) Comparisons between ρ_k

Last, we show that $\rho_N = \rho^{pool.}$ and $\rho_{k+1} \geq \rho_k$ for $k = 1, ..., N-1$. The former is trivial as when $k = N$, the characterization equations for ρ_N and ρ^* coincides with each other. To see the latter, notice that

$$\begin{aligned}
\pi(\rho_{k+1}, \alpha^*) &= -\delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{N-k-2}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
&\geq -\delta[1 + (\delta\gamma) + ... + (\delta\gamma)^{N-k-1}] \cdot [\gamma\pi(\rho_h, \alpha^*) - \gamma\pi(\rho_h, 1)] \\
&= \pi(\rho_k, \alpha^*).
\end{aligned}$$

Therefore, $\rho_{k+1} \geq \rho_k$ comes from the monotonicity of $\pi(\cdot, \alpha)$.

∎

Proof of Proposition 1.8.9. The proof aligns closely with that of Proposition 1.8.8. We adopt the same construction and present sufficient conditions for it to form a PBE. We then show the binding constraint and characterize $\hat{\rho}_k$. Next, the necessity follows the same argument as in Claim 1.12.2 in the proof of Proposition 1.8.8, that is, we already implemented the harshest deviation punishment on buyer's deviations. Finally, we show comparison relationships between different bounds.

The proof of "if part" is same as that in Proposition 1.8.8 with only one exception. After period k, for the history where trade has never been terminated so far (so the buyer has not revealed her private information about θ), sellers posterior beliefs will no longer be α^*. Dealer' rule implies that,

$$\begin{aligned}
\alpha_n &= \alpha^*, n = 1, ..., k+1; \\
\alpha_n &= \frac{\alpha_{n-1}}{\alpha_{n+1} + (1-\alpha_{n-1})\gamma}, \quad n = k+2, ..., N.
\end{aligned} \tag{12.16}$$

The intuition is as follows. According to our construction, after period k, trade occurs under two circumstances. The first case is in good times, when trade always occurs, no matter whether there comes a low-type or high-type seller. The second case is in bad times and when there comes a high-type seller. Due to this probabilistic asymmetry, observing trade occurs, even without any further knowledge about trade details such as transaction price, is informative for future sellers, who will Bayesian update her posterior belief. Since it is more likely for the trade to happen in good times than in bad times, future sellers will tend to have more optimistic beliefs about θ if they keep observing the trade.

Therefore, $\{\alpha_n\}_n$ is a monotonically increasing sequence and $\pi(\rho_h, \alpha_n)$ monotonically decreasing in n.

The dynamics of V_n now becomes

$$\begin{aligned}
V_1 &= \gamma\pi(\rho_h, \alpha_N) \\
V_n &= \gamma\pi(\rho_h, \alpha_{N+1-n}) + \gamma\delta V_{n-1} + (1-\gamma)\sum_{i=1}^{n-1}\delta^i\gamma\pi(\rho_h, 1), \quad n = 2, ..., N-k \\
V_n &= \gamma\pi(\rho_h, \alpha_{N+1-n}) + (1-\gamma)\pi(\rho_\ell, \alpha_{N+1-n}) + \delta V_{n-1} \\
&= \gamma\pi(\rho_h, \alpha^*) + (1-\gamma)\pi(\rho_\ell, \alpha^*) + \delta V_{n-1}, \ n = N-k+1, ..., N
\end{aligned}$$

and the auxiliary sequence $\{\underline{V}_n\}_{n=1}^N$ defined as below,

$$\begin{aligned}
\underline{V}_1 &= \gamma\pi(\rho_h, \alpha_N) \\
\underline{V}_n &= \gamma\pi(\rho_h, \alpha_{N+1-n}) + \gamma\delta\underline{V}_{n-1} + (1-\gamma)\sum_{i=1}^{n-1}\delta^i\gamma\pi(\rho_h, 1), \quad n = 2, ..., N.
\end{aligned}$$

Notice that $\underline{V}_n$ takes the same value as those in Proposition 1.8.8.

We will show that Claim 1.12.1 still holds.

Proof of Claim 1.12.1. Here we check that inductive steps still hold. It is obvious that $A_1, ..., A_{N-k}$ and B_{N-k} holds by construction and the assumption of the theorem.

(I) From A_{n-1} and B_{n-1} to A_n, where $n > N-k$,

$$\begin{aligned}
V_n &= \gamma\pi(\rho_h, \alpha^*) + (1-\gamma)\pi(\rho_\ell, \alpha^*) + \delta V_{n-1} \\
&\overset{(i)}{\geq} \gamma\pi(\rho_h, \alpha^*) + \gamma\delta V_{n-1} + (1-\gamma)\sum_{i=1}^{n-1}\delta^i\gamma\pi(\rho_h, 1) \\
&\overset{(ii)}{\geq} \gamma\pi(\rho_h, \alpha^*) + \gamma\delta\underline{V}_{n-1} + (1-\gamma)\sum_{i=1}^{n-1}\delta^i\gamma\pi(\rho_h, 1) \\
&\overset{(iii)}{=} \underline{V}_n,
\end{aligned}$$

Again, inequality (i) is from B_{n-1} and inequality (ii) is from A_{n-1}. Since $\underline{V}_n$ has the same value as that in the proof of Proposition 1.8.8, equality (iii) still holds. Hence, statement A_n is true.

(II) From A_n and B_{n-1} to B_n, $n > N-k$:

One can show by induction that

$$
\begin{aligned}
\underline{V}_n = \; & \sum_{i=0}^{n-1}(\delta\gamma)^i[\gamma\pi(\rho_h,\alpha_{N+1-n+i})-\gamma\pi(\rho_h,1)]+(1+...+\delta^{n-1})\gamma\pi(\rho_h,1) \\
\Longrightarrow \quad & \pi(\rho_\ell,\alpha^*)+\delta V_n-\sum_{i=1}^{n}\delta^i\gamma\pi(\rho_h,1) \\
\overset{(iv)}{\geq} \quad & \pi(\rho_\ell,\alpha^*)+\delta\underline{V}_n-\sum_{i=1}^{n}\delta^i\gamma\pi(\rho_h,1) \\
\overset{(v)}{=} \quad & \pi(\rho_\ell,\alpha^*)+\delta\sum_{i=0}^{n-1}(\delta\gamma)^i[\gamma\pi(\rho_h,\alpha_{N+1-n+i})-\gamma\pi(\rho_h,1)] \\
& +(\delta+...+\delta^n)\gamma\pi(\rho_h,1)-\sum_{i=1}^{n}\delta^i\gamma\pi(\rho_h,1) \\
= \quad & \pi(\rho_\ell,\alpha^*)+\delta\sum_{i=0}^{n-2}(\delta\gamma)^i[\gamma\pi(\rho_h,\alpha_{N+1-n+i})-\gamma\pi(\rho_h,1)]+\delta(\delta\gamma)^{n-1}\cdot[\gamma\pi(\rho_h,\alpha_N)-\gamma\pi(\rho_h,1)]
\end{aligned}
$$

$$
\begin{aligned}
\overset{(vi)}{\geq} \quad & \pi(\rho_\ell,\alpha^*)+\delta\sum_{i=0}^{n-2}(\delta\gamma)^i[\gamma\pi(\rho_h,\alpha_{N+1-(n-1)+i})-\gamma\pi(\rho_h,1)] \\
& +\delta(\delta\gamma)^{n-1}\cdot[\gamma\pi(\rho_h,\alpha_N)-\gamma\pi(\rho_h,1)] \\
\overset{(vii)}{=} \quad & \pi(\rho_\ell,\alpha^*)+\delta\underline{V}_{n-1}-\sum_{i=1}^{n-1}\delta^i\gamma\pi(\rho_h,1)+\delta(\delta\gamma)^{n-1}\cdot[\gamma\pi(\rho_h,\alpha_N)-\gamma\pi(\rho_h,1)] \\
\overset{(viii)}{\geq} \quad & \delta(\delta\gamma)^{n-1}\cdot[\gamma\pi(\rho_h,\alpha_N)-\gamma\pi(\rho_h,1)] \\
\overset{(ix)}{\geq} \quad & 0
\end{aligned}
$$

where inequality (iv) still comes from statement A_n and inequality $(viii)$ is because of statement B_{n-1}. Equalities (v) and (vii) are due to the explicit form of $\underline{V}_n$. Inequalities (vi) and (ix) are due to the monotonicity of $\pi(\rho_h,\cdot)$ and $\{\alpha_n\}_n$. ∎

The proof of the necessity then follows the same step as those in Proposition 1.8.8. In fact, one can easily show that the only deviation that provides current gain is to decline the trade with a low-type seller in bad times. For this reason, order information itself is a sufficient statistics for buyer's deviation. In other words, by monitoring whether there is a trade or not in a previous period, future sellers can identify a previous deviation and Claim 1.12.2 guarantees that the harshest punishment is triggered in our construction.

Finally we will prove that $\hat{\rho}_k \geq \rho_k$. Denote V_n^{public} as dealer's bad times continuous payoff under public history, and she offers opaque prices in all previous periods with n sellers left. That is, these are the V_n we derived in the proof of Proposition 1.8.8. As $\{\alpha_n\} \geq \alpha^*$ and $\pi(\rho_h,\cdot)$ decreases in α, it can be easily shown by induction that $V_n^{\mathsf{public}} \geq V_n$. From the proof of Proposition 1.8.8 and Proposition 1.8.9, we learn that the binding IC constraints are at

$n = N - k$ and thence $\hat{\rho}_k$ and ρ_k are characterized by the following conditions

$$\begin{aligned} \pi(\rho_k, \alpha^*) + \delta V^{\mathsf{public}}_{N-k} &= \sum_{i=1}^{N-k} \delta^i \gamma \pi(\rho_h, 1), \\ \pi(\hat{\rho}_k, \alpha^*) + \delta V_{N-k} &= \sum_{i=1}^{N-k} \delta^i \gamma \pi(\rho_h, 1). \end{aligned} \tag{12.17}$$

Then $V^{\mathsf{public}}_{N-k} \geq V_{N-k}$ implies that $\hat{\rho}_k \geq \rho_k$. The intuition is as follows. Post-price transparency disables sellers' posterior updating and increases buyer's on-path payoff. This makes IC constraints more slack and opaque pricing strategies easier to sustain. ■

Proof of Proposition 1.8.10. We prove each bullet point one by one.

I: private history

To see this, we only need to show that the opaque price $\hat{p}^{\mathsf{opaque}}_t$ in equilibrium is weakly lower than buyer's worst case evaluation $\mathbb{E}(a(b, z_t))$. Notice that the opaque price now is uniquely determined by the bargaining process. Specifically, when $\theta = b$, $\hat{p}^{\mathsf{opaque}}_t$ divides the surplus from the trade between two parties by a ratio of $\psi_k : (1 - \psi_k)$. Thence, if and only if dealer's evaluation of the risky asset is weakly higher than that of the seller (trader), the dealer can collect non-negative profit from the transaction and afford not revealing her private information. As a result, the IC constraint stays the same and we have exactly same cutoff $\rho^{pool.}$ for the opaque (pooling) pricing to be sustained.

II: public history and order history

Without loss of generality, we can focus on a case with a more general notations of $\{\alpha_n\}$ where $\alpha_1 = ... = \alpha_{k+1} = \alpha^*$. To see that statements are true for the public history case, we can then restrict $\alpha_n = \alpha^*$ for all $n = 1, ..., N$. Similarly to see that statements hold under the order history, we can restrict that $\alpha_1 = ... = \alpha_{k+1} = \alpha^*$ and $\{\alpha_n\}$ follows the evolving rule in (12.16) for $n = k + 2, ..., N$.

First of all, we adopt the same construction as that in the proof of Proposition 1.8.8. Similarly, we now characterize dealer's IC constraint where she is willing to hide her private information. We adopt similar notations and let $V^{\mathsf{bargaining}}_n$ denotes her continuation payoff when $\theta = b$ and when she has not revealed her private information (either by pricing differentiately under the public history or declining trade under the order history). The dynamics of $V^{\mathsf{bargaining}}_n$ is as follows.

$$\begin{aligned} V^{Barg.}_1 &= \psi_{N-1} \gamma \pi(\rho_h, \alpha_N) \\ V^{Barg.}_n &= \gamma[\psi_{N-n} \pi(\rho_h, \alpha_{N+1-n}) + \delta V^{Barg.}_{n-1}] + (1 - \gamma) \sum_{i=1}^{n-1} \delta^i \psi_{N-n+i} \gamma \pi(\rho_h, 1), \quad n = 2, ..., N - k. \\ V^{Barg.}_n &= \psi_{N-n}[\gamma \pi(\rho_h, \alpha_{N+1-n}) + (1 - \gamma) \pi(\rho_\ell, \alpha_{N+1-n})] + \delta V^{Barg.}_{n-1} \\ &= \psi_{N-n}[\gamma \pi(\rho_h, \alpha^*) + (1 - \gamma) \pi(\rho_\ell, \alpha^*)] + \delta V^{Barg.}_{n-1}, \quad n = N - k + 1, ..., N. \end{aligned}$$

Dealer's IC constraints, which are also sufficient conditions for the constructed strategies and beliefs to form a PBE, are

$$\psi_{N-n}\pi(\rho_\ell,\alpha^*)+\delta V_{n-1}^{Barg.}\geq 0+\sum_{i=1}^{n-1}\delta^i\cdot\psi_{N-n+i}\gamma\pi(\rho_h,1),\quad n=N-k+1,...,N. \tag{12.18}$$

Each inequality in condition (12.18) provides a lower bound of ρ_ℓ. Then $\rho_k^{Barg.}$ and $\hat{\rho}_k^{Barg.}$ are characterized by the greatest lower bound.

The necessity proof follows exactly from the fact that we have already implemented the harshest deviation punishment possible (min-max payoffs) in our construction. The comparison between $\rho_k^{Barg.}$ and $\hat{\rho}_k^{Barg.}$ also follows the similarly step as in the proof of Proposition 1.8.9. In fact, under the order history sellers' (traders') posterior belief start increasing after period $k+1$, implying less continuation payoff $V_n^{Barg.}$ than under the public history. Thence any lower bound in equation (12.18) is weakly greater under the order history than that under the public history. As a result, the maximum of those lower bounds under the order history is also weakly higher than that under the public history, which is exactly $\hat{\rho}_k^{Barg.}\geq\rho_k^{Barg.}$.

It is then obvious that $\rho_k^{Barg.}$ and $\hat{\rho}_k^{Barg.}$ are no greater than $\rho^{pool.}$. To see this, notice that if $\rho_\ell\geq\rho^{pooling}$, then $\pi(\rho_\ell,\alpha^*)\geq 0$, and hence inequalities (12.18) are all satisfied. In other words, when $\rho_\ell\geq\rho^{pool.}$, there exists a fully pooling equilibrium and the opaque pricing strategies can always be sustained under the public history or the order history.

Finally, we want to show that $\rho_k^{Barg.}\geq\rho_k$ and $\hat{\rho}_k^{Barg.}\geq\hat{\rho}_k$. In fact, we can show that if

$$P_k\equiv\psi_{N-k}\pi(\rho_\ell,\alpha^*)+\delta V_{k-1}^{Barg.}-\sum_{i=1}^{k-1}\delta^i\cdot\psi_{N-k+i}\gamma\pi(\rho_h,1)\geq 0,$$

then

$$Q_k\equiv\pi(\rho_\ell,\alpha^*)+\delta V_{k-1}-\sum_{i=1}^{k-1}\delta^i\cdot\gamma\pi(\rho_h,1)\geq 0.$$

We prove this by induction. One can prove by induction that for $n=2,...,k$

$$V_n^{Barg.}=\sum_{i=0}^{n-1}(\delta\gamma)^i\psi_{N-n+i}[\gamma\pi(\rho_h,\alpha_{N+1-n+i})-\gamma\pi(\rho_h,1)]+\sum_{i=0}^{n-1}\delta^i\psi_{N-n+i}\gamma\pi(\rho_h,1)$$

Then we have

$$\begin{aligned}
Q_k &= \pi(\rho_\ell, \alpha^*) + \delta V_{k-1} - \sum_{i=1}^{k-1} \delta^i \cdot \gamma\pi(\rho_h, 1) \\
&= \pi(\rho_\ell, \alpha^*) + \delta \sum_{i=0}^{k-1} (\delta\gamma)^i [\gamma\pi(\rho_h, \alpha_{N+1-k+i}) - \gamma\pi(\rho_h, 1)] \\
&\geq \pi(\rho_\ell, \alpha^*) + \delta \sum_{i=0}^{k-1} (\delta\gamma)^i \frac{\psi_{N-k+i}}{\psi_{N-k}} [\gamma\pi(\rho_h, \alpha_{N+1-k+i}) - \gamma\pi(\rho_h, 1)] \\
&= \pi(\rho_\ell, \alpha^*) + \frac{1}{\psi_{N-k}} \delta V_{k-1}^{Barg.} - \frac{1}{\psi_{N-k}} \sum_{i=1}^{k-1} \delta^i \cdot \psi_{N-k+i} \gamma\pi(\rho_h, 1) \\
&= \frac{1}{\psi_{N-k}} P_k \geq 0,
\end{aligned}$$

which completes the proof.

■

Proof of Proposition 1.8.11. We first characterize the restriction of q for the constructed equilibrium to hold. We then prove the welfare result.

(I) if part

First, it is obvious that the constructed belief follows the Bayes' rule whenever it applies, and hence consistency is satisfied. Next, we check that dealer does not want to deviate at all times. Given $\rho \leq \rho^{\mathsf{ISTE}}$, we have $\varphi a_{t-1} + J_b \leq \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. In good times, by offering $\varphi a_{t-1} + J_b$, the dealer is making a positive profit, hence it does not make sense to lower the price and decline the offer. It also does not make sense to increase the price and earn less. In bad times, offering $\varphi a_{t-1} + J_b$ gives the dealer zero profit, while declining the trade also gives her zero. Hence she is indifferent. Plus offering a price weakly above $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ leads to non-positive profit and is always sub-optimal. Finally the only thing left to check is trader's IR constraint:

$$\varphi a_{t-1} + \hat{\alpha} J_g + (1-\hat{\alpha}) J_b - \frac{\rho}{2}[\hat{\alpha}(1-\hat{\alpha})(J_g - J_b)^2 + \hat{\alpha}\sigma_g^2 + (1-\hat{\alpha})\sigma_b^2] \leq \varphi a_{t-1} + J_b, \quad (12.19)$$

where $\hat{\alpha} = \frac{\alpha_{\theta_{t-1}}}{\alpha_{\theta_{t-1}} + (1-\alpha_{\theta_{t-1}})q}$ for some $q \in (0,1)$.

Equivalently, trader's IR holds if and only if there exists $\hat{\alpha} \in (\alpha_{\theta_{t-1}}, 1)$ such that the inequality 12.19 holds:

$$h(\hat{\alpha}) \equiv \hat{\alpha}(J_g - J_b) - \frac{\rho}{2}[\hat{\alpha}(1-\hat{\alpha})(J_g - J_b)^2 + \hat{\alpha}\sigma_g^2 + (1-\hat{\alpha})\sigma_b^2] \leq 0. \quad (12.20)$$

$h(\cdot)$ is a quadratic function and $h(1) = J_g - J_b - \frac{\rho}{2}\sigma_g^2 \geq 0$ for $\rho \leq \rho^{\mathsf{ISTE}}$. Whence, there exists a $\hat{\alpha} \in [\alpha_{\theta_{t-1}}, 1]$ such that $h(\hat{\alpha}) \leq 0$, if and only if

$$h(\alpha_{\theta_{t-1}}) < 0$$

$$\Leftrightarrow \quad \rho > \frac{2\alpha_{\theta_{t-1}}(J_g - J_b)}{\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2}, \quad \forall \theta_{t-1} = g, b$$

$$\Leftrightarrow \quad \rho > \rho^{\mathsf{OSTE}}.$$

(II) only if part

We just showed that $\rho > \rho^{\mathsf{OSTE}}$ is necessary for the existence of semi-opaque equilibria. To see that ρ cannot exceed ρ^{ISTE}, we prove by contradiction. Suppose $\rho > \rho^{\mathsf{ISTE}}$, then $\varphi a_{t-1} + J_b > \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$, then in good times the dealer can deviate and offer a price at $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. Then the dealer will deviate to offer $\varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$. By doing so, the trader will accept the offer and the dealer earns more. As a result, the constructed equilibrium fails. In other words, we must have $\rho \in (\rho^{\mathsf{OSTE}}, \rho^{\mathsf{ISTE}}]$.

(III) welfare result

The social surplus in this semi-opaque static trading equilibrium in period t is

$$\begin{aligned} W^{\mathsf{semi}} &= \alpha_{\theta_{t-1}}[\varphi a_{t-1} + J_g - \varphi a_{t-1} - \hat{\alpha} J_g - (1-\hat{\alpha})J_b + \frac{\rho}{2}(\hat{\alpha}(1-\hat{\alpha})(J_g - J_b)^2 + \hat{\alpha}\sigma_g^2 + (1-\hat{\alpha})\sigma_b^2)] \\ &\quad + (1-\alpha_{\theta_{t-1}})q[\varphi a_{t-1} + J_b - \varphi a_{t-1} - \hat{\alpha} J_g - (1-\hat{\alpha})J_b + \frac{\rho}{2}(\hat{\alpha}(1-\hat{\alpha})(J_g - J_b)^2 + \hat{\alpha}\sigma_g^2 \\ &\quad + (1-\hat{\alpha})\sigma_b^2)] + (1-\alpha_{\theta_{t-1}})(1-q)\cdot 0 \\ &= \frac{\rho}{2}[\frac{\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})q}{(1-\alpha_{\theta_{t-1}})q + \alpha_{\theta_{t-1}}}(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})q\sigma_b^2] \end{aligned}$$

It then follows immediately that W^{semi} increases in the trading probability in bad times q. Therefore, we have:

$$W^{\mathsf{semi}} > W^{\mathsf{ISTE}} = \frac{\rho}{2}\alpha_{\theta_{t-1}}\sigma_g^2$$

and

$$W^{\mathsf{semi}} < W^{\mathsf{OSTE}} = \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2].$$

As $q = 0$ corresponds to the ISTE and $q = 1$ corresponds to an OSTE. ■

Proof of Proposition 1.11.1. We construct the following set of strategies and beliefs. At an on-path history h^t, dealer will offer a price $p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ in good times and a price p_t^b in bad times. Suppose $\varphi a_{t-1} + J_b - p_t^b$ is a constant. Therefore, at observing the price $p^{\tilde{\theta}_t}$, traders will update their posterior and believe the current economic state is $\tilde{\theta}_t$. If dealer deviates and offer something else, trader will believe that the current period is in good times. Since all previous offers and economic states are publicly observed, the dealer and all future traders can determine whether there is deviation or not. If there is at least one deviation in the previous periods, then the game moves into the punishment stage and both players choose the ISTE strategies and form the corresponding beliefs.

We now check all individual rationality constraints:

$$\begin{array}{rcll} p_t^g & \geq & \varphi a_{t-1} + J_g - \dfrac{\rho}{2}\sigma_g^2 & \text{trader's IR in good times} \\ p_t^b & \geq & \varphi a_{t-1} + J_b - \dfrac{\rho}{2}\sigma_b^2 & \text{trader's IR in bad times} \\ p_t^g & \leq & \varphi a_{t-1} + J_g & \text{dealer's IR in good times} \\ p_t^b & \leq & \varphi a_{t-1} + J_b & \text{dealer's IR in bad times} \end{array}$$

$p_t^g = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$ implies that the IR constraints in good times are satisfied. The IR constraints in bad times are equivalent to $\varphi a_{t-1} + J_b - \frac{\rho}{2}\sigma_b^2 \leq p_t^b \leq \varphi a_{t-1} + J_b$. Denote π^{θ_t} as dealer's static payoff at period t if she follows the on-path strategy. Then in our construction $\pi^g = \frac{\rho}{2}\sigma_g^2$ and IR constraints are equivalent to $\pi^b \in [0, \frac{\rho}{2}\sigma_b^2]$.

We now look at dealer's IC constraints. In good times, if dealer deviates to offer a price below p_t^g other than p_t^b, then according to our belief construction, trader will interpret this price as a signal of good times and will reject this offer. By deviating dealer loses current period profit and triggers the punishment. If dealer increases her price offer, then she loses part of her current period profit and triggers the punishment. Clearly both kinds of strategies are dominated by p_t^g. The only deviation left to check is to pretend now is in bad times and offer p_t^b. If she offers p_t^b, then the trader will believe the current economic state is bad and accept this offer. Dealer then collects a static profit of $\varphi a_{t-1} + J_g - p_t^b = \varphi a_{t-1} + J_b - p_t^b + J_g - J_b = \pi^b + J_g - J_b$. In all future periods, the punishment is triggered.

$$U_t^{\mathsf{IDTE}}(g) \geq \underbrace{(1-\delta)(\pi^b + J_g - J_b)}_{\substack{\text{dealer's static payoff after} \\ \text{pretending it's in bad times}}} + \underbrace{\delta[\alpha_g U_t^{\mathsf{off}}(g) + (1-\alpha_g)U_t^{\mathsf{off}}(b)]}_{\substack{\text{dealer's continuous} \\ \text{off-path payoff}}}, \tag{12.21}$$

where $\forall \theta_t$,

$$\begin{array}{rcl} U_t^{\mathsf{IDTE}}(\theta_t) & = & (1-\delta)\pi^{\theta_t} + \delta[\alpha_{\theta_t} U_{t+1}^{\mathsf{IDTE}}(g) + (1-\alpha_{\theta_t})U_{t+1}^{\mathsf{IDTE}}(b)] \\ U_t^{\mathsf{off}}(\theta_t) & = & (1-\delta)\mathbf{1}(\theta_t = g) \cdot \dfrac{\rho}{2}\sigma_g^2 + \delta[\alpha_{\theta_t} U_t^{\mathsf{off}}(g) + (1-\alpha_{\theta_t})U_t^{\mathsf{off}}(b)]. \end{array}$$

are dealer's on-path and off-path continuous payoffs respectively, given the current economic state is θ_t.

Similarly, in bad times, deviating to any price below p_t^g will decline the transaction and result in the loss of current profit; deviating to any price above p_t^g will let the trader agree to trade, but reduce the current profit. Hence, any deviation other than p_t^g is clearly not profitable. The only IC constraint left to check becomes

$$U_t^{\mathsf{IDTE}}(b) \geq \underbrace{(1-\delta)(\pi^g - J_g + J_b)}_{\substack{\text{dealer's static payoff after} \\ \text{pretending it's in good times}}} + \underbrace{\delta[\alpha_b U_t^{\mathsf{off}}(g) + (1-\alpha_b)U_t^{\mathsf{off}}(b)]}_{\substack{\text{dealer's continuous} \\ \text{off-path payoff}}}. \tag{12.22}$$

By the similar argument as in the proof of Proposition 1.5.1, one can show that $U_t^{\mathsf{IDTE}}(\theta_t)$

is stationary. Denote it as $U_t^{\mathsf{IDTE}}(\theta_t) = U^{\mathsf{IDTE}}(\theta_t)$.

$$\begin{bmatrix} 1-\delta\alpha_g & -\delta(1-\alpha_g) \\ -\delta\alpha_b & 1-\delta(1-\alpha_b) \end{bmatrix} \cdot \begin{bmatrix} U^{\mathsf{IDTE}}(g) \\ U^{\mathsf{IDTE}}(b) \end{bmatrix} = (1-\delta) \cdot \begin{bmatrix} \pi^g \\ \pi^b \end{bmatrix},$$

and $U_t^{\mathsf{off}}(\theta_t) = U^{\mathsf{off}}(\theta_t)$ is as calculated in the proof of Proposition 1.5.1.

Whence, there exists such an informative dynamic trading equilibrium where dealer offers distinct prices and trade occurs in both good times and bad times if equation (12.21) and equation (12.22) hold. Equivalently, solving $U^{\mathsf{IDTE}}(\theta_t)$ and $U^{\mathsf{off}}(\theta_t)$ shows that IDTE exists if there exists $\pi^b \in [0, \frac{\rho\sigma_b^2}{2}]$ such that

$$(1 - \frac{\delta(1-\alpha_g)}{(1-\delta)(1-\delta(\alpha_g-\alpha_b))})\pi^b \le \frac{\rho\sigma_g^2}{2} - J_g + J_b \le [1 + \delta\frac{1-\delta\alpha_g-(1-\delta)\alpha_b}{(1-\delta)(1-\delta(\alpha_g-\alpha_b))})]\pi^b.$$

■

Proof of Corollary 1.11.1. Consider the equilibrium σ^{ODTE} constructed in the proof of Proposition 1.5.1, we show that given there exists an IDTE, then σ^{ODTE} is also an equilibrium. From the proof we know that we only need to check that condition (12.4) is satisfied:

$$U^{\mathsf{ODTE}}(b, \theta_{t-1}) \ge (1-\delta) \times 0 + \delta[\alpha_b U^{\mathsf{off}}(g) + (1-\alpha_b)U^{\mathsf{off}}(b)]$$

where $\forall \theta_t, \theta_{t-1}$,

$$\begin{aligned} U^{\mathsf{ODTE}}(\theta_t, \theta_{t-1}) &= (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p_t^{*\mathsf{UI}}(\theta_{t-1})) + \delta[\alpha_{\theta_t} U^{\mathsf{ODTE}}(g, \theta_t) + (1-\alpha_{\theta_t})U^{\mathsf{ODTE}}(b, \theta_t)] \\ U^{\mathsf{off}}(\theta_t) &= (1-\delta)\mathbf{1}(\theta_t = g) \cdot \frac{\rho}{2}\sigma_g^2 + \delta[\alpha_{\theta_t} U^{\mathsf{off}}(g) + (1-\alpha_{\theta_t})U^{\mathsf{off}}(b)]. \end{aligned}$$

are dealer's on-path and off-path continuous payoffs respectively, given the current and last economic states are θ_t and θ_{t-1}.

In fact, suppose σ is such an IDTE equilibrium. At any on-path history h^t, denote dealer's continuous payoff in this IDTE as $\hat{U}_t(\theta_t, h^{t-1})$. If the current period is in good times, she should have no incentive to deviate to pretend it is in bad times and offer a bad time price $p_t^b(h^{t-1})$. That is,

$$\hat{U}_t(g, h^{t-1}) \ge (1-\delta)(\varphi a_{t-1} + J_g - p_t^b(h^{t-1})) + \delta[\alpha_g U^{\mathsf{off}}(g) + (1-\alpha_g)U^{\mathsf{off}}(b)]$$

where $\forall \theta_t$

$$\hat{U}_t(\theta_t, h^{t-1}) = (1-\delta)(\varphi a_{t-1} + J_{\theta_t} - p_t^{\theta_t}(h^{t-1})) + \delta[\alpha_{\theta_t}\hat{U}_t(g, h^t) + (1-\alpha_{\theta_t})\hat{U}_t(b, h^t)]$$

are dealer's on-path continuous payoffs for equilibrium σ and $U_t^{\mathsf{off}}(\theta_t)$ are the same off-path continuous payoff as those in the ODTE, given the current economic state is θ_t.

From the proof of Proposition 1.11.1 we learned that ρ^{IDTE} reaches its lower bound when the following condition is binding:

$$U^{\mathsf{IDTE}}(g) \geq (1-\delta)(\pi^b + J_g - J_b) + \delta[\alpha_g U^{\mathsf{off}}(g) + (1-\alpha_g)U^{\mathsf{off}}(b)],$$

where $\pi^b = \varphi a_{t-1} + J_b - p_t^b(h^{t-1})$ and $p_t^g(h^{t-1}) = \varphi a_{t-1} + J_g - \frac{\rho}{2}\sigma_g^2$, so $\hat{U}_t(g, h^{t-1}) = U^{\mathsf{IDTE}}(g)$.

Therefore, we only need to show that

$$\begin{aligned} & U^{\mathsf{ODTE}}(b, \theta_{t-1}) - \delta[\alpha_b U^{\mathsf{off}}(g) + (1-\alpha_b)U^{\mathsf{off}}(b)] \\ \geq \; & U^{\mathsf{IDTE}}(g) - (1-\delta)(\pi^b + J_g - J_b) - \delta[\alpha_g U^{\mathsf{off}}(g) + (1-\alpha_g)U^{\mathsf{off}}(b)] \end{aligned} \quad (12.23)$$

In fact,

$$\begin{aligned} & U^{\mathsf{ODTE}}(b, \theta_{t-1}) - \delta[\alpha_b U^{\mathsf{off}}(g) + (1-\alpha_b)U^{\mathsf{off}}(b)] \\ - \; & U^{\mathsf{IDTE}}(g) + (1-\delta)(\pi^b + J_g - J_b) + \delta[\alpha_g U^{\mathsf{off}}(g) + (1-\alpha_g)U^{\mathsf{off}}(b)] \\ = \; & U^{\mathsf{ODTE}}(b, \theta_{t-1}) - U^{\mathsf{IDTE}}(g) + (1-\delta)(\pi^b + J_g - J_b) + \delta(\alpha_g - \alpha_b)(U^{\mathsf{off}}(g) - U^{\mathsf{off}}(b)) \\ = \; & (1-\delta)\Lambda^{b,\theta_{t-1}} + \delta[\alpha_b U^{\mathsf{ODTE}}(g, \theta_{t-1}) + (1-\alpha_b)U^{\mathsf{ODTE}}(b, \theta_{t-1})] + (1-\delta)(\pi^b + J_g - J_b) \\ & -(1-\delta)\pi^b - \delta[\alpha_g U^{\mathsf{IDTE}}(g, \theta_{t-1}) + (1-\alpha_g)U^{\mathsf{IDTE}}(b, \theta_{t-1})] + \delta(\alpha_g - \alpha_b)(U^{\mathsf{off}}(g) - U^{\mathsf{off}}(b)) \\ = \; & (1-\delta)(\Lambda^{b,\theta_{t-1}} - \pi^b + \pi^b + J_g - J_b) + \delta(\alpha_g - \alpha_b)[(U^{\mathsf{off}}(g) - U^{\mathsf{off}}(b)) - (U^{\mathsf{IDTE}}(g) - U^{\mathsf{IDTE}}(b))] \\ & +\delta[\alpha_b U^{\mathsf{ODTE}}(g, \theta_{t-1}) + (1-\alpha_b)U^{\mathsf{ODTE}}(b, \theta_{t-1}) - \alpha_b U^{\mathsf{IDTE}}(g) - (1-\alpha_b)U^{\mathsf{IDTE}}(b)] \quad (12.24) \\ \geq \; & 0 \end{aligned}$$

To see the last inequality, notice that

$$\begin{aligned} & \Lambda^{b,\theta_{t-1}} - \pi^b + \pi^b + J_g - J_b \\ = \; & -\alpha_{\theta_{t-1}}(J_g - J_b) + \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2] + J_g - J_b \\ = \; & (1-\alpha_{\theta_{t-1}})(J_g - J_b) + \frac{\rho}{2}[\alpha_{\theta_{t-1}}(1-\alpha_{\theta_{t-1}})(J_g - J_b)^2 + \alpha_{\theta_{t-1}}\sigma_g^2 + (1-\alpha_{\theta_{t-1}})\sigma_b^2] \\ \geq \; & 0; \end{aligned}$$

and

$$\begin{aligned} & U^{\mathsf{off}}(g) - U^{\mathsf{off}}(b) - U^{\mathsf{IDTE}}(g) + U^{\mathsf{IDTE}}(b) \\ = \; & (1-\delta)\left\{[1,-1]\cdot\begin{bmatrix} 1-\delta\alpha_g & -\delta(1-\alpha_g) \\ -\delta\alpha_b & 1-\delta(1-\alpha_b) \end{bmatrix}^{-1}\cdot\begin{bmatrix} \pi^g \\ 0 \end{bmatrix} - [1,-1]\cdot\begin{bmatrix} 1-\delta\alpha_g & -\delta(1-\alpha_g) \\ -\delta\alpha_b & 1-\delta(1-\alpha_b) \end{bmatrix}^{-1}\cdot\begin{bmatrix} \pi^g \\ \pi^b \end{bmatrix}\right\} \\ = \; & \frac{1-\delta}{(1-\delta)(1-\delta(\alpha_g - \alpha_b))}\left\{[1,-1]\cdot\begin{bmatrix} 1-\delta(1-\alpha_b) & \delta(1-\alpha_g) \\ \delta\alpha_b & 1-\delta\alpha_g \end{bmatrix}\cdot\begin{bmatrix} 0 \\ -\pi^b \end{bmatrix}\right\} \\ = \; & \frac{-\pi^b}{1-\delta(\alpha_g - \alpha_b)}\left\{[1,-1]\cdot\begin{bmatrix} \delta(1-\alpha_g) \\ 1-\delta\alpha_g \end{bmatrix}\right\} \\ = \; & \frac{(1-\delta)\pi^b}{1-\delta(\alpha_g - \alpha_b)} \\ \geq \; & 0; \end{aligned}$$

and the non-negativity of the last part in (12.24) is due to the fact that in opaque equilibrium the dealer can extract more information rent and gain a higher ex-ante expected payoff along the equilibrium path. Therefore, we just show that condition (12.23) holds and complete the proof. ■

Proof of Proposition 1.11.2. As in the proof of Proposition 1.7.2, in any informative equilibrium, as long as there is trade in bad times, trader's ex-post utility becomes $p_t^b - \varphi a_{t-1} - J_b + \frac{\rho}{2}\sigma_b^2$, while dealer's ex-post utility is $\varphi a_{t-1} + J_b - p_t^b$. Therefore, in any informative equilibrium, the social welfare in bad times is $\frac{\rho}{2}\sigma_b^2$, providing there is trade. Therefore, IDTE provides the maximal social welfare in all informative equilibria (the ones where the dealer prices discriminatingly), which is $\frac{\rho}{2}\sigma_g^2$ in good times and $\frac{\rho}{2}\sigma_b^2$ in bad times. ■

Proof of Proposition 1.11.3. Similarly as in the proof of Proposition 1.7.1, we can rewrite the off-path continuation payoffs as $U_t^{\mathsf{off}}(\theta_t; q)$, given the current economic state is θ_t and the auditing intensity is q. The recursive formula for this off-path continuation payoff is

$$U_t^{\mathsf{off}}(\theta_t; q) \quad = \quad (1-\delta)\mathbf{1}(\theta_t = g) \cdot \frac{\rho}{2}\sigma_g^2 + \delta(1-q)[\alpha_{\theta_t} U_t^{\mathsf{off}}(g; q) + (1-\alpha_{\theta_t})U_t^{\mathsf{off}}(b; q)].$$

Since if the dealer follows the IDTE strategy, she will never be punished, the on-path continuation payoff remains the same:

$$U_t^{\mathsf{IDTE}}(\theta_t) \quad = \quad (1-\delta)\pi^{\theta_t} + \delta[\alpha_{\theta_t} U_{t+1}^{\mathsf{IDTE}}(g) + (1-\alpha_{\theta_t})U_{t+1}^{\mathsf{IDTE}}(b)].$$

The IC constraint becomes

$$\begin{aligned} U_t^{\mathsf{IDTE}}(g) &\geq (1-\delta)(\pi^b + J_g - J_b) + \delta(1-q)[\alpha_g U_t^{\mathsf{off}}(g; q) + (1-\alpha_g)U_t^{\mathsf{off}}(b; q)]; \\ U_t^{\mathsf{IDTE}}(b) &\geq (1-\delta)(\pi^b + J_b - J_g) + \delta(1-q)[\alpha_b U_t^{\mathsf{off}}(g; q) + (1-\alpha_b)U_t^{\mathsf{off}}(b; q)]. \end{aligned}$$

A similar argument can show that the off-path continuation payoff $U_t^{\mathsf{off}}(\theta_t; q)$ is decreasing in q. Therefore, with a higher auditing intensity, the IC constraint is relaxed, which completes the proof. ■

Chapter 2

Information Choice and Amplification of Financial Crises

2.1 Introduction

Financial investors devote substantial economic activity to acquiring information. In particular, information acquisition played an important role in the most recent financial crisis. In section 2.2, we review evidence about the importance of such information acquisition in the run-up of several recent crisis episodes. In this paper, we show how the acquisition of information by investors amplifies the probability of a financial crises.

Financial crises were historically explained by either weak fundamentals or panics.[1] The global games literature pioneered by Carlsson and van Damme (1993) reconciles both of these views, since weak fundamentals cause the self-fulfilling beliefs about a financial crisis. In particular, global coordination games of regime change are used to study bank runs,

This chapter is coauthored with Toni Ahnert. This chapter is published in the *Review of Financial Studies*, vol 30. We thank Itay Goldstein (the editor) and two anonymous referees for constructive comments that significantly improved the paper. We are also grateful to Daron Acemoglu, Jason Allen, Kartik Anand, Marios Angeletos, Ricardo Caballero, Piotr Danisewicz (discussant), Christian Eufinger, Janet Hua Jiang, Leonid Kogan, Andrey Malenko, Asu Ozdaglar, Andy Postlewaite, Matthew Prittsker (discussant), Josef Schroth, Francesco Trebbi, Russell Wong for very helpful comments.

[1] A self-fulfilling crisis can be caused by a panic among bank depositors, as in Bryant (1980) and Diamond and Dybvig (1983), or among currency speculators, as in Obstfeld (1996). By contrast, fundamental-based crises are analyzed by Chari and Jagannathan (1988), Jacklin and Bhattacharya (1988), and Allen and Gale (1998) for bank runs, and by Krugman (1979) for currency crises. See also Goldstein (2012).

currency attacks, and debt crises.[2] A crisis occurs if sufficiently many depositors withdraw funds from a bank, currency speculators attack a peg, or creditors do not roll over debt. In these models, financial investors base their decisions on an exogenous endowment of public and private information about an unobserved fundamental that measures the profitability of a bank, the foreign reserves of a central bank, or the solvency of a debtor.

In section 3.2, we offer a parsimonious model of information choice in a standard global coordination game of regime change. Investors choose ex ante whether to improve, at a cost, the quality of their private information about the fundamental, for example by hiring analysts or purchasing data. A heterogeneous information cost ensures the existence of a unique equilibrium, in which investors with a sufficiently low cost acquire information.[3]

In section 2.4, we explain how the information choices of investors amplify the probability of a financial crisis. To illustrate the mechanism, consider a debt rollover game. Each investor wishes to roll over debt whenever the debtor is solvent, which is more likely when other investors also roll over. Suppose that public information about debtor solvency is strong, for example the debtor's credit rating is high. What happens after a rating downgrade or an earnings warning, either of which weakens the public information about debtor solvency?

After a downgrade, debt holders have a higher incentive to acquire private information, which is more valuable for two reasons. First, since the public signal casts more doubt about the solvency of the debtor, acquiring private information helps a debt holder determine the solvency of the debtor. Second, acquiring private information helps a debt holder anticipate the rollover decisions of other debt holders. For these fundamental and strategic reasons, a debt holder with more precise private information is more likely to roll over debt when the debtor is solvent, and more likely to withdraw when a debt run occurs and the debtor is insolvent. As a result, more investors acquire private information after a downgrade.

A larger proportion of informed investors increases the probability of a debt crisis. Uninformed investors have less precise private information about debtor solvency and therefore rely more on public information. Since the initial credit rating is high, uninformed investors tend to roll over debt. By contrast, informed investors have more precise private informa-

[2]See Rochet and Vives (2004) and Goldstein and Pauzner (2005) for bank runs; Morris and Shin (1998) and Corsetti et al. (2004) for currency attacks; Morris and Shin (2004) and Corsetti, Guimaraes and Roubini (2006) for debt crises. See Bebchuk and Goldstein (2011) for credit freezes and also Vives (2005, 2014).

[3]Endogenous information in global coordination games can lead to multiple equilibria. In Angeletos and Werning (2006), a public market price aggregates dispersed private information, similar to Grossman and Stiglitz (1980). See also Hellwig, Mukherji and Tsyvinski (2006) and Angeletos, Hellwig and Pavan (2006). For other liquidity papers see Bosshardt and Kakhbod (2020*c*), Bosshardt and Kakhbod (2020*d*) and Kakhbod and Song (2020*b*).

tion and therefore tend to disregard the initially favorable public information, the high credit rating. As a result, informed investors refuse to roll over debt for a larger range of private signals. Hence, a larger proportion of informed investors increases the solvency threshold below which a debt crisis occurs and the ex-ante probability of a debt crisis.[4]

The case of strong but deteriorating public information (e.g., a downgrade of an initially high rating) is arguably empirically relevant for many recent financial crisis episodes. However, amplification also occurs for low and deteriorating public information, though the workings of the mechanism differ.[5] After a downgrade of a debtor with an initially low credit rating, the value of private information is lower and fewer investors acquire private information. Intuitively, there is little doubt about the insolvency of the debtor and the rollover behaviour of other creditors. A smaller proportion of informed investors, however, increases the probability of a debt crisis for a low rating. The intuition is again that uninformed investors are less likely to act against public information, which is unfavorable in this case. The different workings of the amplification mechanism suggest a different policy response.

In section 2.5, we study how a policymaker can enhance financial stability. We consider a deep-pocketed policymaker concerned about reducing the probability of a financial crisis. We start by studying taxes and subsidies on information acquisition, which alter the information choice of investors. However, the appropriate policy response to deteriorating public information depends on the solvency of the debtor. For a creditworthy debtor, taxation is desirable to discourage information acquisition. In contrast, for a less creditworthy debtor, a subsidy is desirable to encourage information acquisition. We show that similar results obtain when the policymaker can tax and subsidize the payoffs of investors instead.

We also study the consequences of an improvement in the quality of public information. A policy that makes public information more precise includes the regulation of credit rating agencies to make future ratings more informative or a commitment to publish the results of future bank stress tests. This policy intervention occurs ex ante, that is, before the public information about debtor solvency or bank asset quality is observed.[6] The first effect is that there is increased coordination on the more informative public signal. The second effect is that fewer investors acquire private information (crowding out). In sum, improving the

[4]This result on the extensive margin of private information (the proportion of informed investors) complements the result on the intensive margin (the precision of private information) in Metz (2002).

[5]The amplification result also holds irrespective of whether public information improves or deteriorates.

[6]In contrast, the tax or subsidy on information acquisition or investor payoffs is ex post, that is after investors have observed the public signal about debtor solvency. It is important to stress that either policy intervention occurs *before* investors decide whether to acquire private information.

quality of public information has an *ambiguous* effect on the probability of a crisis, which is lower whenever the public signal is strong. This policy may have unintended consequences, for example committing to releasing bank stress tests could amplify future banking crises.

Turning to testable implications in section 2.6, we derive and characterize the magnitude of amplification. Our theory implies that the magnitude of amplification is (i) non-monotone in the public information about debtor solvency, such as credit ratings; (ii) higher for less dispersed distributions of the information cost; (iii) higher when the precision improvement in private information is larger. Intuitively, the magnitude is larger the more investors alter their information choice and the more precise private information of informed investors is. To obtain these testable implications, we generalize our model to a generic distribution of information costs and limited precision improvement, which also illustrates the robustness of the amplification mechanism. Throughout this section, we discuss several environments in which these implications about the magnitude of amplification can be tested.

In section 2.7, we study extensions to probe the robustness of the mechanism further. An important feature of the amplification mechanism is how a higher proportion of informed investors affects the probability of a crisis. Analyzing a specification with generalized payoffs proposed by Iachan and Nenov (2015), which generates a more general link between the proportion of informed investors and the probability of a crisis, we state sufficient conditions for amplification to occur. We also consider an extension with a homogeneous information cost that yields multiple equilibria (Hellwig and Veldkamp, 2009). We explain in which cases amplification still obtains. A conclusion and appendices with derivations and proofs follow.

Other theories of amplification have been proposed. Fire sales occur when the natural buyer of an asset experiences financial stress (Shleifer and Vishny, 1992, 1997; Kiyotaki and Moore, 1997). Such sales can be induced by predatory trading (Brunnermeier and Pedersen, 2005). Investors may disengage from markets due to complexity (Caballero and Simsek, 2013) or Knightian uncertainty (Caballero and Krishnamurthy, 2008). Under adverse selection in secondary debt markets (Gorton and Pennacchi, 1990), information production may be destabilizing (Dang, Gorton and Holmström, 2012; Gorton and Ordonez, 2014).

Our model is related to the literature of information choice in coordination games. Hellwig and Veldkamp (2009) first studied the optimal information choice in strategic models. In a beauty contest game, they show that the information choices of investors inherit the underlying strategic motive (complementarity or substitutability).[7] We confirm their "inher-

[7]Myatt and Wallace (2012) and Colombo, Femminis and Pavan (2014) also study information choice in beauty contests.

itance result" in the context of a global coordination game with parsimonious information choice. Furthermore, Hellwig and Veldkamp (2009) show that multiple equilibria arise from a binary information choice and homogeneous information cost. In our model, a unique equilibrium obtains since the information cost is *heterogeneous* across investors. In a regime change game, Szkup and Trevino (2015) study continuous information choice subject to a convex information cost that is *homogeneous* across investors. They investigate the efficiency of equilibrium, when information choices are complements or substitutes, and the trade-off between public and private information, focusing on the *precision* of public information. In contrast, we focus on the *level* of public information, which is crucial for our amplification mechanism. Specifically, we study how changes in the public signal affect both the incentives of investors to acquire information and the probability of a crisis.

Our paper contributes to a literature on dynamic global games of regime change. Dasgupta (2007) studies the option to delay foreign direct investment in emerging markets, where the benefit of more precise information is traded off with a lower return on investment. Angeletos, Hellwig and Pavan (2007) study an infinite-horizon version with the arrival of additional private information over time that generates rich equilibrium dynamics. Like Szkup and Trevino (2015), our focus is on costly information acquisition in a first stage that affects the actions of investors in a second stage, such as their decision to roll over debt. In Yang (2015), the information cost is proportional to the implied reduction in entropy, which generates a coordination motive in information choices and multiple equilibria.[8]

He and Manela (2016) study the acquisition of information about bank liquidity and the dynamic withdrawal decisions of investors. Building on the asynchronous awareness model of Abreu and Brunnermeier (2003), their model yields rich time-series implications about run behaviour. In contrast, our static coordination game is a very different framework and we emphasize the acquisition of information about bank *solvency*.[9]

2.2 Information acquisition and financial crises

In this section, we review existing evidence about the importance of information acquisition in several recent crisis episodes. We also state a noisy but direct measure of information

[8]Kendall (2015) offers another alternative specification of information costs. Modelling a financial market, information acquisition has a time cost due to the expected adverse price movements induced by other traders.

[9]Nikitin and Smith (2008) study costly verification of solvency in a Diamond and Dybvig (1983) setup. See also Kakhbod et al. (2018) for cheap talk and Kakhbod and Loginova (2018) for verifiable communications.

acquisition before the crises at Bear Stearns, Lehman, and Greek debt. Private information acquisition seems to occur before each crisis event once some bad public news arrived.

Kacperczyk and Schnabl (2010) argue that information acquisition was an important factor in the run on commercial paper in 2007–08. Public information about the safety of commercial paper deteriorated due to both concerns about the quality of the underlying collateral and counterparty risk. Consistent with our model, they state on page 45:

> [B]efore the financial crisis, most investors believed that commercial paper almost never defaults and therefore had little incentive to invest in information gathering about issuers of commercial paper. [...] However, during the crisis, investors decided to invest more resources in information gathering activities [...]

According to the Financial Crisis Inquiry Report (2011), the hedge fund investors of Bear Stearns acquired information after Bear Stearns reported its first quarterly loss, but before the eventual run in March 2008. Consistent with our model's prediction of more private information acquisition after public information about the financial health of Bear Stearns deteriorated, the Financial Crisis Inquiry Report (2011) states on page 286:

> The hedge funds that were clients of Bear's prime brokerage services were particularly concerned that Bear would be unable to return their cash and securities. Lou Lebedin, the head of Bear's prime brokerage, told the FCIC that hedge fund clients occasionally inquired about the bank's financial condition in the latter half of 2007, but that such inquiries picked up at the beginning of 2008.

He and Manela (2016) argue that information acquisition played a role in the run on the U.S. commercial bank Washington Mutual and the run by U.S. money market mutual funds on European banks, specifically those with exposure to Greek debt. According to Reuters (2011), information acquisition differed across money market funds, which yielded different conclusions by fund managers, whereby some managers rolled over bank debt, while others withdrew.

Finally, we report an imperfect but direct measure of investor attention and information acquisition for three crises episodes. Da, Engelberg and Gao (2011) suggests Google search frequency as a real-time measure that can be accessed via Google Trends. Figures 2-1, 2-2 and 2-3 report this measure for the crises of Bear Stearns, Lehman Brothers, and Greek debt. It shows an increase in search volumes before each crisis event after bad public news arrived.

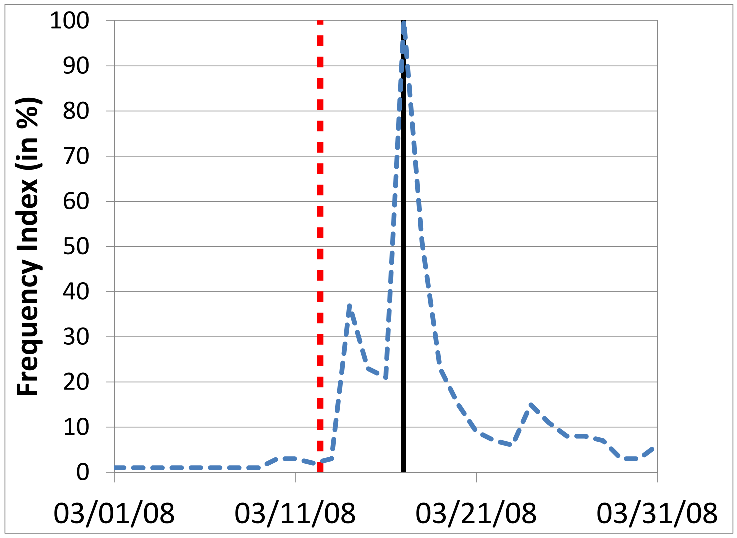

Figure 2-1: Bear Stearns

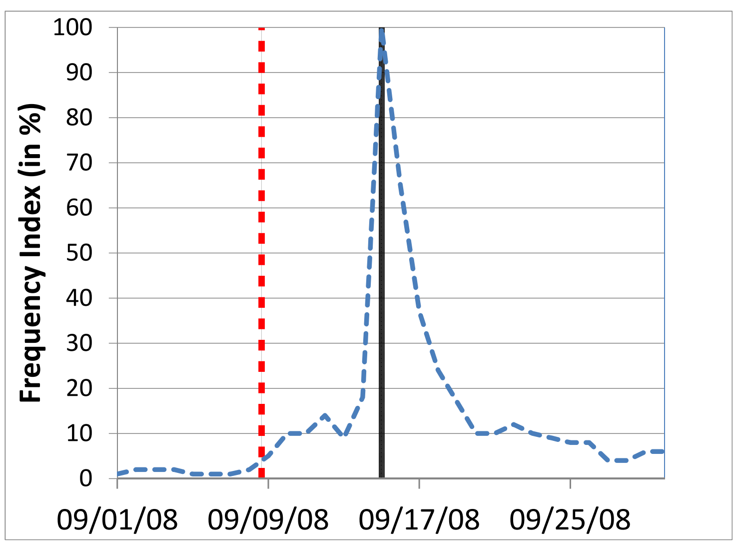

Figure 2-2: Lehman Brothers

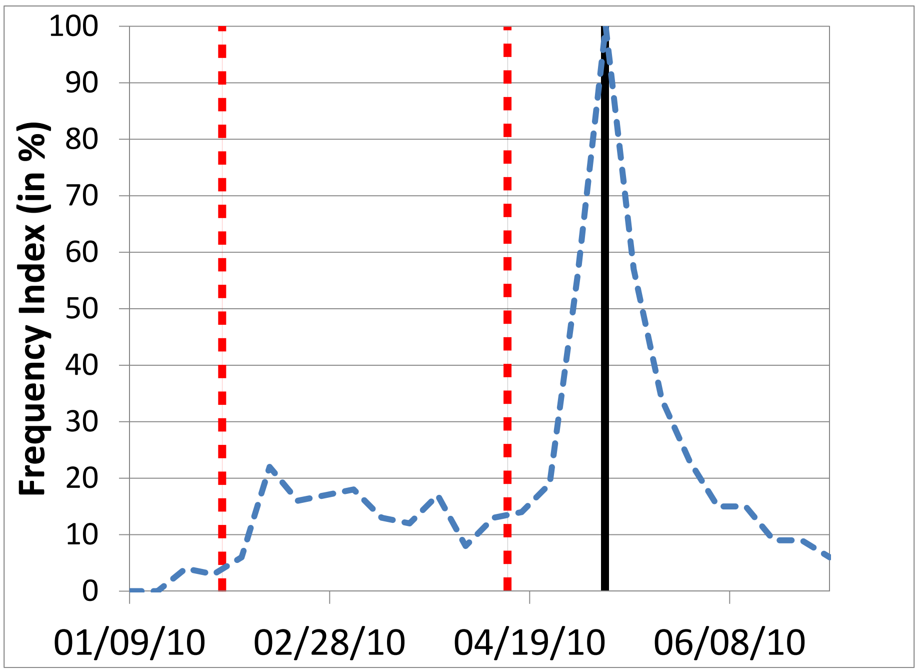

Figure 2-3: Greek sovereign debt crisis. Worldwide web search interest on Google for three crisis episodes, normalized to 100%. Figure 2-1 plots the key word 'Bear Stearns' during March 2008, Figure 2-2 plots 'Lehman Brothers' during September 2008, and this Figure plots 'Greek crisis' during the first half of 2010 as a single dashed blue line. The solid black line represents the crisis event: March 16, 2008 (sale of Bear Stearns to JP Morgan), September 15, 2008 (failure of Lehman Brothers), and May 2, 2010 (political agreement of first Greek bailout package). The dashed red line represents the arrival of bad public news: March 14, 2008 (New York Times, frontpage: 'Who could buy Bear Stearns?'), September 10, 2008 (New York Times, frontpage: 'Wall Street Fears on Lehman Batters Markets'), and February 3, 2010 (The Economist, frontpage: 'A Greek bailout, and soon?') and April 22, 2010 (Wall Street Journal, frontpage: 'Investors Desert Greek Bond Market'). Sources: GoogleTrends.

2.3 A Stylized Crisis Model With Information Choice

We propose a model of parsimonious information choice in a global coordination game of regime change. There is a unit continuum of risk-neutral investors $i \in [0,1]$ who simultaneously decide whether to attack the regime ($a_i = 1$) or not ($a_i = 0$). Regime change occurs if enough investors attack; that is, if the aggregate attack size $A \equiv \int_0^1 a_i \, di$ exceeds a fundamental $\theta \in \mathbb{R}$ that measures the strength of the regime. Building on Vives (2005), attacking yields a benefit $b \in (0,1)$ if regime change occurs, and a loss $\ell \in (0,1)$ otherwise:

$$u(a_i = 1, A, \theta) = b \ \mathbf{1}_{\{A \geq \theta\}} - \ell \ \mathbf{1}_{\{A < \theta\}}, \tag{3.1}$$

where $\mathbf{1}$ is the indicator function. An investor's incentive to attack is measured by the conservativeness ratio $\kappa \equiv \frac{\ell}{b+\ell} \in (0,1)$. The constant payoff from not attacking is normalized to zero. Therefore, the differential payoff from attacking increases in the attack size A and decreases in the fundamental θ. As a result, there is a coordination motive among investors whose attack decisions exhibit global strategic complementarity.

Our preferred interpretation of a regime change is a financial crisis, such as a currency crisis, a bank run, or a sovereign debt crisis. The fundamental θ can be interpreted as the ability of a monetary authority to defend its currency (Morris and Shin, 1998; Corsetti et al., 2004), as the measure of investment profitability (Rochet and Vives, 2004; Goldstein and Pauzner, 2005; Corsetti, Guimaraes and Roubini, 2006) or a sovereign's taxation power or willingness to repay. Investors can be interpreted as currency speculators, as retail or wholesale bank creditors who withdraw funds, or as sovereign debt holders who refuse to roll over.

There is incomplete information about the fundamental that is drawn from an improper uniform prior. Investors receive a public and a private signal (Morris and Shin (2003)):

$$\mu \equiv \theta + \nu, \quad \nu \sim \mathbb{N}\left(0, \alpha^{-1}\right) \tag{3.2}$$

$$x_i \equiv \theta + \epsilon_i, \quad \epsilon_i \sim \mathbb{N}\left(0, \gamma^{-1}\right), \tag{3.3}$$

where the aggregate noise ν is normally distributed with zero mean and precision $\alpha \in (0, \infty)$, and is independent of the fundamental. Idiosyncratic noise ϵ_i is identically and independently distributed across investors and independent of both the fundamental and the aggregate noise. The idiosyncratic noise is normally distributed with zero mean and an *endogenous* precision γ, described below. The information structure is common knowledge.

Table 2.1 summarizes the timeline of events. At the information stage, investors receive the public signal and simultaneously make a costly binary information choice, $n_i \in \{I, U\}$. At the subsequent coordination stage, informed investors ($n_i = I$) receive more precise private information:

$$\gamma_U < \gamma_I. \tag{3.4}$$

Unless stated otherwise, we focus on vanishing noise for informed investors, $\gamma_I \to \infty$, which provides a parsimonious benchmark and maintains uniqueness at the coordination stage.[10]

An information cost captures the resources required to acquire and process information. Skill differences in generating and processing information are captured by heterogeneity in the information cost that is uniformly distributed over a unit interval:[11]

$$c_i \sim \mathcal{U}[0, 1]. \tag{3.5}$$

I. Information stage	**II. Coordination stage**
1. Public signal μ about fundamental θ	1. Private signal x_i about fundamental θ • signal more precise if informed
2. Simultaneous information choice • binary action $n_i \in \{I, U\}$ • heterogeneous information cost c_i	2. Simultaneous attack decision • binary action $a_i \in \{0, 1\}$
	3. Outcome of regime and payoffs

Table 2.1: Timeline

2.3.1 Equilibrium

We start with a definition of the equilibrium concept. We focus on symmetric equilibria.

Definition 2.3.1. *A pure-strategy perfect Bayesian equilibrium comprises an information choice, $n_i^* = z \in \{I, U\}$ of each investor $i \in [0, 1]$, an aggregate proportion of informed investors, $n^* \in [0, 1]$, attack rules, $a_I^*(\cdot) \in \{0, 1\}$ and $a_U^*(\cdot) \in \{0, 1\}$ for informed and uninformed investors, respectively, and an aggregate attack size, $A^* \in [0, 1]$, such that:*

[10] We analyze an extension with limited precision improvement, $\gamma_I \in (\gamma_U, \infty)$, in section 2.6.4.

[11] We analyze an extension with a generic distribution of the information cost, $f(c)$, in section 2.6.3.

1. *For a given proportion of informed investors n^* and a given information choice $n_i^* = z$, the attack rule specifies an optimal behavior for each investor i at the coordination stage:*

$$a_z^*(\cdot) = \arg\max_{a_i} a_i \left[b \Pr(A^* \geq \theta | n^*, n_i^* = z; \mu, x_i) - \ell \Pr(A^* < \theta | n^*, n_i^* = z; \mu, x_i) \right]. \tag{3.6}$$

2. *For a given proportion of informed investors n^*, the aggregate attack size is consistent with the individually optimal attack behavior:*

$$A^* = n^* \int_0^1 a_I^*(\cdot)\, di + (1 - n^*) \int_0^1 a_U^*(\cdot)\, di. \tag{3.7}$$

3. *The aggregate proportion of informed investors is consistent with the individually optimal information choices:*

$$n^* = \int_0^1 \mathbf{1}\{n_i^* = I\}\, di. \tag{3.8}$$

4. *For a given proportion of informed investors n^* and attack rules $a_I^*(\cdot)$ and $a_U^*(\cdot)$, the private information choice n_i^* is optimal for each investor i:*

$$n_i^* = \arg\max_{n_i \in \{I, U\}} \mathbf{1}\{n_i = I\}[EU^I - c_i] + \mathbf{1}\{n_i = U\} EU^U, \tag{3.9}$$

where $EU^z = EU^z(n^*, a_I^*(\cdot), a_U^*(\cdot); \mu)$ is the expected utility of an investor who chooses $n_i^* = z$ at the information stage.

Proposition 2.3.1. ***Existence of a unique equilibrium.***
If private information is sufficiently precise, $\gamma_U > \underline{\gamma} \equiv \left(\frac{\alpha}{\sqrt{2\pi} - 2}\right)^2$, then there generically exists a unique pure-strategy monotone perfect Bayesian equilibrium. It is characterized by (i) a threshold information cost, $\bar{c}$, and (ii) signal thresholds for informed investors, $\overline{x}_I(\bar{c})$, and uninformed investors, $\overline{x}_U(\bar{c})$, and a fundamental threshold, $\overline{\theta}(\bar{c})$. At the information stage, investors acquire information if and only if their individual information cost is below the threshold, $n_i^* = I \Leftrightarrow c_i < \bar{c}$. The proportion of informed investors is $n^* = \bar{c}$. At the coordination stage, each investor attacks the regime if and only if it receives a private signal below the signal threshold specific to its information choice, $x_i < \overline{x}_z(\bar{c})$ for $n_i^* = z \in \{I, U\}$, and a regime change occurs if and only if the fundamental is below the threshold, $\theta < \overline{\theta}(\bar{c})$.

All the proofs are in the appendix.

The expected utility of an investor with information choice $n_i = z$ comprises two terms. An investor receives the benefit b when attacking ($x_i < \overline{x}_z$) if regime change occurs ($\theta < \overline{\theta}$), and the loss ℓ when attacking ($x_i < \overline{x}_z$) if no regime change occurs ($\theta > \overline{\theta}$):

$$\begin{aligned} EU^z &\equiv b\int_{-\infty}^{\overline{\theta}}\int_{-\infty}^{\overline{x}_z} f^z(x|\theta)dx\,dG(\theta) - \ell\int_{\overline{\theta}}^{\infty}\int_{-\infty}^{\overline{x}_z} f^z(x|\theta)dx\,dG(\theta) \\ &= bG(\overline{\theta}) - \ell\int_{\overline{\theta}}^{\infty}\int_{-\infty}^{\overline{x}_z} f^z(x|\theta)dx\,dG(\theta) - b\int_{-\infty}^{\overline{\theta}}\int_{\overline{x}_z}^{\infty} f^z(x|\theta)dx\,dG(\theta), \quad (3.10) \end{aligned}$$

where $G(\theta) \equiv \Phi\Big(\sqrt{\alpha}\,[\theta-\mu]\Big)$ is the cumulative distribution function (cdf) of the fundamental, $f^z(x) \equiv \sqrt{\gamma_z}\,\phi\Big(\sqrt{\gamma_z}\,[x-\theta]\Big)$ is the probability density function (pdf) of the private signal conditional on the fundamental and the information choice, and $\phi(\cdot)$ and $\Phi(\cdot)$ are the pdf and cdf of the standard Gaussian random variable, respectively.

The first term in equation (3.10) is the gain from attacking the regime if the investor had perfectly precise information about the fundamental. The second and third terms measure the mistakes of an investor due to imprecise information. The second term is the type-I error, whereby an investor attacks ($x_i < \overline{x}_z$) but no regime change occurs ($\theta > \overline{\theta}$). The third term is the type-II error, whereby an investor does not attack although a regime change occurs.

The value of more precise private information is therefore $D \equiv EU^I - EU^U$:

$$\begin{aligned} D &= \ell\int_{\overline{\theta}}^{\infty}\Gamma(\theta,\overline{x}_I,\overline{x}_U)dG(\theta) - b\int_{-\infty}^{\overline{\theta}}\Gamma(\theta,\overline{x}_I,\overline{x}_U)dG(\theta) \quad (3.11) \\ \Gamma(\theta,\overline{x}_I,\overline{x}_U) &\equiv \int_{-\infty}^{\overline{x}_U} f^U(x|\theta)dx - \int_{-\infty}^{\overline{x}_I} f^I(x|\theta)dx \equiv \Phi_U(\theta) - \Phi_I(\theta), \quad (3.12) \end{aligned}$$

where $\Gamma(\theta,\overline{x}_I,\overline{x}_U)$ measures, for any given fundamental, the difference in the probability of attacking between an uninformed and informed investor. This difference is generically non-zero, since informed investors receive more precise information than uninformed investors and therefore use a different signal threshold. Informed investors receive a precise private signal and do not make errors, since they attack if and only if regime change occurs. The value of private information is therefore the sum of the errors made by uninformed investors, weighted by the appropriate payoff parameter: uninformed investors sometimes attack although no regime change occurs and sometimes do not attack although a regime change occurs.

An investor with the threshold information cost $\overline{c}$ is indifferent between acquiring and not acquiring information, $\overline{c} = D(\overline{c})$. Figure 2-4 illustrates the unique solution to this fixed-point problem, which determines the threshold information cost.

We establish bounds on the value of private information, $D \in (0,1)$, in Lemma 2.9.3 in

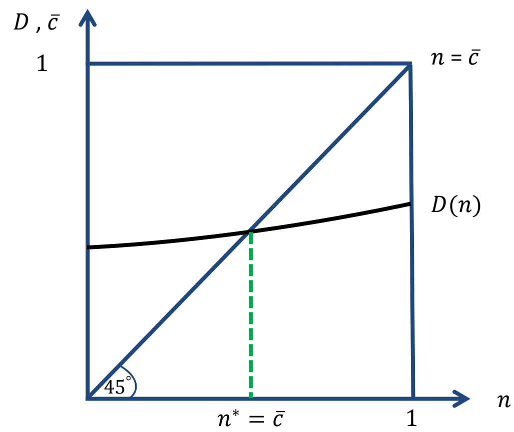

Figure 2-4: The value of more precise private information and the proportion of informed investors: a unique fixed point of the threshold information cost exists. Existence follows from the bounds on the value of information, $D \in (0,1)$ (see Lemma 2.9.3). Uniqueness follows from the bounds on the slope of the value of information, $\frac{dD}{dn^*} \in (0,1)$ (see Lemma 2.9.2).

Appendix 2.9.2. These bounds ensure the existence of a threshold information cost. The value of private information is positive for both fundamental and strategic reasons. First, informed investors form a more precise and more accurate belief about the fundamental than uninformed investors:

$$\theta|n_j = z, x_j \sim \mathbb{N}\left(\frac{\alpha\mu + \gamma_z x_j}{\alpha + \gamma_z}, \frac{1}{\alpha + \gamma_z}\right), \tag{3.13}$$

where $z \in \{I, U\}$ and $\lambda_z^j \equiv \frac{\alpha\mu+\gamma_z x_j}{\alpha+\gamma_z}$ is the posterior mean. Second, as a result, informed investors form a more accurate and more precise belief about the private information received by other investors. For example, the posterior belief of an informed investor j about an informed investor i is

$$x_i^I | x_j^I \sim \mathbb{N}\left(\lambda_I^j, \frac{1}{\alpha + \gamma_I} + \frac{1}{\gamma_I}\right), \tag{3.14}$$

while the posterior belief of an uninformed investor j about an informed investor i is

$$x_i^I | x_j^U \sim \mathbb{N}\left(\lambda_U^j, \frac{1}{\alpha + \gamma_U} + \frac{1}{\gamma_I}\right). \tag{3.15}$$

We show in Lemma 2.9.2 in Appendix 2.9.2 that the information choices of investors are strategic complements, $\frac{dD}{dn^*} > 0$. To obtain intuition, consider the two effects of more informed investors on the value of private information. First, the precision of the belief about the fundamental is independent of the proportion of informed investors. Second, the precision of the belief about the aggregate attack size increases in the proportion of informed investors. Mechanically, the attack behavior of informed investors matters more for the aggregate attack size as more investors are informed. Moreover, an informed investor has a more precise belief about the private information of informed investors than about the private information of uninformed investors:

$$Var\left(x_i^I|x_j^I\right) = \frac{1}{\alpha+\gamma_I} + \frac{1}{\gamma_I} < \frac{1}{\alpha+\gamma_I} + \frac{1}{\gamma_U} = Var\left(x_i^U|x_j^I\right). \quad (3.16)$$

Uniqueness rests on heterogeneous information costs and precise private information. First, for a homogeneous information cost, multiple equilibria arise for a binary information choice (Hellwig and Veldkamp, 2009). Despite the strategic complementarity in information choices, the amount of ex-ante heterogeneity suffices for uniqueness. The support of the information cost includes all possible values of private information, resulting in dominance regions at the information stage. Specifically, there exists a lower dominance region $[0, D(0))$ in which acquiring information is a dominant action, and an upper dominance region $(D(1), 1]$ in which not acquiring information is a dominant action. Second, sufficiently precise private information of uninformed investors ensures that $\frac{dD}{dn^*} < 1$, whereby a larger threshold information cost raises the proportion of informed investors and, via strategic complementarity in information choice, the value of private information at a sufficiently low rate (Lemma 2.9.2). The required lower bound on the precision of private information, $\underline{\gamma}$, is more restrictive than in a stand-alone global coordination game, such as Morris and Shin (2003).[12]

Definition 2.3.2. *Let $\hat{\mu} \equiv 1 - \kappa - \frac{\sqrt{\alpha+\gamma_U}-\sqrt{\gamma_U}}{\alpha}\Phi^{-1}(\kappa)$ denote a threshold level of the public signal. The public signal about the fundamental is **strong** if $\mu > \hat{\mu}$, and it is **weak** if $\mu < \hat{\mu}$.*

A unique equilibrium exists generically. The inequality $\mu \neq \hat{\mu}$ ensures that the information and contagion stages are linked, whereby changes in the aggregate proportion of informed investors affect the fundamental threshold, $\frac{d\overline{\theta}}{dn^*} \neq 0$, as shown in Lemma 2.9.1 in Appendix 2.9.1. All of the subsequent results also hold generically.

[12]We further analyze the lower bound on the precision of private information required for uniqueness in the case of a generic distribution of the information cost in section 2.6.3.

2.4 Amplification

A unique equilibrium is a solid foundation for comparative statics. Changes in the public signal, $d\mu$, give rise to amplification based on the information choice of investors. We study the ex-ante probability of a financial crisis determined by the fundamental threshold:

$$\Pr\{\theta \leq \overline{\theta}\} = \Phi\left(\sqrt{\alpha}\left[\overline{\theta}(\overline{c}) - \mu\right]\right). \tag{4.1}$$

To illustrate the amplification mechanism, we compare endogenous and exogenous information. Let $\tilde{\theta}$ be the fundamental threshold if information is exogenous. A well-known result in the literature is that a weaker public signal raises the fundamental threshold:

$$\left.\frac{d\tilde{\theta}}{d\mu}\right|_{\overline{n},\overline{\theta}} = \frac{A_\mu}{1 - A_{\overline{\theta}}} < 0, \tag{4.2}$$

where A_μ and $A_{\overline{\theta}}$ are the partial derivatives of the aggregate attack size with respect to the public signal and the fundamental threshold, respectively, derived in Appendices 2.10.1 and 2.9.1. All partial derivatives are evaluated at the fundamental threshold $\tilde{\theta}$ and an (exogenous) proportion of informed investors $\overline{n}$. To ensure comparability, this proportion is set to the equilibrium proportion of informed investors in the case of endogenous information, $\overline{n} = n^*$.

A novel effect arises if information is endogenous. Changes in the public signal now also affect the incentives to acquire information (captured by D_μ). In turn, changes in the proportion of informed investors affect the aggregate attack size (captured by A_{n^*}). In Appendix 2.10.1, we derive the total effect of a change in the public signal on the fundamental threshold by totally differentiating the equilibrium conditions of both stages:

$$\left.\frac{d\overline{\theta}}{d\mu}\right|_{n^*,\tilde{\theta}} = \frac{A_\mu + A_{n^*}D_\mu}{1 - A_{\overline{\theta}} - A_{n^*}D_{\overline{\theta}}} < 0, \tag{4.3}$$

where A_{n^*} is the partial derivative of the aggregate attack size with respect to the proportion of informed investors, while D_μ and $D_{\overline{\theta}}$ are the partial derivatives of the value of private information with respect to the public signal and the fundamental threshold. We derive $D_{\overline{\theta}}$ in the proof of Lemma 2.9.2 in Appendix 2.9.2, and D_μ in Appendix 2.10.1.

Proposition 2.4.1. ***Amplification.*** *If private information is sufficiently precise, $\gamma_U > \underline{\gamma}$, then the information choices of investors amplify the impact of changes in public information* *on both the fundamental threshold and on the probability of a financial crisis:*

$$- \left. \frac{d\overline{\theta}}{d\mu} \right|_{n^*, \overline{\theta}} > - \left. \frac{d\tilde{\theta}}{d\mu} \right|_{\overline{n}=n^*, \overline{\theta}}. \quad (4.4)$$

The condition sufficient for uniqueness also suffices for amplification. The lower bound on the precision of private information ensures a positive denominator of equation (4.3).

To provide intuition for the amplification result, we take a closer look at the constituting forces of the mechanism in three steps. First, we describe how a change in the public signal affects the incentives of investors to acquire private information. In equilibrium, there is a non-monotonic relationship between the public signal about the fundamental and the value of private information, $D_\mu (\mu - \hat{\mu}) < 0$, shown in Figure 2-5 and proven in Appendix 2.10.

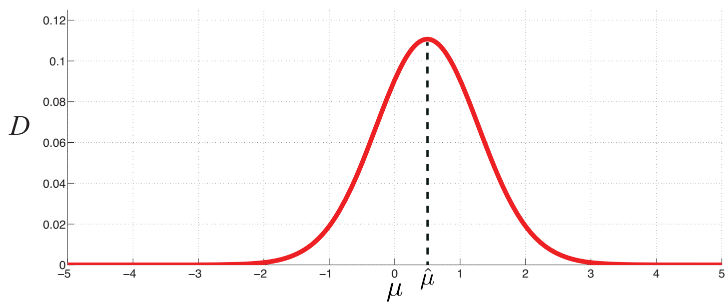

Figure 2-5: The relationship between the public signal μ and the value of private information D is non-monotonic. For low values of the public signal, $\mu < \hat{\mu}$, the value of private information increases in the public signal. The maximum value of private information is reached at $\mu = \hat{\mu}$. For high values of the public signal, $\mu > \hat{\mu}$, the value of private information decreases in the public signal.

Using the signal threshold of uninformed investors, $\overline{x}_U = \overline{\theta} + \frac{\alpha}{\gamma_U}\left[\overline{\theta} - \mu\right] - \frac{\sqrt{\alpha+\gamma_U}}{\gamma_U}\Phi^{-1}(\kappa)$, the value of private information can be expressed as follows for $\gamma_I \to \infty$:

$$\begin{aligned} D \;\to\; & \ell \int_{\overline{\theta}}^{\infty} \Phi\Big(\sqrt{\gamma_U}\,\left[\overline{\theta} - \theta\right] + \frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta} - \mu\right] - \sqrt{1 + \frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\Big)\; \phi\Big(\sqrt{\alpha}\,\left[\theta - \mu\right]\Big) d\theta + \quad (4.5) \\ & b \int_{-\infty}^{\overline{\theta}} \Big[1 - \Phi\Big(\sqrt{\gamma_U}\,\left[\overline{\theta} - \theta\right] + \frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta} - \mu\right] - \sqrt{1 + \frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\Big)\Big]\; \phi\Big(\sqrt{\alpha}\,\left[\theta - \mu\right]\Big) d\theta. \end{aligned}$$

To obtain intuition for the non-monotonicity of the value of private information in the public signal, we build on Szkup and Trevino (2015). If the public signal is weak, the fundamental threshold is high and a crisis is likely from an ex-ante perspective. Hence, there is a low probability of a type-I error ex post, which is attacking the regime when no regime change occurs. As shown by the first term in equation (4.5), the high realized levels of the fundamental required for no regime change to occur are unlikely (the term $\phi(\cdot)$ is small for all $\theta > \overline{\theta}$). There is also a low probability of a type-II error, which is not attacking the regime although regime change occurs, as the second term in equation (4.5) shows. The reason is different. Uninformed investors are likely to attack for a weak public signal (the term $1 - \Phi_U(\cdot)$ is small). Taken together, the value of private information is low for a weak public signal.

As the public signal strengthens, the fundamental threshold falls. The value of private information increases because both those fundamentals consistent with a type-I error are now more likely and the probability of an uninformed investor attacking the regime falls, which increases the expected utility loss due to type-II errors. Since precise private information allows an investor to avoid errors of either type, the value of private information increases. Similarly, the value of private information is low but increasing for a strong but weakening public signal. The maximum value of private information is reached at $\mu = \hat{\mu}$, which implies a fundamental threshold of $\hat{\theta} = 1 - \kappa$. At this point, the fundamental threshold is insensitive to changes in the proportion of informed investors, $\left.\frac{d\overline{\theta}}{dn^*}\right|_{\mu=\hat{\mu}} = 0$, shown in Lemma 2.9.1.

In a second step, a change in the value of private information directly affects the equilibrium proportion of informed investors, as shown in Figure 2-6. A unique solution to the fixed-point problem $\overline{c} = D(n^* = \overline{c})$ exists. An increase in the value of information from D_1 to D_2, for example, raises the equilibrium proportion of informed investors from n_1^* to n_2^*. The magnitude of this change is affected by the degree of strategic complementarity in information choices. If these strategic complementarities are stronger, such that the slope of the value of private information is steeper (right panel), then the increase is larger.

In a third step, changes in the proportion of informed investors affect the fundamental threshold below which a crisis occurs and, as a result, the probability of a financial crisis. For a strong public signal, $\mu > \hat{\mu}$, a larger proportion of informed investors leads to a higher threshold, while the converse result holds for a weak public signal, $\mu < \hat{\mu}$. Metz (2002) examines the dependence of the threshold on the precision of private and public information. Analyzing the intensive margin of private information, Metz (2002) finds that the effect of more precise private information on the fundamental threshold depends on the public signal. More specifically, the threshold increases (decreases) in the precision of private information if the public signal is strong (weak). In contrast, we focus on the proportion of informed investors. We show in Lemma 2.9.1 in Appendix 2.9 that the result of Metz (2002) also holds for the extensive margin of private information, $\frac{d\overline{\theta}}{dn^*}(\mu - \hat{\mu}) > 0$, as shown in Figure 2-7.

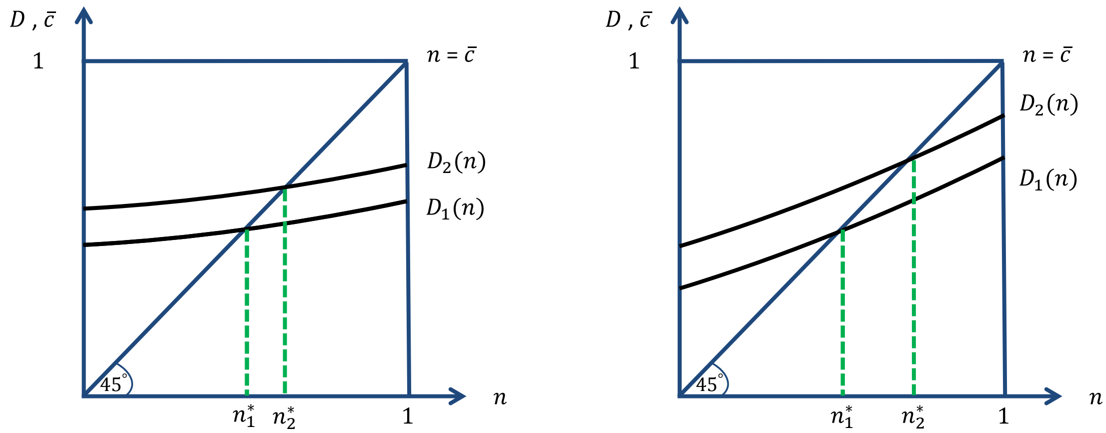

Figure 2-6: A higher value of private information increases the equilibrium proportion of informed investors. This effect is stronger the larger the degree of strategic complementarity in information choices. These strategic complementarities are stronger in the right panel.

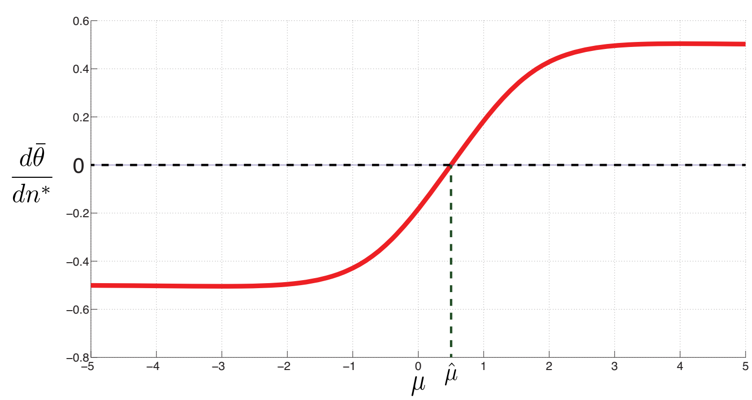

Figure 2-7: The impact of the proportion of informed investors on the fundamental threshold depends on the public signal. For $\mu > \hat{\mu}$, a larger proportion of informed investors leads to a higher fundamental threshold and to a higher probability of a financial crisis. The converse result holds for a weak public signal. Parameter values are $b = \ell = 0.75$, $\alpha = 1$, $\gamma_U = 4$.

The fundamental threshold increases in the proportion of informed investors whenever $\Gamma(\overline{\theta}, \overline{x}_I, \overline{x}_U) < 0$. This condition, evaluated at the threshold value of the fundamental, holds if an informed investor is more likely to attack than an uninformed investor, $\Phi_I(\overline{\theta}) > \Phi_U(\overline{\theta})$, which occurs for a strong public signal (Lemma 2.9.1). Since their private signal is more informative about the fundamental, informed investors put a larger weight on private information and therefore a smaller weight on public information than uninformed investors, as shown by the posterior about the fundamental in condition (3.13). Therefore, informed investors are more likely to disregard a strong public signal and attack the regime than are uninformed investors. In other words, informed investors attack the regime for a larger range of private signals than uninformed investors when the public signal is strong. As a result, the fundamental threshold increases in the proportion of informed investors.[13] Similarly, for a weak public signal, informed investors are less likely to attack the regime than uninformed investors, so the threshold decreases in the proportion of informed investors.

2.4.1 Amplification for both weak and strong public signals

Amplification occurs independently of the strength of the public signal and the direction of its change, as can be seen by combining the previous three steps. However, the mechanism works differently for a strong versus a weak public signal. This different mechanism will have consequences for the design of policy interventions discussed in section 2.5. To illustrate the amplification mechanism, we consider a deterioration in the public signal about the fundamental, which is purely expositional since the opposite result holds for an improvement.

Strong public signal The deterioration of an initially strong public signal raises the value of private information, so more investors choose to acquire information. Since informed investors place a lower weight on public information, they are more likely to attack than uninformed investors for a strong public signal. Therefore, the fundamental threshold and the probability of a financial crisis rise further after the deterioration of the public signal. In short, amplification arises from the information choices of investors.

Weak public signal Likewise, as an initially weak public signal deteriorates, the value of private information falls and fewer investors acquire information. Since informed investors place a lower weight on public information, they are less likely to attack than uninformed

[13] We revisit the link between the proportion of informed investors and the fundamental threshold below which a crisis occurs in section 2.7.1.

investors for a weak public signal. As fewer investors are informed, the fundamental threshold and the probability of a financial crisis increase further, yielding amplification.

2.5 Policy

To explore potential policy interventions, we analyze a policymaker concerned about financial instability, defined as the ex ante probability of a financial crisis given by equation (4.1). Financial instability is determined by the fundamental threshold below which a crisis occurs.

We consider two interventions. First, we study taxes and subsidies on information acquisition. (Similar results obtain for taxing payoffs.) These interventions are ex post, that is once the public signal is observed. Second, we study an improvement of the quality of public information. This intervention is ex ante, that is before the public signal is observed. It captures the regulation of credit rating agencies to induce more informative ratings or the commitment of of a regulator to publish the results of future bank stress tests.[14]

2.5.1 Taxes and subsidies

Suppose that the policymaker can tax or subsidize the information choice of investors, where we assume that the policymaker is deep-pocketed. The information cost of investor i changes to $c_i' \equiv (1-\tau)c_i$ for some $\tau < 1$, where a tax corresponds to $\tau < 0$ and a subsidy to $\tau > 0$. This policy affects optimality at the information stage only, where investor i acquires more precise private information if and only if $(1-\tau)c_i \leq \overline{c}$. Therefore, the proportion of informed investors is $n^* = \Pr\{c_i' \leq \overline{c}\} = \frac{\overline{c}}{1-\tau}$, where a subsidy raises the proportion of informed investors, while a tax lowers it. The fixed-point problem becomes $n^*(1-\tau) = \overline{c} = D$. Proposition 2.5.1 summarizes how changes in the tax or subsidy affect the fundamental threshold.

Proposition 2.5.1. ***Tax or subsidy on information acquisition.*** *A change in the tax or subsidy on information acquisition affects the fundamental threshold according to:*

$$\frac{d\overline{\theta}}{d\tau} = -\frac{\overline{c}}{(1-\tau)^2(1-A_{\overline{\theta}}) - (1-\tau)(b+\ell)g(\overline{\theta})\Gamma(\overline{\theta})^2}\,\Gamma(\overline{\theta}), \tag{5.1}$$

[14] Since the policymaker is uninformed when choosing the quality of public information, there is no signaling to investors. Angeletos, Hellwig and Pavan (2006) study signaling in a game of regime change without information choice.

which is positive (negative) if and only if the public signal is strong (weak), $\frac{d\bar{\theta}}{d\tau}(\mu - \hat{\mu}) > 0$.[15]

As a corollary of Proposition 2.5.1, the policymaker can use these taxes or subsidies to enhance financial stability after a reduction in the public signal, such as after a downgrade. However, the appropriate response (tax or subsidy) depends on the strength of the public signal, since the amplification mechanism works differently for a weak or strong public signal.

Corollary 2.5.1. ***Tax or subsidy policy.*** *To reduce the probability of a crisis, a policymaker's adequate response to a reduced public signal is contingent on the public signal. Taxation ($\tau < 0$) is desirable for a strong signal, while a subsidy ($\tau > 0$) is desirable for a weak signal.*

Figure 2-8 shows the case of a strong public signal, $\mu > \hat{\mu}$. A deterioration in the public signal, $d\mu < 0$, increases the value of private information, $D_\mu < 0$. Taxing information acquisition limits the increase in the proportion of informed investors and, as a result, the increase in the probability of a financial crisis. In contrast, for a weak public signal (not depicted), $\mu < \hat{\mu}$, a deterioration in the public signal decreases the value of private information, $D_\mu > 0$. Subsidizing information acquisition limits the decrease in the proportion of informed investors and the resulting increase in the probability of a financial crisis.

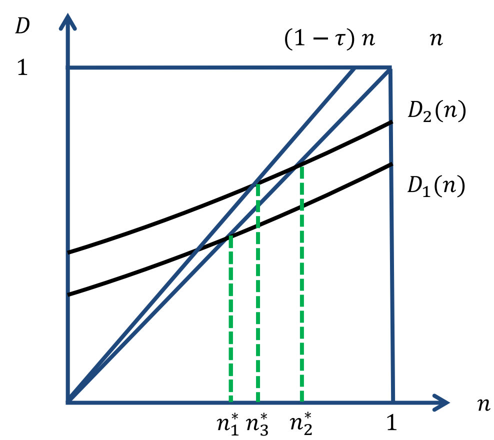

Figure 2-8: The impact of taxing the information acquisition of investors ($\tau < 0$) on the proportion of informed investors for a strong public signal, $\mu > \hat{\mu}$. A reduction in the public signal increases the value of private information from D_1 to D_2, increasing the proportion of informed investors from n_1^* to n_2^*. Taxing information acquisition raises the slope to $1-\tau > 1$ and limits the increase in the proportion of informed investors to $n_2^* < n_3^*$.

[15] An analogous result holds for limited precision improvement with $\hat{\mu}$ replaced by $\tilde{\mu}$. See section 2.6.4.

Similar results obtain if the policymaker can tax and subsidize the payoffs of investors. To illustrate this point, consider a tax on the benefit of attacking the regime changes, $b' \equiv (1-t)b$ for some $0 < t < 1$, combined with tax credit on the losses of attacking in the case of no regime change, $\ell' \equiv (1-t)\ell$. Since this policy leaves the conservatism ratio unchanged, $\kappa' = \kappa$, optimality at the coordination stage for any given proportion of informed investors is unaffected. Both the benefits and losses are reduced proportionally, so the value of private information shrinks accordingly, $D' \equiv (1-t)D$. Intuitively, acquiring more precise private information is less attractive because making errors of either type is now less costly. As a result, fewer investors acquire private information. The results of Corollary 2.5.1 carry over in that higher (lower) taxes are desirable for a strong (weak) public signal.

2.5.2 Improving the quality of public information

Since the information choice of investors may be difficult to verify (and thus difficult to tax), we study an improvement of the quality of public information as an alternative intervention. Suppose the policymaker can affect the precision of the public signal α. To simplify the exposition, we assume $\kappa = \frac{1}{2}$ and a precise private signal of informed investors, $\gamma_I > \underline{\gamma}_I$.

There are two consequences of improving the quality of public information. First, there is a *coordination effect*, whereby investors rely more on the public signal and less on the private signal at the coordination stage (see equation (3.13)). As a result, fewer investors attack the regime, $A_\alpha < 0$, if the public signal is sufficiently strong. For $\kappa = \frac{1}{2}$, this condition simplifies to $\mu > \hat{\mu} = \frac{1}{2}$. Second, there is an *information choice effect*, whereby more precise public information reduces the value of private information and thus crowds out the acquisition of private information, $D_\alpha < 0$. Taken together, Proposition 2.5.2 states the total effect of improving the quality of public information.

Proposition 2.5.2. ***Improving the quality of public information.*** *Suppose $\kappa = \frac{1}{2}$ and private information of informed investors is precise, $\gamma_I > \underline{\gamma}_I$. Improving the quality of public information affects the fundamental threshold according to:*

$$\frac{d\bar{\theta}}{d\alpha} = \frac{A_\alpha + A_{n^*} D_\alpha}{1 - A_{\bar{\theta}} - A_{n^*} D_{\bar{\theta}}}, \tag{5.2}$$

which is negative (positive) if and only if the public signal is strong (weak), $\frac{d\bar{\theta}}{d\alpha}(\mu - \hat{\mu}) < 0$.

As a corollary of Proposition 2.5.2, the effect of improving the quality of public information ex ante on the probability of a financial crisis ex post is ambiguous. The consequences of this policy depend on the strength of the public signal. Specifically, more informative credit ratings reduce the probability of a crisis if the public signal is strong. First, there is increased coordination on a more informative public signal. Second, the more informative public signal crowds out private information acquisition, which may have led to receiving bad news based on which an investor would have attacked. Since the reverse holds for a weak public signal, committing to releasing future stress tests could amplify future banking crises.

2.6 Magnitude of Amplification and Testable Implications

Having studied the qualitative aspects of the amplification mechanism, we now define and characterize its magnitude. Second, we study generalizations of our model to show the robustness of the mechanism and, more importantly, derive testable implications about the magnitude of amplification with respect to (i) the public signal, (ii) the distribution of information costs, and (iii) the precision of private information. Throughout this section, we suggest some environments in which these implications can be tested.

2.6.1 Magnitude of amplification

We define the magnitude of the amplification effect as $MoA \equiv \frac{\frac{d\bar{\theta}}{d\mu}\Big|_{n^*,\bar{\theta}}}{\frac{d\tilde{\theta}}{d\mu}\Big|_{\overline{n},\overline{\theta}}} - 1$ and derive a lower bound on this magnitude in Appendix 2.12.2.

Proposition 2.6.1. ***Magnitude of amplification.*** *For any precision improvement, the magnitude of amplification is*

$$MoA = \frac{\frac{dD}{dn^*}}{(1 - \frac{dD}{dn^*})A_{\overline{\theta}}} > 0. \tag{6.1}$$

It increases in the degree of strategic complementarity in information choices, $\frac{dD}{dn^*}$, and decreases in the sensitivity of the aggregate attack size to the fundamental threshold, $A_{\overline{\theta}}$.

In the following subsections, we derive three specific implications about the magnitude of amplification and discuss how these can be tested. The first result is that the magnitude of amplification is non-monotonic in the public signal (section 2.6.2). This result is immediately testable in public debt markets where unexpected changes in ratings or earnings have implications for the magnitude of amplification. The second result highlights the dispersion in information costs for the magnitude of amplification (section 2.6.3). We offer an empirical strategy based on exploiting differences in the design of over-the-counter versus centralized markets. The third result states that the magnitude of amplification is larger if informed investors have a more precise signal (section 2.6.4). This result can be tested by using differences in the sophistication of institutional investors or in firm and bank characteristics.

2.6.2 Public signal

Next, we study how the magnitude of amplification depends on the public signal.

Proposition 2.6.2. *For any precision improvement, the magnitude of amplification is non-monotonic in the public signal.*

Figure 2-9 shows how the magnitude of amplification is non-monotonic in the public signal, such as a credit rating. For a weak or strong public signal, the changes in the value of private information after changes in the public signal are small (Figure 2-5), so the magnitude of amplification is small. Moreover, for values close to $\hat{\mu}$, the sensitivity of the aggregate attack size to the fundamental threshold, $A_{\overline{\theta}}$, is high and the fundamental threshold changes little with the proportion of informed investors. As a result, the magnitude of amplification is also small. At $\mu = \hat{\mu}$, the magnitude is zero, while it is strictly positive elsewhere. For other values of the public signal, however, the value of private information is quite sensitive to changes in the public signal and the fundamental threshold is sensitive to changes in the proportion of informed investors, resulting in a large magnitude of amplification.

The non-monotonicity result can be tested in a number of markets, including the market for corporate debt. Consider a firm with debt to be rolled over by investors. A reduction in the public signal would correspond to an unexpected rating downgrade or earnings warning. Using the criteria stated in the next paragraph, an empiricist can separate firms according to whether information acquisition is likely or hard to occur. Our theory predicts that (i) amplification occurs for any given initial credit rating, and that (ii) the magnitude of this effect depends on the initial rating. For example, within the class of creditworthy firms (e.g.,

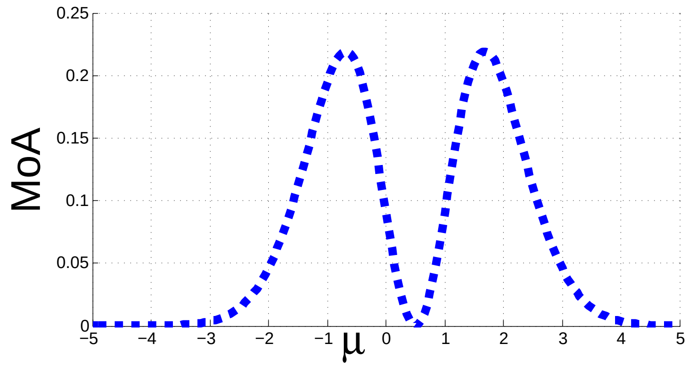

Figure 2-9: The non-monotonic relationship between the magnitude of amplification (MoA) and the public signal. Parameter values are $b = \ell = 0.75$, $\alpha = 1$, $\gamma_U = 4$.

those with investment-grade debt), our theory predicts the magnitude of amplification to be hump-shaped in the initial rating.

There are several potential proxies an empiricist may use for the extent of information investors are able to acquire. First, some firms are publicly listed and their *equity* is traded, so these firms are subject to disclosure requirements imposed by the Securities and Exchange Commission. In contrast, many other firms are privately owned and therefore disclosure requirements do not apply. As a result, information acquisition by debt holders is more relevant in the former case. Second, firm size may proxy for information, since small firms are followed by fewer analysts than large firms. Third, the industry or the age of the firm form other proxies. Young and innovative companies tend to have a higher share of intangible assets, which are more difficult to evaluate (e.g. technology firms versus consumer goods). In addition, their business models and growth perspectives may be less tested and more difficult to evaluate than those of established firms, for which ample data exists.

Kisgen (2006) studies the impact of credit ratings on capital structure decisions. Since rating categories are broad (e.g., AA+, AA, and AA-), a downgrade by one notch has little effect if still in the same category, but it has a large impact if the firm will be in a lower category (i.e., A+). Firms close to an upgrade or downgrade are shown to issue less net debt relative to net equity. Similar to Kisgen (2006), net debt issuance can be constructed for each corporate firm. Beyond corporate debt, similar tests could be conducted for bank commercial paper. It is short-term (typically, with a maturity below 270 days) and is therefore rolled over frequently. Apart from a downgrade or an earnings warning, a reduction in the public signal could come from a downward revision of a bank's asset quality by its supervisor.

Again, an empiricist would need to separate circumstances under which information about the fundamentals of the bank is easy to acquire from those where it may be difficult. The distinction between privately owned and publicly listed companies also applies to banks. Second, the opacity of the bank's assets is a proxy for investor capacity to acquire information. For example, a bank invested in traditional and marketable assets is more transparent than a bank invested in complicated and perhaps illiquid structured products (e.g., exposure to products based on asset-backed securities). Also, the complexity of a bank in terms of its legal and organizational structure is another proxy (Cetorelli and Goldberg, 2014; Goldberg, 2016). It affects the ease with which investors can determine the profitability of the bank or, conversely, whether exposures or outright losses may be hidden.

2.6.3 Generic distribution of information cost

Consider a generic distribution of the information cost given by the probability density function $f(c)$ with support $[c_{min}, c_{max}]$, where $0 \leq c_{min} < c_{max}$. We offer sufficient conditions for a unique equilibrium to exist and show that the amplification result continues to hold. We also describe how the information cost distribution affects the magnitude of amplification.

Proposition 2.6.3. ***Generic distribution of information cost.*** *Let $f(c)$ be the density function of the information cost with support $[c_{min}, c_{max}]$. If $c_{min} < D(0) < D(1) < c_{max}$ and $f(c) < \sqrt{\frac{\pi}{2}}$, then there exists a precision level $\tilde{\gamma} < \infty$ such that a unique equilibrium exists for any $\gamma_U \in (\tilde{\gamma}, \infty)$. Amplification via the information choice occurs in this equilibrium.*

The first sufficient condition ensures that the information cost is sufficiently heterogeneous relative to the bounds on the value of private information. As a result, the required dominance regions at the information stage are ensured and the existence of a threshold information cost follows. The second condition limits, after a small increase in the threshold information cost $\bar{c}$, the increase in the proportion of investors who acquire information. Because of strategic complementarity, it also limits the value of private information, ensuring that a unique threshold information cost exists.

Proposition 2.6.4. ***Information cost distributions and the magnitude of amplification.*** *Suppose the conditions of Proposition 2.6.3 hold. Then the following results for the magnitude of amplification hold uniformly, that is for any given public signal μ:*

1. *If $f(c) > 1$, then the magnitude of amplification is larger than in the case of $c_i \sim U[0,1]$.*

2. *Consider two distributions of the information costs, f^1 and f^2, that satisfy the heterogeneity and slope requirements of Proposition 2.6.3. Let f^2 be a transformation of f^1 in the sense that $c^1_{min} \leq c^2_{min} < D(0) < D(1) < c^2_{max} \leq c^1_{max}$, where the probability mass between $[c^1_{min}, c^2_{min}]$ and $[c^2_{max}, c^1_{max}]$ is evenly distributed over $[c^2_{min}, c^2_{max}]$. Then, the magnitude of amplification for f^2 is higher than for f^1.*

3. *Consider two uniform distributions, $f^3 \sim U[c_{min}, c_{max}]$ and $f^4 \sim U[c_{min}+\varphi, c_{max}-\varphi]$, that satisfy the heterogeneity and slope requirements. Then, the magnitude of amplification for f^4 is higher than for f^3. The difference in the magnitude increases in φ.*

In case (a), if $1 < f(c) < \sqrt{\frac{\pi}{2}}$ (and sufficiently precise private information), changes in the public signal that affect the value of information and the threshold information cost lead to a larger change in the proportion of informed investors, resulting in a larger magnitude of the amplification effect. In cases (b) and (c), a larger mass of the distribution in the relevant range of $[D(0), D(1)]$ implies a larger change in the proportion of informed investors as the threshold information cost changes. Therefore, the magnitude of the amplification effect after a change in the public signal is larger, as shown in Figure 2-10.

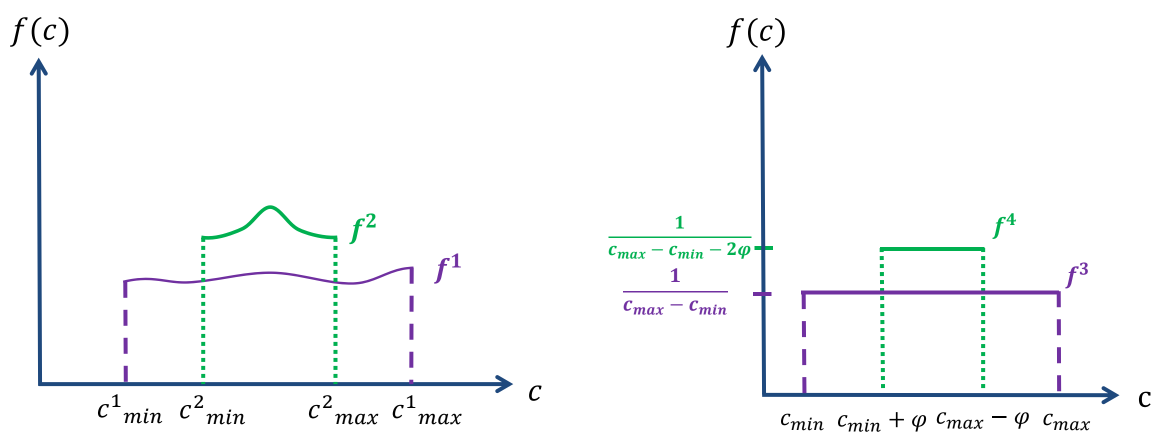

Figure 2-10: The distribution of information costs and the magnitude of amplification. The left panel shows case (b), while the right panel shows case (c).

These results suggest that the magnitude of amplification tends to be lower for more dispersed skills in information acquisition, since a more dispersed distribution of information costs tends to induce smaller changes in the proportion of informed investors.

To test these implications for the magnitude of amplification, one can compare shocks to markets or asset classes that have different distributions of information costs. For example, the dispersion of information costs is higher in over-the-counter markets than in centralized markets. Ang, Shtauber and Tetlock (2013) show that markets in which investors trade over-the-counter are more opaque than centralized markets. Other proxies for the dispersion in information cost are the size of investors (larger investors have lower costs than smaller investors) or the sophistication of investors (retail investors have higher costs than professional investors). Garriott and Walton (2016) provide evidence about the lower informativeness of retail investors. Another proxy could be the (voluntary) disclosure policy of corporate firms and banks, which reduces the level and dispersion of information costs. Lang and Lundholm (1996) show that such disclosures increase the accuracy of earning forecasts of analysts and reduces their dispersion.

2.6.4 Limited precision improvement

Consider next the case of limited precision improvement, $\gamma_I \in (\gamma_U, \infty)$. Investors who acquire information still receive a more precise signal than investors who acquire no information. However, the signal of informed investors is noisy, sometimes resulting in mistakes when (not) attacking the regime. We generalize our result on the existence of a unique equilibrium and show that the amplification result continues to hold.

Proposition 2.6.5. ***Limited precision improvement.*** *Let $\gamma_I \in (\gamma_U, \infty)$. If private information is sufficiently precise, $\gamma_U > \underline{\gamma}$, then there generically exists a unique equilibrium. Amplification via the information choices of investors occurs in this equilibrium.*

Next, we study numerically how the degree of precision improvement affects the magnitude of amplification. Figure 2-11 shows the magnitude of amplification for various precision levels of private information of informed investors. The relationship between the private precision and the magnitude of amplification is monotonic, whereby greater precision raises the magnitude.

To test the implications of different degrees of precision improvement for the magnitude of amplification, the composition of institutional investors can be exploited. Hedge funds may be better at extracting information than mutual or pension funds, so the share of hedge funds among institutional investors is a proxy for the precision of private information. Specifically,

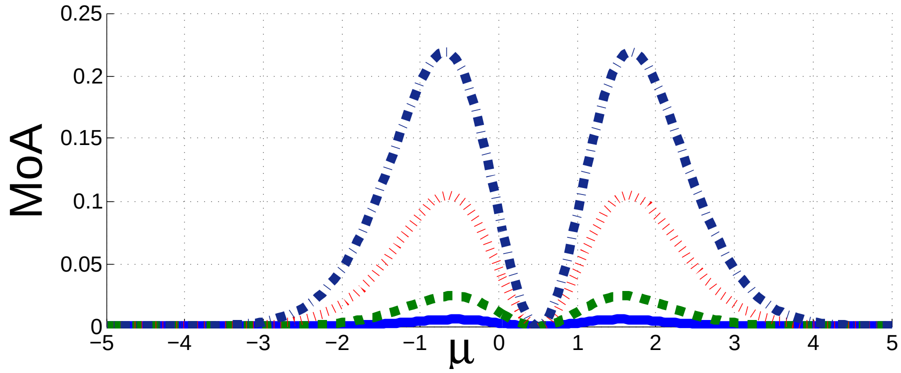

Figure 2-11: The precision of private information and the magnitude of amplification. More precise private information monotonically increases the magnitude of amplification. Parameter values are $b = \ell = 0.75$, $\alpha = 1$, $\gamma_U = 4$, and $\gamma_I \in \{6, 10, 50, \infty\}$.

following a negative shock, our model predicts a larger magnitude of amplification for a larger proportion of hedge fund investors. For the debt of a corporate firm, another proxy could be the industry and the age of the firm (as described above). The improvement in precision is small when much information about the firm is classified or protected (e.g., defense) or when little data or experience exists to evaluate a new business model (e.g., technology). In terms of bank debt, more complex banks and those invested in more opaque assets or with more opaque counterparties allow for smaller precision improvement.

2.7 Extensions

2.7.1 Payoff sensitivity

We have so far considered exogenous payoff parameters. However, Iachan and Nenov (2015) allow the payoffs of investors to depend directly on the fundamental, which allows for a more general link between the proportion of informed investors and the fundamental threshold. To capture this idea in the context of our model, which has both a public signal and endogenous information choice, suppose that $b(\theta) \in (0, 1)$ and $\ell(\theta) \in (0, 1)$, where $b'(\theta) \in (-\infty, 0)$ and $\ell'(\theta) \in (0, \infty)$. The payoff from attacking becomes

$$u(a_i = 1, A, \theta) = b(\theta)\ \mathbf{1}_{\{A \geq \theta\}} - \ell(\theta)\ \mathbf{1}_{\{A < \theta\}}. \tag{7.1}$$

We maintain a uniformly distributed information cost, $c_i \sim \mathcal{U}[0,1]$, and consider limited precision improvement in this section.[16] Some tractability is lost when payoffs are sensitive to fundamentals but we can offer a condition sufficient for amplification to occur.

Proposition 2.7.1. ***Payoff sensitivity and amplification.*** *Let the payoffs be sensitive to the fundamental, whereby $b(\theta) > 0$ and $\ell(\theta) > 0$ with $b'(\theta) < 0$ and $\ell'(\theta) > 0$. If*

$$-g(\overline{\theta})\Big(b(\overline{\theta}) + \ell(\overline{\theta})\Big)\Gamma(\overline{\theta})^2 - \Gamma(\overline{\theta})\Big[\int_{\overline{\theta}}^{\infty} \ell'(\theta)\Gamma(\theta)dG(\theta) - \int_{-\infty}^{\overline{\theta}} b'(\theta)\Gamma(\theta)dG(\theta)\Big] \leq 0, \quad (7.2)$$

amplification via the information choice of investors occurs after changes in the public signal.

In order to obtain more specific results, we study the linear case of $b(\theta) = b_0 - b_1\theta$ and $\ell(\theta) = \ell_0 + \ell_1\theta$, where all coefficients are strictly positive. Since $\overline{\theta} \in (0,1)$, we impose $b_0 > b_1$ to ensure that $b(\theta) > 0$ over the relevant range. Likewise, $\ell_0 > 0$ ensures that $\ell(\theta) > 0$ over the relevant range. We also assume that the slope is identical, $b_1 = \ell_1 \equiv \lambda$.

Proposition 2.7.2. ***A linear case: Existence, uniqueness, and amplification.*** *Consider the linear specification with identical slope coefficient, $b(\theta) = b_0 - \lambda\theta$ and $\ell(\theta) = \ell_0 + \lambda\theta$. If private information is sufficiently precise, $\gamma_U > \underline{\gamma}$, and payoff sensitivity is sufficiently low, $\lambda < \overline{\lambda} > 0$, then there exists a unique equilibrium. Amplification via the information choice of investors occurs in this equilibrium if*

$$b_0 + \ell_0 \geq \frac{\lambda}{g(\overline{\theta})\left|\Gamma(\overline{\theta})\right|}. \quad (7.3)$$

The sensitivity of the payoffs to the fundamental plays a crucial role for both the existence of a unique equilibrium and amplification. We show that a unique equilibrium in the overall game exists as long as the payoffs are not too sensitive to the fundamental. We also offer a simpler condition sufficient for amplification to occur in this equilibrium. Again, an upper bound on the sensitivity of payoffs suffices. These conditions are always met in the baseline case because payoffs are not sensitive to the fundamental, $\lambda = 0$.

In the spirit of Iachan and Nenov (2015), we next consider the special cases in which only one payoff variable is sensitive to the fundamental. These cases may be interpreted as

[16] Iachan and Nenov (2015) also study a threshold function $k(\theta)$, where $A \geq k(\theta)$ is required for regime change. Since the impact of this feature is minor, we set $k(\theta) = \theta$ to focus on the role of payoff sensitivity.

a stylized bank run ($b'(\theta) = 0$) or a currency attack ($\ell'(\theta) = 0$). For simplicity, we continue to consider a constant slope for the sensitive payoff, $\lambda > 0$. Iachan and Nenov (2015) show that more precise private information has different implications for the probability of regime change in these two cases. Therefore, we are interested in analyzing the magnitude of amplification in the cases where only one payoff variable is sensitive to the fundamental.

Proposition 2.7.3. ***One-sided payoff sensitivity.*** *Consider the linear specification in which one payoff coefficient is insensitive to the fundamental, $b(\theta) = b_0 - \lambda\theta$ and $\ell(\theta) = \ell_0$ or $b(\theta) = b_0$ and $\ell(\theta) = \ell_0 + \lambda\theta$. If the private information of an informed investor is sufficiently precise, $\gamma_I > \underline{\gamma}_I$, the magnitude of amplification increases in $\Gamma(\overline{\theta})\, B(\overline{\theta})$, where*

$$B(\overline{\theta}) = \begin{cases} \lambda \int_{\overline{\theta}}^{\infty} \Gamma(\theta) dG(\theta) > 0 & \quad b'(\theta) = 0 \\ & if \\ \lambda \int_{-\infty}^{\overline{\theta}} \Gamma(\theta) dG(\theta) < 0 & \quad \ell'(\theta) = 0 \end{cases} \tag{7.4}$$

Proposition 2.7.3 states that the magnitude of amplification is also different across the different cases of one-sided payoff sensitivity. If informed investors are more likely to attack than uninformed investors, $\Gamma(\overline{\theta}) < 0$, Proposition 2.7.3 suggests that the magnitude of amplification is higher in currency attacks than in bank runs. In contrast, if informed investors are less likely to attack, $\Gamma(\overline{\theta}) > 0$, the magnitude of amplification is higher in bank runs.

2.7.2 Homogeneous information cost

We study an extension with a homogeneous information cost for robustness. Such an information cost structure in a discrete information choice setup yields multiple equilibria (Hellwig and Veldkamp, 2009). We show that our amplification mechanism can obtain in this setup. First, Proposition 2.7.4 states the equilibrium in this environment without proof.

Proposition 2.7.4. ***Homogeneous information costs and multiple equilibria.*** *Consider a homogeneous cost, $c_i \equiv c$. If private information is sufficiently precise, $\gamma_U > \underline{\gamma}$, the information cost determines the number of equilibria in the overall game.*

If the information cost is low, $c < D(0)$, there exists a unique stable equilibrium in which all investors acquire information, $n^* = 1$. Similarly, if $c > D(1)$, there also exists a unique *stable equilibrium without information acquisition, $n^* = 0$. In contrast, for an intermediate information cost, $D(0) \leq c \leq D(1)$, there exist three equilibria. Apart from the previous two equilibria, there is also an asymmetric and unstable equilibrium, in which investors are indifferent in their information choice and the aggregate proportion of informed investors is determined to ensure this indifference of the marginal investor: $n^* = D^{-1}(c)$.*

The optimal behavior of investors at the coordination stage is again uniquely pinned down for a given proportion of informed investors, n^*. It is characterized by the thresholds $\overline{x}_I(n^*)$, $\overline{x}_U(n^*)$, $\overline{\theta}(n^*)$.

Figure 2-12 shows the link between the information cost and the number of equilibria. Given the strict monotonicity of the value of private information in the proportion of informed investors, the latter is uniquely determined in the asymmetric equilibrium. While the amplification effect does not occur in this equilibrium, we exclude the asymmetric equilibrium based on its instability. That is, we focus on the two stable equilibria in which the information choices of investors are symmetric.

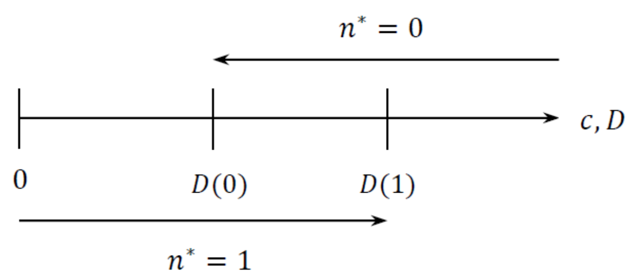

Figure 2-12: A homogeneous information cost: multiple equilibria exist for $D(0) \leq c \leq D(1)$.

We assume a sunspot variable $s \in \{0, 1\}$ with $\Pr\{s = 1\} = q \in (0, 1)$. Whenever both symmetric equilibria exist, investors coordinate on one equilibrium according to the sunspot. That is, if $D(0) \leq c \leq D(1)$, then the equilibrium with information acquisition, $n^* = 1$, occurs with probability q and the equilibrium without information acquisition, $n^* = 0$, occurs with probability $1 - q$. In what follows, we focus on a deteriorating public signal, but the analysis for an improving signal is analogous.

Figure 2-13 shows the case of a strong public signal, where the value of private information increases as the public signal deteriorates (denoted by $\tilde{D}(n)$). There are three areas of inaction with no change in the equilibrium proportion of informed investors and therefore no amplification: $[0, D(0)]$, $[\tilde{D}(0), D(1)]$, and $[\tilde{D}(1), \infty)$. Consider next the range $(D(0), \tilde{D}(0))$, where before the deterioration in the public signal, information acquisition was the equilibrium with probability q, and it is now the unique equilibrium. Therefore, there is no change in the equilibrium proportion of informed investors with probability q but a change, $\Delta n^* = +1$, with probability $1-q$. Hence, with probability $1-q$, amplification occurs, since the probability of a financial crisis increases as more investors are informed for a strong public signal. Likewise for the range $(D(1), \tilde{D}(1))$, where $n^* = 0$ was the unique equilibrium before but $n^* = 1$ is now an equilibrium with probability q. Hence, $\Delta n^* = +1$ and therefore amplification occurs with probability q. The analogous case of a weak public signal yields $\Delta n^* = -1$ and is skipped for brevity.

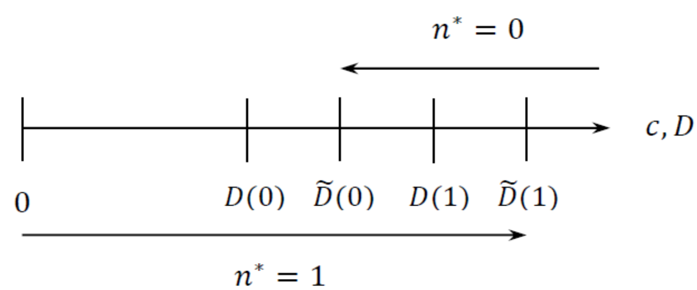

Figure 2-13: A homogeneous information cost and equilibrium selection via a sunspot variable: amplification occurs with positive probability for $D(0) < c < \tilde{D}(0)$ and $D(1) < c < \tilde{D}(1)$.

2.8 Conclusion

We propose an amplification mechanism of financial crises based on the information choice of investors. In a debt rollover game, for instance, an investor wishes to roll over debt whenever the debtor is solvent, which is more likely when other investors also roll over. Adverse news about an initially creditworthy debtor, for example a rating downgrade or an earnings warning, raises the value of private information and more investors acquire information about debtor solvency. In turn, informed investors are more likely to refuse to roll over debt than uninformed investors, which amplifies the probability of a debt crisis.

We show how a policymaker can enhance financial stability. Taxes and subsidies, either on payoffs to investors or on information acquisition, alter the information choice of investors and therefore reduce the probability of a crisis. However, the optimal policy response to deteriorating public information depends on the solvency of the debtor. For an initially creditworthy debtor, taxes reduce the value of private information relative to the effective information cost. By discouraging information acquisition by investors, the probability of a debt crisis is reduced. For an initially less creditworthy debtor, subsidies to encourage information acquisition reduce the probability of a debt crisis. We also study the effects of an improvement in the quality of public information on financial stability.

We derive testable implications about the magnitude of amplification. First, the magnitude of amplification is non-monotone in the public information about debtor solvency, such as credit ratings. Second, we characterize the effect of the distribution of information costs across investors. Amplification is larger the more investors change their information choice as a result of deteriorating public information, which tends to occur for less dispersed distributions. Third, amplification is larger when informed investors have more precise private information. We discuss several environments in which these implications can be tested.

Appendix

2.9 Derivations and Proof of Proposition 2.3.1

We construct the equilibrium by working backwards. In section 2.9.1, we study the optimal attack behavior of investors at the coordination stage for any given information choices $\{n_i^*\}$, and the optimal information choice of investors in section 2.9.2. We state all conditions for a general precision of informed investors, γ_I, to use these conditions also for the case of limited precision improvement. We also state specific results for vanishing private noise, $\gamma_I \to \infty$.

2.9.1 Coordination stage

Given the information choices $\{n_i^*\}$, the global coordination game is standard. An investor's expected utility from attacking conditional on the public signal μ, the private signal x_i, the information choice n_i^*, and the aggregate proportion of informed investors n^* is:

$$E[u(a_i = 1)|n_i^*, n^*; \mu, x_i] = -\ell + (b + \ell) \Pr[A > \theta|n_i^*, n^*; \mu, x_i]. \tag{9.1}$$

Optimality for investor i at the coordination stage requires that his strategy maximizes the conditional expected utility, taking all other investors' strategies as given. Since each investor is atomistic, the aggregate attack size is unaffected by the individual attack decision.

Without loss of generality, we focus on symmetric monotone equilibria at the coordination stage throughout (Morris and Shin (2003); Frankel, Morris and Pauzner, 2003). For a given proportion of informed investors n^*, the equilibrium is fully characterized by an **signal threshold** for informed and uninformed investors, $\overline{x}_I(n^*;\mu)$ and $\overline{x}_U(n^*;\mu)$, and an **fundamental threshold**, $\overline{\theta}(n^*;\mu)$. Investor i attacks the regime if and only if the private signal is below an signal threshold specific to her information choice $n_i^* = z \in \{I, U\}$:

$$a_i^* = 1 \Leftrightarrow x_i \leq \overline{x}(n_i^* = z, n^*; \mu) \equiv \overline{x}_z(n^*;\mu). \tag{9.2}$$

Regime change occurs whenever the realized fundamental is below the fundamental threshold:

$$\theta < \overline{\theta}(n^*;\mu). \tag{9.3}$$

These thresholds are determined by a critical mass condition at the aggregate level and indifference conditions at the individual level. An investor i uses both signals to form a posterior about the unobserved fundamental, where normality is preserved, the posterior mean is a weighted average of the public signal and the private signal, and the posterior precision is the sum of the precisions of the public and private signals (DeGroot, 1970):

$$h_z(\theta, x_i) : \theta | n_i^* = z; \mu, x_i \sim \mathbb{N}\left(\frac{\alpha\mu + \gamma_z x_i}{\alpha + \gamma_z}, \frac{1}{\alpha + \gamma_z}\right), \tag{9.4}$$

where we use the distribution $h_z \equiv \sqrt{\alpha + \gamma_z}\,\phi\left(\sqrt{\alpha+\gamma_z}\left[\theta - \frac{\alpha\mu+\gamma_z x_i}{\alpha+\gamma_z}\right]\right)$ extensively in what follows. Each investor assigns the following probability to regime change, $\Pr\{\theta \leq \overline{\theta} | n_i^* = z, n^*; \mu, x_i\} = \Phi\left(\sqrt{\alpha+\gamma_z}\left[\overline{\theta} - \frac{\alpha\mu+\gamma_z x_i}{\alpha+\gamma_z}\right]\right)$. An investor with information choice $n_i^* = z$ who receives the threshold signal $x_i = \overline{x}_z(n^*;\mu)$ is indifferent between attacking and not attacking.

This **indifference condition** states that the probability of regime change evaluated at the fundamental threshold equals the conservativeness ratio for both informed and uninformed investors, $\Pr\{\theta \leq \overline{\theta} | n^*, n_i^* = z; \mu, x_i = \overline{x}_z\} \equiv \kappa$, and it yields the signal thresholds:

$$\overline{x}_z(n^*;\mu) = \overline{\theta}(n^*;\mu) + \frac{\alpha}{\gamma_z}\Big[\overline{\theta}(n^*;\mu) - \mu\Big] - \frac{\sqrt{\alpha+\gamma_z}}{\gamma_z}\Phi^{-1}(\kappa). \tag{9.5}$$

Since all investors play the threshold strategy with thresholds $\overline{x}_I(n^*;\mu)$ if informed, and $\overline{x}_U(n^*;\mu)$ if uninformed, the aggregate attack size for any fundamental θ is:

$$\begin{aligned} A(n^*;\theta,\overline{x}_I,\overline{x}_U) &= \int_0^1 \mathbf{1}\{x_i \leq \overline{x}_z | n_i^* = z, \theta\} di \equiv n^*\Phi_I(\theta) + (1-n^*)\Phi_U(\theta) \\ &= n^*\Phi\left(\sqrt{\gamma_I}\left[\overline{x}_I - \theta\right]\right) + (1-n^*)\Phi\left(\sqrt{\gamma_U}\left[\overline{x}_U - \theta\right]\right). \end{aligned}$$

The **critical mass condition** states that the aggregate attack size is just sufficient for regime change when the fundamental equals the fundamental threshold:

$$\overline{\theta}(n^*;\mu) = A\left(n^*; \overline{\theta}(n^*;\mu), \overline{x}_I(n^*;\mu), \overline{x}_U(n^*;\mu)\right). \tag{9.6}$$

Inserting the indifference conditions in the critical mass condition yields the fundamental threshold for any given proportion of informed investors, $\overline{\theta} = \overline{\theta}(n^*; \mu)$, implicitly defined by:

$$\begin{aligned}\overline{\theta} &= n^*\Phi_I(\overline{\theta}) + (1-n^*)\Phi_U(\overline{\theta}) \\ &= n^*\Phi\left(\frac{\alpha[\overline{\theta}-\mu]}{\sqrt{\gamma_I}} - \sqrt{1+\frac{\alpha}{\gamma_I}}\Phi^{-1}(\kappa)\right) + (1-n^*)\Phi\left(\frac{\alpha[\overline{\theta}-\mu]}{\sqrt{\gamma_U}} - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right).\end{aligned} \tag{9.7}$$

A unique solution to equation (9.7) for any given proportion of informed investors n^* is ensured by a sufficiently precise private signal of the uninformed investor, $\gamma_U > \underline{\gamma}' \equiv \frac{\alpha^2}{2\pi}$ (Morris and Shin (2003)). Under this condition, the slope of the left-hand side of equation (9.7) exceeds the slope of the right-hand side, $1 > A_{\overline{\theta}} \equiv \frac{\partial A}{\partial \overline{\theta}} > 0$, which is evaluated at the equilibrium values $(\overline{\theta}, n^*)$. To see this, observe that

$$A_{\overline{\theta}} = n^*\frac{\alpha}{\sqrt{\gamma_I}}\phi\Big(\frac{\alpha[\overline{\theta}-\mu]}{\sqrt{\gamma_I}} - \sqrt{1+\frac{\alpha}{\gamma_I}}\Phi^{-1}(\kappa)\Big) + (1-n^*)\frac{\alpha}{\sqrt{\gamma_U}}\phi\Big(\frac{\alpha[\overline{\theta}-\mu]}{\sqrt{\gamma_U}} - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\Big).$$

Since $\phi(\cdot) \leq \frac{1}{\sqrt{2\pi}}$ and $\gamma_U < \gamma_I$, we have $A_{\overline{\theta}} \leq \frac{\alpha}{\sqrt{2\pi\gamma_U}}$, resulting in the stated lower bound $\underline{\gamma}'$. Thus, there exists at most one solution. Since $A \in [0,1]$ and $A_{\overline{\theta}} > 0$, $\forall\, n^*$, there exists a unique fixed point $\overline{\theta}(n^*; \mu)$ in the interval $[0,1]$. Once the unique fundamental threshold is obtained, the signal thresholds $\overline{x}_z(n^*; \mu)$ are backed out from the indifference conditions.

In the limiting case of vanishing private noise of informed investors, $\gamma_I \to \infty$, informed investors base their posterior about the unobserved fundamental completely on their private signal, whereby $f^I(x|\theta)$ is the Dirac delta function at $x = \theta$, and the signal threshold of informed investors reduces to the fundamental threshold, $\overline{x}_I(n^*) \to \overline{\theta}(n^*)$. As a result, an informed investor attacks if and only if regime change occurs, whereby $\Phi_I(\theta) = \mathbf{1}_{\theta \leq \overline{\theta}}$ is a step function. Hence, the fundamental threshold converges to:

$$\overline{\theta} \to n^*(1-\kappa) + (1-n^*)\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta}-\mu\right] - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right). \tag{9.8}$$

Lemma 2.9.1 summarizes the responsiveness of the fundamental threshold to changes in the proportion of informed investors. It states a condition sufficient for a link between the coordination and information stages. This condition is generically satisfied.

Lemma 2.9.1. *Suppose $\gamma_U > \underline{\gamma}'$. If $\mu \neq \hat{\mu}$, then the fundamental threshold at the coordination stage responds to changes in the proportion of informed investors:*

$$\frac{d\overline{\theta}}{dn^*} \neq 0. \tag{9.9}$$

Furthermore, the fundamental threshold increases (decreases) in the proportion of informed investors if the public signal is strong (weak):

$$\frac{d\overline{\theta}}{dn^*}\Big(\mu - \hat{\mu}\Big) > 0. \tag{9.10}$$

Proof. If the private information of uninformed investors is sufficiently precise, then $\overline{\theta}(n^*;\mu)$ is unique for any $n^* \in [0,1]$. The proof is in three steps. First, differentiating the fundamental threshold with respect to the proportion of informed investors yields

$$\begin{aligned}\frac{d\overline{\theta}}{dn^*} &= \frac{\Phi\left(\frac{\alpha}{\sqrt{\gamma_I}}\left[\overline{\theta}-\mu\right]-\sqrt{1+\frac{\alpha}{\gamma_I}}\Phi^{-1}(\kappa)\right)-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta}-\mu\right]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)}{1-A_{\overline{\theta}}} \qquad (9.11)\\ &\to \frac{1-\kappa-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta}-\mu\right]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)}{1-A_{\overline{\theta}}},\end{aligned}$$

for $\gamma_I \to \infty$. Second, recall that we have $0 < A_{\overline{\theta}} < 1$, where the second inequality follows from the sufficient condition for uniqueness at the coordination stage, $\gamma_U > \underline{\gamma}'$. Third, $\mu \neq \hat{\mu}$ ensures a non-zero numerator of the derivative $\frac{d\overline{\theta}}{dn^*}$. This can be proven by contradiction: suppose that the numerator is zero. Since $0 = \Phi^{-1}(\kappa) + \Phi^{-1}(1-\kappa)$ for all $\kappa \in (0,1)$, we have $\hat{\theta} \equiv \mu + \frac{\sqrt{\alpha+\gamma_U}-\sqrt{\gamma_U}}{\alpha}\Phi^{-1}(\kappa)$. Inserting $\hat{\theta}$ in the defining equation of the fundamental threshold, equation (9.7), yields $\mu = \hat{\mu}$. Using the same argument, the numerator is positive (negative) if and only if the public signal is strong (weak). ∎

2.9.2 Information stage

Next, we evaluate the incentive of an investor to acquire information. This is achieved by comparing the expected utility of an informed investor (EU^I) with that of an uninformed investor (EU^U), as defined in equation (3.10). The value of private information stated in equation (3.11) is expressed using the difference in the probability of attacking between an uninformed and informed investor, $\Gamma(\theta, \overline{x}_I, \overline{x}_U)$. This difference is generically non-zero since informed investors receive more precise information than uninformed investors and therefore use a different signal threshold. For $b = \ell$, for example, $\overline{x}_I \neq \overline{x}_U$ whenever $\mu \neq \frac{1}{2}$. If the information advantage of informed investors vanishes, then both types of investors invest with the same probability conditional on the fundamental, that is $\Gamma(\theta) \to 0$ as $\gamma_U \to \gamma_I$.

For the limiting case of vanishing private noise, $\gamma_I \to \infty$, we have:

$$D \to \ell \int_{\overline{\theta}}^{\infty} \Phi_U(\theta)\, dG(\theta) + b \int_{-\infty}^{\overline{\theta}} \left[1-\Phi_U(\theta)\right] dG(\theta). \qquad (9.12)$$

Lemma 2.9.2. *If $\gamma_U > \underline{\gamma}'$, there is strategic complementarity in information choices:*

$$\frac{dD}{dn^*} = (b+\ell)\left[1-A_{\overline{\theta}}\right] g(\overline{\theta}) \left(\frac{d\overline{\theta}}{dn^*}\right)^2 \geq 0, \qquad (9.13)$$

with strict inequality if $\mu \neq \hat{\mu}$. Furthermore, the more-restrictive lower bound on the precision *of private information, $\gamma_U > \underline{\gamma}$, ensures that:*

$$\frac{dD}{dn^*} < 1. \tag{9.14}$$

Proof. First, $\gamma_U > \underline{\gamma}'$, so $\overline{\theta}(n^*;\mu)$ is unique for any n^*. Total differentiation yields $\frac{dD}{dn^*} = \frac{\partial D}{\partial \overline{\theta}}\frac{d\overline{\theta}}{dn^*} + \frac{\partial D}{\partial \overline{x}_U}\frac{d\overline{x}_U}{d\overline{\theta}}\frac{d\overline{\theta}}{dn^*}$. We prove below that $\frac{\partial D}{\partial \overline{x}_z} \equiv D_{\overline{x}_z} = 0$. By Leibniz rule, we have that $\frac{\partial D}{\partial \overline{\theta}} = g(\overline{\theta})[b-(b+l)\Phi_U(\overline{\theta})]$. Inserting $\overline{x}_U$ from equation (9.5) and rewriting Lemma 2.9.1 to obtain an expression for $\Phi_U(\overline{\theta})$, we have:

$$\frac{\partial D}{\partial \overline{\theta}} \equiv D_{\overline{\theta}} = (b+\ell)(1-A_{\overline{\theta}})\frac{d\overline{\theta}}{dn^*}g(\overline{\theta}). \tag{9.15}$$

Therefore, by taking together, we can state $\frac{dD}{dn^*} = (b+\ell)\left[1-A_{\overline{\theta}}\right]g(\overline{\theta})\left(\frac{d\overline{\theta}}{dn^*}\right)^2 \geq 0$. This derivative has four terms. The first three terms are strictly positive because $b>0$, $\ell>0$, $A_{\overline{\theta}}>0$, $A_{\overline{\theta}}<1$ if $\gamma_U > \underline{\gamma}'$, and $g>0$ (pdf of a standardized Gaussian). The fourth term is a square and thus non-negative. It is strictly positive if $\mu \neq \hat{\mu}$ by Lemma 2.9.1. Thus $\mu \neq \hat{\mu}$ suffices for the derivative to be strictly positive.

Next, $\frac{\partial D}{\partial \overline{x}_U} = 0$ by an envelope theorem argument. The threshold $\overline{x}_U$ is chosen by a first-order condition that balances the marginal cost of attacking absent regime change with the marginal benefit of attacking under regime change (see Appendix 2.10.1 for more details):

$$\frac{\partial D}{\partial \overline{x}_U} = -b\int_{-\infty}^{\overline{\theta}} f^U(\overline{x}_U|\theta)g(\theta)d\theta + \ell\int_{\overline{\theta}}^{\infty} f^U(\overline{x}_U|\theta)g(\theta)d\theta = 0. \tag{9.16}$$

Using Lemma 2.9.1, the derivative of the value of private information with respect to the proportion of informed investors is:

$$\frac{dD}{dn^*} = \frac{(b+\ell)g(\overline{\theta})\left[1-\kappa-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta}-\mu\right]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)\right]^2}{1-A_{\overline{\theta}}}, \tag{9.17}$$

where $g \leq \frac{1}{\sqrt{2\pi}}$, $b+\ell < 2$, and $[1-\kappa-\Phi(\cdot)]^2 \leq \max\{\kappa^2,(1-\kappa)^2\} < 1$. Next, $A_{\overline{\theta}} = (1-n^*)\,\phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\left[\overline{\theta}-\mu\right]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)\frac{\alpha}{\sqrt{\gamma_U}} \leq \frac{\alpha}{\sqrt{2\pi}\sqrt{\gamma_U}}$ since $n^* \geq 0$. Hence, $\frac{dD}{dn^*} < 1$ is ensured by $\gamma_U > \underline{\gamma} > \underline{\gamma}'$. ∎

Since $\mu \neq \hat{\mu}$ excludes a parameter space of zero measure, the value of private information increases strictly in the proportion of informed investors generically.

Lemma 2.9.3. *If $\gamma_U > \underline{\gamma}'$, then the value of private information satisfies $D \in (0,1)$.*

Proof. First, we show $D(n^*) > 0$ for all $n^* \in [0,1]$. Formally, $F^U(\theta) \in (0,1)$ and $g(\theta) > 0$ for all $\theta \in (-\infty,\infty)$ and $b>0$ and $\ell>0$, so $D>0$. Intuitively, since $\gamma_U < \infty$, there is positive probability mass on the type-I and type-II errors of an uninformed investor. Second, we show $D(n^*) < 1$ for all n^*. Since $F^U(\theta) \leq 1$ for all $\theta \in [\overline{\theta}, \infty)$ and $F^U(\theta) \geq 0$ for all $\theta \in (-\infty, \overline{\theta}]$, we have $D \leq \ell + (b - \ell)G(\overline{\theta})$. Since $G(\overline{\theta}) \in [0, 1]$, we have that $D \leq \max\{b, \ell\} < 1$. ∎

We are now ready to complete the construction of equilibrium. Consider an investor's optimal information choice. Given the binary action, an investor acquires information whenever the individual information cost is no larger than the value of private information:

$$n_i^* = I \Leftrightarrow c_i < D(n^*). \tag{9.18}$$

Since each investor is atomistic and has no effect on the aggregate proportion of informed investors, the value of private information depends on the proportion of informed investors only. The perfect Bayesian equilibrium is constructed by combining these individual optimality conditions, Lemmas 2.9.2 and 2.9.3, and the consistency between individually optimal information acquisition choices and the aggregate proportion of informed investors.

The optimal information choice is characterized by a threshold strategy with some $\overline{c}$. An investor acquires information if and only if the information cost is below the threshold, $c_i < \overline{c}$, so the proportion of informed investors is $n^* = \overline{c}$. The marginal investor is indifferent between the information choices, $c_i = \overline{c}$, so this threshold is any solution of

$$\overline{c} = D(\overline{c}). \tag{9.19}$$

Uniqueness requires that $D(\overline{c})$ has exactly one fixed point. First, the left-hand side of this equation is continuous, within $[0, 1]$, and has a unit slope. Second, the right-hand side is continuous and strictly positive by Lemma 2.9.3, and its slope lies strictly within $(0, 1)$ by Lemma 2.9.2. Therefore, if a solution exists, it is unique. Next, $D < 1$ by Lemma 2.9.3 ensures existence. This completes the proof of the existence of a unique interior solution $\overline{c}$.

2.10 Amplification and Proof of Proposition 2.4.1

We prove the amplification result in three steps. The first two steps do not resort to the limiting case of vanishing private noise of informed investors. As a result, these two steps directly apply to the case of limited precision improvement studied in section 2.6.4.

2.10.1 Derivation of equation (4.3) and intermediate steps

We derive the total effect of changes in the public signal on the fundamental threshold, $\frac{d\overline{\theta}}{d\mu}$. The equilibrium is given by the following equations:

$$\begin{aligned} \overline{\theta} &= A(n^*, \overline{\theta}, \mu), \\ n^* &= \overline{c} = D(\overline{\theta}, \mu, \overline{x}_z), \end{aligned} \tag{10.1}$$

where the value of private information depends on the proportion of informed investors only indirectly. As for notation, m_x denotes the partial derivative $\frac{\partial m}{\partial x}$. Total differentiation yields:

$$\frac{d\overline{\theta}}{d\mu} = A_{n^*}\frac{dn^*}{d\mu} + A_{\overline{\theta}}\frac{d\overline{\theta}}{d\mu} + A_{\mu},$$
$$\frac{dn^*}{d\mu} = D_{\overline{\theta}}\frac{d\overline{\theta}}{d\mu} + D_{\mu}, \tag{10.2}$$

since $D_{\bar{x}_z} = 0$ for $z \in \{I, U\}$ (see equation (9.16)). Rewriting yields equation (4.3). To evaluate the total derivative in equation (4.3), we first obtain the partial derivatives of the aggregate attack size, which are all evaluated at the equilibrium quantities $(n^*, \overline{\theta})$:

$$A_{\mu} = -\frac{\alpha}{\sqrt{\gamma_I}}n^*\phi_I(\overline{\theta}) - (1-n^*)\frac{\alpha}{\sqrt{\gamma_U}}\phi_U(\overline{\theta}) < 0 \tag{10.3}$$
$$A_{n^*} = \Phi_I(\overline{\theta}) - \Phi_U(\overline{\theta}) = [1 - A_{\overline{\theta}}]\frac{d\overline{\theta}}{dn^*}, \tag{10.4}$$

where $\phi_z(\theta)$ is the probability density function associated with $\Phi_z(\theta)$.

Next, we determine the partial derivatives associated with the value of private information D. A change in the public signal μ affects it via (i) the fundamental threshold $\overline{\theta}$, (ii) the signal thresholds $\overline{x}_z$, and (iii) the distribution of fundamentals $g(\theta)$. First, for the fundamental threshold, we have already derived $D_{\overline{\theta}}$ in the proof of Lemma 2.9.2. Note that equation (9.15) generalizes to the case of limited precision improvement. Second, for the signal thresholds, we have already shown that $\frac{\partial D}{\partial \overline{x}_z} = 0$. Third, we consider the distribution of fundamentals. We use the following result (e.g., Bromiley, 2003).

Corollary 2.10.1. *The product of two normal probability density functions is a scaled normal probability density function. That is, if $f \sim \mathbb{N}(\mu_f, \sigma_f^2)$ and $g \sim \mathbb{N}(\mu_g, \sigma_g^2)$, then $f \circ g \sim S \times \mathbb{N}(\mu_{fg}, \sigma_{fg}^2)$, where*

$$\sigma_{fg} = \sqrt{\frac{\sigma_f^2\sigma_g^2}{\sigma_f^2+\sigma_g^2}}, \quad \mu_{fg} = \frac{\mu_f\sigma_g^2 + \mu_g\sigma_f^2}{\sigma_f^2+\sigma_g^2} \quad \textit{and} \quad S = \frac{1}{\sqrt{2\pi(\sigma_f^2+\sigma_g^2)}}\exp\left(-\frac{(\mu_f-\mu_g)^2}{2(\sigma_f^2+\sigma_g^2)}\right).$$

To obtain the partial effect via a change in the distribution, D_μ, note that $\frac{\partial g}{\partial \mu} = \alpha(\theta - \mu)g(\theta) = -\frac{\partial g}{\partial \theta}$. The partial differentiation of the value of private information is:

$$D_{\mu} = \ell\int_{\overline{\theta}}^{\infty}\Big[\Phi_U(\theta) - \Phi_I(\theta)\Big]\alpha(\theta-\mu)g(\theta)d\theta - b\int_{-\infty}^{\overline{\theta}}\Big[\Phi_U(\theta) - \Phi_I(\theta)\Big]\alpha(\theta-\mu)g(\theta)d\theta. \tag{10.5}$$

We proceed by integrating by parts. We set $u(\theta) \equiv \Phi_U(\theta) - \Phi_I(\theta)$ and $v'(\theta) \equiv \alpha(\theta - \mu)g(\theta)$, so $v(\theta) = -g(\theta)$ and $u'(\theta) = f^U(\overline{x}_U) - f^I(\overline{x}_I)$. From Result 2.10.1, $f^z(\overline{x}_z)(\theta)g(\theta) = S_z\, h_z(\theta)$, where $S_z \equiv \sqrt{\frac{\alpha\gamma_z}{\alpha+\gamma_z}}\phi\left(\sqrt{\frac{\alpha\gamma_z}{\alpha+\gamma_z}}[\mu - \overline{x}_z]\right) > 0$ is a constant and h_z is the probability density function of a normal random variable $\theta \sim \mathbb{N}\left(\frac{\alpha\mu+\gamma_z\overline{x}_z}{\alpha+\gamma_z}, \frac{1}{\alpha+\gamma_z}\right)$, with associated cumulative distribution function H_z. Using the first-order condition for the signal threshold $\overline{x}_z$, which can be written as $\kappa = H_z(\overline{\theta})$, the partial derivative simplifies to:

$$D_\mu = (b+\ell)\, g(\overline{\theta}) \Big[\Phi_U(\overline{\theta}) - \Phi_I(\overline{\theta})\Big]. \tag{10.6}$$

2.10.2 Proving the inequality

Equipped with these preliminaries, we can now prove the inequality that constitutes the amplification result. Specifically, we need to show that

$$-\frac{A_\mu + A_{n^*} D_\mu}{1 - A_{\overline{\theta}} - A_{n^*} D_{\overline{\theta}}} > -\frac{A_\mu}{1 - A_{\overline{\theta}}}. \tag{10.7}$$

To simplify this expression, in a first step we show that both denominators are positive. Consider $\gamma_I \to \infty$. As shown in the proof of Lemma 2.9.1, the denominator on the right-hand side of inequality (10.7) is positive if $\gamma_U > \underline{\gamma}'$. Using the expression for A_{n^*}, we rewrite the denominator on the left-hand side of inequality (10.7):

$$\begin{aligned} 1 - A_{\overline{\theta}} - A_{n^*} D_{\overline{\theta}} &= 1 - A_{\overline{\theta}} - (1 - A_{\overline{\theta}}) \frac{d\overline{\theta}}{dn^*} D_{\overline{\theta}} \\ &= (1 - A_{\overline{\theta}}) \left(1 - D_{\overline{\theta}} \frac{d\overline{\theta}}{dn^*}\right) \\ &= (1 - A_{\overline{\theta}}) \left(1 - \frac{dD}{dn^*}\right) > 0, \end{aligned} \tag{10.8}$$

where the last line follows from the proof of Lemma 2.9.2 and the last inequality follows from equation (9.14), which requires $\gamma_U > \underline{\gamma}$ and $\mu \neq \hat{\mu}$. In sum, both denominators are positive.

In a second step, we rewrite the inequality, using $A_{\overline{\theta}} = -A_\mu$ and $D_{\overline{\theta}} = -D_\mu$. Hence, amplification obtains if and only if

$$A_{n^*} D_\mu = -(b+\ell)\, g(\overline{\theta})\, (1 - A_{\overline{\theta}})^2 \left(\frac{d\overline{\theta}}{dn^*}\right)^2 < 0, \tag{10.9}$$

which holds generically, that is if and only if $\mu \neq \hat{\mu}$. To see this, observe that $b + \ell > 0$, $g > 0$, $A_{\overline{\theta}} \in (0,1)$ if $\gamma_U > \underline{\gamma}'$, and $\left(\frac{d\overline{\theta}}{dn^*}\right)^2 \geq 0$, with strict inequality if $\mu \neq \hat{\mu}$ by Lemma 2.9.1. Collecting the sufficient conditions, we require $\gamma_U > \underline{\gamma}$ and $\mu \neq \hat{\mu}$. This completes the proof of the amplification result regarding the fundamental threshold. Due to equation (4.1), it extends directly to the probability of a financial crisis.

2.11 Policy

2.11.1 Tax or subsidy on information acquisition

To prove Proposition 2.5.1, note that the equilibrium in the case of of tax or subsidy on information choice is given by the following two equations:

$$\begin{aligned} \overline{\theta} &= A(n^*,\overline{\theta}) = n^*\Phi_I(\overline{\theta}) + (1-n^*)\Phi_U(\overline{\theta}) && (11.1)\\ n^*(1-\tau) &= \overline{c} = D(\overline{\theta}). && (11.2)\end{aligned}$$

Note that the tax or subsidy only enters on the left-hand side of the second equation. As a result, total differentiation with respect to τ yields equation (5.1). Finally, we note that $\Gamma(\overline{\theta}) < 0$ if and only if the public signal is strong, as shown in Lemma 2.9.1.

2.11.2 Improving the quality of public information

A higher precision of public information, $d\alpha > 0$, affects the equilibrium conditions at both the coordination stage and the information stage. Total differentiation yields equation (5.2). Since the denominator of $\frac{d\overline{\theta}}{d\alpha}$ is positive, we need to evaluate its numerator. First, for $\kappa = \frac{1}{2}$, we have

$$A_\alpha = (\overline{\theta} - \mu)\left(\frac{n^*}{\sqrt{\gamma_I}}\phi_I(\overline{\theta}) + \frac{1-n^*}{\sqrt{\gamma_U}}\phi_U(\overline{\theta})\right), \tag{11.3}$$

whose sign only depends on $(\overline{\theta} - \mu)$, since the second term is positive. Thus, $A_\alpha < 0$ if and only if $\overline{\theta} < \mu$, which holds if and only if $\mu > \frac{1}{2} = \hat{\mu}(\kappa = \frac{1}{2})$. Second, $A_{n^*} = -\Gamma(\overline{\theta}) > 0$ if and only if $\mu > \frac{1}{2}$. Thus, if $D_\alpha < 0$, then the claim $\frac{d\overline{\theta}}{d\alpha}(\mu - \hat{\mu}) < 0$ follows (generically).

Thus, it remains to show that $D_\alpha < 0$. The value of information depends on the precision of public information only via $g(\theta)$. Note that $2\frac{dg}{d\alpha} = g(\theta)\left[\frac{1}{\alpha} - (\theta-\mu)^2\right] \equiv m(\theta)$. One can prove, for example by differentiation, that $M(\theta) \equiv \int m(\theta)d\theta = g(\theta)\frac{\theta-\mu}{\alpha}$. Thus:

$$2D_\alpha = \ell\int_{\overline{\theta}}^{\infty}\Gamma(\theta)m(\theta)d\theta - b\int_{-\infty}^{\overline{\theta}}\Gamma(\theta)m(\theta)d\theta \tag{11.4}$$

Using partial integration and Result 2.10.1, we observe that

$$\begin{aligned}\int\Gamma(\theta)m(\theta)d\theta &= \Gamma(\theta)M(\theta) - \frac{S_U}{\alpha}\int h_U(\theta)(\theta-\mu)d\theta + \frac{S_I}{\alpha}\int h_I(\theta)(\theta-\mu)d\theta && (11.5)\\ &= \Gamma(\theta)M(\theta) - \frac{S_U}{\alpha}\left[\frac{\gamma_U(\overline{x}_U-\mu)}{\alpha+\gamma_U}H_U(\theta) - \frac{h_U(\theta)}{\alpha+\gamma_U}\right] + \frac{S_I}{\alpha}\left[\frac{\gamma_I(\overline{x}_I-\mu)}{\alpha+\gamma_I}H_I(\theta) - \frac{h_I(\theta)}{\alpha+\gamma_I}\right],\end{aligned}$$

where h_z is evaluated at $\overline{x}_z$ and we define $\breve{\mu}_z \equiv \frac{\alpha\mu+\gamma_z\overline{x}_z}{\alpha+\gamma_z}$ as the mean of h_z, which allows us to rewrite and simplify the integral as follows: $\int h_z(\theta)(\theta-\mu)d\theta = \int h_z(\theta)(\theta-\breve{\mu}_z)d\theta + (\breve{\mu}_z -$

$\mu)H_z(\theta) = -\frac{h_z(\theta)}{\alpha+\gamma_z} + \frac{\gamma_z(\overline{x}_z-\mu)}{\alpha+\gamma_z}H_z(\theta)$. (Recall that $\int y\phi(y)dy = -\phi(y)$.) It follows that:

$$2D_\alpha = -(b+\ell)\left(\Gamma(\overline{\theta})m(\overline{\theta}) + \frac{S_U}{\alpha}\frac{h_U(\overline{\theta})}{\alpha+\gamma_U} - \frac{S_I}{\alpha}\frac{h_I(\overline{\theta})}{\alpha+\gamma_I}\right). \tag{11.6}$$

By continuity, there exists a $\underline{\gamma}_I$ such that $D_\alpha < 0$ for all $\gamma_I > \underline{\gamma}_I$. To see this, note that the third term of equation (11.6) vanishes as $\gamma_I \to \infty$. Also note that $\Gamma(\overline{\theta})m(\overline{\theta}) > 0$ (generically), since $\Gamma(\overline{\theta}) < 0$ and $m(\overline{\theta}) < 0$ for $\mu > \frac{1}{2}$ and $\Gamma(\overline{\theta}) > 0$ and $m(\overline{\theta}) > 0$ for $\mu < \frac{1}{2}$.

2.12 Magnitude of amplification

2.12.1 Proof of Propositions 2.6.1 and 2.6.2

Using equations 4.2 and 4.3, the magnitude of amplification (MoA) is given by:

$$\begin{aligned} \text{MoA}+1 &= \frac{\frac{d\bar{\theta}}{d\mu}\Big|_{n^*,\bar{\theta}}}{\frac{d\tilde{\theta}}{d\mu}\Big|_{\overline{n},\bar{\theta}}} = \frac{(A_\mu + D_\mu A_{n^*})(1-A_{\bar{\theta}})}{(1-A_{\bar{\theta}}-A_{n^*}D_{\bar{\theta}})A_\mu} = \frac{(A_\mu + D_\mu A_{n^*})}{(1-\frac{dD}{dn^*})A_\mu} && (12.1)\\ &= \frac{(A_{\bar{\theta}} - D_\mu A_n)}{(1-\frac{dD}{dn^*})A_{\bar{\theta}}} = \frac{A_{\bar{\theta}} + (1-A_{\bar{\theta}})\frac{dD}{dn^*}}{(1-\frac{dD}{dn^*})A_{\bar{\theta}}} = 1 + \frac{\frac{dD}{dn^*}}{(1-\frac{dD}{dn^*})A_{\bar{\theta}}}, && (12.2)\end{aligned}$$

since $A_\mu = -A_{\bar{\theta}}$ and $-D_\mu A_{n^*} = (b+\ell)g(\bar{\theta})(1-A_\theta)^2(\frac{d\bar{\theta}}{dn^*})^2 = (1-A_{\bar{\theta}})\frac{dD}{dn^*}$.

Public signal Consider first $\gamma_I \to \infty$. It follows that $\text{MoA} = \frac{\frac{dD}{dn^*}}{(1-\frac{dD}{dn^*})}\frac{1}{A_{\bar{\theta}}}$, which increases in $\frac{dD}{dn^*}$ and decreases in $A_{\bar{\theta}}$. We have $MoA(\mu = \hat{\mu}) = 0$ because of Lemma 2.9.1 but $MoA > 0$ for all $\mu \in (-\infty, \hat{\mu})$ or $\mu \in (\hat{\mu}, \infty)$. Moreover, $\lim_{\mu\to+\infty} MoA(\mu) = 0 = \lim_{\mu\to-\infty} MoA(\mu)$. This establishes the non-monotonicity in the public signal. Finally, all of these results also hold for limited precision improvement, where $\hat{\mu}$ is replaced by $\tilde{\mu}$ (see also section 2.6.4).

2.12.2 Lower bound on MoA

Next, we characterize a lower bound on the magnitude of amplification. We first consider the case of vanishing private noise and subsequently study the case of limited precision improvement. In the latter case, the lower bound increases in the precision of private information of informed investors.

Proposition 2.12.1. ***Lower bounds on the magnitude of amplification.***

1. *If $\kappa > \frac{1}{2}$ and $\mu > \bar{\theta}$ (or $\kappa < \frac{1}{2}$ and $\mu < \bar{\theta}$), then*

$$0 < \frac{\zeta\sqrt{2\pi\gamma_U}}{\alpha(1-\zeta)} < MoA, \tag{12.3}$$

where $\zeta \equiv (b+\ell)\left[\Phi\left(\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - \kappa\right]^2 \min\{g(0), g(1)\}$.

2. *If $\kappa = \frac{1}{2}$ and $1 < |\bar{\theta} - \mu|$, then we have two cases:*

   (i) For vanishing private noise, $\gamma_I \to \infty$, we have $0 < \frac{\zeta_0\sqrt{2\pi\gamma_U}}{\alpha(1-\zeta_0)} < MoA$, where $\zeta_0 \equiv 2b\left[\frac{1}{2} - \Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\right)\right]^2 \min\{g(0), g(1)\}$. The lower bound increases in b.

   (ii) For $\gamma_I < \infty$, and $|\bar{\theta} - \mu| < \sqrt{\frac{\ln(\gamma_I/\gamma_U)\gamma_I\gamma_U}{(\gamma_I-\gamma_U)\alpha^2}}$, we have $0 < \frac{\zeta_1\sqrt{2\pi\gamma_U}}{\alpha(1-\zeta_1)} < MoA$, where $\zeta_1 \equiv 2b\left[\Phi(\frac{\alpha}{\sqrt{\gamma_U}}) - \Phi(\frac{\alpha}{\sqrt{\gamma_I}})\right]^2 \min\{g(0), g(1)\}$, which increases in both b and γ_I.

Proof. To determine a non-trivial lower bound, we assume $\kappa > \frac{1}{2}$ and $\mu > \bar{\theta}$. Due to equations (9.5) and (9.17), we have $\frac{dD}{dn^*} = (b+\ell)\frac{g(\bar{\theta})}{1-A_{\bar{\theta}}}[1-\kappa-\Phi(\frac{\alpha}{\sqrt{\gamma_U}}[\bar{\theta}-\mu] - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa))]^2$. Thus we have

$$1-\kappa-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}[\bar{\theta}-\mu] - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) > 1-\kappa-\Phi\left(-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)$$
$$= \Phi\left(\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - \kappa > 0.$$

As a result:

$$\begin{aligned}\frac{dD}{dn^*} &> (b+\ell)\frac{g(\bar{\theta})}{1-A_{\bar{\theta}}}\left[\Phi\left(\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - \kappa\right]^2, \\ &> (b+\ell)g(\bar{\theta})\left[\Phi\left(\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - \kappa\right]^2, \\ &> (b+\ell)\left[\Phi\left(\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - \kappa\right]^2 \min\{g(0), g(1)\} \equiv \zeta,\end{aligned} \tag{12.4}$$

where the second inequality follows from $0 < A_{\bar{\theta}} < 1$, and the third inequality follows from $\bar{\theta} \in [0,1]$ and since $g(\cdot)$ is unimodal at μ, so $g(\bar{\theta}) > \min\{g(0), g(1)\}$. Finally, since MoA increases in $\frac{dD}{dn^*}$ and since $A_{\bar{\theta}} < \frac{\alpha}{\sqrt{2\pi\gamma_U}}$, we have $\frac{\zeta\sqrt{2\pi\gamma_U}}{(1-\zeta)\alpha} < MoA$. Note that $\bar{\theta} = \mu$ obtains for $\mu(\bar{c}) = \bar{c}(1-\kappa) + (1-\bar{c})\left(1-\Phi\left(\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)\right)$, where $\bar{c}$ is the fixed point of $\bar{c} = \ell\int_{\mu(\bar{c})}^{\infty}\left[\mu(\bar{c}) - \theta - \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right]dG(\theta) + b\int_{-\infty}^{\mu(\bar{c})}\left[1-\mu(\bar{c}) + \theta + \sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right]dG(\theta)$.

The case of $\mu < \bar{\theta}$ and $\kappa < \frac{1}{2}$ follows similar steps and also results in the lower bound ζ.

We obtain $\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}[\bar{\theta}-\mu]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) > \Phi\left(-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)$, and $\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}[\bar{\theta}-\mu]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) > 1-\kappa$. Therefore, $\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}[\bar{\theta}-\mu]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - (1-\kappa) > \Phi\left(-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right) - (1-\kappa) > 0$.

Finally, consider the case of $\kappa = \frac{1}{2}$, so $\Phi^{-1}(\kappa) = 0$. If $|\bar{\theta}-\mu| > 1$, we have

$$\left[1-\kappa-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}[\bar{\theta}-\mu]-\sqrt{1+\frac{\alpha}{\gamma_U}}\Phi^{-1}(\kappa)\right)\right]^2 \geq \left[\frac{1}{2}-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\right)\right]^2, \tag{12.5}$$

which leads to $\frac{dD}{dn^*} > 2b\left[\frac{1}{2}-\Phi\left(\frac{\alpha}{\sqrt{\gamma_U}}\right)\right]^2 \min\{g(0), g(1)\} = \zeta_0$.

A lower bound for limited precision improvement Suppose that $\gamma_I < \infty$. We first note a useful Lemma.

Lemma 2.12.1. *Let $a > b > 0$. Then, $\Phi(a\delta) - \Phi(b\delta)$ increases in δ if $|\delta| < \sqrt{\frac{2\ln(a/b)}{a^2-b^2}}$.*

Proof. Let $\Delta(\delta) = \Phi(a\delta) - \Phi(b\delta)$. Then, $\frac{\partial\Delta(\delta)}{\partial\delta} = a\phi(a\delta) - b\phi(b\delta) > 0$ if and only if $\delta^2 < \frac{2\ln(a/b)}{a^2-b^2}$, completing the proof. ∎

Let $\kappa = \frac{1}{2}$ and $|\bar{\theta}-\mu| > 1$. To apply the above lemma, let $a = \frac{\alpha}{\sqrt{\gamma_U}}$, $b = \frac{\alpha}{\sqrt{\gamma_I}}$, and $\delta = \bar{\theta}-\mu$. Then, for $1 < |\bar{\theta}-\mu| < \sqrt{\frac{\ln(\gamma_I/\gamma_U)\gamma_I\gamma_U}{(\gamma_I-\gamma_U)\alpha^2}}$, we have $|\Phi_U(\bar{\theta}) - \Phi_I(\bar{\theta})| = |\Phi(\frac{\alpha}{\sqrt{\gamma_U}}(\bar{\theta}-\mu)) - \Phi(\frac{\alpha}{\sqrt{\gamma_I}}(\bar{\theta}-\mu))| > |\Phi(\frac{\alpha}{\sqrt{\gamma_U}}) - \Phi(\frac{\alpha}{\sqrt{\gamma_I}})|$. Equipped with this lower bound, equations (9.5) and (12.15) imply:

$$\frac{dD}{dn^*} = 2b\frac{g(\bar{\theta})}{1-A_{\bar{\theta}}}[\Phi_U(\bar{\theta}) - \Phi_I(\bar{\theta})]^2 \geq 2b[\Phi(\frac{\alpha}{\sqrt{\gamma_U}}) - \Phi(\frac{\alpha}{\sqrt{\gamma_I}})]^2 \min\{g(0), g(1)\} \equiv \zeta_1. \tag{12.6}$$

Thus, we obtain a lower bound that increases in both γ_I and b:

$$0 < \frac{\zeta_1\sqrt{2\pi\gamma_U}}{\alpha(1-\zeta_1)} < MoA < \infty. \tag{12.7}$$

∎

2.12.3 Generic information cost and Proofs of Proposition 2.6.3 and 2.6.4

The density function $f(c)$ has support $[c_{min}, c_{max}]$ and let $F(c)$ be the associated cumulative distribution function. Investors use the threshold strategy in equilibrium, whereby an investor acquires information, $n_i^* = I$, if and only if the individual information cost is below the threshold information cost $\overline{c}$. As a result, the equilibrium proportion of informed investors is $n^* = F(\overline{c})$, and the equilibrium threshold information cost solves the fixed-point equation

$$\overline{c} = D(F(\overline{c})). \tag{12.8}$$

Two conditions ensure the existence of a unique threshold information cost $\overline{c}$. First, the information cost has to be sufficiently heterogeneous relative to the bounds on the value of private information to ensure the required dominance regions at the information stage:

$$c_{min} < D(0) < D(1) < c_{max}, \tag{12.9}$$

Existence of a threshold information cost follows. Second, uniqueness of the threshold information cost requires the slope of the left-hand side of condition (12.8) to exceed the slope of the right-hand side:

$$1 > \frac{dD}{dn^*} f(\overline{c}), \tag{12.10}$$

where a higher threshold information cost raises the equilibrium proportion of informed investors and, because of the strategic complementarity in information acquisition, the value of private information. Paralleling the previous proof, an upper bound on the probability density function suffices:

$$f(c) \leq \frac{\sqrt{2\pi}}{2} - \frac{\alpha}{2\sqrt{\gamma_U}}, \tag{12.11}$$

for $\gamma_U \in (\frac{\alpha^2}{2\pi}, \infty)$, where the lower bound ensures uniqueness at the coordination stage. If the distribution of information costs satisfies $f(c) < \sqrt{\frac{\pi}{2}}$, then there exists a sufficiently large but finite precision of an uninformed investor's private signal to support uniqueness.

For the uniform distribution over the interval $[0, 1]$ considered in the main text, for example, the slope condition is satisfied, $f(c) = 1$, and implies a range of precision that ensure uniqueness, $\gamma_U \in (\underline{\gamma}, \infty)$. Moreover, since $D \in (0, 1)$ by Lemma 2.9.3 in Appendix 2.9.2, the heterogeneity condition is also satisfied.

We generalize the amplification result to a generic distribution of the information cost:

$$\left.\frac{d\overline{\theta}}{d\mu}\right|_{n^*,\overline{\theta}} = \frac{A_\mu + A_{n^*} D_\mu f(\overline{c})}{1 - A_{\overline{\theta}} - A_{n^*} D_{\overline{\theta}} f(\overline{c})} < 0. \tag{12.12}$$

Amplification always occurs since $A_{n^*} D_\mu f(\overline{c}) < 0$. Moreover, the magnitude of amplification is larger than in the main text if and only if $f(\overline{c}) > 1$, for which $f(c) > 1$ suffices.

More generally, the magnitude of amplification increases in the proportion of investors who change their information acquisition choice after a change in the public signal. The three specification considered in Proposition 2.6.4 are designed to ensure that one density function is above the other for any relevant information cost within $[D(0), D(1)]$. See also Figure 2-10.

2.12.4 Limited precision improvement and Proof of Proposition 2.6.5

For the extension of limited precision improvement, $\gamma_I \in (\gamma_U, \infty)$, we prove the existence and uniqueness of equilibrium and generalize the amplification result. The signal thresholds are given by equation (9.5) and the fundamental threshold is implicitly given by equation (9.7). Lemma 2.12.2 extends Lemma 2.9.1 and its proof parallels that of Lemma 2.9.1 closely and is omitted for brevity. To generalize the definition of a strong public signal, we replace $\hat{\mu}$ with $\tilde{\mu} \equiv \Phi\left(-\frac{\sqrt{\alpha+\gamma_I}-\sqrt{\alpha+\gamma_U}}{\sqrt{\gamma_I}-\sqrt{\gamma_U}}\Phi^{-1}(\kappa)\right) - \frac{\sqrt{\gamma_I(\alpha+\gamma_U)}-\sqrt{\gamma_U(\alpha+\gamma_I)}}{\alpha(\sqrt{\gamma_I}-\sqrt{\gamma_U})}\Phi^{-1}(\kappa)$, where $\tilde{\mu} \to \hat{\mu}$ as $\gamma_I \to \infty$. Moreover, $\hat{\theta}$ generalizes to $\tilde{\theta} \equiv \mu + \frac{\sqrt{\gamma_I(\alpha+\gamma_U)}-\sqrt{\gamma_U(\alpha+\gamma_I)}}{\alpha(\sqrt{\gamma_I}-\sqrt{\gamma_U})}\Phi^{-1}(\kappa)$.

Lemma 2.12.2. *Suppose $\gamma_U > \underline{\gamma}'$. If $\mu \neq \tilde{\mu}$, then the fundamental threshold at the coordination stage responds to changes in the proportion of informed investors:*

$$\frac{d\overline{\theta}}{dn^*} \neq 0 \tag{12.13}$$

Furthermore, the fundamental threshold increases in the proportion of informed investors if and only if the public signal is strong, $\frac{d\overline{\theta}}{dn^*}\Big(\mu - \tilde{\mu}\Big) > 0$.

The value of private information is given by equations (3.12) and (3.11). Lemma 2.9.2 and Lemma 2.9.3 also generalize to the case of limited precision improvement, as we show below.

Lemma 2.12.3. *If $\gamma_U > \underline{\gamma}'$, there is strategic complementarity in information choices:*

$$\frac{dD}{dn^*} = (b+\ell)\left[1 - A_{\overline{\theta}}\right] g(\overline{\theta}) \left(\frac{d\overline{\theta}}{dn^*}\right)^2 \geq 0, \tag{12.14}$$

with strict inequality if $\mu \neq \tilde{\mu}$. Furthermore, the more-restrictive lower bound on the precision of private information, $\gamma_U > \left(\frac{\alpha}{\sqrt{2\pi}-2}\right)^2 = \underline{\gamma}$, ensures $\frac{dD}{dn^*} < 1$.

Proof. The proof follows the same steps as the proof of Lemma 2.9.2. If $\gamma_U > \underline{\gamma}'$, then $\overline{\theta}(n^*; \mu)$ is unique for any $n^* \in [0,1]$. One can show $\frac{dD}{dn^*} = (b+\ell)\left[1 - A_{\overline{\theta}}\right] g(\overline{\theta}) \left(\frac{d\overline{\theta}}{dn^*}\right)^2 \geq 0$, with strict inequality if and only if $\mu \neq \tilde{\mu}$, where $\frac{\partial D}{\partial \overline{x}_z} = 0$ by an envelope theorem argument. Using Lemma 2.12.2, rewriting yields:

$$\frac{dD}{dn^*} = \frac{(b+\ell) g(\overline{\theta}) \left[\Phi_I(\overline{\theta}) - \Phi_U(\overline{\theta})\right]^2}{1 - A_{\overline{\theta}}}. \tag{12.15}$$

Note that $g \leq \frac{1}{\sqrt{2\pi}}$, $b+\ell < 2$, and $[\Phi_I(\overline{\theta}) - \Phi_U(\overline{\theta})]^2 \leq 1$, and $A_{\overline{\theta}} \leq \frac{\alpha}{\sqrt{2\pi\gamma_U}}$ because $n^* \geq 0$ and $\phi \leq \frac{1}{\sqrt{2\pi}}$. Therefore, $\frac{dD}{dn^*} \leq \frac{2\frac{1}{\sqrt{2\pi}}}{1-\frac{\alpha}{\sqrt{2\pi\gamma_U}}}$. Hence, $\frac{dD}{dn^*} < 1$ is ensured by $\gamma_U > \underline{\gamma} > \underline{\gamma}'$. ■

Lemma 2.12.4. *If $\gamma_U > \underline{\gamma}'$, then the value of private information satisfies $D \in (0,1)$.*

Proof. If $\gamma_U > \underline{\gamma}'$, then $\overline{\theta}(n^*;\mu)$ is unique for any $n^* \in [0,1]$. First, we show $D < 1$ for all n^*:

$$D \leq \ell \int_{\overline{\theta}}^{\infty} \Phi_U(\theta)\, dG(\theta) + b \int_{-\infty}^{\overline{\theta}} \Phi_I(\theta)\, dG(\theta) \leq \ell[1-G(\overline{\theta})] + bG(\overline{\theta}) \leq \max\{b,\ell\} < 1, \quad (12.16)$$

where, in the first line, we dropped $-\Phi_I(\theta)$ from the first term and $-\Phi_U(\theta)$ from the second term and, in the second line, used that $\Phi_z(\theta) \leq 1$ and $G(\overline{\theta}) \in [0,1]$.

Second, we show $D > 0$ for all n^*. Because of the strategic complementarity in information acquisition (Lemma 2.12.3), it suffices to show $D > 0$ at $n^* = 0$. At the lower bound, as $\gamma_I \to \gamma_U$, $\Gamma \to 0$ and $D \to 0$. At the upper bound, we have shown that $D > 0$ as $\gamma \to \infty$ (Lemma 2.9.3). Thus, a sufficient condition for $D > 0$ in case of limited precision improvement is $\left.\frac{dD}{d\gamma_I}\right|_{n^*=0} > 0$. In what follows, we use the result of $D_{\bar{x}_I} = 0$ and the result of $\left.\frac{d\overline{\theta}}{d\gamma_I}\right|_{n^*=0} = 0$, since changes in the precision of informed investors do not affect the fundamental threshold if a zero mass of investors acquires information. Hence, by total differentiation:

$$\begin{aligned}
\left.\frac{dD}{d\gamma_I}\right|_{n^*=0} &= \frac{\partial D}{\partial \gamma_I} = \frac{1}{2\sqrt{\gamma_I}}\left[-\ell\int_{\overline{\theta}}^{\infty}[\overline{x}_I - \theta]\phi_I(\theta)dG(\theta) + b\int_{\infty}^{\overline{\theta}}[\overline{x}_I - \theta]\phi_I(\theta)dG(\theta)\right] && (12.17)\\
&> \frac{1}{2\sqrt{\gamma_I}}\left[-\ell\int_{\overline{\theta}}^{\infty}[\overline{x}_I - \overline{\theta}]\phi_I(\theta)dG(\theta) + b\int_{\infty}^{\overline{\theta}}[\overline{x}_I - \overline{\theta}]\phi_I(\theta)dG(\theta)\right] \\
&= \frac{[\overline{x}_I - \overline{\theta}]}{2\sqrt{\gamma_I}}\left[-\ell\int_{\overline{\theta}}^{\infty}\phi_I(\theta)dG(\theta) + b\int_{\infty}^{\overline{\theta}}\phi_I(\theta)dG(\theta)\right] = 0, && (12.18)
\end{aligned}$$

where the first line uses $\phi_z(\theta)$ as the probability density function associated with $\Phi_z(\theta)$, the inequality of the second line arises from different weights used. Specifically, the expression on this line has a higher weight $[\overline{x}_I - \overline{\theta}] > [\overline{x}_I - \theta]$ for $\theta \geq \overline{\theta}$ on the loss and a lower weight $[\overline{x}_I - \overline{\theta}] < [\overline{x}_I - \theta]$ for $\theta \leq \overline{\theta}$ on the benefit. The third line is zero, since this is the first-order condition for the optimal threshold $\overline{x}_I$. This completes the proof. ■

As in Appendix 2.9, existence and uniqueness follows from Lemma 2.12.2 - 2.12.4. The final step is to prove amplification in the case of limited precision improvement. The proof in Appendix 2.10 generalizes. The first two steps did not require vanishing private noise of informed investors. The third step generalizes as well, where the relevant inequality becomes to $\mu \neq \tilde{\mu}$.

2.13 Payoff sensitivity and Proof of Propositions 2.7.1 – 2.7.3

In case of payoffs sensitive to the fundamental, where we follow Iachan and Nenov (2015), Bayesian updating is unchanged but the **indifference conditions** for $z \in \{I, U\}$ becomes:

$$M^z(\overline{\theta}, \overline{x}_z) \equiv \int_{-\infty}^{\overline{\theta}} b(\theta) h_z(\theta, \overline{x}_z) d\theta - \int_{\overline{\theta}}^{\infty} \ell(\theta) h_z(\theta, \overline{x}_z) d\theta \equiv 0 \tag{13.1}$$

where the signal threshold $\overline{x}_z$ enters implicitly via $h_z = h_z(\theta, \overline{x}_z)$. Without payoff sensitivity, equation (13.1) would again yield $\kappa = H_z(\overline{\theta}, \overline{x}_z)$, which implies an explicit expression of $\overline{x}_z$.

For future reference, the partial derivatives of $M^z(\overline{\theta}, \overline{x}_z; \mu)$ are:

$$M^z_{\overline{\theta}} = h_z(\overline{\theta})\Big(b(\overline{\theta}) + \ell(\overline{\theta})\Big) > 0 \tag{13.2}$$

$$M^z_{\overline{x}_z} = -\frac{\gamma_z}{\alpha + \gamma_z}\left[h_z(\overline{\theta})\Big(b(\overline{\theta}) + \ell(\overline{\theta})\Big) - \int_{-\infty}^{\overline{\theta}} b'(\theta) dH_z(\theta) + \int_{\overline{\theta}}^{\infty} \ell'(\theta) dH_z(\theta)\right] < 0 \tag{13.3}$$

$$M^z_{\mu} = \frac{\alpha}{\gamma_z} M^z_{\overline{x}_z} < 0, \tag{13.4}$$

where we used $\frac{\partial h_z}{\partial \overline{x}_z} = h_z(\theta)\gamma_z\left[\theta - \frac{\alpha\mu + \gamma_z \overline{x}_z}{\alpha + \gamma_z}\right] = \frac{\gamma_z}{\alpha}\frac{\partial h_z}{\partial \mu}$ and partial integration for $M^z_{\overline{x}_z}$ and M^z_{μ}.

The **critical mass condition** can be stated as:

$$\overline{\theta} \equiv n^* \, \Phi\left(\sqrt{\gamma_I}\left[\overline{x}_I - \overline{\theta}\right]\right) + (1 - n^*)\, \Phi\left(\sqrt{\gamma_U}\left[\overline{x}_U - \overline{\theta}\right]\right) \equiv A(\overline{\theta}, n^*, \overline{x}_U, \overline{x}_I). \tag{13.5}$$

For future reference, the partial derivatives of $A(\overline{\theta}, n^*, \overline{x}_U, \overline{x}_I)$ are:

$$A_{n^*} = \Phi_I(\overline{\theta}) - \Phi_U(\overline{\theta}) = -\Gamma(\overline{\theta}) \tag{13.6}$$

$$A_{\overline{x}_I} = n^* \sqrt{\gamma_I} \phi_I(\overline{\theta}) > 0 \tag{13.7}$$

$$A_{\overline{x}_U} = (1 - n^*)\sqrt{\gamma_U} \phi_U(\overline{\theta}) > 0 \tag{13.8}$$

$$A_{\overline{\theta}} = -A_{\overline{x}_I} - A_{\overline{x}_U} < 0 \tag{13.9}$$

$$A_{\mu} = 0. \tag{13.10}$$

These partial derivatives do not correspond to the main text directly, since one cannot solve for the signal thresholds $\overline{x}_z$ explicitly in case of payoff sensitivity. For example, A_μ captures only the direct effect of the public signal on the aggregate attack size (which is zero), without taking into account the indirect effect via the indifference condition of the marginal investor.

The value of private information can be stated as:

$$D \equiv \int_{\overline{\theta}}^{\infty} \ell(\theta)\Gamma(\theta) dG(\theta) - \int_{-\infty}^{\overline{\theta}} b(\theta)\Gamma(\theta) dG(\theta), \tag{13.11}$$

where the payoffs now depend directly on the fundamental. The partial derivatives are:

$$D_{\overline{x}_z} = 0 \tag{13.12}$$
$$D_{\overline{\theta}} = -g(\overline{\theta})\Big(b(\overline{\theta})+\ell(\overline{\theta})\Big)\Gamma(\overline{\theta}) \tag{13.13}$$
$$D_{\mu} = g(\overline{\theta})\Big(b(\overline{\theta})+\ell(\overline{\theta})\Big)\Gamma(\overline{\theta}) + \int_{\overline{\theta}}^{\infty}\ell'(\theta)\Gamma(\theta)dG(\theta) - \int_{-\infty}^{\overline{\theta}} b'(\theta)\Gamma(\theta)dG(\theta), \tag{13.14}$$

where the first line again follows from the optimality of $\overline{x}_z$, and the derivation of D_μ parallels the derivation in Appendix 2.10.1 and again uses partial integration as well as the indifference condition in equation (13.1). Observe that D_μ now has additional terms that depend on the slopes of the payoff parameters with respect to the fundamental. Thus $D_\mu \neq -D_{\overline{\theta}}$ in general. Since $c_i \sim \mathcal{U}[0,1]$ as in the main text, the equilibrium proportion of informed investors and the **threshold information cost** is again implicitly given by the fixed point $n^* = \overline{c} = D(\overline{c})$.

Exogenous information If the proportion of informed investors is exogenous, the equilibrium is given by the following set of equations:

$$M^I(\overline{\theta}, \overline{x}_I; \mu) = 0 \tag{13.15}$$
$$M^U(\overline{\theta}, \overline{x}_U; \mu) = 0 \tag{13.16}$$
$$\overline{\theta} = A(\overline{\theta}, \overline{x}_U, \overline{x}_I, \overline{n}). \tag{13.17}$$

Total differentiation of this system of equations with respect to μ and evaluating the resulting expression at the equilibrium values ($\overline{n} = n^*$, $\overline{\theta}$) yields:

$$-\left.\frac{d\tilde{\theta}}{d\mu}\right|_{\overline{n}=n^*,\overline{\theta}} = \frac{\frac{\alpha}{\gamma_I}A_{\overline{x}_I} + \frac{\alpha}{\gamma_U}A_{\overline{x}_U}}{1 - A_{\overline{\theta}} + A_{\overline{x}_I}\frac{M^I_{\overline{\theta}}}{M^I_{\overline{x}_I}} + A_{\overline{x}_U}\frac{M^U_{\overline{\theta}}}{M^U_{\overline{x}_U}}} \equiv \frac{\delta_0}{\delta_1} > 0, \tag{13.18}$$

since $\delta_0 > 0$ and $\delta_1 > 0$ because of the uniqueness of equilibrium. We abstract from deriving sufficient conditions for the existence of a unique equilibrium in the general case but fully describe a special case with linear payoffs below.

Endogenous information choice If the proportion of informed investors is endogenous, the equilibrium is given by the following set of equations:

$$M^I(\overline{\theta}, \overline{x}_I; \mu) = 0 \tag{13.19}$$
$$M^U(\overline{\theta}, \overline{x}_U; \mu) = 0 \tag{13.20}$$
$$\overline{\theta} = A(\overline{\theta}, \overline{x}_U, \overline{x}_I, n^*) \tag{13.21}$$
$$n^* = \overline{c} = D(\overline{\theta}, \overline{x}_I, \overline{x}_U; \mu). \tag{13.22}$$

Total differentiation of this system of equations with respect to μ and evaluating the resulting expression at the equilibrium values (n^*, $\overline{\theta}$) yields:

$$-\left.\frac{d\overline{\theta}}{d\mu}\right|_{n^*,\overline{\theta}} = \frac{\frac{\alpha}{\gamma_I}A_{\overline{x}_I} + \frac{\alpha}{\gamma_U}A_{\overline{x}_U} - A_{n^*}D_\mu}{1 - A_{\overline{\theta}} + A_{\overline{x}_I}\frac{M^I_{\overline{\theta}}}{M^I_{\overline{x}_I}} + A_{\overline{x}_U}\frac{M^U_{\overline{\theta}}}{M^U_{\overline{x}_U}} - A_{n^*}D_{\overline{\theta}}} \tag{13.23}$$

Amplification Note that $-\left.\frac{d\overline{\theta}}{d\mu}\right|_{n^*,\overline{\theta}} > -\left.\frac{d\tilde{\theta}}{d\mu}\right|_{\overline{n}=n^*,\overline{\theta}}$, or amplification via the information choice of investors, occurs whenever:

$$-\delta_0 A_{n^*}D_{\overline{\theta}} + \delta_1 A_{n^*}D_\mu < 0, \tag{13.24}$$

We have $A_{n^*}D_{\overline{\theta}} = g(\overline{\theta})\Big(b(\overline{\theta}) + \ell(\overline{\theta})\Big)\Gamma(\overline{\theta})^2 > 0$ generically. Thus, $A_{n^*}D_\mu \leq 0$ suffices for amplification, which yields the stated condition, as can be seen from $A_{n^*}D_\mu \equiv \delta_3 + \delta_4$:

$$\delta_3 \equiv -g(\overline{\theta})\Big(b(\overline{\theta}) + \ell(\overline{\theta})\Big)\Gamma(\overline{\theta})^2 \leq 0 \tag{13.25}$$

$$\delta_4 \equiv -\Gamma(\overline{\theta})\Big[\int_{\overline{\theta}}^{\infty} \ell'(\theta)\Gamma(\theta)dG(\theta) - \int_{-\infty}^{\overline{\theta}} b'(\theta)\Gamma(\theta)dG(\theta)\Big]. \tag{13.26}$$

2.13.1 A special case with linearity

Consider the linear case of $b(\theta) = b_0 - b_1\theta$ and $\ell(\theta) = \ell_0 + \ell_1\theta$, where all coefficients are strictly positive. Since $\overline{\theta} \in (0,1)$, we impose $b_0 > b_1$ to ensure that $b(\theta) > 0$ over the relevant range $(-\infty, 1)$. Likewise, $\ell_0 > 0$ ensures that $\ell(\theta) > 0$ over the relevant range $(0, \infty)$. We also assume identical slope coefficients $b_1 = \ell_1 \equiv \lambda$.

As a result, the partial derivatives of the indifference condition simplify to:

$$M^z_{\overline{\theta}} = h_z(\overline{\theta})\Big(b_0 + \ell_0\Big) > 0 \tag{13.27}$$

$$M^z_{\overline{x}_z} = -\frac{\gamma_z}{\alpha + \gamma_z}\Big[M^z_{\overline{\theta}} + \lambda\Big] < 0. \tag{13.28}$$

These results allow us to simplify the expressions for δ_0 and δ_1. First, $\delta_0 > 0$ since $\phi(\cdot) \leq \frac{1}{2\pi}$, so $\delta_0 = n^*\phi_I(\cdot)\frac{\alpha}{\sqrt{\gamma_I}} + (1-n^*)\phi_U(\cdot)\frac{\alpha}{\sqrt{\gamma_U}} \leq \frac{\alpha}{\sqrt{2\pi}}\Big[\frac{n^*}{\sqrt{\gamma_I}} + \frac{1-n^*}{\sqrt{\gamma_U}}\Big] \leq \frac{\alpha}{\sqrt{2\pi\gamma_U}}$. Thus, $\delta_0 < 1$ is again ensured by $\gamma_U > \underline{\gamma}'$. Second, one can show that $\delta_1 > 0$ under certain conditions. To see this, insert $A_{\overline{\theta}}$ into δ_1 to obtain:

$$\delta_1 = 1 + A_{\overline{x}_I}\left(1 + \frac{M^I_{\overline{\theta}}}{M^I_{\overline{x}_I}}\right) + A_{\overline{x}_U}\left(1 + \frac{M^U_{\overline{\theta}}}{M^U_{\overline{x}_U}}\right) \tag{13.29}$$

$$= 1 + n^*\sqrt{\gamma_I}\phi_I(\cdot)\frac{\lambda - \frac{\alpha}{\gamma_I}M^I_{\overline{\theta}}}{\lambda + M^I_{\overline{\theta}}} + (1-n^*)\sqrt{\gamma_U}\phi_U(\cdot)\frac{\lambda - \frac{\alpha}{\gamma_U}M^U_{\overline{\theta}}}{\lambda + M^U_{\overline{\theta}}}. \tag{13.30}$$

If $\frac{d}{d\lambda}\frac{\lambda - \frac{\alpha}{\gamma_z}M^z_{\overline{\theta}}}{\lambda + M^z_{\overline{\theta}}} > 0$, then $\delta_1 > 0$ is guaranteed by $\delta_1(\lambda = 0) \equiv \delta_1^{min} > 0$. As before, this is ensured by $\gamma_U > \underline{\gamma}'$. Lemma 2.13.1 states a sufficient condition for the positive total derivative.

Lemma 2.13.1. *There exists a constant $\rho \in (0, \infty)$ such that $\frac{d}{d\lambda}\frac{\lambda - \frac{\alpha}{\gamma_z}M^z_{\overline{\theta}}}{\lambda + M^z_{\overline{\theta}}} > 0$ if $\gamma_I\lambda < \rho$.*

Proof. Taking a derivative with respect to λ implies

$$\frac{d}{d\lambda}\left(\frac{\lambda - \frac{\alpha}{\gamma_z}M^z_{\overline{\theta}}}{\lambda + M^z_{\overline{\theta}}}\right) = \frac{1 + \frac{\alpha}{\gamma_z}}{(\lambda + M^z_{\overline{\theta}})^2}\left(M^z_{\overline{\theta}} - \lambda\frac{dM^z_{\overline{\theta}}}{d\lambda}\right) > 0$$

if and only if $\frac{1}{\lambda}M^z_{\overline{\theta}} > \frac{dM^z_{\overline{\theta}}}{d\lambda}$. Thus, it suffices to show

$$\frac{1}{\lambda}M^z_{\overline{\theta}} > |\frac{dM^z_{\overline{\theta}}}{d\lambda}|. \tag{13.31}$$

Thus, chain rule implies $|\frac{dM^z_{\overline{\theta}}}{d\lambda}| \leq |\frac{\partial M^z_{\overline{\theta}}}{\partial\overline{\theta}}|\ |\frac{d\overline{\theta}}{d\lambda}|$. By following similar steps as in derivation of (13.23), we have

$$-\frac{d\overline{\theta}}{d\lambda}\bigg|_{n^*,\overline{\theta}} = \frac{\frac{M^I_\lambda}{M^I_{\overline{x}_I}}A_{\overline{x}_I} + \frac{M^U_\lambda}{M^U_{\overline{x}_I}}A_{\overline{x}_U} - A_{n^*}D_\lambda}{1 - A_{\overline{\theta}} + A_{\overline{x}_I}\frac{M^I_{\overline{\theta}}}{M^I_{\overline{x}_I}} + A_{\overline{x}_U}\frac{M^U_{\overline{\theta}}}{M^U_{\overline{x}_U}} - A_{n^*}D_{\overline{\theta}}}, \tag{13.32}$$

where due to (13.1) and (13.11):

$$M^z_\lambda = -\frac{\alpha\mu + \gamma_z\overline{x}_z}{\alpha + \gamma_z}, \tag{13.33}$$

$$D_\lambda = E[\theta\Gamma(\theta)]. \tag{13.34}$$

Given the characterizations in (13.6)-(13.10),(13.12)-(13.14), (13.33) and (13.34), there exists a finite ρ such that $|\left[\bar{\theta} - \frac{\alpha\mu+\gamma_z\overline{x}_z}{\alpha+\gamma_z}\right]|\ |\frac{d\bar{\theta}}{d\lambda}| < \rho^{-1}$. The existence of ρ follows by the finiteness of $\bar{\theta}$, $\frac{\alpha\mu+\gamma_z\overline{x}_z}{\alpha+\gamma_z}$ and $\frac{d\bar{\theta}}{d\lambda}$. Recall that $M^z_{\bar{\theta}} = h_z(\bar{\theta})(b_0 + \ell_0)$, thus $-\frac{\partial M^z_{\bar{\theta}}}{\partial\bar{\theta}} = \left[\bar{\theta} - \frac{\alpha\mu+\gamma_z\overline{x}_z}{\alpha+\gamma_z}\right]\ h_z(\bar{\theta})\ (\alpha + \gamma_z)\ (b_0 + \ell_0)$. Finally, plugging $\frac{\partial M^z_{\bar{\theta}}}{\partial\bar{\theta}}$ and (13.32) into (13.31) implies that the inequality holds when $\gamma_z \lambda < \rho$. Since $\gamma_I > \gamma_U$, thus it is sufficient to have $\gamma_I \lambda < \rho$. Let $\overline{\lambda} \equiv \frac{\rho}{\gamma_I}$ denote the upper bound on the sensitivity to the payoff, which completes the proof. ■

We wish to establish the existence of a unique equilibrium. Paralleling the proof in Appendix 2.9, we need to generalize Lemma 2.9.2 and Lemma 2.9.3.

Lemma 2.13.2. *If $\gamma_U > \underline{\gamma}$ and $\lambda < \overline{\lambda}$, then strategic complementarity generically obtains:*

$$\frac{dD}{dn^*} = (b_0 + \ell_0) g(\overline{\theta}) \frac{\Gamma(\overline{\theta})^2}{\delta_1} > 0. \tag{13.35}$$

Furthermore, $\frac{dD}{dn^*} < 1$.

Proof. The expression for $\frac{dD}{dn^*}$ is derived in the same way as in the main text. Using $g(\overline{\theta}) \leq \frac{1}{\sqrt{2\pi}}$, $b_0 + \ell_0 = b + \ell < 2$, $\Gamma^2 \leq 1$, and $\delta_1 \geq \delta_1^{min}$ yields the lower bound on γ_U. ■

Lemma 2.13.3. *If $\gamma_U > \underline{\gamma}$ and $\lambda < \overline{\lambda}$, the value of private information satisfies $D \in (0, 1)$.*

Proof. First, $D > 0$ follows by a direct generalization of Lemma 2.12.4, where the generalized condition for the optimality of $\overline{x}_I$ is used. Second, since $b \in (0, 1)$ and $\ell \in (0, 1)$ and $\Gamma(\theta) \in [-1, 1]$ for all θ, it follows that $D = \int_{\overline{\theta}}^{\infty} \ell(\theta)\Gamma(\theta)dG(\theta) - \int_{-\infty}^{\overline{\theta}} b(\theta)\Gamma(\theta)dG(\theta) < \int_{\overline{\theta}}^{\infty} dG(\theta) - \int_{-\infty}^{\overline{\theta}} (-1)dG(\theta) = 1$. ■

Following the steps of the proof in the baseline model, a unique equilibrium exists. Simplifying the previous sufficient condition, amplification occurs whenever:

$$g(\overline{\theta})(b_0 + \ell_0)\Gamma(\overline{\theta})^2 \geq -\lambda \Gamma(\overline{\theta}) \int_{-\infty}^{\infty} \Gamma(\theta) dG(\theta). \tag{13.36}$$

Using $\Gamma(\theta) \leq 1$ for all θ, a simple condition sufficient for amplification is:

$$b_0 + \ell_0 \geq \frac{\lambda}{g(\overline{\theta}) \left|\Gamma(\overline{\theta})\right|}, \tag{13.37}$$

which places another upper bound on the sensitivity of payoffs.

2.13.2 One payoff sensitive to the fundamental

We derive the magnitude of amplification (MoA) in the case when payoffs are sensitive to the fundamental. Using the previous results, we can state:

$$MoA = \frac{\delta_0 A_{n^*} D_{\overline{\theta}} - \delta_1 A_{n^*} D_{\mu}}{\delta_0[\delta_1 - A_{n^*} D_{\overline{\theta}}]} \tag{13.38}$$

$$= \frac{(\delta_0 + \delta_1) A_{n^*} D_{\overline{\theta}} + \delta_1 \Gamma(\overline{\theta}) B(\overline{\theta})}{\delta_0[\delta_1 - A_{n^*} D_{\overline{\theta}}]}, \tag{13.39}$$

where $B(\overline{\theta}) = \int_{\overline{\theta}}^{\infty} \ell'(\theta)\Gamma(\theta)dG(\theta) - \int_{-\infty}^{\overline{\theta}} b'(\theta)\Gamma(\theta)dG(\theta)$.

Thus, higher values of $\Gamma(\overline{\theta})B(\overline{\theta})$ lead to a higher magnitude of amplification. Under the assumption of one payoff being sensitive and the other depending linearly on the fundamental, we obtain the expression for $B(\overline{\theta})$ stated in the main text. Moreover, the signs follow from continuity and the observation that, for $\gamma_I \to \infty$, we have $\Gamma(\theta) \to F^U(\theta) > 0$ if $\theta > \overline{\theta}$ (perfectly informed investors never attack when no regime change occurs), while $\Gamma(\theta) \to F^U(\theta) - 1 < 0$ if $\theta < \overline{\theta}$ (perfectly informed investors always attack when regime change occurs). This completes the proof.

Chapter 3

Dynamic Pricing in Social Networks: The Word of Mouth Effect

3.1 Introduction

How should a monopolist offering a product in a social network price its product over time? Does the profit-maximizing strategy always keep the prices monotone? Is there a steady state price? This paper introduces a new framework to investigate these questions by considering the mechanism by which information about a product diffuses in networks. In particular, our goal is to investigate the role of word of mouth (WOM) communication of consumers[1], in the optimal pricing policy of the monopoly firm. The multibillion-dollar growing market for smartphone applications, where word of mouth is often the only means of spread of information about the product, is a great real world example of such a scenario.

During the last decade, there has been significant growth in the market for smartphone applications. These applications (apps) are typically cheap, and often the only low-budget[2]

This chapter is coauthored with Amir Ajorlou and Ali Jadbabaie. This chapter is published in the *Management Science* vol 64. I am grateful to Daron Acemoglu, Dirk Bergemann, Matt Elliott, Sanjeev Goyal, Leonid Kogan, Andrey Malenko, Asu Ozdaglar, Hamid Sabourian, Alvaro Sandroni, Xuanming Su, Alireza Tahbaz-Salehi, Juuso Toikka, Rakesh Vohra, and Alexander Wolitzky for helpful comments.

[1] "Word of mouth communication involves the passing of information between a non-commercial communicator (i.e. someone who is not rewarded) and a receiver concerning a brand, a product, or a service", Dichter (1966).

[2] A recent survey by *AppFlood* (McCloughan (2013)) over 1000 independent small, medium and large app

means by which many of these apps spread is the word of mouth communication of their users. Many apps ask users for permission to send notifications about the product to contacts in their address books or to post a message on their online social media, when they purchase or start using the app. A good evidence for the effectiveness of word of mouth is the smartphone application *WhatsApp* which was sold to *facebook* in early 2014 for $19 billion. WOM was the key to *WhatsApp* popularity. As noted by Bloomberg (Satariano (2014)), "They [*WhatsApp* management] eschewed marketing and did not employ a public relations person, relying on the word of mouth recommendations of its users instead".[3]

The price an app developer offers for its product is also a big driver for spreading the information. As such, posting time-varying prices is a common marketing tool for spreading information about the existence of a new app among the users. Figure 3-1 depicts the price history for *Tadaa SLR*, an iPhone photo and video application since its release.[4] An interesting observation from this chart is the *frequent* drops of the price to zero. The same pattern can be seen in the price trends of many other smartphone applications (e.g., *XnShape*, *The Curse*, *Equalizer PRO^{TM}*, *Color Vacuum*, *ContactFlow*, *Coyn*, and *IBSnap*, only to name a few[5]).

Motivated by the above observations, we study the problem of optimal dynamic pricing of a profit-maximizing firm selling a product in a large social network where agents can only get informed about the product via Word of Mouth of previously informed friends. A key feature of this work is the explicit modeling of the effect of the price on the information diffusion via WOM. The (dynamic) price is a control variable by which the firm directly affects the information diffusion of its new product through the underlying social network. Firm's problem is then to decide, at each time step, between optimally exploiting the existing informed network or charging a lower price in favor of a higher spread of information.

The main contribution of this paper is to study pricing in social networks through the channel of information diffusion. We show that when the spread of a durable product is only via word of mouth, the optimal pricing policy is neither monotone nor reaches a steady state. Rather, the optimal policy fluctuates, dropping the price to zero infinitely often, essentially giving away the immediate profit in full to expand the informed network in order to exploit

developers shows that the majority (78%) of developers surveyed had a per app marketing budget of $5000 or less.

[3]This example is solely meant to show the effectiveness of WOM in app marketing.

[4]The price data is gathered from "www.appshopper.com".

[5]These examples are chosen from various app categories of Photos & Videos, Games, Music, Education, Finance, and Utilities. See Appendix 3.9 for the price plots of some of these apps.

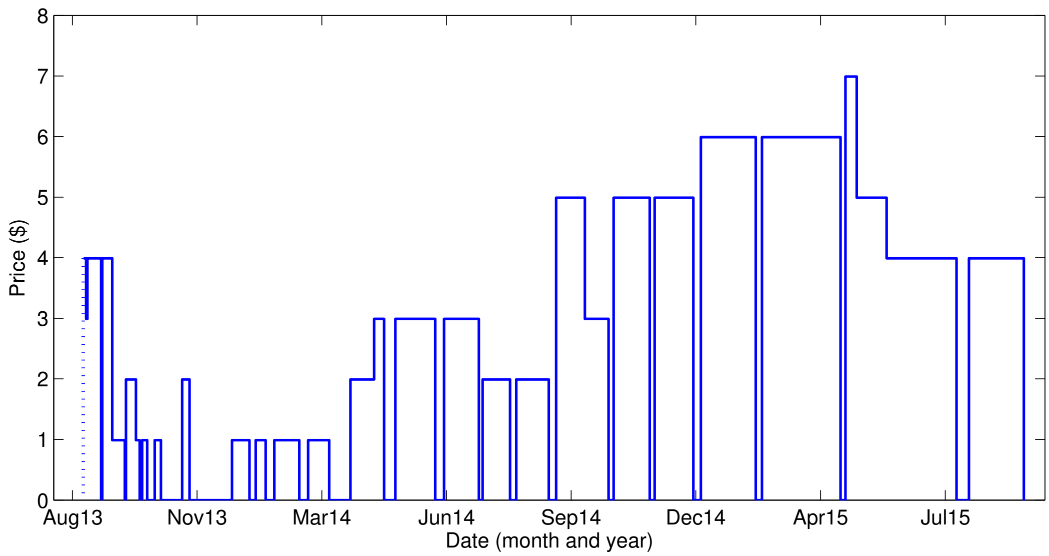

Figure 3-1: Price history for the iPhone application Tadaa SLR since its debut on Aug. 16, 2013.

it in future. This is consistent with the real world evidence from smartphone applications[6] described above.

The key intuition behind this result is that frequent zero-price sales allow the firm to attract consumers who would not buy the product unless offered for free. Giving the product for free to these low-valuation agents and getting them more engaged in the spread, the firm is able to reach sizeable parts of the network that would remain otherwise uninformed. By properly timing the drop in price, the firm can ensure that the marginal gain in future profit by selling the product in this previously unexplored part of the network prevails the loss in the immediate profit caused by offering the product for free, making the drop of the price to zero a profitable course of action. We also show that, although the optimal policy drops the price to zero infinitely often, price will not get trapped near zero.

More importantly, we show that the results remain valid in face of forward-looking agents, homophily-based[7] engagement in word-of mouth, network externalities, and consumer inattention to price changes, under surprisingly mild assumptions. Beside the WOM nature of information diffusion, we further show that the durability of the product is also a key driver for these frequent price drops. For a nondurable product, although the firm may initially

[6]We deal with smartphone applications as durable products. This is because when a user buys an app, she usually does not need to buy the same product any more.

[7]Homophily refers to a tendency of various types of individuals to associate with others who are similar to themselves (Golub and Jackson (2012)).

make some free offers to expand its network, after a finite time it will fix the price at a level that extracts the maximum profit from the already informed population.

3.1.1 Literature

This paper makes contributions to four bodies of literature: WOM marketing, dynamic pricing, zero pricing, and strategic information diffusion in networks.

Since the landmark paper of Katz and Lazarsfeld (1955) in which the authors show that people rarely act on mass-media information unless it is also transmitted through personal ties, WOM has been a major focus of research in the marketing literature (Bass (1969), Godes and Mayzlin (2004), Chevalier and Mayzlin (2006), Besbes and Scarsini (2013)). WOM communication strategies are appealing because they combine the prospect of overcoming consumer resistance with significantly lower costs and fast delivery – especially through the Internet social networking (Trusov, Bucklin and Pauwels (2008)). Although much of the work studying WOM communications focused on behavioral factors affecting information transfer (Herr, Kardes and Kim (1991), Chevalier and Mayzlin (2006), Berger and Schwartz (2011)), the rapid development of Internet social networking and communication technologies, e.g. smartphones, have boosted WOM research that is primarily concerned with network effects of social influence (see e.g., Watts and Dodds (2007), Goldenberg et al. (2009), Stephen and Toubiaz (2010), Katona, Zubcsek and Sarvary (2011), Campbell, Mayzlin and Shin (2013), Hervas-Drane (2015)). Our work also belongs to this latter category of WOM marketing literature. Recently, Campbell, Mayzlin and Shin (2013) study how far information about a product eventually diffuses through the population when consumers desire to signal their type to others. By focusing on a signaling equilibrium, they show that increasing asymmetry among consumers by restricting information of low-typed agents can boost the spread of WOM.[8] As a consequence, they find that advertising may crowd out the incentives for consumers to engage in WOM. In contrast to all of these works, we provide an analytical tractable model to analyze the impact of dynamic pricing as a marketing tool to control information diffusion via WOM.

Dynamic pricing has a rich history in economics and operation research.[9] In general,

[8]In this vein, Hervas-Drane (2015) presents a model of customer search to explain the impact of product recommendations on customer product discovery and the concentration of sales. He shows while recommendations benefit mainstream customers, when recommender systems are based on WOM and social filtering, there is a positive effect in the tail of the sales distribution on customers interested in niche products.

[9]Talluri and van Ryzin (2004) and Phillips (2005) provide an extensive review of this topic. See also Acemoglu, Kakhbod and Ozdaglar (2017) and Kakhbod, Ozdaglar and Schneider (2021) for pricing in renewable

varying prices over time may have different causes. It might be because of the inability of the firms to commit to future actions (e.g. Conlisk, Gerstner and Sobel (1984), Sobel (1991)), or due to learning new experience goods (e.g. Bergemann and Välimäki (1997, 2000), Ifrach, Maglaras and Scarsini (2011, 2013))[10] or the result of the inability of boundedly rational buyers to pay immediate attention to price changes (e.g. Radner, Radunskaya and Sundararajan (2013)). Scarcity of the products with regard to the number of buyers (e.g. Gallego and van Ryzin (1994), Gershkov and Moldovanu (2009)), network externalities (e.g. Cabral, Salant and Woroch (1999)), stochastic incoming demand (e.g. Board (2008)), and time-varying values of buyers (e.g. Garrett (2013)) are among other causes suggested in the literature for varying prices over time. In particular, Garrett (2013) studies profit-maximizing prices in an environment where buyers arrive over time and have values for the good which evolve stochastically. The author shows that for a range of parameter values, optimal prices fluctuate over time. Prices gradually fall up to sales dates and jump thereafter, mainly due to the inter-temporal price discrimination effect introduced by Stokey (1979). In contrast to the present paper, none of these works relate pricing to the extent of the information diffusion.

How can firms profitably give away free products? Several branches of literature yield insight for this phenomenon. For example, multi-product pricing in two-sided markets (Rochet and Tirole (2003), Parker and Alstyne (2005)), forward-looking consumers and durable-goods monopolies with zero marginal cost (Coase (1972), Stokey (1981) and Gul, Sonnenschein and Wilson (1981), and Çelebi, Kakhbod and Li (2020) with preference shocks, and Kakhbod and Lanzani (2020) with dynamic learning with heterogenous speeds), and to a lesser extent bundle pricing (Hanson and Martin (1990)) are potential causes for this phenomenon. In our model the optimal price scheme sets the price to zero infinitely many times. However, the principal probing factor, in contrast to all the aforementioned works, is to control the extent of information diffusion.

The main focus here is durable products with zero or negligible marginal cost. Markets for digital goods and, in particular, smartphone applications are real world examples of this scenario. Market for smartphone applications (apps) is quite large and still rapidly growing.

energy markets.

[10]These models are typically either two sided or one sided. Bergemann and Välimäki (1997, 2000), Ifrach, Maglaras and Scarsini (2013), and Yu, Debo and Kapuscinski (2013) consider two-sided learning models where buyers and sellers both learn the true value of a new product through consumer experiences. Papanastasiou, Bakhshi and Savva (2013) and Ifrach, Maglaras and Scarsini (2011) analyze one-sided learning models when firm knows the product quality, buyers report their experiences and subsequent customers learn from these reports.

Before 2013 mobile apps had an economy with a market size of $25 billion. It is estimated that one billion smartphones will be sold by the end of 2015. Given that %46 of app users report having paid for their apps, the app market is expected to have 268 billion downloads that generate $77 billion worth of revenue by 2017.[11] Frequent zero-pricing of apps is already noticed by the app industry. In fact there are websites that provide lists of the paid apps that become free on a daily basis.[12] Our work is the first analytical approach to this phenomenon in the app market which suggests controlling the extent of information diffusion via dynamic prices as a potential cause for the zero price drops in this market.

This paper is also related to the growing literature on strategic interactions in social networks (e.g. Ballester, Calvó-Armengol and Zenou (2006*b*), Bramoullé and Kranton (2007*b*), Galeotti et al. (2010)).[13] For the most part, in these models prices are static. Su (2007), Nocke and Peitz (2007), and Hörner and Samuelson (2011) consider dynamic pricing with strategic customers, however, in contrast to our work, the price paths are found to be monotone.

The closest result in literature to ours is the work in Campbell (2012), where the author studies pricing for a nondurable product under word of mouth communications. Campbell shows price fluctuations for a nondurable product during the introductory stages. This is in line with our result for the nondurable product where we show that the firm may use price drops to zero during the early stages in order to expand the informed network. However, we show that fluctuations disappear as the size of the spread gets sufficiently large. Campbell models the word of mouth as a branching process, in which each new buyer informs a fixed number of new agents in average. The model neither considers the overlap among the friends of new buyers nor the fact that some of their friends may have already heard about the product. As Campbell emphasizes in his paper, such a model is valid only at the early stages of introducing a product to the network, when the size of the informed population is still very small. Our model, on the other hand, captures both the diminishing marginal contribution of new buyers to diffusion caused by the overlap among their friends, and the slowdown of diffusion caused by the growing size of the agents who have already heard about the product.

[11]See "entrepreneur.com/article/236832" for more statistics on the app market.

[12]For instance, see "www.appsliced.co", "www.appaddict.net" and "www.appspy.com".

[13]Other relevant studies include: strategic information exchange in social networks (e.g. Acemoglu, Bimpikis and Ozdaglar (2014)), optimal static pricing under presence of local network effect (Sundararajan (2008*b*), Hartline, Mirrokni and Sundararajan (2008), Candogan, Bimpikis and Ozdaglar (2012*b*), Jadbabaei and Kakhbod (2019)), and optimal advertising strategies in social networks (Galeotti and Goyal (2009*b*), Galeotti and Mattozzi (2011)).

This work is also related to the literature on diffusion dynamics in social networks. One of the main challenges in information diffusion in social networks is developing tractable models. The combinatorial nature of networks with heterogeneity often makes analysis prohibitively difficult. Several modeling approaches have been developed to reduce the inherent complexity. An early model of diffusion is the Bass (1969) model. Although the proposed model does not capture any explicit social network structure, it still incorporates imitation from others. Some recent models use the concept of mean field theory to model diffusion over the network. The main idea of this approach is to replace all interactions to an agent with an average or effective interaction. These tractable models have been used by Jackson and Rogers (2007*b*) to relate stochastic dominance properties of the degree distribution of the network to the depth of diffusion, by Young (2009) to provide methodologies for characterizing different models of social influence by the time path of adoption, by Jackson and Rogers (2007*a*) to infer how the formation process affects average utility in the network, and by Jackson and Yariv (2007) and López-Pintado (2008) to evaluate strategic adoption decisions of individuals.[14]

From the methodological point of view, our work is also related to the literature on random graph theory. The theory of random graph has been used as a convenient modeling abstraction, which can facilitate modeling and analysis of the information diffusion in networks. Random networks find their origin in studies of random graph by Rapoport (1957) and Erdős and Rényi (1959, 1960, 1961). Random graph theory is also widely used by network scientists. For instance, it has been used by Watts and Strogatz (1998) to present their seminal small-world idea by creating highly clustered networks with small diameters, by Newman, Strogatz and Watts (2001) to model the world-wide web and collaboration networks of company directors and scientists, and by Watts (2002) to model collective actions and the diffusion of norms and innovations. The current paper develops an endogenous network model by using selling prices as dynamic controls for the information diffusion in a social network whose structure is captured by a Poisson random graph (Bollobás (2001)). The proposed model allows a firm to strategically affect the information diffusion about the existence of a new product within the social network by means of dynamic prices.

The rest of the paper is organized as follows. Section 3.2 presents a tractable model for strategic information diffusion via WOM in a large social network whose structure is represented by a Poisson random graph. Section 3.3 discusses the main challenge of the firm

[14]Mean field theory is also used in revenue management, in particular, to study and model complex dynamic demand systems with the objective of maximizing performance, e.g. Gallego and van Ryzin (1994).

as deciding between spreading and exploiting, and presents the main results of the paper. In Section 3.4, we examine the robustness of the frequent zero-price drops in the face of forward-looking agents, homophily-based engagement in WOM, and network externalities. We unravel the key role of the type of the products in the price drops by studying the problem for a nondurable product in Section 3.5. Finally, our conclusions are presented in Section 3.6.

3.2 Model

3.2.1 General Description

The economy consists of a unit measure continuum of agents indexed by $i \in I = [0, 1]$. Agents form a social network, the structure of which is captured by an undirected random graph G with Poisson degree distribution with mean λ. More precisely, each agent $i \in I$ has a total of $d_i \sim \text{Poiss}(\lambda)$ friends uniformly distributed in I.[15] We denote the set of the friends of i in G with N_i. For every $i \in I$, the set of her friends N_i forms a Poisson process in I.

At each time step $t = 0, 1, 2, \cdots$, a firm offers a product to the continuum of agents in the network at price $u(t) \in \mathcal{U}$, where $\mathcal{U}$ is a finite set of admissible prices. The set of admissible prices $\mathcal{U}$ can represent any set of quantized price levels in $[0, 1]$. In particular, we assume $0 \in \mathcal{U}$ to allow for the *free offering* of the product. We denote the set of admissible prices as $\mathcal{U} = \{p_0 = 0 < p_1 < \ldots < p_m \leq 1\}$, where $m \geq 1$ is the number of nonzero price levels.

Each agent has a private valuation $\theta \in [0, 1]$ of the product, distributed according to a cumulative distribution function $F(\theta)$. We assume that F corresponds to a non-atomic PDF and is strictly increasing on $[0, 1]$. The valuations of the agents are time-invariant and for now we assume they are independent of their degrees and the valuations of their friends.[16] Moreover, agents' valuations and their positions in the network are their private information, and hence, not known to the firm.

In order for an agent to buy the product, she should first be informed about its existence.

[15] This can be thought of as a limit case of the well-known Erdős-Rényi graph (Erdős and Rényi (1959)), keeping the mean degree equal to λ and with $I = [0, 1]$ as the limit vertex set. This network model inherits independence of the edges from the Erdős-Rényi graph which proves very convenient in analyzing network behavior. A similar network model is used by Oberfield (2012), Larson (2013), and Galeotti and Goyal (2009*b*).

[16] Later on, we will relax this assumption.

At $t = 0$ to initiate the spread of information, a uniformly randomly subset of the population becomes informed about the product directly by the firm. Later on, at any time $t \geq 1$, other agents can only get informed via word of mouth from a friend who already knows about the product. We make a distinction between the rate of engagement of buyers and non-buyers in the spread of the information about the product: When an agent buys the product, she engages in word of mouth with each friend with some probability $0 < p_B \leq 1$, informing her about the product. If an agent learns about the product but does not make a purchase, she informs each friend with a lower probability $0 \leq p_{\bar{B}} < p_B$. Assigning a nonzero probability to the engagement of non-buyers in word of mouth is motivated by some recent work which provide evidence for the significant role of non-adopters in the spread of information (Banerjee et al. (2013)). It is to be noted that an informed agent buys the product if the offered price does not exceed her valuation, i.e., $u(t) \leq \theta$ where $u(t)$ is the price offered by the firm at time t.

In this framework, firm's objective is to devise an optimal dynamic pricing policy maximizing its accumulated discounted profit over an infinite time horizon. We first study this problem for the case of a *durable* product, such as many smartphone applications, in order to justify the behavior pointed out in the previous section. We then verify the robustness of the price drops to several key model assumptions. We discuss the validity of the results in face of strategic forward-looking agents, homophily-based engagement in word of mouth, network externalities, and consumer inattention to price changes. Finally, we investigate the role of the type of the product in the price drops to zero by studying the problem for a *nondurable* product. It is to be noted that an informed agent may buy a nondurable product at each time step, given that its price is lower than her valuation. However, if she buys a durable product at some time, she will not buy it thereafter.

3.2.2 WOM Diffusion Dynamics

In this subsection, we first present a few notations, definitions, and observations that will be used later to derive the dynamics of the information diffusion in the network. We denote the set of informed agents at time t by $X(t)$ and its size by $x(t)$. $X(0)$ is therefore the set of those agents directly informed by the firm, with $x(0) = x_0$ denoting the size of this set. Considering that we are dealing with a unit measure continuum of agents, an informal use of the strong law of large numbers implies that $x(t) = \text{Prob}(i \in X(t))$. As we will see in the sequel, this will prove very convenient in deriving the dynamics of the information diffusion.[17] The set of informed agents $X(t)$ is increasing, that is $X(t-1) \subseteq X(t)$. $Y(t) = X(t) - X(t-1)$ represents the set of freshly informed agents at time t whose size is denoted by $y(t)$.

We partition the set of freshly informed agents in $Y(t)$ into two subsets: those who buy the product, denoted by $B_Y(t)$, and those who do not buy, denoted by $\bar{B}_Y(t)$. Agents in both subsets contribute to information diffusion by informing a fraction of their friends about the product, which constitutes part of $Y(t+1)$. Noting that θ has the same distribution in $Y(t)$ as it has in I, the fraction of agents from $Y(t)$ that buy the product when offered the price $u(t) = p_r$ is $(1 - F(p_r))y(t)$. This yields

$$\begin{aligned} b_Y(t) &= (1 - F(p_r))y(t), \\ \bar{b}_Y(t) &= F(p_r)y(t), \end{aligned} \tag{2.1}$$

with p_r being the price offered at time t ($u(t) = p_r$) and lowercases denoting the size of the corresponding sets.

Another contribution to diffusion comes from the set of agents previously informed about the product who have not yet purchased. For such agents, the price has not fallen below their valuation since the time they were informed about the product. We denote the set of such agents at time t by $Z(t)$. An agent in this set may buy the product at time t and thus inform some of her friends, if the offered price at time t is below her valuation. Unlike $Y(t)$, the distribution of θ is not given by $F(\cdot)$ for the agents in $Z(t)$ and depends on the price history. However, we can use a stack of m variables (recall that m is the number of nonzero price levels) to fully describe the distribution of θ in $Z(t)$. We can partition $Z(t)$ as $\bigcup_{j=1}^{m} Z_j(t)$, where $Z_j(t)$ is the set of those agents in $Z(t)$ whose valuations lie between price levels p_{j-1} and p_j, that is, $Z_j(t) = \{i \in Z(t) | p_{j-1} \leq \theta_i < p_j\}$. Then, the distribution of θ in $Z(t)$ is fully determined by the sizes of these sets, denoted by $z(t) = [z_1(t) \dots z_m(t)]^T$. If the firm chooses $u(t) = p_r$ as the price at time t, then all the agents in $B_Z(t) = \bigcup_{j=r+1}^{m} Z_j(t)$ which has a size of

$$b_Z(t) = \sum_{j=r+1}^{m} z_j(t), \tag{2.2}$$

will buy the product and will subsequently engage in word of mouth with their friends, while the rest of the agents in $Z(t)$ will be carried over to $Z(t+1)$. Those freshly informed agents in $Y(t)$ with valuations below p_r, which we denoted earlier with $\bar{B}_Y(t)$, constitute another

[17]Following the same logic, we may interchangeably use the words size, fraction, and probability in the paper.

part of $Z(t+1)$. Summarizing the above, we arrive at the following update rule for the size of the set of agents whose valuations are between p_{j-1} and p_j and have not yet bought the product:

$$z_j(t+1) = \begin{cases} z_j(t) + (F(p_j) - F(p_{j-1}))y(t), & 1 \leq j \leq r \\ 0, & \text{otherwise} \end{cases} \tag{2.3}$$

assuming that the offered price at time t is p_r ($u(t) = p_r$). The size of the fresh buyers $B(t) = B_Y(t) \cup B_Z(t)$ is given by

$$b(t) = (1 - F(p_r))y(t) + \sum_{j=r+1}^{m} z_j(t). \tag{2.4}$$

In order to find the size of the informed agents at time $t+1$, we take a closer look at the three subsets involved in the information diffusion: $B_Y(t)$, $\bar{B}_Y(t)$, and $B_Z(t)$. Agents in $B_Y(t)$ are those who were just informed about the product at time t and bought it. Upon buying the product, they may inform each of their friends about the product with some probability p_B. Using the *stationary increments property*[18] of Poisson processes, the number of friends an uninformed agent $i \notin X(t)$ has in $B_Y(t)$ is a Poisson random variable with mean $\lambda b_Y(t)$. Since each such friend may inform i with probability p_B, thus the number of friends in $B_Y(t)$ from which i may hear about the product has a Poisson distribution with mean $\lambda p_B b_Y(t)$.

Agents in $\bar{B}_Y(t)$ are those just informed about the product at time t but found the price too high to buy. These agents still may inform friends about the product with some probability $p_{\bar{B}} < p_B$ (we may also have $p_{\bar{B}} = 0$). The number of friends an uninformed agent $i \notin X(t)$ has in $\bar{B}_Y(t)$ is a Poisson random variable with mean $\lambda \bar{b}_Y(t)$. Therefore, the number of friends in $\bar{B}_Y(t)$ from which i may hear about the product has a Poisson distribution with mean $\lambda p_{\bar{B}} \bar{b}_Y(t)$.

The last contribution to diffusion comes from agents in $B_Z(t)$, who have previously heard about the product but have not made a purchase as of time t. Any such agent has already informed some of her friends when hearing about the product for the first time. Upon buying the product, they once again get engaged in word of mouth with more friends, informing them about the product. Since these agents have already informed each friend with probability $p_{\bar{B}}$, thus the number of friends an uninformed agent $i \notin X(t)$ has in $B_Z(t)$ is a Poisson

[18]According to the stationary increments property of Poisson processes, the probability distribution of the number of occurrences (herein friends) in any subset only depends on the size of the subset (Billingsley (1995)).

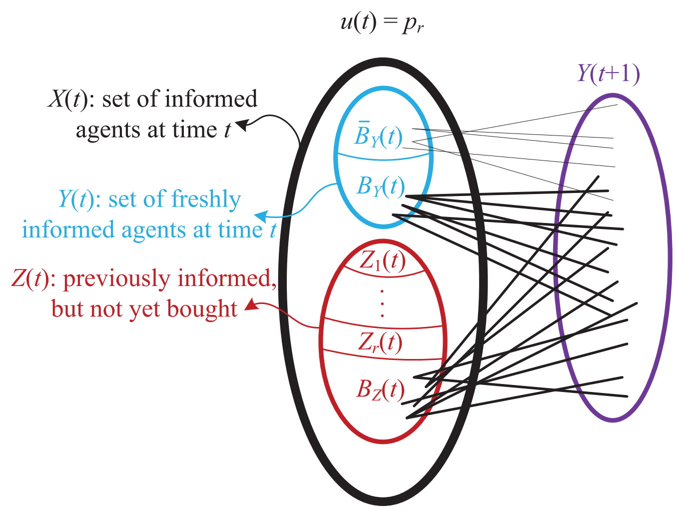

Figure 3-2: An overall view of the information diffusion via word of mouth. Agents in $B_Y(t) \cup B_Z(t)$ buy the product and inform their friends with probability p_B. Fresh non-buyers $\bar{B}_Y(t)$ also inform some friends but with a lower probability (thinner edges reflect the lower likelihood of getting informed via a non-buyer friend).

random variable with mean $\lambda(1-p_{\bar{B}})b_Z(t)$. Therefore, the number of friends in $B_Z(t)$ from which i may hear about the product has a Poisson distribution with mean $\lambda p_B(1-p_{\bar{B}})b_Z(t)$.

Putting the above three cases together, it is easy to see that the number of friends an uniformed agent $i \notin X(t)$ may hear from about the product at time t has a Poisson distribution with mean

$$\lambda p_B b_Y(t) + \lambda p_{\bar{B}} \bar{b}_Y(t) + \lambda p_B(1-p_{\bar{B}})b_Z(t). \tag{2.5}$$

An agent $i \in I$ will be uninformed at time $t+1$ if and only if she is neither informed nor hears from a friend at time t. We can now write the dynamics of the informed population $x(t)$ as

$$1 - x(t+1) = (1-x(t))e^{-\lambda(p_B b_Y(t) + p_{\bar{B}}\bar{b}_Y(t) + p_B(1-p_{\bar{B}})b_Z(t))}, \tag{2.6}$$

where $b_Y(t)$, $\bar{b}_Y(t)$, and $b_Z(t)$ are given by (2.1) and (2.2), $z(t)$ is updated using (4.27), and $y(t+1) = x(t+1) - x(t)$. Moreover, $y(0) = x(0) = x_0$ and $z_j(0) = 0$ for $1 \leq j \leq m$. The overall structure of the information diffusion via WOM is depicted in Figure 3-2.

Comment 3.2.1 (Diminishing Information Diffusion and the Slowdown of Spread)**.** We can use the dynamics derived above to verify our claim in Section 3.1 about the diminishing marginal contribution of new buyers (non-buyers) to diffusion, and the slowdown of spread due to the growing size of the informed population. Using (2.6), the growth in the spread can be written as

$$y(t+1) = (1-x(t))(1-e^{-\lambda(p_B b_Y(t)+p_{\bar{B}}\bar{b}_Y(t)+p_B(1-p_{\bar{B}})b_Z(t))}). \tag{2.7}$$

From this, the marginal contribution of newly informed buyers $B_Y(t)$ to diffusion is

$$\frac{\partial y(t+1)}{\partial b_Y(t)} = \lambda p_B(1-x(t))e^{-\lambda(p_B b_Y(t)+p_{\bar{B}}\bar{b}_Y(t)+p_B(1-p_{\bar{B}})b_Z(t))}. \tag{2.8}$$

When only a few contribute to spread (that is $p_B b_Y(t) + p_{\bar{B}}\bar{b}_Y(t) + p_B(1-p_{\bar{B}})b_Z(t)$ is small), information diffuses at a rate of λp_B from $B_Y(t)$ to the set of uninformed agents $I - X(t)$. However, the marginal contribution to diffusion decays exponentially with the size of those participating in the spread, thus lowering the average rate of diffusion to uninformed agents; the larger the set of spreaders, the higher the chance of having friends in common, hence lowering the average rate of diffusion. The model also clearly captures the slowdown effect of the agents who have already heard about the product on diffusion as $y(t+1)$ is proportional to $1 - x(t)$. Similar arguments hold for the contribution to the spread by buyers $B_Z(t)$ and by new non-buyers $\bar{B}_Y(t)$.

3.3 Firm's Decision Problem: To Spread or to Exploit?

The profit of the firm for a durable product is given by

$$\Pi^D(u(\cdot)) = \sum_{t=0}^{\infty} \beta^t u(t) b(t), \tag{3.1}$$

where $0 < \beta < 1$ is the discount factor, $b(t)$ is the size of the buyers at time t given by (4.28), and the marginal cost of the product is assumed to be zero. Firm's objective is to find a pricing policy that maximizes the above profit, which we denote as $u^D(\cdot)$.

Given the dynamics of the information diffusion for the WOM model developed in previous section and the profit of the firm given by (3.1), the firm's problem is to decide at each time step, between optimally exploiting the network it already has by offering a price that results in the maximum immediate profit, or offering a lower price in favor of a higher spread.

A related problem is to find the maximum achievable size of the informed network via WOM. For any price function $u(\cdot)$, $x(t)$ is bounded and increasing and therefore has a limit as $t \to \infty$.

Define $q(x_0, p_B, p_{\bar{B}}; u(\cdot)) = \lim_{t\to\infty} x(t)$ as the asymptotic size of the population that can be informed about the product via WOM, starting from a uniformly randomly chosen informed population of size x_0 and following a given pricing policy $u(\cdot)$, for given values of p_B and $p_{\bar{B}}$. It is easy to see that for $x_0 < 1$ this asymptotic size is always less than 1, implying that the product cannot take over the entire population I via only WOM. This is simply due to the fact that there are $e^{-\lambda}$ isolated agents (with no friend) in I, out of which $(1-x_0)e^{-\lambda}$ of them are not in $X(0)$ and therefore will never hear about the product via WOM.

To gain more insights on the endogenous dynamics of diffusion, let us start with the case of zero price, i.e., when the product is given for free, that is $u \equiv 0$. Every agent that is informed about the product will in turn inform her friends with probability p_B. Note that since there are no non-buyers, $p_{\bar{B}}$ does not matter, so we simply choose $p_{\bar{B}} = 0$. In this case, $Z(t) = \emptyset$ and $B(t) = Y(t)$, thus the dynamics of diffusion governed by (2.1), (2.2), (4.27), and (2.6) simplifies to

$$1 - x(t+1) = (1 - x(t))e^{-\lambda p_B y(t)}, \tag{3.2}$$

$$y(t+1) = x(t+1) - x(t), \tag{3.3}$$

where $y(0) = x(0) = x_0$. Using this recursively for $t, t-1, \ldots, 0$, we obtain

$$1 - x(t+1) = (1 - x_0)e^{-\lambda p_B x(t)}. \tag{3.4}$$

The asymptotic size of the informed network for the case of zero price when informed agents engage in WOM with friends with probability p_B can now be obtained, noting that $q(x_0, p_B, 0; 0)$ should satisfy the above relation as well:

$$1 - q(x_0, p_B, 0; 0) = (1 - x_0)e^{-\lambda p_B q(x_0, p_B, 0; 0)}. \tag{3.5}$$

Based on this equation, we present several properties for $q(x_0, p_B, 0; 0)$ in the following proposition.

Proposition 3.3.1. *For every $0 < x_0 \leq 1$, the asymptotic size of the informed population for a free product where informed agents engage in the spread with probability p_B is given by the unique solution of $1 - q(x_0, p_B, 0; 0) = (1 - x_0)e^{-\lambda p_B q(x_0, p_B, 0; 0)}$ in $[0, 1]$. The solution is concave and monotonically increasing in x_0. Moreover, $q(x_0, p_B, 0; 0) > 1 - \frac{1}{\lambda p_B}$.*[19]

[19] $q(x_0, p_B, 0; 0)$ can also be represented in terms of the Lambert W function, which is defined as the

Proof. See the appendix. ■

One interesting consequence of Proposition 3.3.1 is the discontinuity in $q(x_0, p_B, 0; 0)$ at $x_0 = 0$ for $\lambda p_B > 1$. Although $q(0, p_B, 0; 0) = 0$, for any nonzero x_0 and $\lambda p_B > 1$, $q(x_0, p_B, 0; 0)$ is lowerbounded by a positive constant independent of x_0. This implies that no matter how small the size of the initially informed population is, a free product with strong engagement of agents in the spread can take over a large portion of the network via WOM given the typically large average number of friends in the networks.

The zero-price case with full engagement of agents in spread ($p_B = 1$) gives an upperbound on the achievable asymptotic size of the informed population, that is
$q(x_0, p_B, p_{\bar{B}}; u(\cdot)) \leq q(x_0, 1, 0; 0)$. In this case, every agent that is informed about the product will in turn inform all of her friends. The information will then spread throughout the network and all the agents that are reachable from an agent $i \in X(0)$ will eventually learn about the product. This upperbound can be obtained by solving $1 - q^0 = (1 - x_0)e^{-\lambda q^0}$ according to Proposition (3.3.1), where q^0 is the short-note for $q(x_0, 1, 0; 0)$.

As the main objective of this paper, we next show that under the optimal policy for a durable product price should drop to zero infinitely often. This matches the real world evidence from smartphone applications discussed in Section 3.1, where price histories witness frequent drops of the price to zero for many apps. We also present tight bounds for the asymptotic size of the spread under the optimal policy.

Theorem 3.3.1. *Under the optimal pricing policy $u^D(\cdot)$ for a durable product with zero marginal cost, where buyers and non-buyers engage in word of mouth with probabilities $0 \leq p_{\bar{B}} < p_B \leq 1$, the price drops to zero infinitely often. That is, there exists an infinite sequence of time instants $0 \leq t_0 < t_1 < \ldots$ such that $u^D(t_j) = 0$ for $j \in \mathbb{N}_0$. Moreover, the asymptotic size of the spread satisfies*

$$q(x_0, p_B, 0; 0) \leq q(x_0, p_B, p_{\bar{B}}; u^D(\cdot)) \leq q(x_0, p_B + p_{\bar{B}} - p_B p_{\bar{B}}, 0; 0). \tag{3.6}$$

Proof. First we note that for any pricing policy, there exists at least one price level that holds infinitely often. This follows from the finiteness of the set of admissible prices $\mathcal{U}$. Let

---

solution to the equation $W(z)e^{W(z)} = z$ (Corless et al. (1996)). Using this notation, we can easily show that $\lambda p_B(1 - q(x_0, p_B, 0; 0)) = -W(-\lambda p_B e^{-\lambda p_B}(1 - x_0))$. W is known to have two branches. It follows from Proposition 3.3.1 that $\lambda p_B(1 - q(x_0, p_B, 0; 0)) < 1$, requiring $W > -1$. This identifies the principal branch of the Lambert W function, denoted by W_0. Therefore, we can write $\lambda p_B(1 - q(x_0, p_B, 0; 0)) = -W_0(-\lambda p_B e^{-\lambda p_B}(1 - x_0))$. This representation enables us to use the properties of the Lambert W function, if ever needed.

$p_r \in \mathcal{U}$ be the smallest price level which holds infinitely often for $u^D(\cdot)$. Then, any price level below p_r is used finitely in $u^D(\cdot)$. Therefore, there exists $T \geq 0$ such that $u^D(t) \geq p_r$ for all $t \geq T$.

Having infinitely many drops to zero under the optimal policy $u^D(\cdot)$ is clearly equivalent to $p_r = 0$ (that is $r = 0$). Therefore, to prove the theorem, we assume $r \geq 1$ and try to reach contradiction by constructing a new policy with a profit higher than that of $u^D(\cdot)$. For this purpose, we show that by zeroing the price to sell the product to a subset of informed agents that would not buy it otherwise, and by getting them (more) engaged in the spread of information, the monopolist can reach out to a part of the network that would remain unexplored under $u^D(\cdot)$. Dropping the price to zero to access this part of the network at a proper time, we then introduce a new policy yielding a profit higher than that of $u^D(\cdot)$ by exploiting this untouched component of the network.

Let $Y_r^D(T) \subset Y^D(T)$ denote those freshly informed agents at time T whose valuations are below p_r, i.e. $Y_r^D(T) = \{i \in Y^D(T) | 0 \leq \theta_i < p_r\}$, with a size of $y_r^D(T) = F(p_r)y^D(T)$.[20] None of the agents in $[\cup_{j=1}^r Z_j^D(T)] \cup Y_r^D(T)$ will ever buy the product under the pricing policy $u^D(\cdot)$, where $\cup_{j=1}^r Z_j^D(T)$ is the set of those previously informed agents at time T whose valuations are below p_r. Now, consider the set of agents that will remain uniformed under $u^D(\cdot)$. The size of this set is clearly $1 - q^D$, where q^D is the asymptotic size of the informed population under $u^D(\cdot)$, i.e., $q^D = q(x_0, p_B, p_{\bar{B}}; u^D(\cdot))$. Define Δ_r as the subset of these agents who have at least one friend in $\cup_{j=1}^r Z_j^D(T) \cup Y_r^D(T)$, that is

$$\Delta_r = \{i \in I | (\nexists t \in \mathbb{N}_0 : i \in x^D(t)) \wedge d_i(\cup_{j=1}^r Z_j^D(T) \cup Y_r^D(T)) \neq 0\}.^{21} \tag{3.7}$$

The number of friends of an uninformed agent among $\cup_{j=1}^r Z_j^D(T) \cup Y_r^D(T)$ has a Poisson distribution with mean $\lambda(1 - p_{\bar{B}})(F(p_r)y^D(T) + \sum_{j=1}^r z_j^D(T))$, since each of these agents has already informed friends with probability $p_{\bar{B}}$. Therefore, by zeroing the price at any time $t > T$ we can reach out a subset of Δ_r with the size of

$$\delta_r = (1 - q^D)(1 - e^{-\lambda p_B (1-p_{\bar{B}})(F(p_r)y^D(T) + \sum_{j=1}^r z_j^D(T))}), \tag{3.8}$$

that could not be reached under $u^D(\cdot)$. The idea now is to show that after a while there is so little profit left to be made in the future under $u^D(\cdot)$ that it is profitable to zero the price to reach out these agents in Δ_r, as will be elaborated below.

[20] We use superscript D to indicate that the variables correspond to the pricing policy $u^D(\cdot)$.

[21] For any $S \subseteq I$ and $i \in I$, we denote the number of friends of agent i in S with $d_i(S)$.

Let t_k, $k = 1, 2, \ldots$, denote the k-th price drop to p_r after time T under the optimal policy $u^D(\cdot)$. If an agent $i \in X^D(t_k)$ does not buy the product at this time, neither will she buy it in future. This means that agents in $X^D(t_k)$ do not contribute to the set of buyers $B^D(t)$ for $t > t_k$. Therefore, the size of the buyers from time $t_k + 1$ to $t_k + \tau$ for any $\tau \geq 1$ can be upperbounded by $x^D(t_k + \tau) - x^D(t_k)$, that is

$$\sum_{t=t_k+1}^{t_k+\tau} b^D(t) \leq x^D(t_k + \tau) - x^D(t_k). \tag{3.9}$$

Letting $\tau \to \infty$, yields

$$\sum_{t=t_k+1}^{\infty} b^D(t) \leq q^D - x^D(t_k). \tag{3.10}$$

Thus, the contribution of the buyers to the firm's profit after t_k can be upperbounded by

$$\begin{aligned}
\Pi^D_{>t_k}(u^D(\cdot)) &= \sum_{t=t_k+1}^{\infty} \beta^t u^D(t) b^D(t) \\
&\leq \beta^{t_k+1} \sum_{t=t_k+1}^{\infty} u^D(t) b^D(t) \\
&\leq \beta^{t_k+1} p_m (q^D - x^D(t_k)).
\end{aligned} \tag{3.11}$$

Next, consider a new policy $\tilde{u}(\cdot)$ having the same value as $u^D(\cdot)$ at all times except $t_k + 1$ and $t_k + 2$. Let $\tilde{u}(t_k + 1) = 0$ and $\tilde{u}(t_k + 2) = u^*$, where $u^*(1 - F(u^*)) = \max_{u \in \mathcal{U}} u(1 - F(u))$. Note that a subset of agents in Δ_r with size δ_r as in (3.8) are among the freshly informed agents $\tilde{Y}(t_k + 2)$ since the agents in $\cup_{j=1}^{r} Z_j^D(T) \cup Y_r^D(T)$ buy the product at time $t_k + 1$. The distribution of θ in Δ_r is given by $F(\cdot)$, hence the discounted profit made from these newly informed agents in Δ_r at time $t_k + 2$ is $\beta^{t_k+2} u^*(1 - F(u^*))\delta_r$. Considering that $x^D(t) \to q^D$ as $t \to \infty$, we can choose k large enough such that

$$q^D - x^D(t_k) < \beta \delta_r \frac{u^*(1 - F(u^*))}{p_m}, \tag{3.12}$$

in which case the profit resulting from $\tilde{u}(\cdot)$ will be clearly higher than that coming from $u^D(\cdot)$. This contradicts the optimality of $u^D(\cdot)$, hence rejecting the initial assumption of $p_r \neq 0$, which completes the proof.

To bound the size of the spread under the optimal policy, we recursively use (2.6) and let $t \to \infty$ to obtain

$$1 - q^D = (1 - x_0)e^{-\lambda(p_B b_Y + p_{\bar{B}} \bar{b}_Y + p_B(1-p_{\bar{B}})b_Z)}, \tag{3.13}$$

where $b_Y = \sum_{t=0}^{\infty} b_Y(t)$, $\bar{b}_Y = \sum_{t=0}^{\infty} \bar{b}_Y(t)$, $b_Z = \sum_{t=0}^{\infty} b_Z(t)$, and q^D is the asymptotic size of the informed population under the optimal policy $u^D(\cdot)$. Since the optimal policy drops the price to zero infinitely often, every informed agent that is not a fresh buyer (i.e., $i \in \bar{B}_Y$), will buy sometime later on and hence is in B_Z. This implies $\bar{b}_Y = b_Z$. Also, every agent that will eventually get informed (a set with size q^D) will be either a fresh buyer ($i \in B_Y$) or not ($i \in \bar{B}_Y$). Thus, we have shown that $\bar{b}_Y = b_Z = q^D - b_Y$. Substituting this in (3.13), we get

$$1 - q^D = (1 - x_0)e^{-\lambda(p_B b_Y + (p_B + p_{\bar{B}} - p_B p_{\bar{B}})(q^D - b_Y))}, \tag{3.14}$$

which using Proposition 3.3.1 and noting that $p_B q^D \leq p_B b_Y + (p_B + p_{\bar{B}} - p_B p_{\bar{B}})(q^D - b_Y) \leq (p_B + p_{\bar{B}} - p_B p_{\bar{B}})q^D$ yields the bound $q(x_0, p_B, 0; 0) \leq q^D \leq q(x_0, p_B + p_{\bar{B}} - p_B p_{\bar{B}}, 0; 0)$. ■

Comment 3.3.1 (Asymptotic Size of the Informed Population for Buyers-Only Spread)**.** For the case where non-buyers do not contribute to the spread of information (i.e., $p_{\bar{B}} = 0$), (3.13) gives the exact size of the spread as $q^D = q(x_0, p_B, 0; 0)$, which is the same as that of a free product. This means that for the special case where only buyers engage in WOM with their friends, the policy maximizing the profit of the firm also maximizes the asymptotic size of the spread.

We can also easily extend the theorem to the case where the marginal cost is nonzero but sufficiently small. However, a significant marginal cost may shift the drops to a price level away from zero. This level is still below the marginal cost given that the gap between the marginal cost and the closest price level to it from below in $\mathcal{U}$ is sufficiently small.[22]

Theorem 3.3.1 shows infinitely many price drops to zero for a durable product under the optimal policy. A question that arises here is that whether it is possible for the optimal price trajectory to get trapped in a vicinity of zero. Noting that the price cannot stay at zero forever (due to the zero profit from such choice), this question translates to the possibility of getting stuck between price levels 0 and p_1 for small values of p_1. If such thing happens, one may even falsely attribute the price drops to a continuous-valued optimal price trajectory asymptotically converging to zero which manifests itself as a quantized price path bouncing between 0 and p_1. The following proposition rejects the possibility of such a price lockdown.

[22]See the appendix for the extension of the theorem to the case of nonzero marginal cost and the proof.

Proposition 3.3.2. *Under the optimal pricing policy* $u^D(\cdot)$, *the price jumps to a level above* p_1 *infinitely often, when* $p_1(1-F(p_1)) < c$, *where*

$$c = \max_{u \in \mathcal{U} \setminus \{p_1\}} \frac{u(1-F(u))(1-\beta\lambda(p_B + p_{\bar{B}} - p_B p_{\bar{B}})(1-q_B^0))}{1-\beta\lambda(p_B(1-F(u)) + p_{\bar{B}}F(u))(1-q_B^0)}, \tag{3.15}$$

and $q_B^0 = q(x_0, p_B, 0; 0)$ *is the asymptotic size of the informed population under the zero price policy.*

Proof. See the appendix. ■

Based on this result, a questions still remains as to how small should p_1 be to guarantee frequent price jumps to levels above p_1. In fact, as we will see in the following example, p_1 does not need to be very small to satisfy the above condition for a wide range of parameters.

Example 1 (Buyers-Only Spread and Uniform Valuations)**.** Assuming $\theta \sim Unif[0,1]$ and $p_{\bar{B}} = 0$, the condition on p_1 simplifies to $p_1 < \frac{1-\sqrt{1-4c}}{2}$, where

$$c = \max_{u \in \mathcal{U} \setminus \{p_1\}} \frac{u(1-u)(1-\beta\lambda p_B(1-q_B^0))}{1-\beta\lambda p_B(1-u)(1-q_B^0)}. \tag{3.16}$$

It is easy to verify that c is decreasing with both β and $\lambda p_B(1-q_B^0)$. As a result, the less the value of each, the looser the bound on p_1. Although it may look counterintuitive at first, the term $\lambda p_B(1-q_B^0)$ is very helpful in loosening the bound on p_1. This term is less than 1 (from Proposition 3.3.1), and indeed is decreasing with λp_B for $\lambda p_B \geq 1$.[23] We now use this background to study a few cases in order to get an insight on the values of p_1 satisfying the above condition. We assume $0.5 \in \mathcal{U}$ to use it as a (sub)maximizer in (3.16). If $\lambda p_B \geq 2$, then $\lambda p_B(1-q_B^0) < 0.41$,[24] which along with $\beta < 1$ yields $c > 0.18$, for which the condition in the above proposition reduces to $p_1 < 0.24$. Some information on β can loosen up this bound even further. For example, if we also know that $\beta \leq 0.5$, then the condition on p_1 becomes $p_1 < 0.33$. This bound gets closer to 0.5 for larger λp_B (or smaller β), assuring infinitely many price jumps to levels above p_1 even for values of p_1 that are not very small (see Appendix 3.8 for an illustrative plot).

[23] Recall that $\lambda p_B(1-q_B^0) = -W_0(-\lambda p_B e^{-\lambda p_B}(1-x_0))$, where W_0 is the principal branch of the Lambert W function (see Footnote 19). W_0 is an increasing function, and $-\lambda p_B e^{-\lambda p_B}$ is also increasing with λp_B for $\lambda p_B \geq 1$. This implies that $\lambda p_B(1-q_B^0)$ is decreasing with λp_B for $\lambda p_B \geq 1$.

[24] $\lambda p_B(1-q_B^0) \leq -W_0(-2e^{-2}) = 0.4064$, for $\lambda p_B \geq 2$.

3.4 Generalizations

In this section, we examine the robustness of the frequent zero-price drops behavior to several key assumptions that were made while developing the results of the previous section.

3.4.1 Forward-Looking Agents

The results of the previous section are derived based on the assumption that agents are myopic; an informed agent buys the product as soon as the firm offers a price below her valuation. At first, it may seem that the frequent free-offering policy may not be a good idea in face of forward-looking agents; a forward-looking agent may wait for a free offer even if the current offered price is below her valuation. The aim of this subsection is to show that in fact the frequent zeroing of the price still persists even if the agents are forward-looking.

We begin by making the necessary changes to the model. Assume that agents share a common discount factor $0 < \beta_c < 1$. An informed agent with valuation θ who pays price u at time t obtains utility $\beta_c^t(\theta - u)$. Therefore, given the optimal price path $u^D(\cdot)$, an informed agent has to choose a purchasing time τ, maximizing the utility

$$\sup_{\tau} \beta_c^{\tau}(\theta - u^D(\tau)). \tag{4.1}$$

An informed agent has the option of not buying, in which case $\tau = \infty$ and the payoff is zero. It is clear that the update rules for the size of previously informed non-buyers in (4.27) and freshly informed buyers and non-buyers in (2.1) are not valid anymore. However, they can be used to obtain bounds for the case of forward-looking agents:

$$z_j(t+1) \geq z_j(t) + (F(p_j) - F(p_{j-1}))y(t) \quad \text{for } 1 \leq j \leq r, \tag{4.2}$$

and

$$\begin{aligned} b_Y(t) &\leq (1 - F(p_r))y(t), \\ \bar{b}_Y(t) &\geq F(p_r)y(t), \end{aligned} \tag{4.3}$$

where p_r is the price level chosen by firm at time t ($u(t) = p_r$). The size of informed agents still obeys the same update rule as in (2.6).

With this model at hand, we are now ready to unravel the rationale behind frequent zeroing of the price even when the consumers are forward-looking. We sketch the main probing factors here and refer the reader to the appendix for a comprehensive proof.

The first observation is that not everyone waits for a free offer; high-valuation agents pay to buy even knowing that the price can be zero later on. In fact, offering a price $u(t)$, all informed agents with valuations $\theta > \frac{u(t)}{1-\beta_c}$ will buy the product immediately. For such agents, $\theta - u(t) > \beta_c\theta$, thus they are willing to pay price $u(t)$ to buy at time t even if they know that it would be offered for free at time $t+1$. In other words, for consumers with high enough valuations, the myopic and forward-looking behaviors coincide. The second observation is that similar to the myopic case, there will still be a growing pile of low-valuation agents that do not buy the product unless offered for free, including those with valuations below the smallest price level p_1.

From the above discussion, we see that there are two subsets of agents whose actions are the same no matter whether they are myopic or forward-looking: these are agents with very high or very low valuations. As a result, the same mechanism as in the case of myopic agents triggers the frequent dropping of the price to zero; firm drops the price to zero to reach out uninformed high-valuation, willing-to-pay agents via low-valuation free-riders by engaging them in the spread at a higher rate; a profitable component which would remain untouched otherwise. We therefore have the following result.

Proposition 3.4.1. *Consider the same setup as in Theorem 3.3.1, with forward-looking agents sharing a common discount factor* $0 < \beta_c < 1 - p_1$. *Then, the optimal pricing policy* $u^D(\cdot)$ *will drop the price to zero infinitely often*

Proof. See the appendix. ■

3.4.2 Homophily-Based WOM

In our base model in Section 3.2, consumers are indiscriminate in passing the information about the product to their friends. Buyers engage in WOM with each friend with the same probability p_B and non-buyers pass on the information with the same probability $p_{\bar{B}}$ to friends when they hear about the product. In this section we aim to extend our model and results to the case where agents' engagement in WOM is based on homophily: Agents tend to get engaged in WOM about the product with friends they believe to have similar valuations for the product. The rationale of zeroing the price to reach out high-valuation willing-to-pay uninformed agents via WOM of low-valuation agents may at first seem to fail here because low-valuation agents tend to engage in WOM with other low-valuation agents based on homophily. However, as we will see, under some mild assumptions on the homophily functions, the same rationale for the price drops still holds except that here we may need a chain of agents (and price drops) to reach the high-valuation uninformed agents.

The social network structure is same as before. We embed homophily in the model by considering valuation-dependent probabilities for the engagement of the informed agents in the spread. A buyer with valuation θ informs a friend whose valuation is θ' with probability $p_B(\theta, \theta')$. A non-buyer with valuation θ passes the information to a friend with valuation θ' with probability $p_{\bar{B}}(\theta, \theta')$, when she hears about the product. [25]

Assuming the independence of the engagement in WOM from the valuations in our base model of Section 3.2 had the advantage of keeping the distribution of θ invariant among the newly informed agents, and we only needed to keep track of the distribution of θ among non-buyers. Here, we will need to keep track of the distribution of θ in all sets in play. Using $\mu_Y(\theta, t)$ to denote the PDF of θ in $Y(t)$ and $M_Y(\theta, t)$ for the corresponding CDF[26] and assuming an offered price $u(t) = p_r$, we can find the size of the sets $B_Y(t)$ and $\bar{B}_Y(t)$ as

$$\begin{aligned} b_Y(t) &= (1 - M_Y(p_r, t))y(t), \\ \bar{b}_Y(t) &= M_Y(p_r, t)y(t), \end{aligned} \tag{4.4}$$

and the distribution of θ in each as

$$\begin{aligned} \mu_{B_Y}(\theta, t) &= \mathbf{1}(\theta \geq p_r)\frac{\mu_Y(\theta, t)}{1 - M_Y(p_r, t)}, \\ \mu_{\bar{B}_Y}(\theta, t) &= \mathbf{1}(\theta < p_r)\frac{\mu_Y(\theta, t)}{M_Y(p_r, t)}. \end{aligned} \tag{4.5}$$

Similar relations hold for the size of buyers and non-buyers among previously informed agents $Z(t)$:

$$\begin{aligned} b_Z(t) &= (1 - M_Z(p_r, t))z(t), \\ \bar{b}_Z(t) &= M_Z(p_r, t)z(t), \end{aligned} \tag{4.6}$$

[25]This is similar to the multi-type random networks model of Golub and Jackson (2012), where types affect the formation of the links. Here, types (valuations) affect the likelihood of engaging two friends in WOM about the product.

[26]We use the same convention of using $\mu_S(\theta, t)$ and $M_S(\theta, t)$ to denote the PDF and CDF of θ in any set $S(t) \subseteq I$.

with corresponding distributions

$$\mu_{B_Z}(\theta,t) = \mathbf{1}(\theta \geq p_r)\frac{\mu_Z(\theta,t)}{1-M_Z(p_r,t)},$$
$$\mu_{\bar{B}_Z}(\theta,t) = \mathbf{1}(\theta < p_r)\frac{\mu_Z(\theta,t)}{M_Z(p_r,t)}. \tag{4.7}$$

In order to find the dynamics of the diffusion, we pick an uninformed agent $i \in \bar{X}(t)$ with valuation θ' (where $\bar{X}(t) = I \setminus X(t)$) and study the likelihood of her getting informed about the product via a contributor to the spread (that is, agents in $B_Y(t) \cup B_Z(t) \cup \bar{B}_Y(t)$). Following the same steps as in our base model, we can easily see that the number of informed agents that may engage in WOM with this agent has a Poisson distribution in each of these sets, except that the mean values here are time-varying and valuation-dependent:

$$\lambda_{B_Y}(\theta',t) = \lambda b_Y(t)\mathbb{E}_{B_Y}[p_B(\theta,\theta')|\theta'] = \lambda b_Y(t)\int p_B(\theta,\theta')\mu_{B_Y}(\theta,t)d\theta,$$
$$\lambda_{\bar{B}_Y}(\theta',t) = \lambda \bar{b}_Y(t)\mathbb{E}_{\bar{B}_Y}[p_{\bar{B}}(\theta,\theta')|\theta'] = \lambda \bar{b}_Y(t)\int p_{\bar{B}}(\theta,\theta')\mu_{\bar{B}_Y}(\theta,t)d\theta, \tag{4.8}$$
$$\lambda_{B_Z}(\theta',t) = \lambda b_Z(t)\mathbb{E}_{B_Z}[p_B(\theta,\theta')(1-p_{\bar{B}}(\theta,\theta'))|\theta'] = \lambda b_Z(t)\int p_B(\theta,\theta')(1-p_{\bar{B}}(\theta,\theta'))\mu_{B_Z}(\theta,t)d\theta,$$

and the update rule for the size of the informed population $x(t)$ becomes

$$1-x(t+1) = (1-x(t))\mathbb{E}_{\bar{X}}[e^{-\lambda(\theta',t)}] = (1-x(t))\int e^{-\lambda(\theta',t)}\mu_{\bar{X}}(\theta',t)d\theta', \tag{4.9}$$

where $\lambda(\theta',t) = \lambda_{B_Y}(\theta',t) + \lambda_{\bar{B}_Y}(\theta',t) + \lambda_{B_Z}(\theta',t)$ and $\mu_{\bar{X}}(\theta',t)$ is the PDF of θ' in $\bar{X}(t)$. Finally, we need the update rules for $\mu_{\bar{X}}$, μ_Y, and μ_Z:

$$\mu_{\bar{X}}(\theta,t+1) = \frac{1-x(t)}{1-x(t+1)}\mu_{\bar{X}}(\theta,t)e^{-\lambda(\theta,t)},$$
$$\mu_Y(\theta,t+1) = \frac{1-x(t)}{y(t+1)}\mu_{\bar{X}}(\theta,t)(1-e^{-\lambda(\theta,t)}),$$
$$\mu_Z(\theta,t+1) = \frac{\bar{b}_Y(t)\mu_{\bar{B}_Y}(\theta,t)+\bar{b}_Z(t)\mu_{\bar{B}_Z}(\theta,t)}{\bar{b}_Y(t)+\bar{b}_Z(t)}. \tag{4.10}$$

To sum up, the model dynamics for homophily-based WOM diffusion is governed by (4.4)-(4.10). Clearly, this is not as compact as the base model in Section 3.2 where $p_B(\theta,\theta') \equiv p_B$ and $p_{\bar{B}}(\theta,\theta') \equiv p_{\bar{B}}$ but is remarkably still tractable. Extracting a few observations from these equations, we will be able to use a similar proof argument for the profitability of the infinitely often dropping of the price to zero. Using the update rule for $\mu_{\bar{X}}(\theta,t)$ in (4.10) recursively, we can obtain

$$\mu_{\bar{X}}(\theta, t+1) = \frac{1-x_0}{1-x(t+1)} e^{-\sum_{\tau=0}^{t} \lambda(\theta,\tau)} f(\theta), \tag{4.11}$$

where $f(\cdot)$ is the PDF of θ in I. The maximum contribution of an informed agent to the summation $\sum_{\tau=0}^{t} \lambda(\theta, \tau)$ is upperbounded by λ if she ever buys the product and zero otherwise. This gives a lowerbound for the PDF of θ in uninformed regions as

$$\mu_{\bar{X}}(\theta, t+1) \geq \frac{1-x_0}{1-x(t+1)} e^{-\lambda q} f(\theta), \tag{4.12}$$

where q is the asymptotic size of the informed population. Integrating the above relation over any range of valuations $[\underline{\theta}, \bar{\theta}]$, we get

$$(M_{\bar{X}}(\bar{\theta}, t+1) - M_{\bar{X}}(\underline{\theta}, t+1))(1 - x(t+1)) \geq e^{-\lambda q}(F(\bar{\theta}) - F(\underline{\theta}))(1 - x_0). \tag{4.13}$$

We refer to the above property as *valuation diversity preservation of WOM among uninformed agents*: if a range of valuations has an initial nonzero measure among uninformed agents $I - X_0$, it will have a nonzero measure in the unexplored part of the network at all times, at least as large as its initial size in $I - X_0$ scaled by $e^{-\lambda q}$. This property holds for any general functions $p_B(\theta, \theta')$ and $p_{\bar{B}}(\theta, \theta')$, and is due to the nature of WOM and not homophily. This will prove very useful.

To proceed further, we impose two conditions on $p_B(\theta, \theta')$ and $p_{\bar{B}}(\theta, \theta')$. The first condition is a *local homophily condition*, requiring a buyer to engage in WOM with those friends having very similar valuations with a probability bounded away from zero. That is, there exist $\delta, \underline{p} > 0$ such that if $|\theta - \theta'| < \delta$, then $p_B(\theta, \theta') > \underline{p}$. The second condition is that the contribution of a non-buyer to spread should be upperbounded by a positive number less than 1. That is, there exists $\bar{p} < 1$ such that $p_{\bar{B}}(\theta, \theta') < \bar{p}$. We refer to this condition as the *limited engagement of non-buyers in spread.* These two conditions are indeed very general and hold even for many non-homophilous functions. We can use the local homophily condition to show that, no matter the pricing policy, there will always be a nonzero mass of low-valuation yet informed non-buyers around.[27] Moreover, we can show that having a nonzero measure set of agents with valuations in the range $[\theta_0, \theta_0 + \frac{\delta}{2}]$ among non-buyers, the firm can reach out a nonzero mass of higher valuations in range of $[\theta_0 + \frac{\delta}{2}, \theta_0 + \delta]$. It should be now clear that by dropping the price to zero (possibly several times in a row depending

[27]See the proof of Proposition 3.4.2 for the details.

on the radius of homophily δ) the firm can reach out high-valuation willing-to-pay agents from low-valuation agents. We therefore have the following result.

Proposition 3.4.2. *Suppose that the probabilities of the engagement of the buyers and non-buyers in WOM spreading of the product are given by the valuation-dependent functions* $p_B, p_{\bar{B}} : [0,1]^2 \to [0,1]$ *satisfying the following conditions:*
i) there exist $\delta, \underline{p} > 0$ *such that if* $|\theta - \theta'| < \delta$*, then* $p_B(\theta, \theta') > \underline{p}$*, and*
ii) there exists $\bar{p} < 1$ *such that* $p_{\bar{B}}(\theta, \theta') < \bar{p}$*.*
Then, the optimal pricing policy $u^D(\cdot)$ *for a durable product drops the price to zero infinitely often.*

Proof. See the appendix. ■

3.4.3 Network Externalities

A firm offering its product in a social network can leverage the spread of its product from network externalities. An informed agent who does not buy the product at a given price may do so later on, if many of her friends buy the product, even if the firm does not lower the price. This raises another interesting question as to whether the price drops would be still profitable in the presence of network externalities. The aim of this section is to formalize and answer this question.

When the product exhibits network externalities, an informed agent buys the product if the offered price does not exceed the sum of her valuation and the total externalities from her friends whom she knows are already using the product. Denote as $\mathfrak{B}(t) = \cup_{\tau=0}^{t-1} B(\tau)$ the set of all previous buyers at time t and let $0 < \alpha \leq 1$ represent the network externality effect. Then, an informed agent $i \in I$ buys the product at time t if the offered price $u(t)$ does not exceed her *augmented* valuation defined as $\theta_i^a(t) = \theta_i + \alpha d_i^{WOM}(\mathfrak{B}(t))$, where $d_i^{WOM}(\mathfrak{B}(t))$ denotes the number of friends at time t who have already bought the product and have engaged in WOM with agent i about the product.

The first step in the analysis is to identify the set of buyers and non-buyers at time t. As before, new buyers are either among the freshly informed agents $Y(t)$ or among those previously informed non-buyers, denoted as $Z(t)$. Define the set of those agents in $Y(t)$ whose augmented valuations are below θ^a as $Y(\theta^a, t)$ and its size by $y(\theta^a, t)$. Note that $y(\theta^a, t)$ fully characterizes the distribution of the augmented valuation in $Y(t)$.[28] Similarly, we use

[28] In fact, $y(\theta^a, t)$ is the CDF of θ^a in $Y(t)$ multiplied by its size $y(t)$.

$Z(\theta^a, t)$ and $z(\theta^a, t)$ to represent the set of those agents in $Z(t)$ whose augmented valuations are below θ^a and its size, respectively.[29] Having been offered a price $u(t) \in \mathcal{U}$, agents in $\bar{B}_Y(t) = Y(u(t), t)$ and $\bar{B}_Z(t) = Z(u(t), t)$ do not have a high enough augmented valuation to buy the product at this price and will form $Z(t+1)$. For agents in $B_Y(t) = Y(t) \setminus \bar{B}_Y(t)$ and $B_Z(t) = Z(t) \setminus \bar{B}_Z(t)$, augmented valuations are higher than (or equal to) $u(t)$ and therefore they will buy the product.

Upon buying the product, buyers and non-buyers engage in WOM with their friends with probabilities p_B and $p_{\bar{B}}$. We need to distinguish between the agents who only hear about the product from non-buyers and those who also hear from some of the buyers as well. Consider the partition $Y(t+1) = Y_0(t+1) \cup Y_{>0}(t+1)$. Here, $Y_0(t+1)$ are those freshly informed agents who have not heard from any friend in $B_Y(t) \cup B_Z(t)$, if they have any friend among them. $Y_{>0}(t)$, on the other hand, is the set of those freshly informed agents who have heard about the product from a friend in $B_Y(t) \cup B_Z(t)$. Recalling that an agent only receives externality from friends whom she knows are using the product, the augmented valuation for any freshly informed agent will be $\theta_i^a(t+1) = \theta_i + \alpha d_i^{WOM}(B(t))$. Using this we can find the update rule for $y_{>0}(\theta^a, t+1)$ and $y_0(\theta^a, t+1)$ as

$$
\begin{aligned}
y_{>0}(\theta^a, t+1) &= (1 - x(t)) e^{-\lambda_B(t)} \sum_{d=1}^{\infty} \frac{(\lambda_B(t))^d}{d!} F(\theta^a - \alpha d), \\
y_0(\theta^a, t+1) &= (1 - x(t)) e^{-\lambda_B(t)} (1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(t)}) F(\theta^a),
\end{aligned} \tag{4.14}
$$

where $\lambda_B(t) = \lambda p_B(b_Y(t) + (1 - p_{\bar{B}}) b_Z(t))$ is the average number of friends among the buyers $B(t)$ that an uninformed agent hears from.

Finding the update rule for $Z(\theta^a, t+1)$ is much more involved, mainly due to the engagement of non-buyers in the spread. Whether or not an agent is informed about the product via a non-buyer changes the future likelihood of them getting engaged in WOM, if the non-buyer ever buys the product. To include this into the diffusion dynamics we need to keep track of both the time an agent first hears about the product and whether she first learns about it only through non-buyers, or not. We write $Z(t) = Z^0(t) \cup \ldots \cup \; Z^{t-1}(t)$, where we use the superscript to time stamp the moment when a non-buyer first hears about the product. That is, $Z^\tau(t)$, where $0 \leq \tau \leq t-1$, are those agents who heard about the product at time τ and have not yet made a purchase by time t. We also use the subscript to indicate whether an agent has first learned about the product via only non-buyers or not. So, for each

[29] In general, for any $S(t) \subseteq I$, we use the notation $S(\theta^a, t)$ to denote the set of those agents in $S(t)$ whose augmented valuations are below θ^a and $s(\theta^a, t)$ to denote its size.

$0 \leq \tau \leq t-1$, we write $Z^\tau(t) = Z_0^\tau(t) \cup Z_{>0}^\tau(t)$.[30] Similarly we can partition the set of buyers according to the time they first hear about the product: $B(t) = B^0(t) \cup \ldots \cup B^{t-1}(t) \cup B^t(t)$, for which

$$B^\tau(t) = \begin{cases} Z^\tau(t) \setminus Z^\tau(u(t), t), & \text{for } 0 \leq \tau \leq t-1 \\ B_Y(t), & \text{for } \tau = t. \end{cases} \tag{4.15}$$

To complete the dynamics, we note that $Z(t+1) = Z^0(t+1) \cup \ldots \cup Z^t(t+1)$, and

$$Z^\tau(t+1) = \begin{cases} Z^\tau(u(t), t), & \text{for } 0 \leq \tau \leq t-1 \\ \bar{B}_Y(t), & \text{for } \tau = t. \end{cases} \tag{4.16}$$

The augmented valuation of an agent $i \in Z^\tau(t+1)$ may increase by time $t+1$ as some of her friends may buy the product at time t and let her know about it. More precisely, $\theta_i^a(t+1) = \theta_i^a(t) + \alpha d_i^{WOM}(B(t))$. Therefore, in order to find the update rule for $z(\theta^a, t+1)$ we need to determine how many friends an agent $i \in Z^\tau(t+1)$ interacts with among the new buyers $B(t)$. The distribution of interactions varies depending on the time an agent first gets informed and whether or not a buyer was involved in informing her. To see this heterogeneity, note that the likelihood of the engagement of an agent $i \in Z^\tau(t+1)$, is highest for agents in $B_0^{\tau+1}(t)$ and lowest for buyer at time t who were informed before $\tau - 1$. This is because agents in $B_0^{\tau+1}(t)$ were first informed via new non-buyers at time τ and agent i was one of them. This increases the posterior probability of them engaging in WOM. On the other hand, agents informed before $\tau - 1$ engaged in WOM once when they first heard about the product but did not get engaged in WOM with i since she just heard about the product at time τ, lowering the posterior probability of them engaging in WOM when they buy the product.

Using the Bayes update rule and some manipulation we can show that number of buyers in $B(t)$ that agent $i \in Z_{>0}^\tau(t+1)$ may hear from has a Poisson distribution with mean[31]

$$\lambda_{>0}^\tau(t) = \lambda p_B \left((1 - p_{\bar{B}}) b(t) + p_{\bar{B}} (b^{\tau-1}(t) + b^\tau(t) + b^{\tau+1}(t)) + \frac{e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau)}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau)}} p_{\bar{B}} b_0^{\tau+1}(t) \right). \tag{4.17}$$

Using this, the update rule for $z_{>0}^\tau(\theta^a, t+1)$ is

$$z_{>0}^\tau(\theta^a, t+1) = \sum_{d=0}^{\infty} e^{-\lambda_{>0}^\tau(t)} \frac{(\lambda_{>0}^\tau(t))^d}{d!} z_{>0}^\tau(\min(\theta^a - \alpha d, u(t)), t), \tag{4.18}$$

[30] Similarly, for any $S(\cdot) \subseteq I$ we use the notation $S_0(\cdot)$ to represent those in $S(\cdot)$ who were first informed only via non-buyers and $S_{>0}(\cdot)$ to denote the rest.

[31] To improve the readability, the details are moved to the appendix.

for $0 \leq \tau \leq t-1$, and

$$z_{>0}^{t}(\theta^a, t+1) = \sum_{d=0}^{\infty} e^{-\lambda_{>0}^{t}(t)} \frac{(\lambda_{>0}^{t}(t))^d}{d!} y_{>0}(\min(\theta^a - \alpha d, u(t)), t), \tag{4.19}$$

for $\tau = t$. The update rule for $z_0^{\tau}(\theta^a, t+1)$ is even more involved. Agents in $Z_0^{\tau}(t+1)$ were informed via non-buyers in $\bar{B}_Y(\tau-1)$. Therefore, the number of WOM engagement they have with those in $\bar{B}_Y(\tau-1)$ that buy the product at time t also depends on the number of their engagement they have already had with others in $\bar{B}_Y(\tau-1)$ that have made a purchase before t. Therefore, we need to break down $z_0^{\tau}(\theta^a, t+1)$ further. We write $z_0^{\tau}(\theta^a, t+1) = \sum_{d=0}^{\infty} z_0^{\tau}(\theta^a, d, t+1)$, where $z_0^{\tau}(\theta^a, d, t+1)$ denotes the mass of non-buyers at time $t+1$, first informed via non-buyers in $\bar{B}_Y(\tau-1)$ at time τ, who has engaged with d friends in $\bar{B}_Y(\tau-1)$ that have bought the product sometime between τ and t, that is $d_i^{WOM}(\cup_{t'=\tau}^{t} B^{\tau-1}(t')) = d$. We can find the update rule for $z_0^{\tau}(\theta^a, d, t+1)$ as[31]

$$\begin{aligned} z_0^{\tau}(\theta^a, d, t+1) =& (1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(\tau-1) - p_B \sum_{t'=\tau}^{t} b^{\tau-1}(t'))}(1-p_{\bar{B}})^d) \times \\ & \sum_{k=0}^{d} \frac{e^{-\lambda p_B b^{\tau-1}(t)} \frac{(\lambda p_B b^{\tau-1}(t))^{d-k}}{(d-k)!}}{1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(\tau-1) - p_B \sum_{t'=\tau}^{t-1} b^{\tau-1}(t'))}(1-p_{\bar{B}})^k} \times \\ & \sum_{d'=0}^{\infty} e^{-\lambda_0^{\tau}(t)} \frac{(\lambda_0^{\tau}(t))^{d'}}{d'!} z_0^{\tau}(\min(\theta^a - \alpha(d+d'-k), u(t)), k, t), \end{aligned} \tag{4.20}$$

for $0 \leq \tau \leq t-1$, and

$$\begin{aligned} z_0^{t}(\theta^a, d, t+1) =& (1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(t-1) - p_B b^{t-1}(t))}(1-p_{\bar{B}})^d) \times \\ & \frac{e^{-\lambda p_B b^{t-1}(t)} \frac{(\lambda p_B b^{t-1}(t))^d}{d!}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(t-1)}} \sum_{d'=0}^{\infty} e^{-\lambda_0^{t}(t)} \frac{(\lambda_0^{t}(t))^{d'}}{d'!} y_0(\min(\theta^a - \alpha(d+d'), u(t)), t), \end{aligned} \tag{4.21}$$

where

$$\lambda_0^{\tau}(t) = \lambda_{>0}^{\tau}(t) - \lambda p_B b^{\tau-1}(t). \tag{4.22}$$

Despite the above complex dynamics, the rationale for persistence of the zero price drops is rather simple. The first observation is that unless the externality is very strong, the optimal price path will jump above α infinitely often. As we will see, this property is very similar to Proposition 3.3.2. We then use the above dynamics to show that raising the price above α will bring in some non-buyers which will inform a nonzero measure subset of low-valuation agents (with valuations below p_1) given $p_{\bar{B}} > 0$. Due to network externalities, these low-valuation agents may eventually buy the product if many of their friends do so, elevating their augmented valuations above p_1. The next observation is to show that this cannot vanish the set of low-valuation agents. In fact by recursively using (4.20), we can show that[31]

$$z_0^\tau(p_1, t) \geq (1 - p_B)e^{-\lambda(p_B + p_{\bar{B}})} y_0(p_1, \tau), \tag{4.23}$$

for $t > \tau$, that is, a nonzero fraction of those low-valuation agents that get informed about the product at time τ will never buy the product unless we drop the price to zero. Therefore, the same intuition of dropping the price to reach out new parts of the network via these low-valuation agents still holds in the presence of network externalities.

Proposition 3.4.3. *Consider the same setup as in Theorem 3.3.1 with $0 < p_{\bar{B}} < p_B$, and assume that the network externality effect $0 < \alpha < 1$ satisfies*

$$\min_{u \in \mathcal{U}} \frac{\alpha(1 - \beta\lambda p_B(1 - F(u - \alpha))(1 - \hat{q}_B^0))}{u(1 - F(u - \alpha))} < 1 - \beta\lambda p_B(1 - q_B^0), \tag{4.24}$$

where $q_B^0 = q(x_0, p_B, 0; 0)$ and $\hat{q}_B^0 = q(x_0, p_B + p_{\bar{B}} - p_B p_{\bar{B}}, 0; 0)$. Then, the optimal pricing policy should drop the price to zero infinitely often.

Proof. See the appendix. ∎

How small the externality effect α needs to be to satisfy the condition in (4.24)? To get an idea, let us consider the case of uniform valuations $\theta \sim Unif[0, 1]$ and assume $\frac{1+\alpha}{2} \in \mathcal{U}$ to use it as a sub-minimizer in (4.24). After some simple algebra and noting that $q_B^0 \leq \hat{q}_B^0$, the condition reduces to $\alpha < \frac{1-\beta\lambda p_B(1-\hat{q}_B^0)}{1+\beta\lambda p_B(1-\hat{q}_B^0)}$. We can now state the following corollary.

Corollary 3.4.1. *Suppose that $\theta \sim Unif[0, 1]$ and that $\frac{1+\alpha}{2} \in \mathcal{U}$. If the network externality effect satisifies*

$$\alpha < \frac{1 - \beta\lambda p_B(1 - \hat{q}_B^0)}{1 + \beta\lambda p_B(1 - \hat{q}_B^0)}, \tag{4.25}$$

where $\hat{q}_B^0 = q(x_0, p_B + p_{\bar{B}} - p_B p_{\bar{B}}, 0; 0)$, then the optimal pricing policy should drop the price to zero infinitely often.

In fact, the network externality needs to be very strong to fail (4.25). This bound is decreasing with β and $\lambda p_B(1 - \hat{q}_B^0)$. Also, $\lambda p_B(1 - \hat{q}_B^0) \leq \lambda p_B(1 - q_B^0) < 1$ according to Proposition 3.3.1 and as we explained in Example 1, $\lambda p_B(1 - q_B^0)$ is decreasing with λp_B for $\lambda p_B \geq 1$. With this background, let us look at a couple of examples. If $\lambda p_B \geq 2$, that is each buyer passes the information to at least two friends in average, then for α to fail the condition we should have $\alpha > 0.42$. [32] We can improve this bound if we also have some information about β. For example, if we also know that $\beta < 0.5$, then to fail we should have $\alpha > 0.66$. These are extremely strong externalities, as $\alpha > 0.42$ ($\alpha > 0.66$) suggests that everyone with three (two) friends among the buyers will buy the product even if the price is at its maximum ($u(t) = 1$).

3.4.4 Consumer Inattention to Price Changes

The core model of Section 3.2 assumes that consumers have unlimited ability of tracking down the price changes. An agent who knows about the product, is available to make a purchase as soon as the price falls below her valuation. With limited attention to price changes, consumers may check for a better deal every now and then. This section aims to verify the persistence of the frequent zero-price drops when consumers are partially attentive to price changes.

To model this, we assume that each informed non-buyer checks for price changes with an attentiveness probability $0 < p_A \leq 1$ at each subsequent time instant. This results in a geometrical distribution for the next time this agent becomes available to make a purchase.

The extension of the model in this case is quite straightforward. Instead of the whole set $Z(t)$, only a fraction p_A of them will be available to make a purchase. Therefore, the size of the buyers among the previously informed agents at time t in (2.2) changes to

$$b_Z(t) = p_A \sum_{j=r+1}^{m} z_j(t), \tag{4.26}$$

assuming a price $u(t) = p_r$. This consequently changes the update rule for the distribution of valuations among informed non-buyers in (4.27):

$$z_j(t+1) = \begin{cases} z_j(t) + (F(p_j) - F(p_{j-1}))y(t), & 1 \leq j \leq r \\ (1 - p_A)z_j(t), & \text{otherwise} \end{cases} \tag{4.27}$$

implying that $(1 - p_A)$ fraction of agents who would make a purchase if they had noticed the new price $u(t) = p_r$ miss the opportunity to buy the product. The total size of the new

[32]See Example 1 and Footnotes 23-24 on how to derive this.

buyers in (4.28) becomes

$$b(t) = (1 - F(p_r))y(t) + p_A \sum_{j=r+1}^{m} z_j(t), \tag{4.28}$$

and the dynamics of the spread is given by (2.6) as before. The rationale for dropping the price to zero to reach out to high-valuation agents via low-valuation agents is preserved in this case, except that the bridge to the otherwise untouchable profitable component is taken via a fraction p_A of the low-valuation agents that would notice the free-offer. The other point worth mentioning here is that the immediate loss in the profit by dropping the price to zero may have a higher margin in this case; some high-valuation agents may be still around at the time of the drop even if the firm has already used low prices, due to the consumers' limited attention to price changes. We formally present this result in the next proposition.

Proposition 3.4.4. *Consider the same setup as in Theorem 3.3.1 with the exception that consumers pay limited attention to price changes. That is, upon learning about the product, each non-buyers checks back for price changes with some probability* $0 < p_A \leq 1$ *at each subsequent time instant. Then, the optimal pricing policy* $u^D(\cdot)$ *will drop the price to zero infinitely often.*

Proof. See the appendix. ■

3.5 Dynamic Pricing for a Nondurable Product

As we saw in the previous sections, the WOM nature of the information diffusion is a key deriver for dropping the price to zero. If agents (users) are not involved in spreading the information about the product, the firm will not have any incentive to drop the price to zero. In fact, it is easy to show that for the case of full information, in which everybody is directly informed by the firm, the optimal pricing policy is monotone (decreasing) exploiting those who are willing to pay more first and then gradually lowering the price. The aim of this section is to show that beside the WOM nature of the information diffusion, these drops are also caused by the durability of the product.

For a nondurable product, every agent $i \in X(t)$ can buy the product if the offered price is below her valuation. The size of the buyers at time t is $(1 - F(u(t)))x(t)$, and therefore the accumulated discounted profit of the firm over an infinite time horizon is given by

$$\Pi^{ND}(u(\cdot)) = \sum_{t=0}^{\infty} \beta^t u(t)(1 - F(u(t)))x(t). \tag{5.1}$$

Firm's objective is to find the optimal pricing policy $u^{ND}(\cdot)$ that maximizes the above profit. An informed agent can buy a nondurable product as many times as the offered price is below her valuation, while she may buy a durable product only once. Therefore, in order to keep the dynamics of the spread the same for both cases, we focus on the case where $p_B = 1$ and $p_{\bar{B}} = 0$. In this setting, despite the type of the product, an agent informs her friends about the product as soon as she buys it.

Denote by u^* the price level maximizing the immediate profit, that is, $u^* = \text{argmax}_{u \in \mathcal{U}}\, u(1 - F(u))$. If there are two such price levels in $\mathcal{U}$, denote the smaller one with u^*. A useful observation is that $u^{ND}(t) \leq u^*$ for all $t \geq 0$. Otherwise, lowering the price to u^* would increase both the immediate profit and the size of the informed population at future times. Next theorem presents a steady state fixed-price property for the optimal pricing policy of a nondurable product.

Theorem 3.5.1. *Given the optimal pricing policy $u^{ND}(\cdot)$ for a nondurable product, there exists a finite time T after which the price is set to the fixed level u^* maximizing the immediate profit, that is $u^{ND}(t) = u^*$ for $t \geq T$.*

Proof. Denote by q^{ND} the asymptotic size of the informed population under the optimal policy $u^{ND}(\cdot)$, i.e., $q^{ND} = q(x_0, 1, 0; u^{ND}(\cdot))$. We claim that when the size of the informed population gets large enough, then no price other than u^* can be used by the optimal policy. In particular, we show that if $x^{ND}(t) > \gamma q^{ND}$, where

$$\gamma = \max_{\{u \in \mathcal{U} | u < u^*\}} \frac{\beta u^*(1 - F(u^*))}{u^*(1 - F(u^*)) - (1 - \beta)u(1 - F(u))}, \tag{5.2}$$

then $u^{ND}(t) = u^*$. Clearly $\gamma < 1$ since $u(1 - F(u)) < u^*(1 - F(u^*))$ for every $u < u^*$ in $\mathcal{U}$.

In order to prove the above claim, we again use contradiction and assume there is some time t_0 at which $x^{ND}(t_0) > \gamma q^{ND}$, but $u^{ND}(t_0) \neq u^*$, and we try to reach contradiction by constructing a new policy with a higher profit. We construct the new policy $\tilde{u}(\cdot)$, by shifting $u^{ND}(\cdot)$ one step ahead for $t > t_0$, changing the price to u^* at t_0, and keeping the policy unchanged for $t < t_0$. More specifically, we have

$$\tilde{u}(t) = \begin{cases} u^{ND}(t), & t < t_0 \\ u^*, & t = t_0 \\ u^{ND}(t-1), & t > t_0. \end{cases} \tag{5.3}$$

The key observation here is to note that defining $\tilde{u}$ in this way, any agent who is informed about the product under the optimal policy $u^{ND}(\cdot)$ will also be informed under the new policy $\tilde{u}(\cdot)$ with at most one step delay. This assures $X^{ND}(t-1) \subseteq \tilde{X}(t)$ for $t > t_0$, implying that $x^{ND}(t-1) \leq \tilde{x}(t)$ for $t > t_0$. Using this, we can lowerbound the accumulated discounted profit under the new policy $\tilde{u}(\cdot)$ from time t_0 on by the immediate profit under this policy, plus the accumulated discounted profit under the optimal policy $u^{ND}(\cdot)$ from time t_0 on discounted by β to account for the one step delay. This can be written as

$$\begin{aligned} \Pi^{ND}_{\geq t_0}(\tilde{u}(\cdot)) &= \sum_{t=t_0}^{\infty} \beta^t \tilde{u}(t)(1 - F(\tilde{u}(t)))\tilde{x}(t) \\ &\geq \beta^{t_0} u^*(1 - F(u^*))x^{ND}(t_0) + \beta \Pi^{ND}_{\geq t_0}(u^{ND}(\cdot)), \end{aligned} \tag{5.4}$$

where $\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot))$ is the accumulated discounted profit under the optimal policy $u^{ND}(\cdot)$ from time t_0 on.[33]

The profit of the firm for $t < t_0$ is the same under both policies. Therefore, in order to prove that the new policy $\tilde{u}(\cdot)$ results in a higher profit than $u^{ND}(\cdot)$, we need to show that $\Pi^{ND}_{\geq t_0}(\tilde{u}(\cdot)) > \Pi^{ND}_{\geq t_0}(u^{ND}(\cdot))$. Applying (5.4), it thus suffices to show that

$$(1 - \beta)\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot)) < \beta^{t_0} u^*(1 - F(u^*))x^{ND}(t_0). \tag{5.5}$$

We can find an upperbound for $\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot))$ as[34]

$$\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot)) \leq \beta^{t_0} u^{ND}(t_0)(1 - F(u^{ND}(t_0)))x^{ND}(t_0) + \frac{\beta^{t_0+1}}{1-\beta} u^*(1 - F(u^*))q^{ND}. \tag{5.6}$$

Using the above upperbound along with the fact that $x^{ND}(t_0) > \gamma q^{ND}$, where γ is defined in (5.2), and after some simplifications, we can easily show that (5.5) holds. This completes the proof. ■

Although the above theorem assures that there will be no free-offering of the product

[33] See Appendix 3.7 for more details on how to obtain (5.4).

[34] See Appendix 3.7 for more details on how to obtain (5.6) and to use it in proving (5.5).

after some finite time in the nondurable case, it is still possible that the firm drops the price to zero during the early stages in order to expand its network, as shown in the next proposition.

Proposition 3.5.1. *Consider the optimal pricing policy $u^{ND}(\cdot)$ for a nondurable product and assume that $\beta > \frac{1-F(u^*)}{u^F(\lambda^*-1)}$, where $\lambda^* = (1-F(u^*))\lambda$ and*

$$u^F = \frac{\min_{p_1 \le u \le u^*} u(1-F(u))}{\max_{p_1 \le u \le u^*} \frac{u(1-F(u))}{F(u)}}. \tag{5.7}$$

Then, there exists $x^c > 0$ such that for $x^{ND}(t) < x^c$:
i) $u^{ND}(t)u^{ND}(t+1) = 0$,
ii) $x^{ND}(t+1) > \lambda^* x^{ND}(t)$.

Proof. See the appendix. ■

Part i) of the above proposition implies that, as long as the size of the informed population is below a certain threshold, there are no successive nonzero price levels in the optimal policy. This means that the firm should initially offer the product for free at least half of the time in order to expand its network. Part ii) presents a lowerbound on the effectiveness of these free offers. Using this, it is easy to see that for $\lambda^* > 1$, these drops can result in an exponential growth of the informed population.

3.6 Conclusions

In this paper, we analyzed optimal dynamic pricing in social networks from the information diffusion point of view. We developed a tractable, yet rich, model for information diffusion via word of mouth, where an agent can only get informed about a product through a friend who already knows about the product. Both buyers and non-buyers may contribute to the information diffusion except that buyers are more likely to engage in the spread. Firm can hence use dynamic prices as a tool to control the endogenous information diffusion process. Word of mouth is the only means by which many apps spread among smartphone users. Using this model, we showed that the optimal pricing policy for a durable product with zero or negligible marginal cost, such as many smartphone applications, should drop the price to zero infinitely often. The rationale for this behavior is that by dropping the price to zero and selling the product to agents with low valuations of the product and getting them more engaged in the spread, firm can reach out a new part of the network that would remain untouched otherwise. By timing the drop properly, firm can make sure that the marginal growth in future profit by exploiting this new part of the network prevails the loss in the immediate profit caused by dropping the price to zero. We also showed that although the optimal policy drops the price to zero infinitely often, the price trajectory cannot get trapped in a vicinity of zero meaning that it jumps away from this vicinity infinitely often.

We also examined the validity of our results in face of strategic forward-looking agents, homophily-based engagement in word of mouth, network externalities, and consumer inattention to price changes, by generalizing our base model to these cases. Finally, we showed that beside the word of mouth nature of the information diffusion, this behavior is also rooted in the durability of product being offered. For a nondurable product, although the firm may initially make some free offers to expand its network, after a while it will set the price at a fixed level which extracts the maximum profit from the already informed population. When the network gets large, the loss in the immediate profit by dropping the price in favor of a higher spread would become too large to compare with the marginal gain in future resulted from the excess expansion of the network.

Appendix

3.7 Proofs

Proof of Proposition 3.3.1. Let $\lambda_B = \lambda p_B$. We can use (3.5) to write x_0 as a function of q

$$x_0 = 1 - (1-q)e^{\lambda_B q}. \tag{7.1}$$

The existence of a solution $q \in [0,1]$ for any $0 < x_0 \leq 1$ now follows from the continuity of x_0 in q and that $x_0(q=0) = 0$ and $x_0(q=1) = 1$. Taking derivatives from the above equation, we obtain

$$\frac{dx_0}{dq} = (1 - \lambda_B(1-q))e^{\lambda_B q}, \tag{7.2}$$

$$\frac{d^2x_0}{dq^2} = \lambda_B(2 - \lambda_B(1-q))e^{\lambda_B q}. \tag{7.3}$$

It follows from (7.2) that x_0 attains its minimum at $q^* = 1 - \frac{1}{\lambda_B}$ and is strictly increasing (decreasing) for $q \geq q^*$ ($q \leq q^*$). It also follows from (7.3) that x_0 is convex for $q \geq q^*$. Next, we show that for $q \in [0,1]$ the constraint $0 < x_0 \leq 1$ implies $q > q^*$. This is automatically satisfied for the case where $\lambda_B < 1$ since $q^* < 0$. For $\lambda_B \geq 1$, we have $q^* \geq 0$. However,

$x_0(q)$ is decreasing for $0 \le q \le q^*$ resulting in $x_0(q) \le x_0(q=0) = 0$. This shows that also in this case we should have $q > q^*$.

Now, the uniqueness of the solution in $[0,1]$ for $0 < x_0 \le 1$ follows from the fact that x_0 is strictly increasing for $q \ge q^*$. Also, since x_0 is strictly increasing and convex for $q \ge q^*$, thus q is strictly increasing, but is concave in x_0. ■

Extension of Theorem 3.3.1 to a durable product with nonzero marginal cost. Denote the marginal cost with $c > 0$ and let $p_s \in \mathcal{U}$ be some given price level. We claim that if the gap between c and p_s is small enough, then the optimal price trajectory will drop the price to a level below or equal to p_s infinitely often. This specially implies the validity of Theorem 3.3.1 for a nonzero but negligible marginal cost. In particular we claim that if

$$c - p_s < \frac{\beta\lambda p_B(1-p_{\bar{B}})(1-\hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B}(u_s^* - p_s)(1-F(u_s^*))}{1+\beta\lambda p_B(1-p_{\bar{B}})(1-\hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B}(1-F(u_s^*))}, \tag{7.4}$$

where $\hat{q}_B = q(x_0, p_B + p_{\bar{B}} - p_B p_{\bar{B}}, 0; 0)$ and $(u_s^* - p_s)(1-F(u_s^*)) = \max_{u \in \mathcal{U}}(u - p_s)(1-F(u))$, then the optimal policy $u^D(\cdot)$ drops the price to a level below or equal to p_s infinitely often.

The proof follows the same line as of the proof of Theorem 3.3.1, so we here only highlight the differences. Defining p_r as the smallest price level which is used by firm infinitely often and assuming $p_r > p_s$ we try to reach contradiction using (7.4). Let $u(t) \ge p_r$ for $t \ge T$ and t_k denote the k-th drop of the price to p_r after T. To extend the results to a nonzero marginal cost, we need to slightly modify $\tilde{u}(\cdot)$. Let $\tilde{u}(\cdot)$ have the same value as $u^D(\cdot)$ at all times except t_k+1 and t_k+2, for some k that we specify later. We choose $\tilde{u}(t_k+1) = p_s$ and $\tilde{u}(t_k+2) = u_s^*$. Similar to (3.8), we can lowerbound the size of the freshly informed agents at time t_k+2 as

$$\begin{aligned}\tilde{y}(t_k+2) &\ge (1-q^D)(1-e^{-\lambda p_B(\tilde{b}_Y(t_k+1)+(1-p_{\bar{B}})\tilde{b}_Z(t_k+1))}) \\ &\ge (1-q^D)(1-e^{-\lambda p_B(1-p_{\bar{B}})(\tilde{b}_Y(t_k+1)+\tilde{b}_Z(t_k+1))}).\end{aligned} \tag{7.5}$$

The main difference with the proof of Theorem 3.3.1 is that a drop may acquire some cost if $p_s < c$. The cost of the drop is $(c-p_s)(\tilde{b}_Y(t_k+1)+\tilde{b}_Z(t_k+1))$. How much profit can we make from $\tilde{y}(t_k+2)$? Using the simple fact that $\frac{1-e^{-\zeta w}}{w}$ is decreasing with w for $w \ge 0$ and $\zeta > 0$ (and so is $(1-w)\frac{1-e^{-\zeta w}}{w}$ for $0 < w < 1$) and that $q^D \le \hat{q}_B$, we can use (7.5) to obtain

$$\begin{aligned}\frac{\tilde{y}(t_k+2)}{\tilde{b}_Y(t_k+1)+\tilde{b}_Z(t_k+1)} &\ge (1-q^D)\frac{1-e^{-\lambda p_B(1-p_{\bar{B}})q^D}}{q^D} \\ &\ge (1-\hat{q}_B)\frac{1-e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B}}{\hat{q}_B} \\ &\ge \lambda p_B(1-p_{\bar{B}})(1-\hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B},\end{aligned} \tag{7.6}$$

yielding

$$\tilde{y}(t_k+2) \ge \lambda p_B(1-p_{\bar{B}})(1-\hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B}(\tilde{b}_Y(t_k+1)+\tilde{b}_Z(t_k+1)). \tag{7.7}$$

Therefore, the profit from the choice $\tilde{u}(t_k+1) = p_s$ and $\tilde{u}(t_k+2) = u_s^*$ can be lowerbounded by

$$\Pi^D_{>t_k}(\tilde{u}(\cdot)) \geq \beta^{t_k+1}((u_s^*-c)(1-F(u_s^*))\beta\lambda p_B(1-p_{\bar{B}})(1-\hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B}-(c-p_s))(\tilde{b}_Y(t_k+1)+\tilde{b}_Z(t_k+1)). \tag{7.8}$$

Using (7.4), we can easily see that

$$((u_s^* - c)(1 - F(u_s^*))\beta\lambda p_B(1 - p_{\bar{B}})(1 - \hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B} - (c - p_s)) > 0, \tag{7.9}$$

i.e., it is a positive constant. Also, $(\tilde{b}_Y(t_k + 1) + \tilde{b}_Z(t_k + 1)) > \sum_{j=s}^{r} z_j^D(T) + (F(p_r) - F(p_s))y^D(T) > 0$. To have the new strategy yield a higher profit it suffices to choose k such that

$$\begin{aligned}(p_m - c)(q^D - x^D(t_k)) <&\beta((u_s^* - c)(1 - F(u_s^*))\beta\lambda p_B(1 - p_{\bar{B}})(1 - \hat{q}_B)e^{-\lambda p_B(1-p_{\bar{B}})\hat{q}_B} - (c - p_s))\times \\ &(\sum_{j=s}^{r} z_j^D(T) + (F(p_r) - F(p_s))y^D(T)),\end{aligned} \tag{7.10}$$

which can be satisfied noting that $x^D(t) \to q^D$ as $t \to \infty$. ■

Proof of Proposition 3.3.2. In order to prove the proposition, we show that if the optimal policy $u^D(\cdot)$ gets stuck between 0 and p_1 after a finite time, then we should have $p_1(1 - F(p_1)) \geq c$. Suppose that there is $T \geq 0$ such that $u^D(t) \in \{0, p_1\}$ for all $t \geq T$. Denote with t_k the k-th drop of the price to zero after T under $u^D(\cdot)$. Note that there are infinitely many such drops according to Theorem 3.3.1. We first try to find an upperbound for the accumulated profit of the firm under $u^D(\cdot)$ after t_k. Consider a new policy $\hat{u}(\cdot)$ which has the same values as $u^D(\cdot)$ before t_k but is zero afterwards,

$$\hat{u}(t) = \begin{cases} u^D(t), & t \leq t_k \\ 0, & t > t_k \end{cases} \tag{7.11}$$

that is a free product after t_k. Assume also that in this alternative scenario, informed agents spread the word with increased probability $\hat{p}_B = p_B + p_{\bar{B}} - p_B p_{\bar{B}}$. Our first claim is that the size of the informed population $x^D(t)$ for $t \geq t_k$ is upperbounded by $\hat{x}(t)$, that is $x^D(t) \leq \hat{x}(t)$ for $t \geq t_k$. Recursively using (2.6) from t_k to t we get

$$1 - x^D(t+1) = (1 - x^D(t_k))e^{-\lambda b(t_k,t)}, \tag{7.12}$$

for $t \geq t_k$, where

$$\begin{aligned} b(t_k, t) &= \sum_{\tau=t_k}^{t} p_B b_Y(\tau) + p_{\bar{B}}\bar{b}_Y(\tau) + p_B(1 - p_{\bar{B}})b_Z(\tau) \\ &\leq \hat{p}_B(x^D(t) - x^D(t_k - 1)). \end{aligned} \tag{7.13}$$

On the other hand, $1 - \hat{x}(t+1) = (1 - \hat{x}(t_k))e^{-\lambda\hat{p}_B(\hat{x}(t) - \hat{x}(t_k-1))}$. Using a simple induction and the fact that $\hat{x}(t_k - 1) = x^D(t_k - 1)$ and $\hat{x}(t_k) = x^D(t_k)$, we can show that $x^D(t) \leq \hat{x}(t)$ for

$t \geq t_k$.

Our next claim is that the profit made by the firm under $u^D(\cdot)$ after t_k is upperbounded by

$$\Pi^D_{>t_k}(u^D(\cdot)) \leq p_1(1 - F(p_1)) \sum_{t=t_k+1}^{\infty} \beta^t \hat{y}(t). \tag{7.14}$$

To prove this, it apparently suffices to show that $\sum_{t=t_k+1}^{\infty} \beta^t y^D(t) \leq \sum_{t=t_k+1}^{\infty} \beta^t \hat{y}(t)$. This is immediate from the previous result, if we use $y(t+1) = x(t+1) - x(t)$ to rewrite these summations in terms of $x^D(t)$ and $\hat{x}(t)$, with positive coefficients $(\beta^t - \beta^{t+1})$. The dynamics of $\hat{y}(t)$ for $t > t_k$ has the simple form of

$$\hat{y}(t+1) = (1 - \hat{x}(t))(1 - e^{-\lambda \hat{p}_B \hat{y}(t)}), \tag{7.15}$$

where $\hat{y}(t_k+1) = y^D(t_k+1)$ and $\hat{x}(t_k+1) = x^D(t_k+1)$. Using this, we can easily obtain

$$\hat{y}(t+1) \leq \lambda \hat{p}_B(1 - x^D(t_k+1))\hat{y}(t), \tag{7.16}$$

for all $t > t_k$. Using this along with (7.14), the profit of the firm under $u^D(\cdot)$ after t_k can be upperbounded by

$$\Pi^D_{>t_k}(u^D(\cdot)) \leq \frac{\beta^{t_k+1} p_1(1 - F(p_1)) y^D(t_k+1)}{1 - \beta\lambda\hat{p}_B(1 - x^D(t_k+1))}. \tag{7.17}$$

Next, we compare this profit with that of a modified policy $\tilde{u}(\cdot)$ having the same value as $u^D(\cdot)$ for $t \leq t_k$ and a fixed value $u^c > p_1$ for $t > t_k$. The profit of the firm for policy $\tilde{u}(\cdot)$ after t_k is given by

$$\Pi^D_{>t_k}(\tilde{u}(\cdot)) = u^c(1 - F(u^c)) \sum_{t=t_k+1}^{\infty} \beta^t \tilde{y}(t). \tag{7.18}$$

The dynamics of $\tilde{y}(t)$ for $t > t_k$ is given by

$$\tilde{y}(t+1) = (1 - \tilde{x}(t))(1 - e^{-\lambda^c \tilde{y}(t)}), \tag{7.19}$$

where $\lambda^c = \lambda(p_B(1 - F(u^c)) + p_{\bar{B}}F(u^c))$, $\tilde{y}(t_k+1) = y^D(t_k+1)$ and $\tilde{x}(t_k+1) = x^D(t_k+1)$. We next aim to lowerbound $\tilde{y}(t)$ for $t > t_k$ with a geometric sequence, in order to find a closed form lowerbound for the profit of the firm for $\tilde{u}(\cdot)$ after t_k given by (7.18). This can be easily done by rewriting (7.19) as

$$\begin{aligned}
\tilde{y}(t+1) &= (1 - \tilde{x}(t+1))(e^{\lambda^c \tilde{y}(t)} - 1) \\
&\geq \lambda^c(1 - \tilde{x}(t+1))\tilde{y}(t) \\
&> \lambda^c(1 - \tilde{q})\tilde{y}(t),
\end{aligned} \tag{7.20}$$

where $\tilde{q}$ is the asymptotic size of the informed population under $\tilde{u}(\cdot)$. The profit of the firm after t_k for $\tilde{u}(\cdot)$ can thus be lowerbounded by

$$\Pi^D_{>t_k}(\tilde{u}(\cdot)) > \frac{\beta^{t_k+1} u^c(1 - F(u^c)) y^D(t_k+1)}{1 - \beta\lambda^c(1 - \tilde{q})}. \tag{7.21}$$

The asymptotic size $\tilde{q}$ satisfies $1-\tilde{q} = (1-x^D(t_k+1))e^{-\lambda^c(\tilde{q}-x^D(t_k))} > (1-x^D(t_k+1))e^{-\lambda p_B(\tilde{q}-x^D(t_k))}$. Shifting $t \to \infty$ in (7.12) and noting that the optimal policy drops the price to zero infinitely often, we can show that $1-q^D \leq (1-x^D(t_k+1))e^{-\lambda p_B(q^D-x^D(t_k))}$. From this we can get $\tilde{q} < q^D$, and hence

$$\Pi^D_{>t_k}(\tilde{u}(\cdot)) > \frac{\beta^{t_k+1}u^c(1-F(u^c))y^D(t_k+1)}{1-\beta\lambda^c(1-q^D)}. \tag{7.22}$$

Noting that $u^D(\cdot)$ is the optimal policy, we should have $\Pi^D_{>t_k}(u^D(\cdot)) \geq \Pi^D_{>t_k}(\tilde{u}(\cdot))$ for all choices of u^c. This, along with (7.17) and (7.22) yields

$$\frac{p_1(1-F(p_1))}{1-\beta\lambda\hat{p}_B(1-x^D(t_k+1))} > \frac{u^c(1-F(u^c))}{1-\beta\lambda^c(1-q^D)}, \tag{7.23}$$

for any choices of k and u^c. Shifting $k \to \infty$, we can obtain

$$\frac{p_1(1-F(p_1))}{1-\beta\lambda\hat{p}_B(1-q^D)} > \frac{u^c(1-F(u^c))}{1-\beta\lambda^c(1-q^D)}. \tag{7.24}$$

To get rid of q^D in above, we note that $\frac{1-\beta\lambda\hat{p}_B(1-q^D)}{1-\beta\lambda^c(1-q^D)}$ is increasing with q^D since $\lambda\hat{p}_B > \lambda^c$. Thus, (7.24) yields

$$p_1(1-F(p_1)) > \frac{u^c(1-F(u^c))(1-\beta\lambda\hat{p}_B(1-q^0_B))}{1-\beta\lambda^c(1-q^0_B)}, \tag{7.25}$$

for every $u^c \in \mathcal{U}\setminus\{p_1\}$, since $q^D \geq q^0_B$ according to Theorem 3.3.1. This completes the proof. ■

Proof of Proposition 3.4.1. Similar to the proof of Theorem 3.3.1, let $p_r \in \mathcal{U}$ denote the smallest price level in $\mathcal{U}$ which holds infinitely often for the optimal pricing policy $u^D(\cdot)$. Since any price level below p_r is used only finitely by $u^D(\cdot)$, there exists $T \geq 0$ such that $u^D(t) \geq p_r$ for all $t \geq T$. It clearly suffices to show that $p_r = 0$, that is $r = 0$. Therefore, to prove the proposition, we assume $r \geq 1$ and try to reach contradiction by showing that the firm can increase its profit by deviating from the price path $u^D(\cdot)$.

Let $Y^D_r(T) \subset Y^D(T)$ denote those freshly informed agents at time T whose valuations are below p_r, i.e. $Y^D_r(T) = \{i \in Y^D(T) | 0 \leq \theta_i < p_r\}$, with a size of $y^D_r(T) = F(p_r)y^D(T)$. None of the agents in $\cup_{j=1}^r Z^D_j(T) \cup Y^D_r(T)$ will ever buy the product under the pricing policy $u^D(\cdot)$, where $\cup_{j=1}^r Z^D_j(T)$ is the set of those previously informed agents at time T whose valuations are below p_r. Now, consider the set of agents that will remain uniformed under $u^D(\cdot)$. The size of this set is clearly $1-q^D$, where q^D is the asymptotic size of the informed population under $u^D(\cdot)$, i.e., $q^D = q(x_0, p_B, p_{\bar{B}}; u^D(\cdot))$. Define Δ_r as the subset of these agents who have at least a friend in $\cup_{j=1}^r Z^D_j(T) \cup Y^D_r(T)$. As in Theorem 3.3.1, the number of friends of an uninformed agent among $\cup_{j=1}^r Z^D_j(T) \cup Y^D_r(T)$ has a Poisson distribution with mean $\lambda(1-p_{\bar{B}})(F(p_r)y^D(T) + \sum_{j=1}^r z^D_j(T))$. Therefore, by zeroing the price at any time $t > T$ we can reach out a subset of Δ_r with the size of

$$\delta_r = (1 - q^D)(1 - e^{-\lambda p_B(1-p_{\bar{B}})(F(p_r)y^D(T)+\sum_{j=1}^{r} z_j^D(T))}), \tag{7.26}$$

from which clearly $\delta_r > 0$. Now, the idea is to show that after a while there is so little profit left to be made in future under $u^D(\cdot)$ that it is profitable to zero the price to reach out these agents in Δ_r. Let t_k, $k = 1, 2, \ldots$, denote the k-th price drop to p_r after time T under the optimal policy $u^D(\cdot)$. If a strategic agent $i \in X^D(t_k)$ does not buy the product at this time, neither will she buy it in future, as this is the cheapest offer she will ever get. This means that agents in $X^D(t_k)$ do not contribute to the set of buyers $B^D(t)$ for $t > t_k$. Therefore, the size of the buyers from time $t_k + 1$ can be upperbounded by

$$\sum_{t=t_k+1}^{\infty} b^D(t) \leq q^D - x^D(t_k). \tag{7.27}$$

Thus, the contribution of the buyers to the firm's profit after t_k can be upperbounded by

$$\Pi^D_{>t_k}(u^D(\cdot)) \leq \beta^{t_k+1} p_m (q^D - x^D(t_k)). \tag{7.28}$$

Now, consider the following deviation $\tilde{u}(\cdot)$ from the price path $u^D(\cdot)$ only at times $t_k + 1$ and $t_k + 2$. Let $\tilde{u}(t_k + 1) = 0$ and $\tilde{u}(t_k + 2) = \hat{u}$, where $\hat{u}(1 - F(\frac{\hat{u}}{1-\beta_c})) = \max_{u \in \mathcal{U}} u(1 - F(\frac{u}{1-\beta_c}))$. Given that $\beta_c < 1 - p_1$, $\hat{u}$ is always nonzero. Note that a subset of agents in Δ_r with size δ_r as in (7.26) are among the freshly informed agents $\tilde{Y}(t_k + 2)$. Those of them with valuations $\theta > \frac{\hat{u}}{1-\beta_c}$ buy the product since $\theta - \hat{u} > \beta_c \theta$, meaning that the payoff of a purchase today is higher than even a free purchase tomorrow. The discounted profit made from these newly informed agents in Δ_r at time $t_k + 2$ is hence lowerbounded by $\beta^{t_k+2}\hat{u}(1 - F(\frac{\hat{u}}{1-\beta_c}))\delta_r$. Considering that $x^D(t) \to q^D$ as $t \to \infty$, we can choose k large enough such that

$$q^D - x^D(t_k) < \frac{\beta \delta_r}{p_m} \hat{u}(1 - F(\frac{\hat{u}}{1 - \beta_c})), \tag{7.29}$$

in which case the profit resulted from $\tilde{u}(\cdot)$ will be clearly higher than that from $u^D(\cdot)$. This contradicts the optimality of $u^D(\cdot)$, hence rejecting the initial assumption of $p_r \neq 0$, which completes the proof. ■

Proof of Proposition 3.4.2. First, we show that under the local homophily condition, there will always be a nonzero mass of low-valued informed non-buyers around, with valuations below $\delta_1 = \min(\frac{\delta}{2}, p_1)$. More precisely, we claim that

$$M_Y(\delta_1, t)y(t) + M_Z(\delta_1, t)z(t) > 0. \tag{7.30}$$

First note that $M_Y(\delta_1, 0)y(0) = F(\delta_1)x_0 > 0$. As long as the offered price $u(t) \geq p_1$, our claim is apparent. If the price becomes zero ($u(t) = 0$), then for an uninformed agent with $\theta' \leq \delta_1$, we have

$$\lambda(\theta', t) \geq \lambda \underline{p}(1 - \bar{p})(M_Y(\delta_1, t)y(t) + M_Z(\delta_1, t)z(t)). \tag{7.31}$$

Denoting the RHS with $\underline{\lambda}(t) > 0$ along with (4.10) and (4.13), we get

$$M_Y(\delta_1, t+1)y(t+1) \geq e^{-\lambda q}(1 - e^{-\underline{\lambda}(t)})F(\delta_1)(1 - x_0) > 0, \tag{7.32}$$

completing the proof of our claim of having a nonzero measure of low-valued informed non-buyers around at all times. Using a similar approach we can show that if there is a mass w of informed non-buyers with valuations in the range of $[\underline{\theta}, \bar{\theta}]$ at time t, where $|\underline{\theta} - \bar{\theta}| \leq \frac{\delta}{2}$, then by dropping the price to zero we can reach out a set of agents with valuations in the range of $[\bar{\theta}, \bar{\theta} + \frac{\delta}{2}]$ whose size $\tilde{w}$ satisfies

$$\tilde{w} \geq e^{-\lambda q^0}(1 - e^{-\lambda \underline{p}(1-\bar{p})w})(F(\bar{\theta} + \frac{\delta}{2}) - F(\bar{\theta}))(1 - x_0) > 0, \tag{7.33}$$

where we have also used the fact that the asymptotic size of the informed population is upperbounded by q^0 as defined in Section 3.3. With these results at hand, we now use an approach similar to that of Theorem 3.3.1 to complete the proof.

Let $p_r \in \mathcal{U}$ denote the smallest price level in $\mathcal{U}$ which holds infinitely often for the optimal pricing policy $u^D(\cdot)$. Since any price level below p_r is used only finitely by $u^D(\cdot)$, there exists $T \geq 0$ such that $u^D(t) \geq p_r$ for all $t \geq T$. It clearly suffices to show that $p_r = 0$, that is $r = 0$. Therefore, to prove the proposition, we assume $r \geq 1$ and try to reach contradiction by showing that the firm can increase its profit by deviating from the price path $u^D(\cdot)$.

Let t_k, $k = 1, 2, \ldots$, denote the k-th price drop to p_r after time T under the optimal policy $u^D(\cdot)$. The contribution of the buyers to the firm's profit after t_k can be upperbounded by

$$\Pi^D_{>t_k}(u^D(\cdot)) \leq \beta^{t_k+1} p_m (q^D - x^D(t_k)). \tag{7.34}$$

Denote with $W^D(T)$ the set of non-buyer agents at time T whose valuations are below δ_1, and its size with $w^D(T)$. Agents in $W^D(T)$ will never buy the product under the optimal policy $u^D(\cdot)$, and hence $w^D(t) \geq w^D(T)$ for all $t \geq T$. Now, consider a new policy $\tilde{u}(\cdot)$ having the same values as $u^D(\cdot)$ at all times except $t_k + 1 \leq t \leq t_k + s + 1$, where s is the smallest integer for which $\delta_1 + \frac{\delta(s-1)}{2} \geq p_1$. Let $\tilde{u}(t_k + 1) = \ldots = \tilde{u}(t_k + s) = 0$ and $\tilde{u}(t_k + s + 1) = p_1$. From above, we already know that $w^D(T) > 0$. Dropping the price to zero s times in a row, we can reach a nonzero measure population with valuations in $[\delta_1 + \frac{\delta(s-1)}{2}, \delta_1 + \frac{\delta s}{2}]$. By recursively using (7.33), we can obtain a lower bound $\underline{w} > 0$ for the size of this set. [35] The discounted profit made from these newly informed agents at time $t_k + s + 1$ is lowerbounded by $\beta^{t_k+s+1} p_1 \underline{w}$. Considering that $x^D(t) \to q^D$ as $t \to \infty$, we can choose k large enough such that

$$q^D - x^D(t_k) < \frac{\beta^s p_1 \underline{w}}{p_m}, \tag{7.35}$$

in which case the profit resulted from $\tilde{u}(\cdot)$ will be clearly higher than that from $u^D(\cdot)$. This contradicts the optimality of $u^D(\cdot)$, hence rejecting the initial assumption of $p_r \neq 0$, which completes the proof. ■

[35] Note that this lower bound is in terms of $w^D(T)$ and is independent of t_k.

Distribution of WOM Engagement among buyers and non-buyers in Section 3.4.3. To find the update rule for $Z(\theta^a, t+1)$, we need to figure out how many friends an agent $i \in Z^\tau(t+1)$ interacts with among the new buyers $B(t)$.

$$\begin{aligned}
\text{Prob}(d_i^{WOM}(B_0^{\tau+1}(t)) = d | i \in Z^\tau(t+1)) &\sim \text{Poiss}(\lambda p_B(1 - p_{\bar{B}} + \frac{p_{\bar{B}}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau)}}) b_0^{\tau+1}(t)) \\
\text{Prob}(d_i^{WOM}(B_{>0}^{\tau+1}(t)) = d | i \in Z^\tau(t+1)) &\sim \text{Poiss}(\lambda p_B b_{>0}^{\tau+1}(t)) \\
\text{Prob}(d_i^{WOM}(B^{\tau}(t)) = d | i \in Z^\tau(t+1)) &\sim \text{Poiss}(\lambda p_B b^{\tau}(t)) \\
\text{Prob}(d_i^{WOM}(B^{\tau-1}(t)) = d | i \in Z_{>0}^\tau(t+1)) &\sim \text{Poiss}(\lambda p_B b^{\tau-1}(t)),
\end{aligned} \tag{7.36}$$

where $b^{-1}(0) = b^{t+1}(t) = 0$. The number of the interactions between $i \in Z_0^\tau(t+1)$ and new buyers that were informed at $\tau - 1$, also depends on the number of interactions between i and those previous buyers that were informed at the same time $\tau - 1$. More precisely,

$$\begin{aligned}
&\text{Prob}\left(d_i^{WOM}(B^{\tau-1}(t)) = d | i \in Z_0^\tau(t+1) \wedge d_i^{WOM}(\cup_{t'=\tau}^{t-1} B^{\tau-1}(t')) = k\right) = \\
&\frac{e^{-\lambda p_B b^{\tau-1}(t)} \frac{(\lambda p_B b^{\tau-1}(t))^d}{d!} (1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(\tau-1) - p_B \sum_{t'=\tau}^{t} b^{\tau-1}(t'))} (1 - p_{\bar{B}})^{d+k})}{1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(\tau-1) - p_B \sum_{t'=\tau}^{t-1} b^{\tau-1}(t'))} (1 - p_{\bar{B}})^k}.
\end{aligned} \tag{7.37}$$

Finally, for $\tau' \notin \{\tau - 1, \tau, \tau + 1\}$, we have

$$\text{Prob}(d_i^{WOM}(B^{\tau'}(t)) = d | i \in Z^\tau(t+1)) \sim \text{Poiss}(\lambda p_B(1 - p_{\bar{B}}) b^{\tau'}(t)). \tag{7.38}$$

■

Proof of Proposition 3.4.3. We first prove the following lemma:

Lemma 3.7.1. *Consider a price function $u(\cdot)$ with finite number of drops to zero, that is, there exists $T \geq 0$ such that $u(t) \neq 0$ for $t \geq T$. Then, for any $t > \tau \geq T$ we have*

$$z_0^\tau(p_1, t) \geq (1 - p_B) e^{-\lambda(p_B + p_{\bar{B}})} y_0(p_1, \tau). \tag{7.39}$$

Using only the term corresponding to $d' = 0$ for $d = 0$ in (4.21), we can obtain

$$z_0^\tau(p_1, 0, \tau+1) \geq \frac{1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(\tau-1) - p_B b^{\tau-1}(\tau))}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau-1)}} e^{-\lambda_{>0}^\tau(\tau)} y_0(p_1, \tau). \tag{7.40}$$

Using this along with recursive use of (4.20) for $d = 0$ using only the term corresponding to $d' = 0$, we get

$$z_0^\tau(p_1, 0, t+1) \geq \frac{1 - e^{-\lambda p_{\bar{B}}(\bar{b}_Y(\tau-1) - p_B \sum_{t'=\tau}^{t} b^{\tau-1}(t'))}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau-1)}} e^{-\sum_{t'=\tau}^{t} \lambda_{>0}^\tau(t')} y_0(p_1, \tau), \tag{7.41}$$

for $t \geq \tau$. We now note that

$$\sum_{t'=\tau}^{t} b^{\tau-1}(t') < \bar{b}_Y(\tau - 1). \tag{7.42}$$

On the other hand, using (4.17) we can easily show that

$$\lambda^{\tau}_{>0}(t') \leq \lambda p_B b(t') + \lambda p_B p_{\bar{B}} \frac{e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau)}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau)}} y_0^{\tau+1}(t') \leq \lambda p_B b(t') + p_B \frac{y_0^{\tau+1}(t')}{\bar{b}_Y(\tau)}, \tag{7.43}$$

which yields

$$\sum_{t'=\tau}^{t} \lambda^{\tau}_{>0}(t') \leq \lambda p_B q + p_B \frac{y_0(\tau+1)}{\bar{b}_Y(\tau)} \leq \lambda p_B + \lambda p_{\bar{B}}(1 - x(\tau)), \tag{7.44}$$

where we have used the fact that $q \leq 1$, and that $y_0(\tau+1) \leq \lambda p_{\bar{B}} \bar{b}_Y(\tau)(1 - x(\tau))$ (this is quite straightforward using (4.14)). Using (7.41),(7.42), and (7.44), we get

$$\begin{aligned} z_0^{\tau}(p_1, 0, t+1) &\geq \frac{1 - e^{-\lambda p_{\bar{B}}(1-p_B)\bar{b}_Y(\tau-1)}}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau-1)}} e^{-\lambda(p_B + p_{\bar{B}}(1 - x(\tau)))} y_0(p_1, \tau) \\ &\geq \frac{e^{\lambda p_{\bar{B}}(1-p_B)\bar{b}_Y(\tau-1)} - 1}{1 - e^{-\lambda p_{\bar{B}} \bar{b}_Y(\tau-1)}} e^{-\lambda(p_B + p_{\bar{B}}(1 - x(\tau) + \bar{b}_Y(\tau-1)))} y_0(p_1, \tau) \\ &\geq (1 - p_B) e^{-\lambda(p_B + p_{\bar{B}})} y_0(p_1, \tau), \end{aligned} \tag{7.45}$$

which completes the proof of the lemma. We now get back to the proof of the proposition. Following the same line as of the proof of Theorem 3.3.1, we assume a finite number of drops to zero under the optimal policy $u^D(\cdot)$ and try to reach contradiction by constructing a new policy with a profit higher than that of $u^D(\cdot)$.

Let t_0 be the last drop of the price to zero (we study the case that there is no drop to zero at all later). So, $u^D(t_0) = 0$ and $u(t) > 0$ for $t > t_0$. Our first claim is that, under assumption (4.24), there exists $t_1 > t_0$ such that $u^D(t_1) > \alpha$., that is the price path cannot stay below α after the last drop to zero. The proof is very similar to that of Proposition 3.3.2. We assume $u^D(t) \leq \alpha$ for $t > t_0$ and try to reach contradiction by finding a more profitable price function. Let $u^\alpha = \max\{u \in \mathcal{U} | u \leq \alpha\}$. Then, $u^D(t) \leq \alpha$ for $\alpha > t_0$ indeed implies $u^D(t) = u^\alpha$ for $\alpha > t_0$. Note that at this price after a drop to zero every body that hears about the product has at least a friend among the buyers and hence her augmented valuation is at least α. Therefore, there is no need to use a price less than u^α even if we want to keep it below α. Using almost the same approach as in Proposition 3.3.2, we can find an upperbound similar to (7.17) for the profit of the firm from $t' > t_0$ afterwards as

$$\Pi^D_{>t'}(u^D(\cdot)) \leq \frac{\beta^{t'+1} u^\alpha y^D(t'+1)}{1 - \beta \lambda p_B (1 - x^D(t'+1))}. \tag{7.46}$$

Next, we compare this profit with that of a modified policy $\tilde{u}(\cdot)$ having the same value as $u^D(\cdot)$ for $t \leq t'$ and a fixed value u^c for $t > t'$. Considering only the part of the spread by buyers and the effect of only one of the buyer friends in augmented valuation, we can lowerbound the profit of the firm after t', similar to (7.22), by

$$\Pi^D_{>t'}(\tilde{u}(\cdot)) > \frac{\beta^{t'+1}u^c(1-F(u^c-\alpha))y^D(t'+1)}{1-\beta\lambda p_B(1-F(u^c-\alpha))(1-\tilde{q})}. \tag{7.47}$$

Noting that $u^D(\cdot)$ is the optimal policy, we should have $\Pi^D_{>t'}(u^D(\cdot)) \geq \Pi^D_{>t'}(\tilde{u}(\cdot))$ for all choices of u^c and all $t' > t_0$. Therefore,

$$\frac{u^\alpha}{1-\beta\lambda p_B(1-q^D)} \geq \frac{u^c(1-F(u^c-\alpha))}{1-\beta\lambda p_B(1-F(u^c-\alpha))(1-\tilde{q})}. \tag{7.48}$$

Noting that $q^D \geq q^0_B$ and $\tilde{q} \leq \hat{q}^0_B$, and that $u^\alpha \leq \alpha$ we should thus have

$$\frac{\alpha(1-\beta\lambda p_B(1-F(u^c-\alpha))(1-\hat{q}^0_B))}{u^c(1-F(u^c-\alpha))} \geq 1-\beta\lambda p_B(1-q^0_B), \tag{7.49}$$

for all $u^c \in \mathcal{U}$which contradicts (4.24). This proves our first claim, that is, there exists $t_1 > t_0$ such that $u^D(t_1) > \alpha$. Apparently, we can assume that t_1 is the earliest time after t_0 for which $u^D(t_1) > \alpha$. Thus, for $t_0 \leq t < t_1$, the offered price is below α and all the freshly informed agents in $y(t)$ buy the product. Let $u(t_1) = p_r > \alpha$. Consider $Y^D(p_r, t_1)$, the set of freshly informed agents at time t_1 whose augmented valuations are below p_r. We can lowerbound the size of this set using only the term corresponding to $d = 1$ in (4.14), as

$$y^D(p_r, t_1) \geq (1-x^D(t_1-1))e^{-\lambda p_B y^D(t_1-1)}(\lambda p_B y^D(t_1-1))F(p_r-\alpha), \tag{7.50}$$

which implies $y^D(p_r, t_1) > 0$ using the assumption $p_r > \alpha$. These agents will not buy the product but will inform friends with a lower probability $p_{\bar{B}} > 0$. Using (4.14) now yields

$$y^D_0(p_1, t_1+1) = (1-x^D(t_1))e^{-\lambda^D_B(t_1)}(1-e^{-p_{\bar{B}}\lambda \bar{b}^D_Y(t_1)})F(p_1). \tag{7.51}$$

Using this along with $\bar{b}^D_Y(t_1) \geq y^D(p_r, t_1) > 0$, we get $y^D_0(p_1, t_1+1) > 0$, a nonzero measure set of low-valued agents with no buyer friend yet (in case that $u^D(\cdot)$ does not drop the price to zero at all, we can choose $t_1 = 0$ for which we can show similar to the above that $y^D_0(p_1, t_1+1) > 0$). Due to network externalities, agents in this set may eventually buy the product if many friends do so, elevating their augmented valuations above p_1. However, Lemma 7.39 guarantees that a nonzero fraction of them will never buy the product, thus,

$$z^D_0(p_1, t) \geq (1-p_B)e^{-\lambda(p_B+p_{\bar{B}})}y^D_0(p_1, t_1+1), \tag{7.52}$$

for all $t > t_1$.

Now, consider a new policy $\tilde{u}(\cdot)$ having the same value as $u^D(\cdot)$ at all times except t and $t+1$, where we will specify $t > t_1$ later. Let $\tilde{u}(t) = 0$ and $\tilde{u}(t+1) = u^*$, where $u^*(1-F(u^*-\alpha)) = \max_{u\in\mathcal{U}} u(1-F(u-\alpha))$. Agents in $z^D_0(p_1, t)$ will buy the product and inform their friends with probability p_B, giving a lowerbound on the size of the freshly informed agents at time $t+1$ as

$$\tilde{y}(t+1) \geq (1-q^D)(1-e^{-\lambda p_B(1-p_{\bar{B}})z_0^D(p_1,t)}) \\ \geq (1-q^D)(1-e^{-\lambda p_B(1-p_B)(1-p_{\bar{B}})e^{-\lambda(p_B+p_{\bar{B}})}y_0^D(p_1,t_1+1)}), \quad (7.53)$$

leading to a profit of at least $u^*(1-F(u^*-\alpha)) * y(t+1)$ at time $t+1$. Using an approach similar to that used in the proof of Theorem 3.3.1, it is quite straightforward to show that by a proper choice of t, this new policy will yield a profit higher than that of $u^D(\cdot)$. Firm's profit from time t onward under the optimal policy can be upperbounded by

$$\Pi^D_{\geq t}(u^D(\cdot)) = \sum_{\tau=t}^{\infty} \beta^\tau u^D(\tau) b^D(\tau) \\ \leq \beta^t p_m \sum_{\tau=t}^{\infty} b^D(\tau). \quad (7.54)$$

Noting that $\sum_{\tau=0}^{\infty} b^D(\tau) \leq q^D$, we can choose t large enough to ensure that

$$\sum_{\tau=t}^{\infty} b^D(\tau) < \frac{\beta u^*(1-F(u^*-\alpha))y(t+1)}{p_m}, \quad (7.55)$$

in which the profit from $\tilde{u}(\cdot)$ will be clearly higher than that of $u^D(\cdot)$. This contradicts the optimality of $u^D(\cdot)$, hence rejecting the initial assumption of finite drops to zero under the optimal policy. ■

Proof of Proposition 3.4.4. The proof is very similar to that of Theorem 3.3.1 so we only point out the necessary changes here. The first differnece is the size δ_r in (3.8). That is, the size of the agents that can only be reached via low-valuation agents by offering them the product for free. With only a fraction p_A of low-valuation agents noticing the free-offer, the size of this set becomes

$$\delta_r = (1-q^D)(1-e^{-\lambda p_A p_B(1-p_{\bar{B}})(F(p_r)y^D(T)+\sum_{j=1}^{r} z_j^D(T))}). \quad (7.56)$$

The other difference is the lowerbound on the size of the future buyers in (3.10) which becomes

$$\sum_{t=t_k+1}^{\infty} b^D(t) \leq q^D - x^D(t_{k-k'}) + (1-p_A)^{k'+1} x^D(t_{k-k'}) \\ \leq (1+(1-p_A)^{k'+1})q^D - x^D(t_{k-k'}), \quad (7.57)$$

for any $1 \leq k' \leq k$. The rough idea is that while bounding the size of the future buyers, we should keep in mind the extra mass of informed non-buyers that have been inattentive to the previous drops of the price to p_r. Using this, the contribution of the buyers to the firm's profit after t_k can be now upperbounded by

$$\Pi^D_{>t_k}(u^D(\cdot)) \leq \beta^{t_k+1} p_m((1+(1-p_A)^{k'+1})q^D - x^D(t_{k-k'})), \tag{7.58}$$

for any $1 \leq k' \leq k$. To make the deviation $\tilde{u}(\cdot)$, as defined in the proof of Theorem 3.3.1, preferable over $u^D(\cdot)$ it suffices to have

$$(1+(1-p_A)^{k'+1})q^D - x^D(t_{k-k'}) < \beta\delta_r \frac{u^*(1-F(u^*))}{p_m}. \tag{7.59}$$

Considering that $x^D(t) \to q^D$ as $t \to \infty$, and that $(1-p_A)^{k'+1} \to 0$ as $k' \to \infty$, we can choose k' and $k-k'$ large enough to satisfy this condition, hence completing the proof. ■

Proof of Theorem 3.5.1. Below, we provide more details on parts of the proof of Theorem 3.5.1.

i) *Proof of the lowerbound on $\Pi^{ND}_{\geq t_0}(\tilde{u}(\cdot))$ given by* (5.4)*:* For $t > t_0$, we have $\tilde{x}(t) \geq x^{ND}(t-1)$ and $\tilde{u}(t) = u^{ND}(t-1)$. Also, $\tilde{x}(t_0) = x^{ND}(t_0)$ and $\tilde{u}(t_0) = u^*$. Therefore, we can write

$$\begin{aligned}
\Pi^{ND}_{\geq t_0}(\tilde{u}(\cdot)) &= \sum_{t=t_0}^{\infty} \beta^t \tilde{u}(t)(1-F(\tilde{u}(t)))\tilde{x}(t) \\
&\geq \beta^{t_0}\tilde{u}(t_0)(1-F(\tilde{u}(t_0)))\tilde{x}(t_0) + \sum_{t=t_0+1}^{\infty} \beta^t u^{ND}(t-1)(1-F(u^{ND}(t-1)))x^{ND}(t-1) \\
&= \beta^{t_0} u^*(1-F(u^*))x^{ND}(t_0) + \beta\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot)). \qquad (7.60)
\end{aligned}$$

ii) *Proof of the upperbound on $\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot))$ given by* (5.6)*:* Noting that $x^{ND}(t) \leq q^{ND}$, we have

$$\begin{aligned}
\Pi^{ND}_{\geq t_0}(u^{ND}(\cdot)) &= \sum_{t=t_0}^{\infty} \beta^t u^{ND}(t)(1-F(u^{ND}(t)))x^{ND}(t) \\
&\leq \beta^{t_0} u^{ND}(t_0)(1-F(u^{ND}(t_0)))x^{ND}(t_0) + \sum_{t=t_0+1}^{\infty} \beta^t u^*(1-F(u^*))q^{ND} \\
&= \beta^{t_0} u^{ND}(t_0)(1-F(u^{ND}(t_0)))x^{ND}(t_0) + \frac{\beta^{t_0+1}}{1-\beta} u^*(1-F(u^*))q^{ND}. \qquad (7.61)
\end{aligned}$$

iii) *Proof of the inequality given by* (5.5)*:* Applying the upperbound given by (5.6) and simple algebra, we can easily see that this inequality is satisfied if

$$(1-\beta)u^{ND}(t_0)(1-F(u^{ND}(t_0)))x^{ND}(t_0) + \beta u^*(1-F(u^*))q^{ND} < u^*(1-F(u^*))x^{ND}(t_0), \tag{7.62}$$

or equivalently

$$x^{ND}(t_0) > \frac{\beta u^*(1-F(u^*))}{u^*(1-F(u^*)) - (1-\beta)u^{ND}(t_0)(1-F(u^{ND}(t_0)))} q^{ND}, \tag{7.63}$$

where $u^{ND}(t_0) < u^*$ (this follows from the assumption $u^{ND}(t_0) \neq u^*$ and that $u^{ND}(t_0) \leq u^*$).

The above follows then from the fact that $x^{ND}(t_0) > \gamma q^{ND}$ and the definition of γ in (5.2). ■

Proof of Proposition 3.5.1. The proof is by induction. However, in order to use induction, we will need a more accurate but dirtier version of the proposition as follows.
Claim: Choose some λ_1 satisfying $\lambda^* < \lambda_1 < \frac{\lambda(\lambda^*-1)}{\lambda^*} + 1$. Note that RHS is greater than LHS since $\lambda^* < \lambda$. For any such λ_1, there exists $x^c > 0$ such that for $x^{ND}(t) < x^c$:
i) $u^{ND}(t)u^{ND}(t+1) = 0$,
ii) if $u^{ND}(t) = 0$ then $x^{ND}(t+1) > \lambda_1 x^{ND}(t)$. And, if $u^{ND}(t) \neq 0$ then $x^{ND}(t+1) > \lambda^* x^{ND}(t)$. Note that in either case in ii) we have $x^{ND}(t+1) > \lambda^* x^{ND}(t)$ since $\lambda_1 > \lambda^*$.

Define

$$g_1(x) = 1 - (1-x)e^{-\frac{\lambda^*(\lambda_1-1)}{\lambda_1}x} - \lambda^* x, \tag{7.64}$$

$$g_2(x) = 1 - (1-x)e^{-\frac{\lambda(\lambda^*-1)}{\lambda^*}x} - \lambda_1 x, \tag{7.65}$$

$$g_3(x) = (1-x)(e^{-\frac{\lambda(\lambda_1-1)(1-F(u))}{\lambda_1}x} - e^{-\frac{\lambda(\lambda_1-1)}{\lambda_1}x}) - \frac{\lambda(\lambda^*-1)F(u)}{\lambda^*}x, \tag{7.66}$$

where in $g_3(x)$, u is a nonzero price level in $\mathcal{U}$. We can easily verify that the derivatives of these functions at $x = 0$ are all positive. This implies that all the three functions are strictly increasing in a vicinity of $x = 0$. Using this along with the fact that $g_1(0) = g_2(0) = g_3(0) = 0$, we can conclude that there exists $\tilde{x}^c > 0$ such that $g_1(x) > 0, g_2(x) > 0,$ and $g_3(x) > 0$ for all $0 < x < \tilde{x}^c$. Also, define

$$h(y) = e^{-\lambda(1-F(u))y} - e^{-\lambda y}. \tag{7.67}$$

It is quite straightforward to show that h is strictly increasing for $0 < y < \frac{-\ln(1-F(u))}{\lambda F(u)}$, and that $\frac{1}{\lambda} < \frac{-\ln(1-F(u))}{\lambda F(u)}$ for $0 < F(u) < 1$. Thus, h is strictly increasing for $0 < y < \frac{1}{\lambda}$ for any nonzero price level $u \in \mathcal{U}$. We set $x^c = \min\{\tilde{x}^c, \frac{1}{\lambda}\}$ and prove the claim above for this choice of x^c with induction.

We start with the transition part of the induction. Assuming that the claim holds for $t-1$, we try to prove it for t. Suppose that $x^{ND}(t) < x^c$. First of all note that $x^{ND}(t-1) < x^c$ since $x^{ND}(t-1) < x^{ND}(t)$. We first tackle part ii). If $u^{ND}(t) \neq 0$, then $u^{ND}(t-1) = 0$ since, according to the assumption of induction, $u^{ND}(t-1)u^{ND}(t) = 0$. This implies $Z^{ND}(t) = \emptyset$, and hence $b^{ND}(t) = (1 - F(u^{ND}(t)))y^{ND}(t)$ from (4.28). Using (2.6), we have

$$x^{ND}(t+1) = 1 - (1 - x^{ND}(t))e^{-\lambda(1-F(u^{ND}(t)))y^{ND}(t)}. \tag{7.68}$$

Now, we try to lowerbound $y^{ND}(t)$ in terms of $x^{ND}(t)$. From the assumption of induction and that $u^{ND}(t-1) = 0$, we get $x^{ND}(t) > \lambda_1 x^{ND}(t-1)$, which along with $y^{ND}(t) = x^{ND}(t) - x^{ND}(t-1)$ yields

$$y^{ND}(t) > \frac{\lambda_1 - 1}{\lambda_1}x^{ND}(t). \tag{7.69}$$

This, along with (7.68) and that $g_1(x^{ND}(t)) > 0$ since $x^{ND}(t) < x^c$, we find

$$\begin{aligned} x^{ND}(t+1) &> 1-(1-x^{ND}(t))e^{-\frac{\lambda(1-F(u^{ND}(t)))(\lambda_1-1)}{\lambda_1}x^{ND}(t)} \\ &\geq 1-(1-x^{ND}(t))e^{-\frac{\lambda(1-F(u^*))(\lambda_1-1)}{\lambda_1}x^{ND}(t)} \\ &> \lambda^* x^{ND}(t), \end{aligned} \tag{7.70}$$

where we have also used the fact that for the optimal policy $u^{ND}(t) \leq u^*$. This proves that if $u^{ND}(t) \neq 0$ then $x^{ND}(t+1) > \lambda^* x^{ND}(t)$. If, on the other hand, $u^{ND}(t) = 0$, then $b^{ND}(t) \geq y^{ND}(t)$ from (4.28), and we can use (2.6) to obtain

$$x^{ND}(t+1) \geq 1-(1-x^{ND}(t))e^{-\lambda y^{ND}(t)}. \tag{7.71}$$

Using the assumption of induction, $x^{ND}(t) > \lambda^* x^{ND}(t-1)$, which along with $y^{ND}(t) = x^{ND}(t) - x^{ND}(t-1)$ yields

$$y^{ND}(t) > \frac{\lambda^*-1}{\lambda^*}x^{ND}(t). \tag{7.72}$$

Using (7.71) and (7.72), and that $g_2(x^{ND}(t)) > 0$, we get

$$\begin{aligned} x^{ND}(t+1) &> 1-(1-x^{ND}(t))e^{-\frac{\lambda(\lambda^*-1)}{\lambda^*}x^{ND}(t)} \\ &> \lambda_1 x^{ND}(t), \end{aligned} \tag{7.73}$$

which completes the proof of part ii) for t. Now, we get to the proof of part i). Assume that $u^{ND}(t)u^{ND}(t+1) \neq 0$, and construct a new policy $\tilde{u}(\cdot)$ that is obtained from $u^{ND}(\cdot)$ by only changing $u^{ND}(t)$ to 0. We claim that the new policy will result in a profit higher than that of $u^{ND}(\cdot)$. First of all, note that for all times $\tau \geq 0$, $X^{ND}(\tau) \subseteq \tilde{X}(\tau)$, thus $x^{ND}(\tau) \leq \tilde{x}(\tau)$. In particular, we are interested in calculating $\tilde{x}(t+1) - x^{ND}(t+1)$. From the assumption of induction, we should have $u^{ND}(t-1)u^{ND}(t) = 0$. Therefore, $u^{ND}(t-1) = 0$ since $u^{ND}(t) \neq 0$. This implied that $Z^{ND}(t) = \emptyset$. Hence, using (4.28) and (2.6) we get

$$\tilde{x}(t+1) - x^{ND}(t+1) = (1-x^{ND}(t))(e^{-\lambda(1-F(u^{ND}(t)))y^{ND}(t)} - e^{-\lambda y^{ND}(t)}). \tag{7.74}$$

It follows from $u^{ND}(t-1) = 0$ and the assumption of induction that $x^{ND}(t) > \lambda_1 x^{ND}(t-1)$, which in turn implies

$$y^{ND}(t) > \frac{\lambda_1-1}{\lambda_1}x^{ND}(t). \tag{7.75}$$

Now, considering that $h(y)$ defined in (7.67) is strictly increasing for $0 < y < x^c$, (7.74) yields

$$\begin{aligned} \tilde{x}(t+1) - x^{ND}(t+1) &> (1-x^{ND}(t))(e^{-\frac{\lambda(\lambda_1-1)(1-F(u^{ND}(t)))}{\lambda_1}x^{ND}(t)} - e^{-\frac{\lambda(\lambda_1-1)}{\lambda_1}x^{ND}(t)}) \\ &> \frac{\lambda(\lambda^*-1)F(u^{ND}(t))}{\lambda^*}x^{ND}(t), \end{aligned} \tag{7.76}$$

where the last inequality comes from $g_3(x^{ND}(t)) > 0$.

In order to show that the new policy $\tilde{u}(\cdot)$ results in a higher profit, it suffices to show that

$$u^{ND}(t)(1 - F(u^{ND}(t)))x^{ND}(t)+\beta u^{ND}(t+1)(1 - F(u^{ND}(t+1)))x^{ND}(t+1) < \\ \beta u^{ND}(t+1)(1 - F(u^{ND}(t+1)))\tilde{x}(t+1), \quad (7.77)$$

or equivalently,

$$u^{ND}(t)(1-F(u^{ND}(t)))x^{ND}(t) < \beta u^{ND}(t+1)(1-F(u^{ND}(t+1)))(\tilde{x}(t+1)-x^{ND}(t+1)). \quad (7.78)$$

Applying (7.76) and some simplifications, we can see that the above is satisfied if

$$\frac{u^{ND}(t)(1 - F(u^{ND}(t)))}{F(u^{ND}(t))} < \beta u^{ND}(t+1)(1 - F(u^{ND}(t+1)))\frac{\lambda(\lambda^* - 1)}{\lambda^*}. \quad (7.79)$$

Therefore, it suffices to have

$$1 < \beta u^F \frac{\lambda(\lambda^* - 1)}{\lambda^*}, \quad (7.80)$$

which holds if $\beta > \frac{1-F(u^*)}{u^F(\lambda^*-1)}$, noting the definition of u^F in (5.7). This shows that $\tilde{u}(\cdot)$ has a higher profit that $u^{ND}(\cdot)$ which contradicts its optimality. This completes the proof of part i).

The only thing which is left is to verify the base of the induction, that is to prove the claim for $t = 0$. For part ii), it is easy to see that all the relations (7.68)-(7.73) also hold for $t = 0$, noting that $y^{ND}(0) = x^{ND}(0)$ and $Z(0) = \emptyset$. Similar story holds for part i). ■

3.8 Upperbound on p_1 Assuring Infinitely Many Jumps above p_1 in Example 1

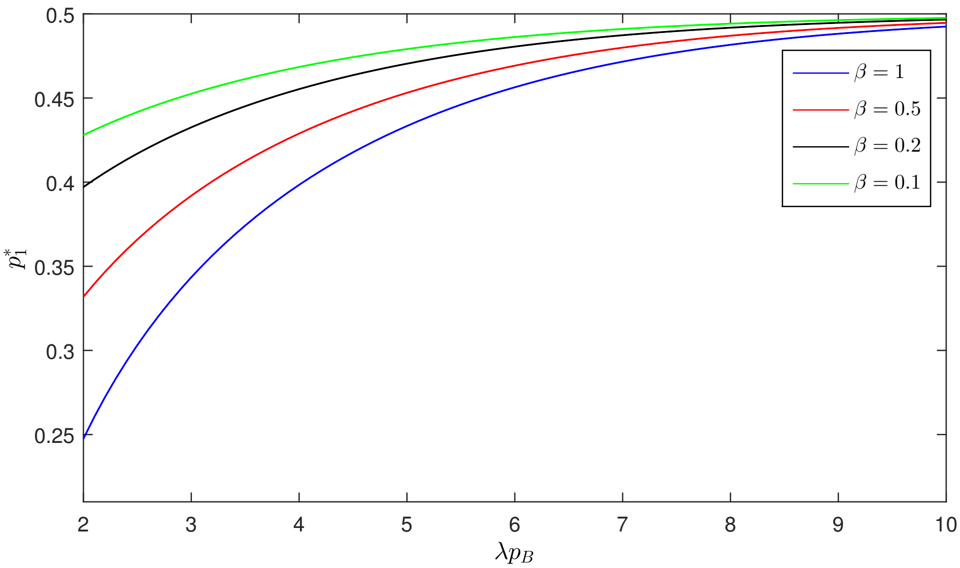

Figure 3-3: Values of p_1 that guarantee frequent jumps above p_1 for different values of β and λp_B, using $u = 0.5$ as a (sub)maximizer. Having $p_1 < p_1^*$ guarantees infinitely many jumps above p_1 under the optimal policy.

3.9 More Evidence on Price Drops from App Market

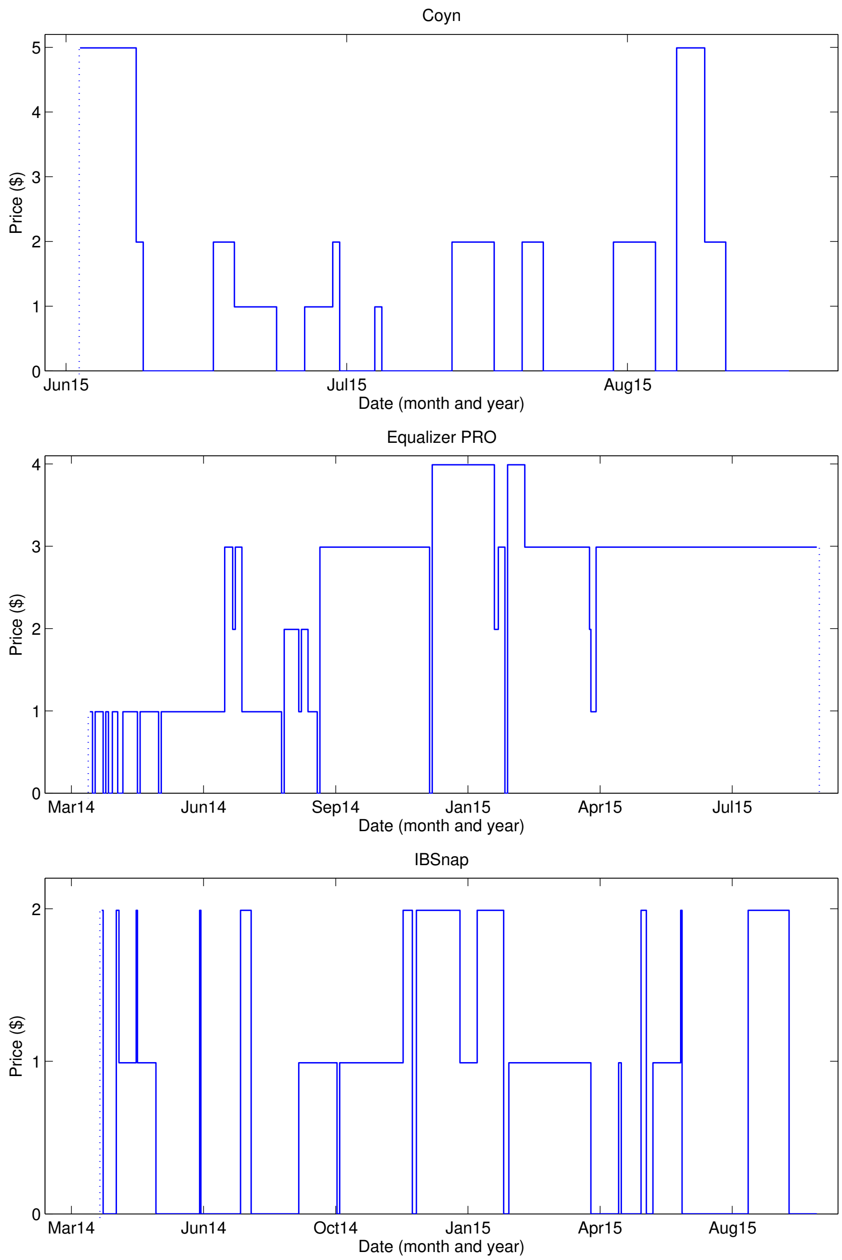

Figure 3-4: Price histories of three other apps (Coyn, Equalizer PRO, and IBSnap) from the time they were debuted to September 2015.

Chapter 4

Optimal Contracting in Networks

4.1 Introduction

Over the past decade, there has been a flurry of research on social networks and monetization of social data. From display advertising to viral marketing, firms have been trying to use the consumers' social interactions to increase their sales revenue leading to exponential growth in revenue of online social networking platforms. While most revenue generated from social media is through advertising, there has been a major push to devise intelligent strategies for pricing products that exhibit network effects. Information about goods and services often spreads in networks as "word of mouth", while in other instances, the product itself has features that induces positive network effects. Examples include Nike+, the technology that tracks data of every run and connects runners around the world, to online music services such as Apple Music and Spotify. In all these examples, firms can utilize the positive externality to sell more goods and services since the decision of each consumer is influenced by the choices of her friends and acquaintances in social networks.

In many instances firms have data on the consumption of the products/services they sell, as well as the social network activity of their consumers. For example, online social

This chapter is coauthored with Ali Jadbabaie. This chapter is published in the *Journal of Economic Theory*, vol. 183. I am grateful to our anonymous referees for their detailed comments. I also thank Daron Acemoglu, Amir Ajorlou, Dirk Bergemann, Ricardo Caballero, Matt Elliott, Kris Iyer, Matt Jackson, Leonid Kogan, Andrey Malenko, Asu Ozdaglar, Hamid Sabourian, Alireza Tahbaz Salehi, Juuso Toikka, Alex Wolitzky and seminar participants at MIT, Harvard, UC Berkeley and USC for helpful comments.

networking communities, such as Facebook, Twitter, Instagram and Pinterest allow firms to target users based on users' social interactions. There are also some companies that provide services to firms based on information on aggregate network effects obtained by quantitative analyses of consumers' online behavior (e.g., Klout, Commun.it, Cloze.com and Kred). Firms spend major effort to exploit the underlying network effects and forces to maximize profit. However, a major difficulty in using this information is uncertainty in how much agents *value* the externalities. The externality value is often private information of the agents. As an example, an agent may become aware of opinions or experience of her friends about their usage of a specific good or service via her social interactions, but how much she cares (i.e., the extent of her attention) about their opinions is often her private information.

A monopoly firm selling a divisible good that demonstrates network effects would naturally face the following questions: How should firms incorporate knowledge of the underlying social network structure of their customers in their selling strategies when there is incomplete information in network effects? Is uncertainty between agents about network effect beneficial to the firm? Does explicit knowledge of the network structure always matter in firms' profit maximizing strategies?

To address these questions, we study optimal contracting (multilateral and bilateral) between a firm and a network of agents (consumers) with two distinguishing features: First, the positive externality that agents receive from their neighbors' consumption is captured by the underlying social structure that is often incomplete (i.e., not an all-to-all graph) with different agents in the network have varying degrees of influence.[1] Second, and more importantly, the extent of network externality to each agent is privately held and is unknown to the firm and other agents. Our aim is to understand how network structure affects optimal contracts and to investigate the nature of resulting distortions and inefficiencies in principal-agent(s) problems, when the above two features are both present.

The goal of this paper is to study optimal contracting in networks with incomplete information about strength of network effects. A major twist in our model, and a point of departure from existing literature, is that the impact of aggregate network effect (externalities) to each agent is her private information, interpreted as the agent's *type*. Network structure and uncertainty in the strength of network effect are essential features of our anal-

[1]This heterogeneity of influence is important as the extent of externality is often determined by the influential agents in the network. These agents are key individuals whose consumption behavior has a large impact on the consumption levels of others. Developing new methods and algorithms to identify influential/key agents from big data sets has recently attracted a lot of attention, see Probst, Grosswiele and Pfleger (2013) for a recent survey of this literature.

ysis leading to the following implications.

4.1.1 Overview of the results

Network-dependent distortion propagation. We start with explicitly specifying optimal multilateral contracts. Our first—rather intuitive—result shows that because of the positive network externality effect, inefficiency in an agent's trade propagates throughout the network, causing a downward network-dependent distortion in other agents' trades. Due to the connectivity in the symmetrized network,[2] inefficiency in only one agent's trade is sufficient to cause a downward distortion in *all* the agents' trades—even trades of efficient agents.[3] Furthermore, we show that the maximum distortion in the whole network is caused by influential agents who are connected to other agents with high influence.

Given the complexities associated with implementation of optimal multilateral contracts and difficulties associated with commitment and enforcement, we specify simpler yet suboptimal bilateral contracts, where additional restrictions are added to contract space to ensure that the maps are diagonal.[4] We explicitly characterize optimal bilateral contracts. Given the explicit characterization, we first consider how uncertainty among consumers (agents) affects firm's profit.

Benefits of uncertainty. We show that the firm benefits from increasing uncertainty among consumers. The intuition is as follows. There are two sources of uncertainty in the model. One uncertainty is among the agents themselves, and the other is between the firm and the agents. The uncertainty among the firm and the agents results in the canonical adverse selection effect that manifests itself in the difference between the first-best and the second-best solutions. However, impact of the uncertainty among the agents themselves is more profound and is captured by the structure of the first-best solution.[5] Given this observation, we show that higher uncertainty about an agent's type increases the expected quantity consumed by *all* agents, and, thus, due to the strategic complement property, this

[2] The symmetrized network is defined as the sum of the network adjacency and its transpose. As we will see later, since the influences are asymmetric, the network adjacency matrix is not (necessarily) a symmetric matrix.

[3] Following Laffont and Martimort (2002), we call an agent *efficient*, only when she reports the maximum type.

[4] In the literature sometimes bilateral contracting problems are called non-linear pricing.

[5] In the first-best solution there is no uncertainty between the firm and the agents, thus uncertainty in the first-best solution comes from the uncertainty among the agents.

effect in turn increases the firm's expected profit. Similar result holds in the multilateral contracts.

Multilateral v.s. Bilateral contracts. Our explicit characterizations of the optimal contracts show that *even a small shock to externality* will make a potentially large change in the first and second best allocations as well as agents' influences[6] in the network. This structural difference particularly manifests itself in the importance of second-hop information in the structure of the bilateral contracts.

We also consider how firm's profit in these contracts, consumers' social influences and network characteristics are related to one another. Surprisingly, when the agents' centralities in the underlying network are similar the resulting bilateral contract is network-independent—network structure has no effect on firm's profit. As a result, in balanced networks, explicit knowledge of network structures has no benefit for the firm. We further derive simple upper and lower bounds on the firm's profit in terms of the network spectral properties (i.e., the maximum and minimum eigenvalues).

Clearly, firm's profit in any multilateral contract upper-bounds the profit in the corresponding bilateral contract. How is this gap related to network structures? Next, we study how network structure affects the firm's profit gap and profit ratio for multilateral and bilateral contracts.

Profit Gap, profit ratio and network structure. We show that the optimality gap is directly related to the centrality of nodes of an auxiliary network in which certain directed paths of length 2 have become edges. Furthermore, the optimality gap monotonically increases as additional links are added and decreases when edge weights are reduced. We also find graph-theoretic upper bound on this profit gap that is increasing in the largest eigenvalue of the underlying network. A lower bound on the profit gap is also provided which is increasing in the maximum in-degree in the network. Hence, the main take away of this result is that the profit gap is larger in most dense networks.

We further analyze the ratio of the profits in multilateral and bilateral contracts. For large balanced networks the ratio remains constant. Therefore, in large economies with limited heterogeneity in in-degrees, firm's profits in multilateral and bilateral contracts are proportional to one another. We can derive an upper bound on the profit ratio in terms of the network spectral properties and relate it to the extent of dispersion in the original

[6] As measured by their centralities.

network. In particular for networks with high dispersion (e.g., networks with large diameter) the proposed upper bound decreases. We also provide a lower bound on this profit ratio in terms of the Frobenius norm of Bonacich centrality in the symmetrized network and relate it to the extent of concentrations in networks. Particularly, the main take away of this result is that in most networks with the mean-preserving spread of centralities, as the standard deviation of the centralities increases the lower bound rises, increasing the profit ratio.

Networks with unbounded or vanishing profit gap. Finally, we focus on particular core-periphery network structures and show that the extent of asymmetry in agents' in-degrees can have a major impact on the profit gap. Specifically, for certain networks (core-periphery structures), the profit gap can grow unbounded as the network size grows. According to this result, in large networks with large influence asymmetries (e.g, stark heterogeneity in in-degrees) firm's restriction to bilateral contracts may result in major loss of profit. We also show that the profit gap can go to zero when externality weights are getting sufficiently small as the network size grows.

4.1.2 Related Literature

The goal of this paper is to study optimal contracting in networks with incomplete information in the strength of interactions. A major twist in our model, and a point of departure from existing literature, is the uncertainty in strength of externalities which is private information of each agent. We consider a model with strategic complementarity where an increase in the trade of others leads a given agent's higher trade to have relatively higher payoff compared to the same agent's lower trade. This effect has been the subject of extensive work (e.g., Farrell and Saloner (1985), Katz and Shapiro (1986)). However in the existing literature, the network effects often correspond to *all to all* or complete graphs, whereas in our set up agents interact locally only with their neighbors according to the underlying network structure.

Strategic interactions under presence of local network effects have been analyzed as *network games* by a series of papers in the past decade. These include Ballester, Calvo-Armengol and Zenou (2006*a*), Bramoulle and Kranton (2007*a*), Candogan, Bimpikis and Ozdaglar (2012*a*), Bloch and Quérou (2013), Corbo, Calvo-Armengol and Parkes (2007), Galeotti and Goyal (2009*a*), Jackson and Zenou (2015), Fainmesser and Galeotti (2016*a*,*b*). Following the common trend in this literature, we assume the payoff function of agents takes the form of a quadratic function. In Ballester, Calvo-Armengol and Zenou (2006*a*), the authors explicitly characterize the Nash equilibrium of a network game when agents choose their effort simultaneously. Furthermore, the authors show that the peer effect game has a unique Nash equilibrium in which each agent's effort is proportional to her Bonacich centrality in the original network. Other authors such as Bramoulle and Kranton (2007*a*), Corbo, Calvo-Armengol and Parkes (2007), and Bramoullé, Kranton and D'Amours (2014) study a similar game in the context of provisioning of public goods. Candogan, Bimpikis and Ozdaglar (2012*a*) extend this framework to include optimal pricing when the firm and the agents have perfect knowledge of the network structure. The authors show that the optimal consumption of each consumer depends on the Bonacich centrality of the network and as a result, optimal pricing strategy may involve offering discounts to agents who have a central position in the network and are a source of extra utility for their neighbors, while others who are the recipient of this extra utility will receive a mark-up in the price. Bloch and Quérou (2013) study a similar pricing setting and also model network effects as a deterministic graph that is commonly known to the firm and agents. They show that the value of lcoation-based price discrimination is dependent on the extent of convexity of the firm's cost function. Other authors such as Galeotti et al. (2010) and Sundararajan (2008*a*) consider the more realistic setting of limited knowledge of the firm about the structure of the original social network, such as knowledge of the degree distribution and thus derive optimal mechanisms that depend on this first degree measure of influence of an agent. Recently, Fainmesser and Galeotti (2016*a*) consider a set up similar to Galeotti et al. (2010) and Sundararajan (2008*a*) which utilizes the knowledge of the degree distribution according to which consumers are aware of their own in-degree and out-degree, and are uncertain about the level of interactions of others. The authors develop price discrimination strategies based on firm's information about consumers' influence and consumers' susceptibility to influence, and evaluate the value of information on network effects for the firm and show the value of information about consumers' influence and/or consumers' susceptibility increases in the dispersion of the in- and out-degree distributions as well as the average level of network effects.[7] In contrast to these works, we consider optimal contracting with local externalities and *incomplete* information in the strength of interactions.[8]

[7]Fainmesser and Galeotti (2016*b*) recently, in an oligopoly markets, study the practice of influencer marketing and its effect on market efficiency. They show firms subsidize (charge premia) consumers whose influence is above (below) average influential consumers. The size of premia/discounts depend on the strength of network effects and the level of information that firms have on consumers' influence. For Pay to bid auction framework see Kakhbod (2013*a*).

[8]Other network related works include: Business cycles (Acemoglu et al. (2012)), learning (Acemoglu et al. (2011), Golub and Jackson (2012), Golub and Jackson (2010)), advertising and targeted pricing (Kakhbod (2013*b*), Bimpikis, Ozdaglar and Yildiz (2016), Bloch (2016), Leduc, Jackson and Johari (2017), Shin (2017)

Our work is also in line with the body of literature on optimal contracting in principal-agent models by Mussa and Rosen (1978); Maskin and Riley (1984); JJ and Tirole (1990), and, in particular, devising optimal contracts with externalities by Segal (1999, 2003); Csorba (2008); Bernstein and Winter (2012).[9] In particular, Segal (1999, 2003) in his seminal papers develops a model of contracting with externalities under complete information with an all-to-all (complete) network structure and characterizes the nature of arising distortions and inefficiencies. When externalities are positive, he shows that each agent's consumption level is smaller in the resulting equilibrium allocation than in the socially efficient one. In contrast to these works, we consider optimal contracting under presence of a general *network structure* with *incomplete information* in the aggregate network effect. We particularly consider how network structures affect multilateral and bilateral contracts.

The rest of the paper is organized as follows. Section 4.2 presents the general model. Sections 4.3 characterizes the optimal multilateral contract and makes a linkage to the Bonacich centrality measure. The optimal bilateral contract is specified in Section 4.4, and the comparison between these contracts is in Section 4.5. Finally, conclusion closes the paper. Proofs are relegated to the appendix.

4.2 Model: optimal (multilateral) contracts

A monopoly firm sells a divisible good that may be used by n consumer agents in varying quantities. The firm's marginal cost of production is normalized to zero. Agents in this market form a social network denoted by $G = (V, E)$, where the vertex set $V = \{1, \cdots, n\}$ corresponds to the set of agents, and the edge set $E \subseteq V^2$ corresponds to social relationships. The corresponding adjacency in the network is captured by a matrix denoted by G (with a slight abuse of notation G denotes the network as well as the adjacency matrix corresponding

Chen, Zenou and Zhou (2018)), dynamic pricing (Ajorlou, Jadbabaei and Kakhbod (2018)), network games (Zhou and Chen (2015), Belhaj, Bervoets and Deroian (2016)). For implementation theory on networks see Kakhbod and Teneketzis (2012*b*), Kakhbod and Teneketzis (2012*c*), Kakhbod and Teneketzis (2012*a*) and Kakhbod et al. (2014), and for interactions between network and information theory see Kakhbod and Yazdi (2010), Mohajer, Pakzad and Kakhbod (2012), Mohajer and Kakhbod (2011) and Kakhbod and Zadimoghaddam (2011).

[9]Bernstein and Winter (2012) set up is in complete information that studies how to subsidies in order to obtain efficient coordination in a setting. Csorba (2008) considers Segal (1999, 2003) works to include incomplete information and demonstrate that the joint presence of asymmetric information and positive externality leads to downward distortion from the welfare-maximizing allocation for all agents. Importantly, the nature of externality in those work are such that the utility of an agent depends on the behavior of the whole set of agents, however, in our set up agents interact locally only with their neighbors, in addition, uncertainty is in network effects.

to it). The ij-th entry of G, denoted by g_{ij}, represents the *strength* of the influence of agent j on i. We assume $g_{ij} \in [0\ 1]$ for all i, j and we set $g_{ii} = 0$ for all i. For any $i, j \in V$, when $g_{ij} > 0$, agent j induces positive externality on agent i. We further assume the underlying undirected network formed by ignoring orientation of edges is connected.

We further assume the payoff function of agents is quadratic. More specifically, each agent i's payoff is specified as follows:

$$u_i(\theta_i, x_i, x_{-i}, t_i) = \underbrace{ax_i - \frac{b}{2}x_i^2}_{\text{direct utility}} + \underbrace{\theta_i \sum_{j \neq i} g_{ij} x_i x_j}_{\substack{\text{indirect utility} \\ \text{(type-dependent network effect)}}} - \underbrace{t_i}_{\text{payment}}, \tag{2.1}$$

where x_i is the amount of the good she consumes, $x_{-i} \triangleq (x_1, x_2, \cdots, x_{i-1}, x_{i+1}, \cdots, x_n)$ is the consumption of other agents excluding agent i and t_i is the disutility charged for x_i by the firm.

In addition, the parameter $\theta_i \in [\underline{\theta}\ \bar{\theta}]$ controls the (aggregate) strength of the network effect for agent i. Each agent's type θ_i is her private information (i.e., neither the firm nor other agents know the type), but it is commonly known that θ_i is drawn independently from a cumulative distribution F, for all i. We assume F is continuous and has a full support, with continuous density f such that $f(\cdot) > 0$. Moreover, $\phi(\tau) \triangleq \frac{1-F(\tau)}{f(\tau)}$, the reciprocal of the hazard rate, is assumed to be non-increasing for all τ.

The payoff function of each agent i can be interpreted as follows. The first two terms of (2.1) represent the direct utility agent i derives from her trade x_i, independent of the trades of her neighbors. The third term represents the (type-dependent) network externality effect when agent i accepts the proposed offer. Note that g_{ij} and g_{ji} might be different from one another. When $g_{ij} > 0$ trading of i and j with the firm has strategic complement property. That is, an increase in agent j's trade triggers an upward shift in agent i's trade. Finally, the last term is the cost of the trade.

In this economy, firm's objective is to devise a menu of optimal incentive quantity-price pairs $\{x_i(\cdot), t_i(\cdot)\}$ for each agent i so as to maximize her total ex-ante profit defined as:

$$\mathbf{E}_{\theta \in [\underline{\theta}\ \bar{\theta}]^n} \left[\sum_{i=1}^{n} t_i(\theta) \right]. \tag{2.2}$$

Notations Before proceeding further, we introduce some notation that will be used throught. We will occasionally write "$-i$" to mean agents other than agent i, and θ to denote a type profile, i.e., $\theta \in [\underline{\theta}\ \bar{\theta}]^n$. If M is a square matrix, then M^T denotes its transpose. Unless indicated otherwise, a bold letter denotes a column vector (e.g. $\mathbf{x}$), its transpose is denoted by $\mathbf{x}^T$ and the i's-th element is denoted by $[\mathbf{x}]_i$. Finally, I denotes the identity matrix.

4.3 Contracting game

The timing of the contracting game is as follows:

- Period 1: Each agent i observes her private type θ_i.
- Period 2: Firm announces to each agent i the menu of quantity-price pairs $\{x_i(\cdot), t_i(\cdot)\}$.
- Period 3: Each agent i based on her private type θ_i determines her optimal announcement type that is

  $$\hat{\theta}_i = \arg\max_{\tau_i \in [\underline{\theta}\ \bar{\theta}]} \mathbf{E}_{\theta_{-i}} \left[u_i(\theta_i, x_i(\tau_i, \theta_{-i}), x_{-i}(\tau_i, \theta_{-i}), t_i(\tau_i, \theta_{-i}))\right],$$

  subject to her participation constraint

  $$\mathbf{E}_{\theta_{-i}} \left[u_i(\theta_i, x_i(\hat{\theta}_i, \theta_{-i}), x_{-i}(\hat{\theta}_i, \theta_{-i}), t_i(\hat{\theta}_i, \theta_{-i}))\right] \geq 0,$$

  (we have assumed the reservation utility of each agent is normalized to zero).
- Period 4: The payoff of all agents and the firm are realized.

We assume each agent i, before making her decision, observes the menus offered to all agents. This assumption essentially rules out the situation where the firm and one agent trade secretly, and removes the case where one agent's trade is contingent on the trades of the others.[10] Appealing to the Revelation Principle, in characterizing optimal quantity-price pairs $\{x_i(\cdot), t_i(\cdot)\}_{i=1}^n$ firm can focus on direct revelation mechanisms in which the agents announce their types and truthful reporting constitutes a Bayes-Nash equilibrium[11]. Thus,

[10] This is common in contracting literature, see Segal (1999, 2003).

[11] By Revelation Principle, without loss of generality, firm can restrict her search to design contracts profiles wherein each agent i always finds it optimal to report her type θ_i truthfully.

the firm's problem is rewritten as:

$$\begin{aligned}
\max_{\{x_i(\cdot),t_i(\cdot)\}_{i=1}^n} \quad & \mathbf{E}\left[\sum_{i=1}^n t_i(\theta)\right] \\
\text{subject to} \quad & \text{(IC):} \;\; \theta_i = \arg\max_{\tau_i \in [\underline{\theta}\ \bar{\theta}]} \mathbf{E}_{\theta_{-i}}\left[u_i(\theta_i, x_i(\tau_i,\theta_{-i}), x_{-i}(\tau_i,\theta_{-i}), t_i(\tau_i,\theta_{-i}))\right] \;\; \forall i, \;\; \theta_i, \quad (3.1) \\
& \text{(PC):} \;\; \mathbf{E}_{\theta_{-i}}\left[u_i(\theta_i, x_i(\theta_i,\theta_{-i}), x_{-i}(\theta_i,\theta_{-i}), t_i(\theta_i,\theta_{-i}))\right] \geq 0 \;\forall i, \;\; \theta_i, \quad (3.2)
\end{aligned}$$

where (3.1) and (3.2) are the corresponding incentive-compatibility (IC) and participation constraints (PC), respectively. As in standard adverse selection problems, the solution of the above program under absence of the IC constraints is called the *first best* solution, denoted by $\{x_i^{FB}(\cdot), t_i^{FB}(\cdot)\}_{i=1}^n$. And when IC constraints are present, the solution is called the *second best* solution and is denoted by $\{x_i^{SB}(\cdot), t_i^{SB}(\cdot)\}_{i=1}^n$.

4.3.1 First best and second best solutions

Since without the IC constraints in the firm's problem all the PC constraints must bind, the first best trade profile $\{x_i^{FB}(\cdot)\}_{i=1}^n$ is indeed (ex ante) total surplus maximizing:

$$\{x_i^{FB}(\cdot)\}_{i=1}^n \in \arg\max_{\{x_i(\cdot)\}_{i=1}^n} \mathbf{E}\sum_{i=1}^n \left[a x_i - \frac{b}{2}x_i^2 + \theta_i x_i \sum_{j\neq i} g_{ij} x_j\right]. \quad (3.3)$$

$\{x_i^{FB}(\cdot)\}_{i=1}^n$ serves as a benchmark. Any discrepancy between the second best trade profile $\{x_i^{SB}(\cdot)\}_{i=1}^n$ and the benchmark solution, for each realization of θ, is called *distortion* which is the source of inefficiency. To ensure the objective function in (3.3) is indeed concave and has an interior solution for each type profile $\theta \in [\underline{\theta}\ \bar{\theta}]^n$ we make the following assumption. This condition stipulates that local complementarities have to be small enough compared to own concavity, which prevents excessive feedback which can lead to the absence of a finite trade profile.

Assumption 4.3.1. *For each $i \in V$, $b > \bar{\theta}\sum_{j\neq i}(g_{ij} + g_{ji})$.*

Before characterizing the second best solution we have the following assumption, to ensure $x_i^{SB}(\theta)$ is an interior solution[12] to the firm's problem, for all $i \in V$ and $\theta \in [\underline{\theta}\ \bar{\theta}]^n$.

[12] Without Assumption 4.3.2 we may need to do bunching.

Assumption 4.3.2. $\psi(\underline{\theta}) \triangleq \underline{\theta} - \frac{1-F(\underline{\theta})}{f(\underline{\theta})} \geq 0.$[13]

$\psi(\theta_i)$ is usually referred as *virtual type*. Next proposition characterizes the first best and the second best trade profiles in terms of the network structure.

Proposition 4.3.1. *For any type profile $\theta \in [\underline{\theta}\ \bar{\theta}]^n$:*

(i) the first best trade profile is given by

$$\mathbf{x}^{FB}(\theta) = a\ [bI - (M_\theta G + G^T M_\theta)]^{-1}\mathbf{1}, \tag{3.4}$$

where $M_\theta \triangleq diag(\theta_1, \theta_2, \cdots, \theta_n)$.

(ii) the second best trade profile is given by

$$\mathbf{x}^{SB}(\theta) = a\ [bI - (M_\psi G + G^T M_\psi)]^{-1}\mathbf{1}, \tag{3.5}$$

where $M_\psi \triangleq diag(\psi(\theta_1), \psi(\theta_2), \cdots, \psi(\theta_n))$.

Proof. See Appendix. ■

The consequences of this explicit characterizations of the first best and the second best trade profiles in terms of the underlying network structure are analyzed in the following sections.

4.3.2 Distortion vs. network structure: Role of influential agents

Given the result in Proposition 4.3.1, we now introduce a distortion vector as a function of the underlying network structure. By characterizing the distortion in terms of Bonacich centrality of agents, we determine how inefficiency in one agent's trade propagates throughout the network. Finally, we show the impact of this distortion propagation on the firm's profit varies depending on the agents' locations in the network, and thus identify influential agents whose inefficiencies result in maximum loss on the firm's profit. Let us first provide the definitions of Bonacich centrality and the distortion vector, for any type profile $\theta \in [\underline{\theta}\ \bar{\theta}]^n$.

[13] As an example, one can easily show that the following class of ($\beta-$parametric) cumulative distribution

$$F_\beta(\tau) = 1 - \left(\frac{\bar{\theta} - \tau}{\bar{\theta} - \underline{\theta}}\right)^{\frac{1}{\beta}}, \quad \beta > 0, \quad \tau \in [\underline{\theta}\ \bar{\theta}],$$

has increasing hazard rate and, in addition, for any $\beta \leq \frac{\underline{\theta}}{\bar{\theta} - \underline{\theta}}$, Assumption 4.3.2 is satisfied.

Definition 4.3.1 (Bonacich centrality)**.** *Given a network with adjacency matrix G, a scalar α, and a vector $\mathbf{v}$ (wighted) Bonacich centrality with parameter α is defined by*[14]

$$K(G, \alpha, \mathbf{v}) = (I - \alpha G)^{-1}\mathbf{v} = [k_1 \ k_2 \ \cdots k_n]^T,$$

and k_i denotes agent i's Bonacich centrality.

Definition 4.3.2. *For any $\theta \in [\underline{\theta} \ \bar{\theta}]^n$, the difference between the first-best and the second best is called distortion and is denoted by $\mathbf{d}(\theta)$ as follows:*

$$\mathbf{d}(\theta) \triangleq \mathbf{x}^{FB}(\theta) - \mathbf{x}^{SB}(\theta). \tag{3.6}$$

Using the first-best and second-best trade profiles, the distortion vector is characterized by the following Lemma.

Lemma 4.3.1. *For any type profile $\theta \in [\underline{\theta} \ \bar{\theta}]^n$:*

(i) $\mathbf{d}(\theta) = K\left(M_\theta G + G^T M_\theta, \frac{1}{b}, \frac{a}{b}\mathbf{1}\right) - K\left(M_\psi G + G^T M_\psi, \frac{1}{b}, \frac{a}{b}\mathbf{1}\right).$

(ii) Distortion is downward*, i.e.,* $\mathbf{d}(\theta) \geq \mathbf{0}$.

(iii) Let $[\mathbf{d}(\theta)]_i$ denote the distortion in agent i's trade with regard to the type profile θ. Then:

$$[\mathbf{d}(\theta)]_i > 0$$

if there exists at least one agent j whose $\theta_j \neq \bar{\theta}$.

Proof. See Appendix. ∎

The key insights of the above lemma are summarized as follows. First, because of the positive network externality in the payoff functions (see (2.1)), distortion is always downward (i.e. $\mathbf{d}(\theta) \geq \mathbf{0}$). Second, because of connectivity (i.e., existence of a path between any pair of agents in the symmetrized network $G + G^T$) inefficiency in only *one* agent's trade is *sufficient* to distort other agents' allocations downward. In other words, positive network externality together with connectivity imply that distortion in one agent's trade propagates throughout the network and distorts downward even the trade of an efficient agent whose

[14] It is assumed that $(I - \alpha G)^{-1}$ is well defined and nonnegative.

type is the highest type $\bar{\theta}$. Moreover, given the explicit characterization in part (i), the way distortion in one agent's trade propagates throughout the network depends on its position in the symmetrized network, captured by the node's Bonacich centrality.[15] This observation immediately gives rise to the following question: Are central agents (in terms of Bonacich centrality) in $G+G^T$ necessarily those key agents whose distortion have the highest effect on firm's profit? By key agents we naturally mean agents whose inefficiency create maximum distortion in the whole network. We show that this is not necessarily the case. To proceed, we first provide the definition of total distortion in a network for any type profile θ. The total distortion is directly related to the total firm's loss caused by propagations of distortions due to inefficiencies in agents' trades. Thus, we also specify agents whose private information (and firm's lack of access to it) can result in the largest *loss* to the firm's profit.

Definition 4.3.3. *For any type profile $\theta \in [\underline{\theta}\ \bar{\theta}]^n$, the total distortion is the sum of the discrepancies in the first and the second best allocations of all agents, given by:*

$$\mathcal{T}(\theta) = \mathbf{1}^T \mathbf{d}(\theta).$$

To show that central agents in $G+G^T$ are *not* necessarily key agents, we use the explicit characterization of the distortion vector in Lemma 4.3.1 and determine the overall distortion as well as the loss that inefficiency of any agent's type can create, when others are all at the efficient type (i.e. $\theta_j = \bar{\theta}$ for all $j \neq i$).[16]

Proposition 4.3.2. *Consider any agent $i \in \{1, \cdots, n\}$ and let $\theta_j = \bar{\theta}$, for all $j \neq i$. Let k_r denote agent r's Bonacich centrality characterized by $(I - \frac{\bar{\theta}}{b}(G+G^T))^{-1}\mathbf{1} = [k_1, k_2, \cdots, k_n]$. Then, the followings hold when $\Delta\theta = \bar{\theta} - \underline{\theta}$ is sufficiently small.*

(i) When $\theta_i \leq \bar{\theta}$, the overall distortion created by agent i in any network G is given by:

$$\mathcal{T}_i(\theta_i, \bar{\theta}_{-i}) = \frac{2a}{b^2}\ |\phi'(\bar{\theta})|\ (\bar{\theta} - \theta_i) \left[\sum_{j=1}^{n} k_i k_j g_{ij}\right]. \tag{3.7}$$

(ii) Let $\theta_i \in \{\underline{\theta}, \bar{\theta}\}$ with $Prob\{\theta_i = \underline{\theta}\} = \mathsf{v} > \frac{\Delta\theta}{\theta}$. Then, the lack of firm's certain knowledge (incomplete information) about agent i's type maximizes expected loss in the firm's

[15] To visualize better the way distortion propagates in a network, see Example B-1 in Appendix 4.7.

[16] The assumption that all except one agent is efficient is to obtain a meaningful normalization. This simplifies the comparisons and serves as a simple benchmark to achieve our goal to show that central agents in $G+G^T$ are *not* necessarily key agents.

profit when $i \in V^\star$ where

$$V^\star \triangleq \arg \max_{i\in\{1,2,\cdots,n\}} \mathbb{E}\left[\underbrace{\Pi^{FB}(\theta_i, \bar{\theta}_{-i}) - \Pi^{SB}(\theta_i, \bar{\theta}_{-i})}_{loss}\right] = \arg \max_{i\in\{1,2,\cdots,n\}} \sum_{j=1}^{n} k_i k_j g_{ij}, \quad (3.8)$$

where Π^{FB} and Π^{SB} denote firm's profit due to the corresponding first best and second best contracts, respectively.

Proof. See Appendix. ∎

The above proposition provides three results that are worth highlighting: First, the overall distortion created by agent i decreases in her type. This is intuitive because distortions created by agent i decreases as agent i becomes more efficient. The second point is about identifying key agents whose inefficiencies create the maximum distortion in the whole network. According to (3.7), the maximum distortion in the whole network is due to agent for whom the product of her centrality and her neighbor's centrality are maximized in the symmetrized network. Finally, the maximum loss on the firm's profit occurs when the firm does not have access to the key agents' type.

Hence, in oder to identify the key agents in G, not only each agent i's centrality in $G + G^T$ (i.e. k_i) but also its neighbors' centralities (i.e., k_j through g_{ij}) have to be taken into account. As a consequence, the notion of a key agent in our model is different from other models with the linear-quadratic payoff functions (e.g., most notably Ballester, Calvo-Armengol and Zenou (2006*a*), Candogan, Bimpikis and Ozdaglar (2012*a*) and Bloch and Quérou (2013)).

In these works key agents are directly central agents in the underlying network G or the symmetrized network $G + G^T$. This structural difference in our model is mainly due to the nature of uncertainty, which is placed in the network effect. In Appendix 4.11 we show that if uncertainty is placed in the *direct* utility, then maximum distortion is due to agents that are central in the symmetrized network $G + G^T$, following the previous works in the literature.

The impact of this observation is highlighted by the next example. In the next example we show that centrality of an agent in the symmetrized network $G + G^T$ is not sufficient for the agent to be a key agent as far as trade is concerned.

Example 2 (Two-star Network)**.** Consider the following Two-star network, $g_{ij} \in \{0, \kappa\}$, for[17] all i, j, capturing the interconnection among the agents. Agent 1 obtains externality from agent 2 but not vice versa. In this network, agents 1 and 2 are both equally central in $G + G^T$, i.e., $k_1 = k_2$. In addition, the peripheral nodes all have the same centralities, lower than the central nodes, i.e., $k_j = k_i < k_1 = k_2$, for all $i, j \in \{3, 4, \cdots, 2d+2\}$. Let $\Lambda(\theta_i) \triangleq \Big(2\phi'(\bar{\theta})(\theta_i - \bar{\theta})\Big)$, for all $\theta_i \in [\underline{\theta} \;\; \bar{\theta}]$. Using Proposition 4.3.2, we show, although agents 1 and 2 both have the same centralities, inefficiency in agent 1's report creates *more* distortion (in the whole network) than agent 2's report. To see this, let $\theta_1 = \theta_2 = \gamma \in [\underline{\theta} \; \bar{\theta}]$, then:

$$\mathcal{T}_1(\gamma, \bar{\theta}_{-1}) - \mathcal{T}_2(\gamma, \bar{\theta}_{-2}) = \Big(\kappa\Lambda(\gamma)\Big) \; k_1 k_2 = \underbrace{\Big(\kappa\Lambda(\gamma)\Big)}_{>0} \; (k_1)^2 > 0,$$

where $(\gamma, \bar{\theta}_{-2}) \triangleq (\bar{\theta}, \gamma, \bar{\theta}, \cdots, \bar{\theta})$ and $(\gamma, \bar{\theta}_{-1}) \triangleq (\gamma, \bar{\theta}, \bar{\theta}, \cdots, \bar{\theta})$. Notice that this discrepancy is strictly increasing in κ and k_1, resulting in a potential drastic difference by increasing κ and/or k_1.

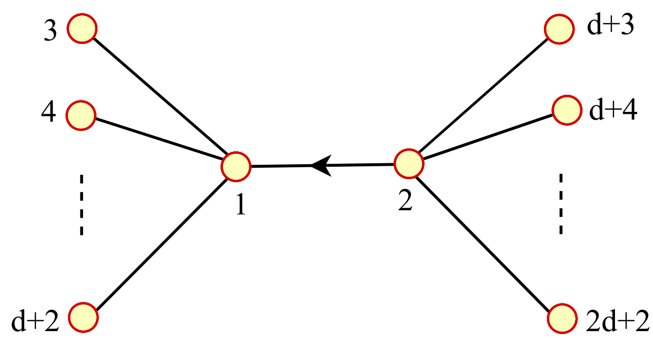

Figure 4-1: Two-star network. Agents 1 and 2 are both equally central in $G + G^T$. But distortion created in the whole network due to inefficiency in agent 1's type is strictly higher than 2's.

4.4 Bilateral contracts

So far we have analyzed the optimal *multilateral* contracts under the presence of network externalities. In this section, we consider a situation where the firm is forced to offer "bilateral contracts" (which can be essentially interpreted as nonlinear pricing). Although a bilateral

[17] $\kappa > 0$ is small enough so that $(I - \bar{\theta}(G + G^T))^{-1}$ is well-defined. Note also that $a = b = 1$.

contract is not profit maximizing,[18] it is practically simpler to implement. In the next section, we compare multilateral and bilateral contracts in terms of network structure. In bilateral contracts firm's objective is to devise a menu of optimal incentive quantity-price pairs $\{x_i(\theta_i), t_i(\theta_i)\}$ for each agent i, where $\theta_i \in [\underline{\theta}\ \bar{\theta}]$, so as to maximize her total ex-ante profit defined as: $\sum_{i=1}^{n} \mathbf{E}\left[t_i(\theta_i)\right]$. In contrast to multilateral contracts, each agent i's menu is *only* a function of her (reported) type and *not* the type profile. Using the Revelation Principle, firm's problem can be recast as follows:

$$\begin{aligned}
&\max_{\{x_i(\cdot),t_i(\cdot)\}_{i=1}^{n}} && \sum_{i=1}^{n} \mathbf{E}\left[t_i(\theta_i)\right] \\
&\text{subject to} && \text{(IC):} \quad \theta_i = \arg\max_{\tau_i \in [\underline{\theta}\ \bar{\theta}]} u_i(\theta_i, x_i(\tau_i), t_i(\tau_i)) \quad \forall i,\ \theta_i, && (4.1)\\
& && \text{(PC):} \quad u_i(\theta_i, x_i(\theta_i), t_i(\theta_i)) \geq 0 \quad \forall i,\ \theta_i, && (4.2)
\end{aligned}$$

where (4.1) and (4.2) are the corresponding incentive-compatibility and participation constraints, respectively. Moreover, the payoff function of each agent i is now updated as:

$$u_i(\theta_i, x_i, t_i) = \underbrace{a x_i - \frac{b}{2} x_i^2}_{\text{direct utility}} + \underbrace{\theta_i x_i \sum_{j \neq i} g_{ij}\ \mathbf{E}[x_j(\theta_j)]}_{\text{type-dependent indirect utility}} - \underbrace{t_i}_{\text{payment}}.$$

Comment 4.4.1. Before characterizing the optimal bilateral contract, we notice that multilateral and bilateral contracts can be different, not only in the magnitude of allocations but also in the induced order of allocations. While having a different allocation is natural, we note that certain network structures might also induce a different rank order on the amount of allocations (see Example B-3 in Appendix 4.7).

We start by a technical following assumption. This assumption ensures that matrices (that appear in Proposition 4.4.1) in the allocations are indeed invertible and, thus, the corresponding first and second best solutions are interior and bounded solutions.[19]

[18]It is not profit maximizing, since these contracts are less constrained and there exists inefficiency even at the highest type profile.

[19]Another meaningful (but more conservative) presentation of this assumption is given by Lemma 4.6.1 in the appendix. Assumption 4.4.1 is ensured if

$$\mu_\tau \left(\sum_{j=1}^{n} g_{ij} + \sum_{j=1}^{n} g_{ji} \right) + \frac{\sigma_\tau^2}{b} \sum_{j=1}^{n} \sum_{k=1}^{n} g_{kj} = \mu_\tau \left(d_i(in) + d_i(out) \right) + \frac{\sigma_\tau^2}{b} \sum_{i=1}^{n} d_i(in) < b,$$

for any i, where $d_i(in)$ is the in-degree and $d_i(out)$ is the out-degree of agent i and τ stands for the corre-

Assumption 4.4.1. *Let $\rho(A)$ denote the spectral radius of A.*[20] *Then*

$$\max\{\rho\left(\frac{1}{b}\left[\mu_\psi(G+G^T)+\frac{\sigma_\psi^2}{b}G^TG\right]\right),\rho\left(\frac{1}{b}\left[\mu(G+G^T)+\frac{\sigma^2}{b}G^TG\right]\right)\}<1,$$

where

$$\mu \triangleq \mathbf{E}[\theta_i], \qquad \sigma^2 \triangleq Var(\theta_i), \qquad \mu_\psi \triangleq \mathbf{E}[\psi(\theta_i)], \qquad \sigma_\psi^2 \triangleq Var(\psi(\theta_i)).$$

Next proposition characterizes the first best and the second best trade profiles in terms of network structures. **E**dskip

Proposition 4.4.1. *For any i and θ_i, the optimal trade profiles are characterized as follows.*

(i) The optimal first best trade profile is given by

$$x_i^{FB}(\theta_i)=\frac{a}{b}(\theta_i-\mu)[GK\mathbf{1}]_i+a[K\mathbf{1}]_i \tag{4.3}$$

where $K \triangleq \left[bI-\mu(G+G^T)-\frac{\sigma^2}{b}G^TG\right]^{-1}$.

(ii) The optimal second best trade profile is given by

$$x^{SB}(\theta_i)=\frac{a}{b}(\psi(\theta_i)-\mu_\psi)[GK_\psi\mathbf{1}]_i+a[K_\psi\mathbf{1}]_i \tag{4.4}$$

where $K_\psi \triangleq \left[bI-\mu_\psi(G+G^T)-\frac{\sigma_\psi^2}{b}G^TG\right]^{-1}$.

Proof. See Appendix. ∎

Given the above characterization, the optimal first-best trade profile explicitly depends on the types, network structure, and the mean and variance of the types. The uncertainty (i.e., the variance of the agents' types) affects the optimal trade profile through the matrix K. We emphasize that matrix K has an important role in the characterization of the optimal trade profiles in terms of agents' location in the underlying network. Similar intuition holds for the second best trade profile with K_ψ.

sponding ψ and θ (the mean and variance).

[20] Let $\{\lambda_1,\cdots,\lambda_n\}$ be the eigenvalues of a matrix $A\in\mathbf{R}^{n\times n}$. Then its spectral radius $\rho(A)$ is defined as: $\rho(A)=\max\{|\lambda_1|,\ldots,|\lambda_n|\}$.

It is important to note that when $[\underline{\theta}\ \bar{\theta}]$ shrinks to a single value θ, the allocation in our model becomes $x(\theta) = a[bI - \theta(G + G^T)]^{-1}\mathbf{1}$, which is, expectedly, the same allocation as in Candogan, Bimpikis and Ozdaglar (2012*a*) and Bloch and Quérou (2013).[21] Thus, the interesting point that the above characterization shows is that even a small uncertainty in θ (captured by σ^2) will make a substantial change in the structure of the allocation.[22] Moreover, somewhat expectedly, a quick comparison between Propositions 4.3.1 and 4.4.1 shows that multilateral and bilateral contracts may *not* necessarily induce the same order of allocations (see Appendix 4.10). Finally, given the explicit characterizations in Proposition 4.4.1, the following proposition shows that increasing uncertainty in an agent's type increases expected consumption in *all* the agents. The intuition is simple. Increasing uncertainty in an agent's type increases expected consumption of the agent.[23] Moreover, $G+G^T$ is a connected network. Thus, due to the strategic complementarity in consumptions, an increase in the expected consumption of the agent lead to an increase in the expected consumptions of other agents. Similar result holds in the second best contract, by considering increasing uncertainty in corresponding virtual types.

Proposition 4.4.2. *Given an agent i, we can state the following:*

(i) The expected first best (second best) consumption of agent i is given by $\mathbf{E}[x_i^{FB}(\tilde{\theta}_i)] = a\,[K\mathbf{1}]_i$ *(respectively,* $\mathbf{E}[x_i^{SB}(\tilde{\theta}_i)] = a\,[K_\psi\mathbf{1}]_i$*).*

(ii) Increasing uncertainty in agent j's type increases the expected consumption of every one (i.e., $\frac{\partial \mathbf{E}[x_i(\tilde{\theta}_i)]}{\partial \sigma_j^2} \geq 0$*). Similar result holds in the second best contract.*

Proof. See Appendix. ■

4.4.1 Firm's profit

In this section we consider firm's profit in bilateral contracts and its relation to network structures, and provide simple upper and lower bounds on the firm's profit in terms of the network's *spectral* properties.

[21] We notice that marginal cost of production in our model is, without loss of generality, normalized to zero.

[22] We further note that in Candogan, Bimpikis and Ozdaglar (2012*a*) and Bloch and Quérou (2013) the authors consider *linear pricing*. We, however, in bilateral contracts consider *non-linear pricing*.

[23] This is due to the nature of $K \propto \sum_{k=0}(\mu(G + G^T) + \frac{\sigma^2}{b}G^TG)^k$ and the linear-quadratic feature of the payoffs.

Proposition 4.4.3. *The firm's optimal expected profit in the first best is given by* $\Pi_G^{bi.} = \frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1}$. *Similar result holds in the second best, i.e.,* $\Pi_G^{bi.,SB} = \frac{a^2}{2}\ \mathbf{1}^T K_\psi \mathbf{1}$. *Consequently, using Definition 4.3.1, the firm's optimal expected profit is also given by* $\Pi_G^{bi.,SB} = \frac{a}{2} \left\| K\left(\mathcal{G} + \mathcal{G}^T, \frac{1}{b}, \frac{a}{b}\mathbf{1}\right) \right\|_1$, *where* $\mathcal{G} \triangleq \mu_\psi G + \frac{\sigma_\psi^2}{2b} G^T G$ *and* $\|\cdot\|_1$ *denotes the* L_1*-norm.*[24]

Proof. See Appendix. ∎

Proposition 4.4.3 characterizes the firm's optimal profit with respect to network structures and, moreover, makes a linkage to the sum of Bonacich centralities (L_1-norm) corresponding to the network $\mathcal{G}$. In other words, Proposition 4.4.3 shows that under presence of uncertainty in the extent of externality, what's important is the centrality in $\mathcal{G}$ (in which, due to $G^T G$, certain paths of length two have become edges). The following example highlights this observation.

Example 3. Consider an economy consisting of a firm and 5 consumer agents.[25] Let the interconnection among the agents be captured by a class of networks denoted by G_τ, $\tau \in [0\ 1]$. Let $G_\tau = \tau G_1 + (1-\tau)G_0$, $\forall \tau \in [0\ 1]$, where G_0 and G_1 denote the star-outward and star-inward networks, respectively (see Fig. 4-2). Thus, $G_0 \rightsquigarrow G_1$ when τ varies from 0 to 1.[26]

Given Proposition 4.4.3, the firm's optimal profit is the sum of Bonacich centralities in $\mathcal{G}_\tau + \mathcal{G}_\tau^T$, where, given the above parameters, for any $\tau \in [0\ 1]$:

$$\begin{aligned} \mathcal{G}_\tau + \mathcal{G}_\tau^T &= \mu_\psi (G_\tau + G_\tau^T) + \frac{\sigma_\psi^2}{b} G_\tau^T G_\tau |_{b=1, \mu_\psi = 1, \sigma_\psi^2 = \frac{1}{3}} = G_\tau + G_\tau^T + \frac{1}{3} G_\tau^T G_\tau \\ &= \underbrace{G_0 + G_1}_{\text{fixed for all } \tau} + \frac{1}{3} G_\tau^T G_\tau, \end{aligned} \tag{4.5}$$

where the last equality follows by the construction of G_τ. Observing this structure allows us to disentangle the effect of second-hop neighbor information via the $G^T G$ term. The impact of changing τ, from 0 to 1, on the firm's profit is pictorially depicted in Fig. 4-2. As numerically shown in this figure, among all star networks, characterized by G_τ, *star-inward*, i.e. G_1, results in the highest (ex-ante) profit. This is due to the fact that the second hop neighbors have the most impact when $\tau = 1$.

[24] That is $\|x\| = \sum_i |x_i|$, the L_1-norm is simply the sum of the absolute values. The distance derived from this norm is called the Manhattan distance or L_1 distance.

[25] Number of agents is set to 5 just for ease of illustration.

[26] *Parameters*: Focusing on the network structure, we assume each agent's type is either low or high, i.e., $\theta_i \in \{\underline{\theta} = 1, \bar{\theta} = 2\}$, with probability $\text{Prob}\{\theta_i = 1\} = 1 - \text{Prob}\{\theta_i = 2\} = \mathsf{v} = \frac{3}{4}$, for all i. Thus, $\mu_\psi = \underline{\theta} = 1$ and $\sigma_\psi^2 = \left(\frac{1-\mathsf{v}}{\mathsf{v}}\right)(\Delta\theta)^2|_{\mathsf{v}=\frac{3}{4}} = \frac{1}{3}$. In addition, $b = 1, a = 10$ and $g_{ij} \in \{0, .1\}$, for all i, j.

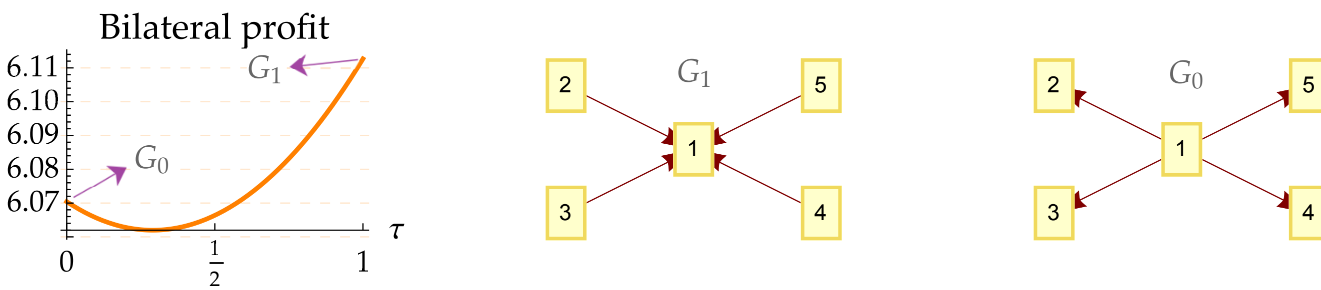

Figure 4-2: Interconnection among the agents.

We wrap up this section by the following proposition characterizing quantitative upper and lower bounds on the firm's bilateral profit in terms of the network's *spectral* properties. More specifically, in terms of the maximum and minimum eigenvalues of the original symmetrized network.

Proposition 4.4.4. *For ease of exposition normalize $b = 1$. Consider an $n-$agent economy with the symmetric*[27] *network structure G. Let λ_{min} and λ_{max} be the smallest and largest eigenvalues*[28] *corresponding to G. Then for small σ^2:*

$$\frac{na^2}{2} f(\lambda_{min}) \leq \Pi_G^{bi.} \leq \frac{na^2}{2} f(\lambda_{\max}),$$

where $f(\lambda) = \frac{1-2\xi\lambda+\sigma^2\lambda^2}{(1-2\xi\lambda)^2}$ is increasing and convex in λ. Similar result holds for the second best contract.

Proof. See Appendix. ∎

The above result shows that aggregate characteristics of networks such as the minimum and maximum eigenvalues are enough to bound the profit. Intuitively, since $\lambda_{\max} \leq d_{\max}$ (where $d_{\max}$ denotes the network's maximum degree[29]), thus when the underlying network structure becomes more dispersed (decreasing $d_{\max}$) the proposed upper bound in most networks falls, decreasing the profit. Dispersed networks with low $d_{\max}$ include: path-like networks, networks with large diameter and low $d_{\max}$.

[27] Symmetric means $G^T = G$.
[28] Note that since G is loop-less (i.e. $g_{ii} = 0$, for all i), thus $\sum_i \lambda_i = 0$, implying that $\lambda_{min} \leq 0 \leq \lambda_{max}$.
[29] For the proof of $\lambda_{\max} \leq d_{\max}$ see Lovász (2007).

4.4.2 Benefits of uncertainty

Given the findings in Proposition 4.4.1 and Proposition 4.4.3, it is important to distinguish between two sources of uncertainty in the model. One source is the uncertainty among the agents, and the other is the uncertainty between the firm and each agent. Under the imperfect information case (the second best) both sources are present. Under the full information case (the first best) only uncertainty among the agents is present. Importantly, the uncertainty between the firm and each agent manifests its effect in the disparity of the first-best and the second-best solutions in which, as in a canonical adverse selection problem, θ changes to $\psi(\theta)$, which is *not* structural. However, the impact of uncertainty among the agents is more profound and is captured by the *structure* of the first-best solution. Similar arguments hold in the multilateral contract. The following remark summarizes.

Comment 4.4.2. We note that similar to multilateral contracts, analyzing the effect of uncertainty in the full information case and imperfect information case are essentially the same. In the imperfect information case the uncertainty is in the virtual type, whereas in the full information case uncertainty is in the type θ.

We next consider the impact of uncertainty on the firm's expected profit.

Proposition 4.4.5. *The optimal second best expected profit of the firm increases as the uncertainty in the agents' (virtual) types increases.*

The preceding proposition states that the *firm prefers the uncertainty in agents' types to be higher*. It is helpful to examine the reasoning behind this statement. As mentioned before, increasing the uncertainty in the agents' types has two effects: it not only increases the adverse selection effect that the firm faces against each agent (uncertainty between firm and agents), but also increases the uncertainty in the beliefs of one agent regarding her neighbors' types (i.e., uncertainty among the agents). The former effect tends to reduce the firms expected profit, as evidenced by the fact that the firm would obtain higher profits in the full information setting. However, from Proposition 4.4.2, we see that higher uncertainty about a neighbor's type increases the quantity consumed by other agents, and hence the latter effect tends to increase the firm's expected profit. Thus, there are two opposing forces at play here and the proposition states that the latter effect is stronger. We attribute this to the positive externality in the network: increasing the consumption profile of any one agent causes consumption of other agents in the network to increase thus earning even higher profits for the firm.[30]

The following corollary shows that in the first best allocation, firm's profit is increasing in the variance of the types.

Corollary 4.4.1. *Let μ be fixed. Then, firm's first best expected profit increases with a greater uncertainty among the agents, i.e.* $\frac{\partial \Pi_G^{bi.}}{\partial \sigma^2} > 0$.[31]

Proof. See Appendix. ∎

4.5 Multilateral vs. Bilateral contracts in networks

In this section, we explore differences in the firm's profit between multilateral and bilateral contracts under the presence of network externality and uncertainty. In particular, we wish to address the following network dependent questions: How and when does network structure matter? How does network structure affect the difference in firm's profit using these contracts? Is it possible for this difference to grow unboundedly or even go to zero? And finally, how does network structure affect the ratio of the profits? In the sequel, for ease of expositions we assume without loss of generality that the curvature term of the utility of each agent to be normalized to 1 (i.e., $b = 1$).

4.5.1 Balanced networks and bilateral contracts

It turns out that when the in (or out)-degree of the symmetrized network is the same constant across agents, the network structure has no impact on firm's profit in the case of bilateral contracts. This means that explicit knowledge of the network structure (beyond the knowledge of uniformity of centralities) has no value for the firm. To highlight this observation, we next define balanced networks to be those networks whose symmetrized graph is degree-regular.

[30] It is also worth noting that employing Taylor expansions for the moments of the function $\psi(\cdot)$ of the random variable θ_i we have: $Var(\psi(\tilde{\theta}_i)) \propto \left(\frac{d}{d\rho}\psi(\rho)|_{\rho=\mathbf{E}[\tilde{\theta}_i]}\right)^2 Var(\tilde{\theta}_i) = (\dot{\psi}(\mathbf{E}[\tilde{\theta}_i]))^2 \; Var(\tilde{\theta}_i)$. That is, increasing $Var(\tilde{\theta}_i)$ increases $Var(\psi(\tilde{\theta}_i))$, when θ is around its mean.

[31] One way to change σ^2 and keep μ fixed is via changing $\Delta\theta = \bar{\theta} - \underline{\theta}$. That is, to ensure that a change in σ^2 does not affect μ, we change $\Delta\theta$ so that μ is kept fixed.

Definition 4.5.1. *An economy with the network structure G is balanced if the symmetrized network $G+G^T$ has constant row sums, i.e., there exists a $\tau > 0$ such that*[32] $\sum_j (g_{ij}+g_{ji}) = \tau$, *for all i.*[33]

Intuitively, balanced economies are those in which no agent is overtly influential in terms of exertion of externality.

Proposition 4.5.1. *Consider an n-agent economy with the balanced structure G. Then, when σ^2 is small, firm's profit using bilateral contract is given by* $\Pi_G^{bi.} = \frac{a^2}{2}\left[n\zeta + \sigma^2\zeta^2 \sum_{i=1}^n \left(\sum_{j=1}^n g_{ij}\right)^2\right]$, *where* $\zeta = \frac{1}{1-\xi\tau}$. *Similar result holds for the second best contract, when θ is replaced with $\psi(\theta)$.*

Proof. See Appendix. ■

The above proposition implies that when a firm uses bilateral contracts, the profit in any balanced network is only a function of the number of agents and the common degree of the symmetrized network (which is the row sum of the symmetrized network matrix). When the number of agents is fixed, all networks corresponding to the same τ result in the *same* (expected) profit. Consequently, explicit knowledge of the network structure beyond the common degree has no value for the firm. This result is particularly useful in practice, as it says when uncertainty among the agents is small and firm uses the bilateral contracts (i.e., non-linear pricing), then the firm's profit does not depend on the details of the interconnection structure. This, however, may not be the case for multilateral contracts (see Fig. 4-3). The following corollary summarizes the above discussed result.

Corollary 4.5.1. *All balanced networks in which each agent has an equal in-degree (e.g. cycle, wheel, regular networks) result in the same (expected) profit when firm uses bilateral contracts.*

---

[32]Note that τ is small enough so that all matrices remain invertible.

[33]We note that the term balanced graphs in graph theory has a slightly different meaning than the one we use here: in graph theory, a directed graph is called balanced if the in-degree and out-degree of each node is the same. Here, on the other hand, we call the graph balanced if the *symmetrized network* is degree-regular graph. This means that the sum of in- and out-degrees is the same constant for all nodes in the network.

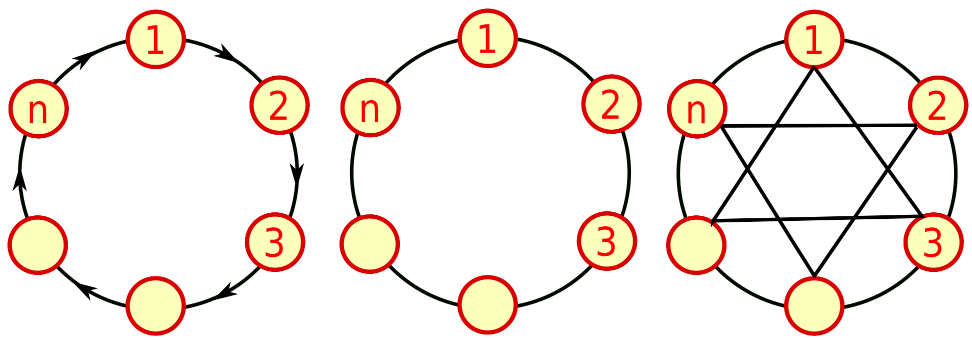

Figure 4-3: All of these networks with bilateral contract result in the same profit, however, for multilateral contract they may behave differently.

4.5.2 Difference in firm's profit (Profit Gap): bilateral vs. multilateral contracts

We now explore the difference in firm's profit between multilateral and bilateral contracts as a function of the network structure. In particular, the following lemma provides a closed form expression characterizing the suboptimality gap.

Lemma 4.5.1 (Profit Gap). *Suppose there exists an $\hat{m} > 0$, such that $\mathbf{E}[(\theta_i - \mu)^k] < (\hat{m}\sigma)^k$, for all i and*[34] *$k \geq 3$. For any network G and for small σ^2, let $S \triangleq (I - \mu(G + G^T))$, and $K \triangleq (I - \mu(G + G^T))^{-1}\mathbf{1} = S^{-1}\mathbf{1}$ be the centrality of the symmetrized graph, and define $R_i \triangleq E_i G + G^T E_i$, where E_i is the matrix with only ith diagonal set as 1 and other entries as zero, then*[35]

$$\Pi_G^{multi.} - \Pi_G^{bi.} = \left(\frac{a^2}{2}\sigma^2\right) K^T \left\{\sum_{i=1}^{n} R_i(S^{-1} - I)R_i + diag[(G \circ G)\mathbf{1}]\right\} K. \tag{5.1}$$

Similar result holds in the second best contract with θ replaced by $\psi(\theta)$.

Proof. See Appendix. ∎

With the above useful lemma we can now characterize the suboptimality of bilateral contracts. Some results are immediate. First, it is immediate[36] that $\Pi_G^{\text{multi.}} \geq \Pi_G^{\text{bi.}}$ and also

[34] One can easily show this constraint is easily satisfied by many distributions like uniform and power distributions, for appropriately chosen parameters.

[35] $\circ$ denotes the elementwise Hadamard product. Given matrices $A = (a_{ij})$ and $C = (c_{ij})$ of order $n \times n$ $A \circ C = (a_{ij}c_{ij})_{ij}$, where $a_{ij}c_{ij}$ is a scalar and $A \circ C$ is of order $n \times n$.

[36] Notice that since S^{-1} is well defined and positive, $S^{-1} - I = \sum_{i \geq 1}(\mu(G + G^T))^k$ is a positive matrix.

the gap increases with the extent of uncertainty, i.e., $\frac{\partial\{\Pi_G^{\text{multi.}} - \Pi_G^{\text{bi.}}\}}{\partial\sigma^2} > 0$. Second, this profit gap is directly related to Bonacich centrality of the symmetrized network.

Using Lemma 4.5.1, the following proposition characterizes several comparative statics on the profit gap in terms of network structure and its spectral properties.

Proposition 4.5.2 (Profit Gap and Network Properties). *Consider an n-agent economy with the interaction matrix G.*

(i) Suppose there is no link (externality) from j to i in G, and let $G + \{g_{ij}\}$ denote the new matrix induced by introducing the link $\{ij\}$. Assuming all the invertibility assumptions are preserved, then $\Pi_G^{multi.} - \Pi_G^{bi.} < \Pi_{G+\{g_{ij}\}}^{multi.} - \Pi_{G+\{g_{ij}\}}^{bi.}$.

(ii) If $0 < \alpha < 1$, then $\Pi_{\alpha G}^{multi.} - \Pi_{\alpha G}^{bi.} < \alpha^2 \left(\Pi_G^{multi.} - \Pi_G^{bi.}\right)$.

(iii) Let G be symmetric[37] and λ_{max} denote the largest eigenvalue of G. Then

$$0 \leq \Pi_G^{multi.} - \Pi_G^{bi.} \leq \left(\frac{na^2}{2}\sigma^2\right) \; h(\lambda_{\max}),$$

where $h(\lambda) \triangleq \lambda^2 \; \frac{3+2\xi\lambda}{(1-2\xi\lambda)^3}$, that is increasing and convex in λ.

(iv) Let $g_{ij} \in \{0, g\}$, for all i, j. Then[38]:

$$\Pi_G^{multi.} - \Pi_G^{bi.} \geq \frac{a^2}{2}\sigma^2\mu^2 \; \left(d_{max}^3 \; g^4\right),$$

where d_{max} denotes the maximum in-degree in G.

Similar results hold in the second best contract with θ replaced by $\psi(\theta)$.

Proof. See Appendix. ∎

Intuitively, independent of network structures, the profit gap monotonically increases when extra links are added and networks become denser. This is because more links amplifies strategic complementarity between agents, increasing the influence of network externality on each agent's payoff, leading to an increase in the profit gap. Similarly, the profit gap deceases with lower weights g_{ij}. In this regard, Part (ii) shows that the influence of weight reduction on the profit gap is at least quadratic. Parts (iii) and (iv) provide quantitative upper and

[37]Symmetric means $G^T = G$.

[38]Note that g is potentially a function of n and small enough preserving all the invertibility assumptions.

lower bounds on the profit gap in terms of the network's *spectral* properties. The upper bound is increasing in maximum eigenvalue of underlying networks. Therefore, when the underlying network structure becomes more dispersed (e.g., path-like networks) the proposed upper bound in most networks falls, decreasing the profit gap. Part (iv) provides a lower bound in terms of the mean and the variance of the types and maximum in-degrees of underlying networks. This result implies that with more concentration in networks (i.e., increasing maximum in-degree of networks) the proposed lower bound increases, suggesting that the profit gap can be large in dense networks.

Can the profit gap explode or vanish? In the following example, we establish that when networks are heterogeneous with high dispersion (as in core-periphery networks) the restriction to bilateral contracts can result in profit losses that grow unboundedly with the size of networks. The converse holds as well. That is, expectedly, when the influencing weights (i.e., $g_{i,j}$) become proportionally small with the size of networks, the profit gap converges to zero. All of these statements are formally proved in Appendix 4.8.

Example 4. Let us start with formally defining star-family networks:

Definition 4.5.2 (Star-family)**.** *Let $G_0(n)$ and $G_1(n)$ denote, respectively, star-outward and star-inward graphs over n nodes. The star family $\mathfrak{G}$ includes all the δ-convex combination of the two star structures defined as $\mathfrak{G}(n) = \{G_\delta(n) : \quad \exists\ \delta \in [0,1]$ s.t. $G_\delta(n) = \delta G_1(n) + (1-\delta)G_0(n)\}$.*[39]

The following figure depicts the profit gap due to using the bilateral contracts in star networks with 50, 100 and 150 nodes.[40] The two extremes are star outward, $\delta = 0$, and star inward, $\delta = 1$, respectively, and the rests are the convex combinations of the two extremes. As the following graph shows, the maximum loss (within the star-family) occurs in star-inwards. We formally prove this in Appendix 4.8.[41] Moreover, the graph highlights the fact that with increasing number of agents, the profit gap may explode. This, of course, is not always true and depends on the network constructions. In Appendix 4.8, we show that

[39]Notice that the star-family of networks is centrality-preserving (i.e., $K = (I - \mu(G + G^T))^{-1}\mathbf{1}$ is the same for all the networks within the core-periphery (star) family). This is because all the networks in $\mathfrak{G}(n)$ have the same symmetrized adjacency (i.e., $G_\delta(n) + G_\delta(n) = G_{\delta'}(n) + G_{\delta'}(n)$ for all $\delta, \delta' \in [0,1]$).

[40]Other parameters: $b = \sqrt{n}$, $\mu = 1/2$, $\sigma^2 = 1/12$, and $a^2/2 = 1$.

[41]It is worth noting that the main driver of this result is the second term of Eq (5.1) in Lemma 4.5.1. Clearly, in the class of core-periphery structures the maximum centrality is in the core node, lets denote it by k_1, and thus by the symmetry in the periphery nodes in $S^{-1}\mathbf{1} = [k_1\ k_2\ \cdots\ k_n]^T$, we must have $k_1 > k_2 = \cdots = k_n$. Next, considering the second term of Eq (5.1) (in Lemma 4.5.1), the Hadamard product

the explosion in the profit gap can happen in networks that have major asymmetry in their in-degrees. When there is an agent dominant in its in-degree (e.g., star-inward) then the loss in the firm's profit because of using the simpler bilateral contracts may become unbounded, and, thus, firm's restriction to the simpler bilateral contracts may result in major loss in firm's profit.

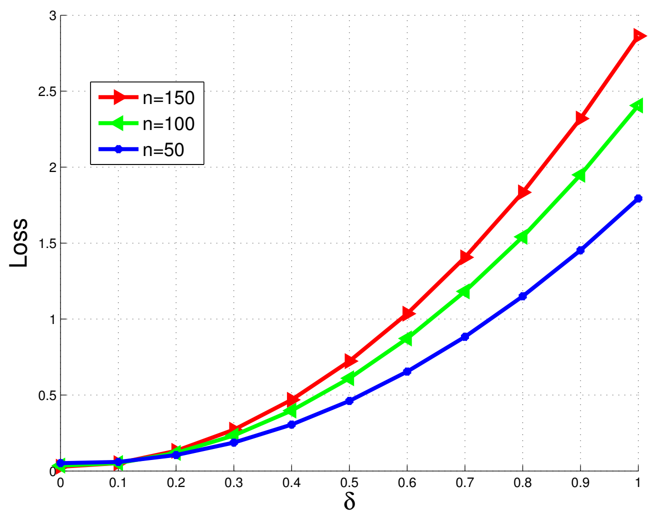

Figure 4-4: The chart visualizes the loss of using bilateral contracts for the family of centrality preserving star networks. The highest amount of loss happens in star (in)ward, when $\delta = 1$. Moreover, as the number of agents, i.e. n, increases, the amount of loss monotonically becomes higher.

4.5.3 Profit Ratio

In this section we consider the impact of heterogeneity of influences in the network on the ratio of the bilateral and multilateral profits. We first consider balanced and ring structures and then find upper and lower bounds on the profit ratio in general networks in terms of various spectral properties of the network. We show that in networks with balanced

$G \circ G$ implies that

$$K^T \text{diag}[(G \circ G)\mathbf{1}]K = \sum_{i=1}^{n} k_i^2 \parallel \deg_{in}(i) \parallel_2^2 = \mathbf{1}^T S^{-1} \text{diag} \left(\sum_{j=1}^{n} g_{1j}^2, \sum_{j=1}^{n} g_{2j}^2, \cdots, \sum_{j=1}^{n} g_{nj}^2 \right) S^{-1}\mathbf{1}.$$

Thus, due to $k_1 > k_2 = \cdots = k_n$, this term is maximized for star-inward, because $\sum_{j=1}^{n} g_{1j}^2$ is maximized in star-inward.

structures (as defined in Definition 4.5.1) firm's profit in these two contracts are of the same order (their ratio is of order 1). In other words, the following proposition asserts in economies with *limited heterogeneity in in-degrees* multilateral and bilateral contracts are (in profit) proportional to one another.

Proposition 4.5.3. *Consider the sequence of balanced networks $\{G(n)\}_{n=2}^{\infty}$ with the same in-degree at each node*[42]*. As n grows:*

$$\lim_{n\to\infty} \frac{\Pi^{multi.}_{G(n)}}{\Pi^{bi.}_{G(n)}} = O(1).$$

Proof. See Appendix. ∎

The following result shows the above observation is also true in cycle (clock-wise) networks that may not necessarily have a balanced structure.

Proposition 4.5.4. *Consider the sequence of (clock-wise) ring/cycle networks $\{G(n)\}_{n=2}^{\infty}$, where weights are either zero or one. Let $2\xi + \sigma^2 < 1$ and $\xi < \frac{1}{4}$. As n grows:*

$$\lim_{n\to\infty} \frac{\Pi^{multi.}_{G(n)}}{\Pi^{bi.}_{G(n)}} = O(1).$$

Proof. See Appendix. ∎

Profit Ratio: Upper and Lower bounds

This subsection aims to characterize upper and lower bounds on the ratio of the firm's profit in multilateral and bilateral contracts. These bounds are in terms of the graph spectral properties.

Proposition 4.5.5 (Upper bound)**.** *Consider an $n-$ agent economy with the network structure G. Let λ_{max} denote the maximum eigenvalue of the symmetrized network $G + G^T$. Then*

$$\frac{\Pi^{multi.}_{G}}{\Pi^{bi.}_{G}} \leq 1 + \sigma^2 g(\lambda_{\max}),$$

where $g(\lambda_{\max}) \triangleq \frac{\lambda^2_{max}}{(1-\xi\lambda_{max})^3}$ is increasing in $\lambda_{\max}$.

[42] Meaning that there exists a $0 < \kappa < 1$ so that $\sum_j g_{ij} = \kappa\tau$, for all i. We further assume $G + G^T$ is k regular, where k is finite.

Proof. See Appendix. ∎

The simple upper bound is in terms of the symmetrized network's *spectral* radius. Let $d_{max} = \max_i \sum_j (g_{ij} + g_{ji})$ denote the maximum degree in $G + G^T$. Since (by Proposition 2.1 in Lovász (2007))

$$\max\{\sqrt{d_{max}}, \text{average degree in } G + G^T\} \leq \lambda_{max} \leq d_{max}$$

and $g(\cdot)$ is an increasing function, thus as the symmetrized network $G + G^T$ becomes more *dispersed* (path-like networks, networks with large diameter, $d_{max} \downarrow$) the proposed upper bound in most networks falls, decreasing the profit ratio.

Next, we provide a lower bound on the profit ratio.

Proposition 4.5.6 (Lower bound)**.** *Consider an $n-$ agent economy with the network structure G. Let $K \triangleq (I - \mu(G + G^T))^{-1}\mathbf{1}$ be the centrality of the symmetrized network. Then*

$$\frac{\Pi_G^{multi.}}{\Pi_G^{bi.}} \geq 1 + \sigma^2 \frac{||diag(K)\ G||_F^2}{||K||_1 + ||K||_2^2},$$

where $||\cdot||_F$ stands for the Frobenius norm, and $diag(K) = diag(k_1, k_2, \cdots, k_n)$. Particularly, if $||d_i||_2^2 = 1$, for all i, (where $d_i = (g_{i1}, g_{i2}, \cdots, g_{ni})$, $i-th$ row of G), then $||diag(K)\ G||_F^2 = ||K||_2^2$. Thus:

$$\frac{\Pi_G^{multi.}}{\Pi_G^{bi.}} \geq 1 + \sigma^2 \frac{\frac{||K||_2^2}{||K||_1}}{1 + \frac{||K||_2^2}{||K||_1}}. \tag{5.2}$$

Proof. See Appendix. ∎

The lower bound in (5.2) is increasing in $||K||_2^2$. Therefore, the above result suggests that in most networks with the mean-preserving spread of centralities (i.e., keep $||K||_1$ fixed), as the standard deviation of the centralities increases (i.e., increasing $||K||_2^2$) the lower bound rises, increasing the profit ratio (Fixing $||K||_1$, the larger $\frac{||K||_2^2}{||K||_1}$, the more disperse the network becomes).

4.6 Conclusion

We studied optimal contracting strategies for a firm selling a divisible good that exhibits positive externality to a finite number of consumers in a social network. A special feature of the model, and a point of departure from existing literature, was the magnitude of network externality being the private information of each agent

We explicitly characterized the firm's optimal multilateral and bilateral contracts as a function of the underlying network structure. Due to the presence of positive network effects, inefficiency in one agent's trade propagates throughout the network and creates an unequal and network-dependent downward distortion in *all* the agents' trades, even agents with the highest type. The distortion vector can be characterized in terms of Bonacich centrality of the symmetrized network. In particular, inefficiencies in the trade of highly central agents who are connected to other central agents create the maximum distortion, causing the maximum loss in the firm's profit. We also explicitly characterized optimal bilateral contracts. The explicit characterization shows that, surprisingly, uncertainty in the value of the network externality among consumers is beneficial to the firm — meaning that the firm will be better off *not* to invest on reducing uncertainty among the agents.

Focusing on the suboptimality of the bilateral contracts, we explicitly characterize the profit gap between the multilateral and bilateral contracts in terms of network structures. We showed that heterogeneity in the centralities of different agents plays an important role in the profit gap. When the network structure is balanced, firm's profit in the bilateral contracts becomes independent of network structures. Furthermore, the optimality gap increases when new links are introduced. Focusing on the core-periphery structures, we also showed that the optimality gap can grow unboundedly when the network size grows. Finally we showed that increasing the standard deviation of the centralities leads to an increase in the profit ratio of bilateral and multilateral contracts.

Our results sheds light on scenarios in which it is profitable for firms to invest in finding the social network of their consumers and when it is not worth making such investments.

Appendix

In the following lemma we present an alternative condition for Assumption 4.4.1 to hold.

Lemma 4.6.1. *In order to obtain* $\rho\left(\frac{1}{b}\left[\mu_\tau(G+G^T)+\frac{\sigma_\tau^2}{b}G^TG\right]\right)<1$*, it suffices to have for* *any i:*

$$\mu_\tau \left(\sum_{j=1}^{n} g_{ij} + \sum_{j=1}^{n} g_{ji} \right) + \frac{\sigma_\tau^2}{b} \sum_{j=1}^{n} \sum_{k=1}^{n} g_{kj} = \mu_\tau \left(d_i(in) + d_i(out) \right) + \frac{\sigma_\tau^2}{b} \sum_{i=1}^{n} d_i(in) < b,$$

where $d_i(in)$ is the in-degree and $d_i(out)$ is the out-degree of agent i, and τ stands for the corresponding ψ and θ (the mean and variance).

Proof. Let us denote $S \triangleq \frac{1}{b}\left[\mu_\tau(G+G^T) + \frac{\sigma_\tau^2}{b} G^T G\right]$. Let $v = [v_1 \ v_2 \ \cdots \ v_n]^T$ be an eigenvector of S and its corresponding eigenvalue be λ. Thus, by definition, we have

$$Sv = \lambda v.$$

Let $v_i \triangleq \max\{|v_1|, |v_2|, \cdots, |v_n|\}$, where $|A|$ denotes the absolute value of A. Hence

$$|\lambda v_i| = |S_i v| = |\sum_{j=1}^{n} s_{ij} v_j| \leq |\sum_{j=1}^{n} s_{ij}||v_i|, \tag{6.1}$$

where S_i denotes the i^{th} row of S. Let $\hat{g}_{ij}$ denote an entry of G^T located at i^{th} row and j^{th} column, thus, $\hat{g}_{ij} = g_{ji}$. Therefore,

$$\begin{aligned}
\sum_{j=1}^{n} s_{ij} &= \frac{1}{b}\left[\mu_\tau \left(\sum_{j=1}^{n} g_{ij} + \sum_{j=1}^{n} g_{ji} \right) + \frac{\sigma_\tau^2}{b} \sum_{j=1}^{n} \left(\sum_{k=1}^{n} \hat{g}_{ik} g_{kj} \right)\right] \\
&= \frac{1}{b}\left[\mu_\tau \left(\sum_{j=1}^{n} g_{ij} + \sum_{j=1}^{n} g_{ji} \right) + \frac{\sigma_\tau^2}{b} \sum_{j=1}^{n} \left(\sum_{k=1}^{n} g_{ki} g_{kj} \right)\right] \\
&\leq \frac{1}{b}\left[\mu_\tau \left(\sum_{j=1}^{n} g_{ij} + \sum_{j=1}^{n} g_{ji} \right) + \frac{\sigma_\tau^2}{b} \sum_{j=1}^{n} \sum_{k=1}^{n} g_{kj}\right],
\end{aligned}$$

where the last inequality follows since $g_{ij} \in [0\ 1]$, thus $g_{ki} g_{kj} \leq g_{kj}$, $\forall\ k, j$ and i.

Now, if $\frac{1}{b}\left[\mu_\tau\left(\sum_{j=1}^{n} g_{ij} + \sum_{j=1}^{n} g_{ji}\right) + \frac{\sigma_\tau^2}{b} \sum_{j=1}^{n} \sum_{k=1}^{n} g_{kj}\right] < 1$, then the above inequality implies that $|\sum_{j=1}^{n} s_{ij}| < 1$. Therefore, due to Eq. (6.1), we have $|\lambda| < 1$, and consequently, $\rho(S) < 1$.

We further note that each eigenvalue of $I - S$ is equal to $1 - \lambda$ where λ is an eigenvalue of S. Thus $\rho(S) < 1$ implies that all the eigenvalues of $I - S$ are non-zero and, therefore, $I - S$ is invertible. ∎

4.7 Appendix: extra examples

Example B-1 (Path Network): Nature of distortion This example considers distortion propagations in a path network when there is uncertainty in the extent of interactions.

In particular, it shows that the downstream propagation and upstream propagation due to inefficiency in one agent's trade are different in nature. The downstream propagation (i.e. distortion in trades) is due to strategic complementarities. However, the upstream propagation is due to both complementarities and the need to reduce information rent.

Consider a market including a firm and three consumers connected via simple chain as depicted in Fig. 4-5. The payoff function of each agent is given as follows:

$$\begin{aligned} u_1(x_1, x_2, x_3, t_1) &= ax_1 - \frac{1}{2}x_1^2 - t_1, \\ u_2(\theta_2, x_1, x_2, x_3, t_2) &= ax_2 - \frac{1}{2}x_2^2 + \theta_2\ x_1\ x_2 - t_2, \\ u_3(x_1, x_2, x_3, t_3) &= ax_3 - \frac{1}{2}x_3^2 + \gamma\ x_2\ x_3 - t_3, \end{aligned}$$

where x_i and t_i are consumption and payment of agent i, respectively.

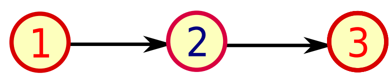

Figure 4-5: Line network. Interconnection among the agents.

In this market, agents 2 and 3 obtain positive externality from the consumptions of their neighbors. The externality from agent 2 to agent 3 is weighted by $\gamma \in (0\ 1)$ that is publicly known. In contrast, the externality from agent 1 to agent 2 is weighted by θ_2 that is agent 2's *private* information. It is, however, commonly known that $\theta_2 \in \{\underline{\theta}, \bar{\theta}\}$, $0 < \underline{\theta} < \bar{\theta} \leq \gamma < \frac{1}{\sqrt{2}}$, with $\text{Prob}\{\theta_2 = \underline{\theta}\} = \mathsf{v} > \frac{\Delta\theta}{\bar{\theta}}$, where $\Delta\theta \triangleq \bar{\theta} - \underline{\theta}$. Firm's objective is to devise a menu of optimal incentive quantity-price pairs $\{(x_i(\theta_2), t_i(\theta_2))\}$ for each agent i, given agent 2's report $\theta_2 \in \{\underline{\theta}, \bar{\theta}\}$, so as to maximize her total ex-ante profit subject to the corresponding participation constraint (PC) of each agent $i \in \{1, 2, 3\}$,[43] and incentive compatibility constraints (IC) of agent 2. Thus, firm's problem is precisely written as:

$$\begin{aligned} \max_{\{x_i(\cdot), t_i(\cdot)\}_{i=1}^3} \quad & \mathsf{v} \sum_{i=1}^{3} t_i(\underline{\theta}) + (1 - \mathsf{v}) \sum_{i=1}^{3} t_i(\bar{\theta}) \\ \text{subject to} \quad & \text{PC}_i(\theta_2), \quad \theta_2 \in \{\underline{\theta}, \bar{\theta}\},\ \text{for all } i \in 1, 2, 3, \\ & \text{IC}_2(\theta_2), \quad \theta_2 \in \{\underline{\theta}, \bar{\theta}\}. \end{aligned} \tag{7.1}$$

In the full information case where θ_2 is commonly known, only PC constraints are active in the firm's problem, thus, the first-best trade profile maximizes the social surplus. In the second-best solution, the incentive constraint is active. Thus, in the firm's problem, the IC constraint of agent 2 binds at her efficient type and firm must give some rent to agent 2 for revealing her type truthfully. Importantly, given the above network structure, this rent is

[43] The reservation utility of each agent is normalized to zero.

controlled not only by agent 2's trade but also agent 1's trade. To understand this better, let us first characterize the first and second best solutions.

Lemma 4.7.1. *The first and second-best trade profiles are given by:*

$$x_1^{SB}(\bar{\theta}) = x_1^{FB}(\bar{\theta}), \quad x_3^{SB}(\bar{\theta}) = x_3^{FB}(\bar{\theta}), \quad x_2^{SB}(\bar{\theta}) = x_2^{FB}(\bar{\theta}),$$

$$x_1^{FB}(\underline{\theta}) = a + \underline{\theta}\; x_2^{FB}(\underline{\theta}), \quad x_3^{FB}(\underline{\theta}) = a + \gamma\; x_2^{FB}(\underline{\theta}), \quad x_2^{FB}(\underline{\theta}) = \frac{a(1 + (\underline{\theta} + \gamma))}{1 - (\underline{\theta}^2 + \gamma^2)}.$$

$$x_1^{SB}(\underline{\theta}) = a + (\underline{\theta} - S)\; x_2^{SB}(\underline{\theta}), \quad x_3^{SB}(\underline{\theta}) = a + \gamma\; x_2^{SB}(\underline{\theta}), \quad x_2^{SB}(\underline{\theta}) = \frac{a(1 + ((\underline{\theta} - S) + \gamma))}{1 - ((\underline{\theta} - S)^2 + \gamma^2)}.$$

where $S \triangleq \left(\frac{1-\mathsf{v}}{\mathsf{v}}\right) \Delta\theta > 0.$

Proof. See Appendix (omitted proofs). ∎

As expected, when agent 2 is efficient, i.e. $\theta_2 = \bar{\theta}$, the first and second best are equal, i.e. $x_i^{SB}(\bar{\theta}) = x_i^{FB}(\bar{\theta})$, $i \in \{1, 2, 3\}$ and when agent 2 is inefficient, i.e. $\theta_2 = \underline{\theta}$, agent 2's allocation is distorted downward, i.e. $x_2^{SB}(\underline{\theta}) < x_2^{FB}(\underline{\theta})$. However, while agents 1's and 3's payoffs are both common knowledge, due to the positive network externality, in contrast with adverse selection environments with no externality,[44] distortion in agent 2's trade yields downward distortion in agents 1's and 3's trades. Most importantly, this distortion has *unequal* effects on agents 1's and 3's trades. To understand the reasoning, we look at agent 2's information rent, that is $R_2 = \Delta\theta\; x_1^{SB}(\underline{\theta}) x_2^{SB}(\underline{\theta})$. Essentially, the downstream propagation and upstream propagations are different in nature. The downstream propagation (i.e. distortion in trades) is only because of *strategic complementarities.* Distortion in the allocation of the low-type agent 2 causes the optimal allocation of agent 3 to change. But this change is only there as a consequence of the distortion to low-type 2. And of course, this can be simply seen in the first order optimality condition for $x_3^{SB}(\underline{\theta})$: any change in $x_2^{SB}(\underline{\theta})$ has to be reflected in $x_3^{SB}(\underline{\theta})$. But, the same is *not* true for agent 1's allocation. Changing 1's trade also distort's high-type 2's *incentive* in mimicking low-type 2. In fact, by distorting 1's trade downward, the firm makes sure that 2's IC constraint is satisfied in *a cheaper way*, as essentially the effective type of agent 2 is now $\theta_2 x_1^{SB}(\underline{\theta})$. Lower $x_1^{SB}(\underline{\theta})$ helps the firm reduces the rent she has to pay to agent 2's high-type.

Example B-2 (Kite Network) Consider the following network:[45] Let us assume all the agents report the efficient type $\bar{\theta}$ except agent 6. Agent 6's report varies from the lowest to the highest type, i.e., $\theta_i = \bar{\theta}$ for $1 \leq i \leq 5$ and $\underline{\theta} \leq \theta_6 \leq \bar{\theta}$. If there was no network there would be no distortion for the efficient agents $1, 2, 3, 4, 5$ and onlyagent 6's allocation would distort downward. However, due to strategic complementarities and the interconnection among the agents, distortion in agent 6's allocation propagates in the *whole* network. More

[44] Since the firm has complete information about agents 1's and 3's payoffs, as it is standard in adverse selection problems, one would expect to achieve *efficiency* in their trades, i.e., equality of the second and the first bests.

[45] Parameters: $g_{ij} = g_{ji} \in \{0, .15\}$, $i, j \in \{1, 2, \cdots, 5\}$, capturing the interrelations among the agents. Let θ_i be uniformly distributed on $[\frac{2}{3}\; 1]$ (b is normalized to 1 and $a = 10$). In this network agent 2 is most central in terms of Bonacich centrality.

importantly, each agent ,depending on its proximity to agent 6 experiences different amount of reduction in its allocation. The following figure shows the corresponding distortion and the way it unequally propagates throughout the network as θ_6 changes from $\underline{\theta}$ to $\bar{\theta}$. Note that when agent 6 is efficient, i.e. $\theta_6 = \bar{\theta}$, distortion vanishes, i.e. the first and second-best for all agents are equal. However, when $\theta_6 < \bar{\theta}$ distortion propagates and all agents experience reduction in their allocations, even though their types are *efficient*.

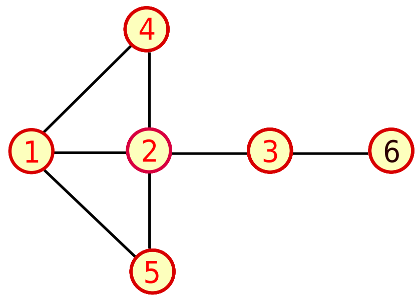

Figure 4-6: Kite network. Interconnection among the agents.

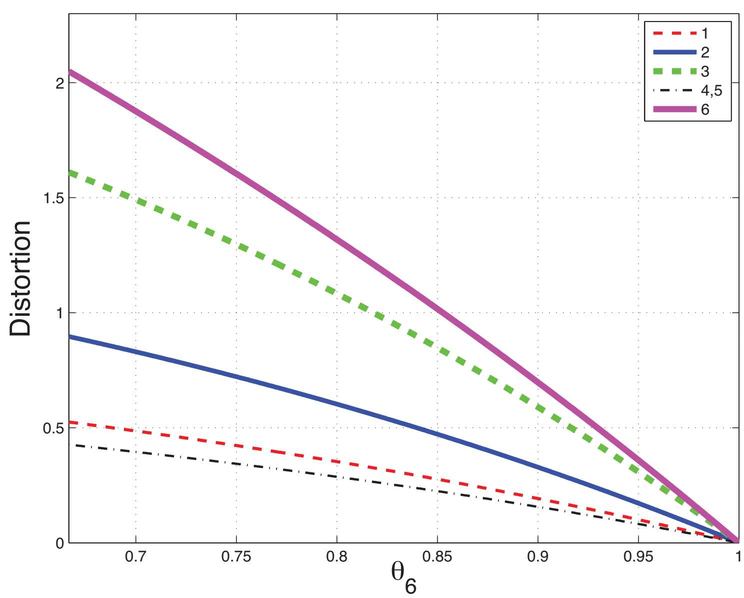

Figure 4-7: Distortion propagates throughout the network.

Example B-3 (different orders in bilateral and multilateral contracts) This example compares bilateral and multilateral contracts in a network with interconnections as in Example B-1 (above). We consider a market including a firm and three consumer agents. Fig. 4-5 visualizes the interconnection among the agents.

Let us first start with the multilateral contract.

Multilateral contract: The payoff function of each agent is as follows:

$$\begin{aligned} u_1(x_1, x_2, x_3, t_1) &= ax_1 - \frac{1}{2}x_1^2 - t_1, \\ u_2(\theta_2, x_1, x_2, x_3, t_2) &= ax_2 - \frac{1}{2}x_2^2 + \theta_2\ x_1\ x_2 - t_2, \\ u_3(x_1, x_2, x_3, t_3) &= ax_3 - \frac{1}{2}x_3^2 + \bar{\theta}\ x_2\ x_3 - t_3. \end{aligned}$$

In the above network, agents 2 and 3 obtain positive externality from the consumptions of their neighbors. The externality from agent 2 to agent 3 is weighed by $\bar{\theta}$ that is publicly known. However, the externality from agent 1 to agent 2 is weighed by θ_2 that is agent 2's *private* information. It is, however, commonly known that $\theta_2 \in \{\underline{\theta}, \bar{\theta}\}$, $0 < \underline{\theta} < \bar{\theta}$, with $\text{Prob}\{\theta_2 = \underline{\theta}\} = \mathsf{v} > \frac{\Delta\theta}{\bar{\theta}}$, where $\Delta\theta \triangleq \bar{\theta} - \underline{\theta}$. Firm's objective is to devise a menu of optimal incentive quantity-price pairs $\{(x_i(\theta_2), t_i(\theta_2))\}$ for each agent i, given agent 2's report $\theta_2 \in \{\underline{\theta}, \bar{\theta}\}$, so as to maximize her total ex-ante profit subject to the corresponding PC of each agent $i \in \{1, 2, 3\}$, and IC constraint of agent 2. Therefore, the firm's problem is written as:

$$\begin{aligned} \max_{\{x_i(\cdot), t_i(\cdot)\}_{i=1}^3} \quad & \mathsf{v}\sum_{i=1}^{3} t_i(\underline{\theta}) + (1-\mathsf{v})\sum_{i=1}^{3} t_i(\bar{\theta}) \\ \text{subject to} \quad & \text{PC}_i(\theta_2), \quad \theta_2 \in \{\underline{\theta}, \bar{\theta}\}, \text{ for all } i \in 1, 2, 3, \\ & \text{IC}_2(\theta_2), \quad \theta_2 \in \{\underline{\theta}, \bar{\theta}\}. \end{aligned}$$

Solving the above program simply implies:

$$\begin{aligned} x_1^{SB}(\bar{\theta}) &= x_3^{SB}(\bar{\theta}), \\ x_1^{SB}(\underline{\theta}) &< x_3^{SB}(\underline{\theta}). \end{aligned}$$

The above implies that (due to agent 2's information rent $R_2(\underline{\theta}) = \Delta\theta x_1(\underline{\theta})x_2(\underline{\theta})$) if agent 2 reports her efficient type, then agents 1 and 3's allocations are *exactly the same*. But agent 1's allocation is distorted (downward) more than 3's when agent 2 reports her inefficient type.

Next, we consider the bilateral contract in the same network

Bilateral contract: We show using bilateral contracts taht the above intuition does not carry over. Here, *independent* of agent 2's report, agent 1's allocation is *always* lower than agent

3's allocation. Focusing on the bilateral contract, each agent i's payoff is written as follows:

$$u_1(x_1, t_1) = ax_1 - \frac{1}{2}x_1^2 - t_1,$$
$$u_2(\theta_2, x_2, t_2) = ax_2 - \frac{1}{2}x_2^2 + \theta_2 \; x_1 \; x_2 - t_2, \quad \theta_2 \in \{\underline{\theta}, \bar{\theta}\}$$
$$u_3(x_3, t_3) = ax_3 - \frac{1}{2}x_3^2 + \bar{\theta} \; x_3 \; \mathbf{E}[x_2(\theta_2)] - t_3,$$

and, thus, the firm's objective is to characterize

$$\{(x_1, t_1), \{(x_2(\underline{\theta}), t_2(\underline{\theta}))), (x_2(\bar{\theta}), t_2(\bar{\theta}))\}, (x_3, t_3)\},$$

maximizing her ex-ante profit with respect to the IC and PC constraints, Eqs. (4.1) and (4.2), respectively. Finally, we have the following.

Lemma 4.7.2. *Using bilateral contracts,* $x_3^{SB} - x_1^{SB} = \Delta\theta \; x_2^{SB}(\underline{\theta}) > 0$.

Therefore, in bilateral contracts independent of agent 2's report, $x_3^{SB} > x_1^{SB}$. This is because agent 2's information rent is $R_2(\underline{\theta}) = \Delta\theta \; x_1 \; x_2(\underline{\theta})$, distorting (downward) more agent 1's allocation.

In summary, bilateral and multilateral contracts may induce different orders on the agents' allocations under presence of network externalities. By similar argument, one can show this effect can be even dramatic (for details, see Fig. 4-8). In the next section by characterizing the optimal bilateral trade profiles we present a more precise comparison for general networks.

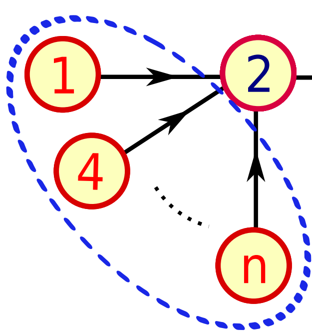

Figure 4-8: Star network. Using multilateral contracts, if agent 2 reports her efficient type, then $x_3(\bar{\theta}) = x_i(\bar{\theta})$, for all $i \in \{1, 4, 5, \cdots, n\}$. However, using bilateral contracts, if agent 2 reports her efficient type, then $x_3(\bar{\theta}) > x_i(\bar{\theta})$, for all $i \in \{1, 4, 5, \cdots, n\}$.

[**Analysis of Example B-1**] **Full information (First-best)**: Let θ_2 be commonly known. Hence, in the firm's problem only PC constraints are active, and, thus, the first best trade profile is characterized as follows.

Lemma 4.7.3 (First best)**.** *The first-best trade profile is given by:*

$$x_1^{FB}(\bar\theta) = a + \bar\theta\ x_2^{FB}(\bar\theta), \qquad x_1^{FB}(\underline\theta) = a + \underline\theta\ x_2^{FB}(\underline\theta),$$
$$x_3^{FB}(\bar\theta) = a + \gamma\ x_2^{FB}(\bar\theta), \qquad x_3^{FB}(\underline\theta) = a + \gamma\ x_2^{FB}(\underline\theta),$$
$$x_2^{FB}(\bar\theta) = \frac{a(1+(\bar\theta+\gamma))}{1-(\bar\theta^2+\gamma^2)}, \qquad x_2^{FB}(\underline\theta) = \frac{a(1+(\underline\theta+\gamma))}{1-(\underline\theta^2+\gamma^2)}.$$

We next characterize the second-best solution.

Sketch of the proof of Lemma 4.7.3*.* In the full information case, since the payoff functions of all agents are commonly known, the $PC_1(\underline\theta), PC_2(\underline\theta), PC_3(\underline\theta), PC_1(\bar\theta), PC_2(\bar\theta)$, and $PC_3(\bar\theta)$ must bind, characterizing $t_1(\underline\theta)$, $t_2(\underline\theta), t_3(\underline\theta), t_1(\bar\theta), t_2(\bar\theta)$, and $t_3(\bar\theta)$, respectively. Then, plugging the characterized payments into the objective function and taking the first optimality condition yields the result.

∎

Sketch of the proof of Lemma 4.7.1*.* Since the payoff functions of agents 1 and 3 are both fully known, thus the corresponding $PC_1(\underline\theta), PC_3(\underline\theta), PC_1(\bar\theta)$ and $PC_3(\bar\theta)$ must bind, characterizing $t_1(\underline\theta)$, $t_3(\underline\theta)$, $t_1(\bar\theta)$, and $t_3(\bar\theta)$, respectively. In addition $PC_2(\underline\theta)$ must bind, characterizing $t_2(\underline\theta)$. In addition, the (downward incentive constraint) $IC_2(\bar\theta)$ must bind, characterizing $t_2(\bar\theta)$, plus the corresponding information rent that is equal to $R = \Delta\theta x_1(\underline\theta)x_2(\underline\theta)$. The result then follows from the first order optimality condition of the objective function after plugging the payments, characterized in the above, into it. ∎

Proof of Lemma 4.7.2*.* [**Analysis of Example B-3**] For ease of exposition, let $\ell \triangleq x_2^{SB}(\underline\theta)$ and $\hbar \triangleq x_2^{SB}(\bar\theta)$, and ignore the superscript SB for the other agents. All the PC constraints for agents 1, 3, and 2 (at her inefficient type) as well as the IC of agent 2 at her efficient type must bind. Thus, the payment are as follows:

$$\begin{aligned}
t_1 &= ax_1 - \frac{1}{2}x_1^2,\\
t_3 &= ax_3 - \frac{1}{2}x_3^2 + \gamma\ x_3\ \mathbf{E}[x_2(\theta_2)],\\
t_2(\underline\theta) &= a\ell - \frac{1}{2}\ell^2 + \underline\theta\ \ell\ x_1,\\
t_2(\bar\theta) &= a\hbar - \frac{1}{2}\hbar^2 + \bar\theta\ \hbar\ x_1 - \underbrace{\Delta\theta\ x_1\ \ell}_{\text{information rent}},
\end{aligned}$$

Firm's objective becomes the following program:

$$\max_{x_1,x_2,\ell,\hbar\in\mathbb{R}_+} \quad t_1 + t_3 + \mathsf{v}t_2(\underline\theta) + (1-\mathsf{v})t_2(\bar\theta)$$

By first order optimality condition, we have:

$$\begin{aligned}
x_1 &= a + \mathsf{v}\underline{\theta}\ell + (1-\mathsf{v})(\gamma\hbar - \Delta\theta\ \ell),\\
x_3 &= a + \gamma\ \mathbf{E}[x_2(\theta_2)],\\
\hbar &= a + \gamma(x_1 + x_3),\\
\ell &= a + \left(\underline{\theta} - \frac{1-\mathsf{v}}{\mathsf{v}}\Delta\theta\right) x_1 + \gamma x_3 > 0.
\end{aligned}$$

Therefore, $x_3 - x_1 = \ell\Delta\theta > 0$. ∎

4.8 Role of influential agents in core-periphery structures: Infinite vs. Vanishing profit gap

In this section we formally prove the statements of Example 4 in Section 4.5.2.

In the next proposition, using Lemma 4.5.1, we show that among all networks within the class of star-family the maximum loss in the firm's profit occurs in the star-inward network.

Proposition 4.8.1. *Among all $G \in \mathfrak{G}(n)$ (see Definition (4.5.2)),* star-inward *maximizes $\Pi_G^{multi.} - \Pi_G^{bi.}$.*

Proof. See the Appendix—proof section. ∎

As the figure in Example 4 shows, the maximum loss (within the star-family) occurs in star-inwards. Moreover, the figure highlights the fact that with increasing number of agents, the firm's loss for using bilateral contracts grows unboundedly. This, of course, is not always true and depends on the network constructions. The following propositions highlight the points.

Proposition 4.8.2. *Consider the sequence of star-inward networks $\{G(n)\}_{n=2}^{\infty}$, where $\alpha_n \triangleq g_{12} = g_{13} = \cdots = g_{1n}$, and the rest of entries are zero. If $\alpha_n = \frac{1}{\sqrt{(n-1)(\mu^2+\sigma^2+\epsilon)}}$, where $\epsilon > 0$ is constant, then as n grows:*

$$\Pi_{G(n)}^{multi.} - \Pi_{G(n)}^{bi.} \ \longrightarrow\ \infty.$$

Proof. See the Appendix—proof section. ∎

The above proposition shows that the marginal benefit of using multilateral contracts may grow unboundedly in networks that have major asymmetry in their in-degrees. When there is an agent dominant in its in-degree (e.g., star-inward) then the loss in the firm's profit because of using the simpler bilateral contracts may become unbounded. In summery, in large networks exhibiting large asymmetry in their in-degrees, firm's restriction to the simpler bilateral contracts may result in major loss in firm's profit.

When will the profit gap vanish? In contrast, the following proposition asserts that in large directed core-periphery economies the difference (in profit) between multilateral and bilateral contracts might be negligible. This is intuitive because, as shown in Proposition 4.5.2 (part *(ii)*), decreasing the influencing weights (i.e., $g_{ij}\ \forall i, j$) reduces the benefit of using multilateral contracts. To prove the result we use Proposition 4.8.1, that is, in the class of star-like core-periphery structures[46] the benefit of using multilateral contacts in star-inward networks is higher than the others.

Proposition 4.8.3. *Let $\mathcal{G}(n)$ denote the set of all economies including n agents with the directed star network*[47] *structure. For any $G(n) \in \mathcal{G}(n)$, if its non-zero weights are equal to $\beta_n \triangleq \frac{1}{(n-1)\sqrt{\mu^2+\sigma^2+\epsilon}}$, where $\epsilon > 0$ is constant, then as n grows:*

$$\Pi^{multi.}_{G(n)} - \Pi^{bi.}_{G(n)} \longrightarrow 0.$$

Proof. See the Appendix—proof section. ∎

4.9 Appendix: Omitted Proofs

Proof of Proposition 4.3.1. We prove each part of the proposition separately:
Proof of Part (i): To find $\{x_i^{FB}(\cdot)\}_{i=1}^n$, we maximize the objective function in (3.3) point-wise for any type profile θ. Thus, for a given θ, to find $(x_1^{FB}(\theta), x_2^{FB}(\theta), \cdots, x_n^{FB}(\theta))$ we faced with the follwing program:

$$\max_x \sum_{i=1}^n \left[a x_i - \frac{b}{2} x_i^2 + x_i \theta_i \sum_{j=1}^n g_{ij} x_j \right]$$
$$\text{subject to:} \quad x_i \geq 0 \quad \forall i \in V. \tag{9.1}$$

The Hessian matrix corresponding to the objective function in (9.1) is given by:

$$\begin{pmatrix} -b & \theta_1 g_{12} + \theta_2 g_{21} & \cdots & \theta_1 g_{1n} + \theta_n g_{n1} \\ \theta_1 g_{12} + \theta_2 g_{21} & -b & \cdots & \theta_2 g_{2n} + \theta_n g_{n2} \\ \vdots & \vdots & \ddots & \vdots \\ \theta_1 g_{1n} + \theta_n g_{n1} & \theta_2 g_{2n} + \theta_n g_{n2} & \cdots & -b \end{pmatrix}.$$

The Hessian matrix is a Hermitian, strictly diagonally dominant matrix[48] (due to Assumption 4.3.1) with real negative diagonal entries, thus it is negative semi-definte (due to

[46] In which $g_{ij} \in \{0, \kappa\}$ for all i, j, $(\kappa > 0)$.
[47] With no parallel links.
[48] A matrix is called strictly diagonally dominant if for every row of the matrix, the magnitude of the diagonal entry in a row is larger than the sum of the magnitudes of all the other (non-diagonal) entries in that row. That is, the matrix $A = [a_{ij}]_{(i,j)\in V^2}$ is diagonally dominant if $|a_{ii}| > \sum_{j\neq i}(|a_{ij}| + |a_{ji}|)$, for all $i \in V$.

Sylvester's criterion, Horn and Johnson (2012)). As a result, the objective function in (9.1) is concave.

Next, assuming the solution to program (9.1) is interior implies that it satisfies the following first order optimality condition:

$$a - bx_i^{FB}(\theta) + \left(\theta_i \sum_{j=1}^{n} g_{ij} x_j^{FB}(\theta) + \sum_{j=1}^{n} g_{ji}\theta_j x_j^{FB}(\theta)\right) = 0, \qquad \forall\; i \in V. \tag{9.2}$$

Recall that $M_\theta = \text{diag}(\theta_1, \theta_2, \cdots, \theta_n)$. Thus, the above equation in its matrix form is written as:

$$a\mathbf{1} + (M_\theta G + G^T M_\theta)\mathbf{x}^{FB}(\theta) = b\mathbf{x}^{FB}(\theta),$$

Hence, we obtain:

$$\mathbf{x}^{FB}(\theta) = a\;[bI - (M_\theta G + G^T M_\theta)]^{-1}\mathbf{1}.$$

Note that since $(bI - (M_\theta G + G^T M_\theta))$ is a strictly diagonally dominant matrix, it is invertible (due to Levy-Desplanques theorem, Horn and Johnson (2012)). Furthermore, notice that $x^{FB}(\theta)$ is in the positive orthant, since $a > 0$ and Assumption 4.3.1. Finally, we show that there is no corner solution, and, consequently, the above solution is unique. Suppose it were not true; that is, there exists a non-interior solution that we denote it by y^*. Let $W \subset V$, $(W \neq \emptyset)$, such that $x_i^* = 0$ if and only if $i \in V/W$ (that is, $x_i^* > 0$ when $i \in W$). Since $a > 0$, thus $\mathbf{0}$ can not be an optimal solution, therefore, $W \neq \emptyset$. Let $y_i^* = 0$, due to the optimality condition, by (9.2) we should have:

$$a - by_i + \left(\theta_i \sum_{j=1}^{n} g_{ij} y_j + \sum_{j=1}^{n} g_{ji}\theta_j y_j\right)\Bigg|_{y=y^*} = a + \left(\theta_i \sum_{j=1}^{n} g_{ij} y_j^* + \sum_{j=1}^{n} g_{ji}\theta_j y_j^*\right) \leq 0.$$

However, since $a > 0$, achieving the last inequality is impossible, which is a contradiction.

Edskip

Proof of Part (ii): Define, for every i and $\theta_i \in [\underline{\theta}\;\bar{\theta}]$,

$$V_i(\theta_i) \triangleq \mathbf{E}_{\theta_{-i}}\left[a x_i(\theta_i, \theta_{-i}) - \frac{b}{2} x_i^2(\theta_i, \theta_{-i})\right] \tag{9.3}$$

$$\gamma_i(\theta_i) \triangleq \mathbf{E}_{\theta_{-i}}\left[x_i(\theta_i, \theta_{-i}) \sum_{j\neq i} g_{ij} x_j(\theta_i, \theta_{-i})\right] \tag{9.4}$$

$$T_i(\theta_i) \triangleq \mathbf{E}_{\theta_{-i}}\left[t_i(\theta_i, \theta_{-i})\right]. \tag{9.5}$$

Thus, the agent i's interim utility for reporting $\hat{\theta}_i$, while her real type is θ_i, is given by:

$$U_i(\theta_i, \hat{\theta}_i) = \theta_i \gamma_i(\hat{\theta}_i) + V_i(\hat{\theta}_i) - T_i(\hat{\theta}_i).$$

By appealing to the revenue equivalence theorem, a direct quantity-price schedule $\{x_i^{SB}(\cdot), t_i^{SB}(\cdot)\}_{i=1}^n$ is IC and PC if and only if for every $i \in V$:

(i) $\gamma_i(\theta_i)$ is increasing;

(ii) For every $\theta_i \in [\underline{\theta}\ \bar{\theta}]$:

$$T_i(\theta_i) = \theta_i \gamma_i(\theta_i) + V_i(\theta_i) - \int_{\underline{\theta}}^{\theta_i} \gamma_i(\tau) d\tau. \tag{9.6}$$

Given the specification of the incentive compatible and individual rational mechanisms the firm's problem is re-written as:

$$\begin{aligned} \max_{\{x_i(\cdot), t_i(\cdot)\}_{i=1}^n} \quad & \sum_{i=1}^n \mathbf{E}_{\theta_i} \Big[T_i(\theta_i) \Big] \\ \text{subject to} \quad & T_i(\theta_i) = \theta_i \gamma_i(\theta_i) + V_i(\theta_i) - \int_{\underline{\theta}}^{\theta_i} \gamma_i(\tau) d\tau \qquad \forall i,\ \theta_i, & (9.7) \\ & \gamma_i(\theta_i) \text{ is increasing} \quad \forall\ i & (9.8) \end{aligned}$$

To find the solution of the above program we first ignore the monotonicity constraints (9.8), later we show (9.8) is indeed satisfied. Thus, plugging (9.7) into the objective function, and using the fact that $\mathbf{E}_{\theta_i} \left[\int_{\underline{\theta}}^{\theta_i} \gamma_i(\tau) d\tau \right] = \mathbf{E}_{\theta_i} \left[\frac{1-F(\theta_i)}{f(\theta_i)} \gamma_i(\theta_i) \right]$ we obtain

$$\max_{\{x_i(\cdot)\}_{i=1}^n} \quad \sum_{i=1}^n \mathbf{E}_{\theta_i} \left[\psi(\theta_i) \gamma_i(\theta_i) + V_i(\theta_i) \right], \tag{9.9}$$

recall that $\psi(\theta_i) = \theta_i - \frac{1-F(\theta_i)}{f(\theta_i)}$ (virtual type). Given the definitions in (9.3)-(9.5), and the fact that $\mathbf{E}[\cdot]$ is a linear operator, (9.9) can be re-written as:

$$\begin{aligned} & \max_{\{x_i(\cdot)\}_{i=1}^n} \quad \mathbf{E}_{\theta} \sum_{i=1}^n \left[\psi(\theta_i) x_i(\theta_i, \theta_{-i}) \sum_{j \neq i} g_{ij} x_j(\theta_i, \theta_{-i}) + a x_i(\theta_i, \theta_{-i}) - \frac{b}{2} x_i^2(\theta_i, \theta_{-i}) \right] \\ \equiv & \max_{\{x_i(\cdot)\}_{i=1}^n} \quad \mathbf{E}_{\theta} \sum_{i=1}^n \left[\psi(\theta_i) x_i(\theta) \sum_{j \neq i} g_{ij} x_j(\theta) + a x_i(\theta) - \frac{b}{2} x_i^2(\theta) \right]. \end{aligned} \tag{9.10}$$

To find the optimal solution to (9.10), we maximize it point-wise. Let $\theta \in [\underline{\theta}\ \bar{\theta}]^n$ be fixed and given, hence $\{x_i^{SB}(\theta)\}_{i=1}^n$ solves the following program:

$$\max_{\{x_i(\theta)\}_{i=1}^n} \sum_{i=1}^n \left[\psi(\theta_i) x_i(\theta) \sum_{j \neq i} g_{ij} x_j(\theta) + a x_i(\theta) - \frac{b}{2} x_i^2(\theta) \right]. \tag{9.11}$$

Note that the objective function in (9.10) is concave. Because its corresponding Hessian matrix,

$$\begin{pmatrix} -b & \psi(\theta_1)g_{12}+\psi(\theta_2)g_{21} & \cdots & \psi(\theta_1)g_{1n}+\psi(\theta_n)g_{n1} \\ \psi(\theta_1)g_{12}+\psi(\theta_2)g_{21} & -b & \cdots & \psi(\theta_2)g_{2n}+\psi(\theta_n)g_{n2} \\ \vdots & \vdots & \ddots & \vdots \\ \psi(\theta_1)g_{1n}+\psi(\theta_n)g_{n1} & \psi(\theta_2)g_{2n}+\psi(\theta_n)g_{n2} & \cdots & -b \end{pmatrix},$$

is Hermitian and strictly diagonally dominant and thus it is negative semi-definte (due to Sylvester's criterion, Horn and Johnson (2012)). The diagonally dominant property of the above matrix is ensured by Assumption 4.3.1, 4.3.2, and the fact that $\psi(\theta_i) \leq \theta_i$ for all $\theta_i \in [\underline{\theta}\ \bar{\theta}]$ and $i \in V$.

The first order optimality condition of (9.11) (assuming the solution is interior) yields

$$a - bx_i^{SB}(\theta) + \left(\psi(\theta_i)\sum_{j=1}^{n} g_{ij}x_j^{SB}(\theta) + \sum_{j=1}^{n} g_{ji}\psi(\theta_j)x_j^{SB}(\theta)\right) = 0, \qquad \forall\, i \in V. \tag{9.12}$$

Recall that $M_\psi = \text{diag}(\psi(\theta_1), \psi(\theta_2), \cdots, \psi(\theta_n))$. Thus, the above equation in its matrix form is written as:

$$a\mathbf{1} + (M_\psi G + G^T M_\psi)\mathbf{x}^{SB}(\theta) = b\mathbf{x}^{SB}(\theta).$$

Since $(bI - (M_\psi G + G^T M_\psi))$ is a strictly diagonally dominant matrix, it is invertible (due to Levy-Desplanques theorem, Horn and Johnson (2012)). Thus,

$$\mathbf{x}^{SB}(\theta) = a\ [bI - (M_\psi G + G^T M_\psi)]^{-1}\mathbf{1}. \tag{9.13}$$

Furthermore, notice that $x^{SB}(\theta)$ is in the positive orthant, since $a > 0$ and Assumption 4.3.1 and 4.3.2.

Notice that, since $\psi(\theta_i) \geq 0$ for all $\theta_i \in [\underline{\theta}\ \bar{\theta}]$ and $i \in V$, similar to the proof of the first-best case, one can easily show $\mathbf{x}^{SB}(\theta)$, characterized in (9.13), is indeed the unique to the program (9.11), for any given θ.

Next, we present a lemma that we make a use of it in the proof later.

Lemma 4.9.1. *Let $K \triangleq [bI - (M_\psi G + G^T M_\psi)]^{-1}$. Then, $\frac{\partial K}{\partial \theta_i}$ is a matrix with non-negative entries, for any $i \in V$.*

Proof. Observe that, by chain rule,

$$\mathbf{0} = \frac{\partial I}{\partial \theta_i} = \frac{\partial K K^{-1}}{\partial \theta_i} = \frac{\partial K}{\partial \theta_i}\ K^{-1} + K\ \frac{\partial K^{-1}}{\partial \theta_i}. \tag{9.14}$$

Furthermore, as G does not depend on θ_i, we have

$$\frac{\partial K^{-1}}{\partial \theta_i} = \frac{\partial}{\partial \theta_i}\Big[bI - (M_\psi G + G^T M_\psi)\Big] = -\frac{\partial}{\partial \theta_i}\Big[(M_\psi G + G^T M_\psi)\Big] = -\frac{\partial M_\psi}{\partial \theta_i}G - G^T\frac{\partial M_\psi}{\partial \theta_i}$$
$$= -(\mathsf{E}_{ii}G + G^T\mathsf{E}_{ii}),$$

where $\mathsf{E}_{ii} \triangleq \frac{\partial M_\psi}{\partial \theta_i}$ is a matrix with $\frac{\partial \psi(\theta_i)}{\partial \theta_i}$ at the ii th entry, and zero otherwise. Notice that $\frac{\partial \psi(\theta_i)}{\partial \theta_i} \geq 0$, (due to the monotonicty assumption in the hazard rate). Thus, using (9.38) we obtain

$$\frac{\partial K}{\partial \theta_i} = -K\ \frac{\partial K^{-1}}{\partial \theta_i}\ K = K(\mathsf{E}_{ii}G + G^T\mathsf{E}_{ii})K.$$

Thus, since the right hand side in the above equation is non-negative, $\frac{\partial K}{\partial \theta_i}$ is a matrix with non-negative entries. ■

This Lemma has two important implications: (i) $\frac{\partial x_i^{SB}(\theta)}{\partial \theta_i} \geq 0$, this property is intuitively immediate due to the IC constraint. (ii) $\frac{\partial x_j^{SB}(\theta)}{\partial \theta_i} \geq 0$, $(j \neq i)$, this property is due to the strategic complement property.

Finally, to wrap up the proof, it is left to show that $\gamma_i(\theta_i)$ is increasing in θ_i, for all $\theta_i \in [\underline{\theta}\ \bar{\theta}]$, and $i \in V$ (recall that monotonicity of $\gamma_i(\theta_i)$ is the constraint must be satisfied to achieve IC).

By the definition of $\gamma_i(\theta_i)$ (see (9.4)) and the above two points ((i) and (ii)), we have:

$$\frac{\partial \gamma_i(\theta_i)}{\partial \theta_i} = \mathbf{E}_{\theta_{-i}}\left[\frac{\partial\{x_i(\theta_i,\theta_{-i})\sum_{j\neq i} g_{ij}x_j(\theta_i,\theta_{-i})\}}{\partial \theta_i}\right]$$
$$= \mathbf{E}_{\theta_{-i}}\left[\underbrace{\frac{\partial x_i(\theta_i,\theta_{-i})}{\partial \theta_i}}_{\geq 0 \text{ by (i)}}\sum_{j\neq i} g_{ij}x_j(\theta_i,\theta_{-i}) + x_i(\theta_i,\theta_{-i})\sum_{j\neq i} g_{ij}\underbrace{\frac{\partial x_j(\theta_i,\theta_{-i})}{\partial \theta_i}}_{\geq 0 \text{ by (ii)}}\right]$$
$$\geq 0.$$

The proof is complete. ■

Proof of Lemma 4.3.1. Proof of Part (i) follows directly from Proposition 4.3.1 and the definition of Bonacich centrality measure (see Def. 4.3.1). We proceed to prove the rest.

Given the definition of distortion vector and make a use of Proposition 4.3.1, we obtain:

$$\mathbf{d}(\theta) = a\left([bI - (M_\theta G + G^T M_\theta)]^{-1} - [bI - (M_\psi G + G^T M_\psi)]^{-1}\right)\mathbf{1}$$
$$= \frac{a}{b}\left([I - \frac{1}{b}(M_\theta G + G^T M_\theta)]^{-1} - [I - \frac{1}{b}(M_\psi G + G^T M_\psi)]^{-1}\right)\mathbf{1}.$$

Next, define $M \triangleq [I - \frac{1}{b}(M_\theta G + G^T M_\theta)]^{-1}$ and $N \triangleq [I - \frac{1}{b}(M_\psi G + G^T M_\psi)]^{-1}$. Thus, their entries are characterized as:

$$m_{ij} = \sum_{k=0}^{\infty} \frac{1}{b^k} f_{ij}^{[k]}$$
$$n_{ij} = \sum_{k=0}^{\infty} \frac{1}{b^k} h_{ij}^{[k]}$$

where $f_{ij}^{[k]}$ is the ij entry of $(M_\theta G + G^T M_\theta)^k$, and $h_{ij}^{[k]}$ is the ij entry of $(M_\psi G + G^T M_\psi)^k$. Recall that $\psi(\tau) = \tau - \frac{1-F(\tau)}{f(\tau)}$, $\tau \in [\underline{\theta}\ \bar{\theta}]$. Thus, $\psi(\tau) \leq \tau$ with equality only at $\tau = \bar{\tau}$. This implies that $f_{ij}^{[k]} \geq h_{ij}^{[k]}$, for any k. Consequently, $m_{ij} \geq n_{ij}$, for all i, j. Since $\theta_j < \bar{\theta}$, thus $\psi(\theta_j) < \theta_j$. Consequently, the result follows. ■

Proof of Proposition 4.3.2. We prove the parts separately as follows. **E**dskip

Part (i): For ease of exposition we consider agent 1, i.e., $i = 1$. Thus, $\theta_1 \leq \bar{\theta}$, while all the other agents are efficient, i.e., $\theta_j = \bar{\theta}$ (for all $j \neq 1$). Let $(\theta_1, \bar{\theta}_{-1})$ denote the corresponding type profile.

Since θ_1 is sufficiently close to $\bar{\theta}$, thus:

$$\mathcal{T}_1(\theta_1, \bar{\theta}_{-1}) = \mathcal{T}_1(\bar{\theta_1}, \bar{\theta}_{-1}) - (\bar{\theta} - \theta_1)\mathcal{T}_1'(\bar{\theta_1}, \bar{\theta}_{-1}), \tag{9.15}$$

where $\mathcal{T}_1'(\bar{\theta_1}, \bar{\theta}_{-1}) = \frac{\partial \mathcal{T}_1'(\theta_1, \bar{\theta}_{-1})}{\partial \theta_1}|_{\theta_1 = \bar{\theta_1}}$.

Notice that, by Definition 4.3.3, $\mathcal{T}_1(\bar{\theta_1}, \bar{\theta}_{-1}) = \mathbf{1}^T \mathbf{d}(\bar{\theta_1}, \bar{\theta}_{-1})$. According to Proposition (4.3.1), $\mathbf{d}(\bar{\theta_1}, \bar{\theta}_{-1}) = 0$, since $M_\theta = M_\psi = \text{diag}(\bar{\theta}, \bar{\theta}, \cdots, \bar{\theta})$, and thus $\mathcal{T}_1(\bar{\theta_1}, \bar{\theta}_{-1}) = 0$.

Therefore, we only need to characterize $(\bar{\theta} - \theta_1)\mathcal{T}_1'(\bar{\theta_1}, \bar{\theta}_{-1})$. To do so, for ease of exposition, define $S^{-1} \triangleq (I - \frac{1}{b}(M_\theta G + G^T M_\theta))^{-1}$ and $T^{-1} \triangleq (I - \frac{1}{b}(M_\psi G + G^T M_\psi))^{-1}$. Therefore, due to Proposition 4.3.1, $\mathbf{d}(\theta) = \frac{a}{b}(S^{-1} - T^{-1})\mathbf{1}$.

Next, since $SS^{-1} = I$, we have, $\frac{\partial S}{\partial \theta_1} S^{-1} + S \frac{\partial S^{-1}}{\partial \theta_1} = 0$, that yields:

$$\frac{\partial S^{-1}}{\partial \theta_1} = -S^{-1} \frac{\partial S}{\partial \theta_1} S^{-1}.$$

By the definition,

$$\begin{aligned} \frac{\partial S}{\partial \theta_1}|_{\theta_1 = \bar{\theta}} &= -\frac{1}{b}[\text{diag}(1, 0, \cdots, 0)G + G^T \text{diag}(1, 0, \cdots, 0)] \\ &= -\frac{1}{b}\left(R_1 + R_1^T\right), \end{aligned}$$

where,

$$R_1 \triangleq \begin{bmatrix} g_{11} & g_{12} & \cdots & g_{1n} \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}. \tag{9.16}$$

Similarly,

$$\begin{aligned}\frac{\partial T}{\partial \theta_1}|_{\theta_1=\bar{\theta}} &= -\frac{1}{b}[\text{diag}(\psi'(\bar{\theta}), 0, \cdots, 0)G + G^T \text{diag}(\psi'(\bar{\theta}), 0, \cdots, 0)] \\ &= -\frac{1}{b}\psi'(\bar{\theta})\left(R_1 + R_1^T\right),\end{aligned}$$

where $\psi'(\bar{\theta}) = 1 - \phi'(\bar{\theta})$. Moreover, since $\psi(\bar{\theta}) = \bar{\theta}$, hence $S^{-1}|_{\theta=\bar{\theta}} = T^{-1}|_{\theta=\bar{\theta}}$. Therefore,

$$\begin{aligned}\frac{\partial \mathbf{d}(\theta_i, \bar{\theta}_{-1})}{\partial \theta_1}|_{\theta_1=\bar{\theta}} &= \frac{a}{b}\left[-S^{-1}\frac{\partial S}{\partial \theta_1}S^{-1} + T^{-1}\frac{\partial T}{\partial \theta_1}T^{-1}\right]\mathbf{1} = \frac{a}{b}\left[-S^{-1}\left(\frac{\partial S}{\partial \theta_1} - \frac{\partial T}{\partial \theta_1}\right)S^{-1}\right]\ \mathbf{1} \\ &= \frac{a}{b^2}\left[-S^{-1}\Big((-1+\psi'(\bar{\theta}))(R_1 + R_1^T)\Big)S^{-1}\right]\ \mathbf{1} \\ &= \frac{a}{b^2}(1-\psi'(\bar{\theta}))\left[S^{-1}\Big(R_1 + R_1^T\Big)S^{-1}\right]\ \mathbf{1} \\ &= \frac{a}{b^2}\phi'(\bar{\theta})\left[S^{-1}\Big(R_1 + R_1^T\Big)S^{-1}\right]\ \mathbf{1}\end{aligned}$$

The above equality along with (9.15) yields

$$\begin{aligned}\mathcal{T}_1(\theta_1, \bar{\theta}_{-1}) &= -(\bar{\theta} - \theta_1)\mathcal{T}_1'(\bar{\theta}_1, \bar{\theta}_{-1}) \\ &= (\theta_1 - \bar{\theta})\ \mathbf{1}^T\ \frac{\partial \mathbf{d}(\theta_i, \bar{\theta}_{-1})}{\partial \theta_1}|_{\theta_1=\bar{\theta}} \\ &= \frac{a}{b^2}\Big((\theta_1 - \bar{\theta})\phi'(\bar{\theta})\Big)\mathbf{1}^T\left[S^{-1}\Big(R_1 + R_1^T\Big)S^{-1}\right]|_{\theta_1=\bar{\theta}}\ \mathbf{1} \\ &\overset{(a)}{=} \frac{a}{b^2}\Big((\theta_1 - \bar{\theta})\phi'(\bar{\theta})\Big)[k_1\ k_2\ \cdots k_n]\Big(R_1 + R_1^T\Big)[k_1\ k_2\ \cdots k_n]^T \\ &\overset{(b)}{=} \frac{a}{b^2}\Big((\theta_1 - \bar{\theta})\phi'(\bar{\theta})\Big)\left[2\sum_j k_1 k_j g_{1j}\right],\end{aligned}$$

where (a) follows since $S^{-1}|_{\theta=\bar{\theta}}\mathbf{1} = T^{-1}|_{\theta=\bar{\theta}}\mathbf{1} = \left(I - \frac{1}{b}\bar{\theta}(G+G^T)\right)^{-1}\mathbf{1} \triangleq [k_1\ k_2\ \cdots\ k_n]^T$, where k_i is agent i's Bonacich centrality in $G + G^T$. And (b) follows since[49]

$$[k_1\ k_2\ \cdots k_n]\Big(R_1 + R_1^T\Big)[k_1\ k_2\ \cdots k_n]^T = [k_1\ k_2\ \cdots k_n]\left[g_{11}k_1 + \sum_j g_{1j}k_j \quad g_{12}k_1 \quad g_{13}k_1 \quad \cdots \quad g_{1n}k_1\right]^T$$
$$= k_1 \sum_j g_{1j}k_j + \sum_j g_{1j}k_jk_1$$
$$= 2\sum_j g_{1j}k_jk_1.$$

The above computations work for all i. Thus

$$\begin{aligned} \mathcal{T}_i(\theta_i, \bar{\theta}_{-i}) &= \frac{2a}{b^2}(\theta_i - \bar{\theta})\phi'(\bar{\theta})\left[\sum_j k_ik_jg_{ij}\right] \\ &= \frac{2a}{b^2}\ (\bar{\theta} - \theta_i)\ |\phi'(\bar{\theta})|\ \left[\sum_j k_ik_jg_{ij}\right]. \end{aligned} \tag{9.17}$$

Notice that $\phi'(\bar{\theta}) < 0$.

Part (ii): Lets first derive $\mathbb{E}[\Pi^{FB}]$. As shown in Proposition 4.3.1, the optimal (first best) trade profile is $\mathbf{x}^{FB}(\theta) = a\ [bI - (M_\theta G + G^T M_\theta)]^{-1}\mathbf{1}$. Thus, for each type profile θ the (ex post) firm's profit is given by the following. For ease of exposition, let x_i denote agent i's (first best) allocation with respect to the type profile θ then

$$\begin{aligned} \sum_{i=1}^n \left(ax_i - \frac{b}{2}x_i^2 + x_i\theta_i\sum_j g_{ij}x_j\right) &= a\mathbf{1}^T\mathbf{x} - \frac{b}{2}\mathbf{x}^T\mathbf{x} + (M_\theta\mathbf{x})^TG\mathbf{x} \\ &= a\mathbf{1}^T\mathbf{x} - \frac{b}{2}\mathbf{x}^T\mathbf{x} + \frac{1}{2}\mathbf{x}^T(M_\theta G + G^TM_\theta)\mathbf{x} \\ &= a\mathbf{1}^T\mathbf{x} - \frac{1}{2}\Big[\mathbf{x}^T\ \ [bI - (M_\theta G + G^TM_\theta)]\ \ \mathbf{x}\Big] \\ &= \frac{a^2}{2}\mathbf{1}^T[bI - (M_\theta G + G^TM_\theta)]^{-1}\mathbf{1}. \end{aligned}$$

Therefore $\mathbb{E}[\Pi^{FB}] = \frac{a^2}{2}\mathbf{1}^T\mathbb{E}\left[[bI - (M_\theta G + G^TM_\theta)]^{-1}\right]\mathbf{1}$.

Next, to derive $\mathbb{E}[\Pi^{SB}]$ we first note that as shown in Proposition 4.3.1 (see (9.7), (9.9), (9.10)) $\mathbb{E}[\Pi^{SB}] = \mathbb{E}\sum_{i=1}^n T(\theta_i) = \mathbb{E}\sum_{i=1}^n\left[\psi(\theta_i)y_i(\theta)\sum_{j\neq i}g_{ij}y_j(\theta) + ay_i(\theta) - \frac{b}{2}y_i^2(\theta)\right]$, where $y_i(\theta) = x_i^{SB}(\theta)$. Thus, similar to the full information case, let y_i denote agent i's (second

[49] Recall that $g_{11} = 0$.

best) allocation with respect to the type profile θ then

$$\begin{aligned}\sum_{i=1}^{n}\left(ay_i - \frac{b}{2}y_i^2 + y_i\psi(\theta_i)\sum_j g_{ij}y_j\right) &= a\mathbf{1}^T\mathbf{y} - \frac{b}{2}\mathbf{y}^T\mathbf{y} + (M_\psi \mathbf{y})^T G\mathbf{y} \\ &= a\mathbf{1}^T\mathbf{y} - \frac{b}{2}\mathbf{y}^T\mathbf{y} + \frac{1}{2}\mathbf{y}^T(M_\psi G + G^T M_\psi)\mathbf{y} \\ &= a\mathbf{1}^T\mathbf{y} - \frac{1}{2}\left[\mathbf{y}^T \quad [bI - (M_\psi G + G^T M_\psi)] \quad \mathbf{y}\right] \\ &= \frac{a^2}{2}\mathbf{1}^T[bI - (M_\psi G + G^T M_\psi)]^{-1}\mathbf{1}.\end{aligned}$$

Thus,

$$\begin{aligned}\mathbb{E}[\Pi^{FB} - \Pi^{SB}] &= \frac{a^2}{2}\mathbb{E}\mathbf{1}^T\left[[bI - (M_\theta G + G^T M_\theta)]^{-1} - [bI - (M_\psi G + G^T M_\psi)]^{-1}\right]\mathbf{1} \\ &= \frac{a^2}{2}\mathbb{E}\ \mathcal{T}(\theta_i, \theta_{-i}).\end{aligned}$$

Next, under the assumption that it is known to the firm that $\theta_j = \bar{\theta}$, for all $j \neq i$. Then, when $\theta_i \in \{\underline{\theta}, \bar{\theta}\}$ with $\text{Prob}\{\theta_i = \underline{\theta}\} = \mathsf{v} > \frac{\Delta\theta}{\bar{\theta}}$ then $\mathbb{E}[\Pi^{FB} - \Pi^{SB}] = \frac{a^2}{2}\text{Prob}\{\theta_i = \underline{\theta}\}\mathcal{T}_i(\theta_i = \underline{\theta}, \bar{\theta}_{-i}) = \frac{a^3}{b^2}\ \mathsf{v}\ \Delta\theta\ |\phi'(\bar{\theta})|\ \left[\sum_j k_i k_j g_{ij}\right]$, where the last equality follows by (9.17) and the fact that $\Delta\theta$ is sufficiently small. Therefore, it is immediate that $V^* = \arg\max_{i\in\{1,2,\cdots,n\}} \mathbb{E}[\Pi^{FB} - \Pi^{SB}] = \arg\max_{i\in\{1,2,\cdots,n\}} \sum_{j=1}^{n} k_i k_j g_{ij}$, completing the proof.

■

Proof of Proposition 4.4.1. We first derive the second best trade profile. Define, for every i and $\theta_i \in [\underline{\theta}\ \bar{\theta}]$,

$$V_i(\theta_i) \triangleq a x_i(\theta_i) - \frac{b}{2}x_i^2(\theta_i) \tag{9.18}$$

$$\gamma_i(\theta_i) \triangleq x_i(\theta_i)\sum_{j\neq i} g_{ij}\mathbf{E}_{\theta_j}[x_j(\theta_j)] \tag{9.19}$$

Thus, the agent i's interim utility for reporting $\hat{\theta}_i$, while her real type is θ_i, is given by:

$$U_i(\theta_i, \hat{\theta}_i) = \theta_i\gamma_i(\hat{\theta}_i) + V_i(\hat{\theta}_i) - t_i(\hat{\theta}_i).$$

By appealing to the revenue equivalence theorem, a direct quantity-price schedule $\{x_i^{SB}(\cdot), t_i^{SB}(\cdot)\}_{i=1}^{n}$ is IC and PC if and only if for every $i \in V$:

(i) $\gamma_i(\theta_i)$ is increasing;

(ii) For every $\theta_i \in [\underline{\theta}\ \bar{\theta}]$:

$$t_i(\theta_i) = \theta_i \gamma_i(\theta_i) + V_i(\theta_i) - \int_{\underline{\theta}}^{\theta_i} \gamma_i(\tau) d\tau. \tag{9.20}$$

Given the specification of the incentive compatible and individual rational mechanisms the firm's problem is re-written as:

$$\begin{aligned}
\max_{\{x_i(\cdot),t_i(\cdot)\}_{i=1}^n} \quad & \sum_{i=1}^n \mathbf{E}_{\theta_i}\Big[t_i(\theta_i)\Big] \\
\text{subject to} \quad & t_i(\theta_i) = \theta_i\gamma_i(\theta_i) + V_i(\theta_i) - \int_{\underline{\theta}}^{\theta_i} \gamma_i(\tau) d\tau \qquad \forall i,\ \theta_i, & (9.21)\\
& \gamma_i(\theta_i) \text{ is increasing} \quad \forall\ i & (9.22)
\end{aligned}$$

To find the solution of the above program we first ignore the monotonicity constraints (9.22), later we show (9.22) is indeed satisfied. Note that $\mathbf{E}_{\theta_i}\left[\int_{\underline{\theta}}^{\theta_i}\gamma_i(\tau)d\tau\right] = \mathbf{E}_{\theta_i}\left[\frac{1-F(\theta_i)}{f(\theta_i)}\gamma_i(\theta_i)\right]$.

Hence, we obtain

$$\max_{\{x_i(\cdot)\}_{i=1}^n} \quad \sum_{i=1}^n \mathbf{E}_{\theta_i}\left[\psi(\theta_i)\gamma_i(\theta_i) + V_i(\theta_i)\right], \tag{9.23}$$

recall that $\psi(\theta_i) = \theta_i - \frac{1-F(\theta_i)}{f(\theta_i)}$ (virtual type). Given the definitions in (9.18)-(9.19), and the fact that $\mathbf{E}[\cdot]$ is a linear operator, (9.23) can be re-written as:

$$\max_{\{x_i(\cdot)\}_{i=1}^n} \quad \mathbf{E}\sum_{i=1}^n \left[ax_i(\theta_i) - \frac{b}{2}x_i^2(\theta_i) + \psi(\theta_i)x_i(\theta_i)\sum_{j\neq i} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\right] \tag{9.24}$$

To find the optimal solution to (9.24), we maximize it point-wise.

Note that $\sum_{i=1}^n \mathbf{E}\left[\psi(\theta_i)x_i(\theta_i)\Big[\sum_{j\neq i} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big]\right]$ can be rewritten as

$$\begin{aligned}
&= \mathbf{E}\left[\psi_1(\theta_1)x_1(\theta_1)\Big[\sum_{j\neq 1} g_{1j}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big]\right] + \sum_{i=2}^n \mathbf{E}\left[\psi_i(\theta_i)x_i(\theta_i)\Big[\sum_{j\neq i} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big]\right] \\
&= \mathbf{E}\left[\psi_1(\theta_1)x_1(\theta_1)\Big[\sum_{j\neq 1} g_{1j}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big]\right] + \mathbf{E}\left[x_1(\theta_1)\Big[\sum_{j\neq 1} g_{j1}\mathbf{E}[\psi_j(\tilde{\theta}_j)x_j(\tilde{\theta}_j)]\Big]\right] \\
&\quad + \sum_{i=2}^n \mathbf{E}\left[\psi_i(\theta_i)x_i(\theta_i)\Big[\sum_{\substack{j\neq i\\ j\neq 1}} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big]\right]
\end{aligned}$$

$$= \mathbf{E}\left[\psi_1(\theta_1)x_1(\theta_1)\Big[\sum_{j\neq 1} g_{1j}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big] + x_1(\theta_1)\Big[\sum_{j\neq 1} g_{j1}\mathbf{E}[\psi_j(\tilde{\theta}_j)x_j(\tilde{\theta}_j)]\Big]\right]$$
$$+ \sum_{i=2}^{n}\mathbf{E}\left[\psi_i(\theta_i)x_i(\theta_i)\Big[\sum_{\substack{j\neq i\\ j\neq 1}} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big]\right] \tag{9.25}$$

Now, plugging (9.25) in the firm's problem, i.e. Eq. (9.24), and decoupling agent 1 from the rest, we obtain

$$\sum_{i=1}^{n}\mathbf{E}\left[ax_i(\theta_i) - \frac{b}{2}x_i^2(\theta_i) + \psi(\theta_i)x_i(\theta_i)\sum_{j\neq i} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\right]$$
$$= \sum_{i=2}^{n}\mathbf{E}\left[ax_i(\theta_i) - \frac{b}{2}x_i^2(\theta_i) + \psi(\theta_i)x_i(\theta_i)\sum_{\substack{j\neq i\\ j\neq 1}} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)]\right]$$
$$+ \mathbf{E}\left[a_1x_1(\theta_1) - \frac{b}{2}x_1^2(\theta_1) + \psi(\theta_1)x_1(\theta_1)\Big[\sum_{j\neq 1} g_{1j}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big] + x_1(\theta_1)\sum_{j\neq 1} g_{j1}\mathbf{E}[\psi_j(\tilde{\theta}_j)x_j(\tilde{\theta}_j)]\right].$$

FOC with respect to $x_1(\theta)$, by keeping $\mathbf{x}_{-1}$ fixed, gives

$$a - bx_1(\theta_1) + \psi(\theta_1)\Big[\sum_{j\neq 1} g_{1j}\mathbf{E}[x_j(\tilde{\theta}_j)]\Big] + \sum_{j\neq 1} g_{j1}\mathbf{E}[\psi_j(\tilde{\theta}_j)x_j(\tilde{\theta}_j)] = 0. \tag{9.26}$$

a similar equality can be obtained for any agent $i = 1, 2, \cdots, n$, that is for any i, we have

$$0 = a - bx_i(\theta_i) + \psi(\theta_i)\sum_{j=1}^{n} g_{ij}\mathbf{E}[x_j(\tilde{\theta}_j)] + \sum_{j=1}^{n} g_{ji}\mathbf{E}[\psi(\tilde{\theta}_j)x_j(\tilde{\theta}_j)]. \tag{9.27}$$

Equation (9.27) is rewritten in its matrix form as

$$b\mathbf{x}(\theta) = a\mathbf{1} + \psi(\theta)G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})]. \tag{9.28}$$

Taking expectation, we obtain

$$\begin{aligned} b\mathbf{E}[\mathbf{x}(\tilde{\theta})] &= a\mathbf{1} + \mathbf{E}[\psi(\tilde{\theta})]G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] \\ &= a\mathbf{1} + \mu_\psi G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})]. \end{aligned} \tag{9.29}$$

Thus, we have

$$G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] = (bI - \mu_\psi G)\mathbf{E}[\mathbf{x}(\tilde{\theta})] - a\mathbf{1}. \tag{9.30}$$

Next, multiplying (9.28) by $\psi(\theta)$ and taking expectation[50], we obtain

$$\begin{aligned} b\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] &= \mathbf{E}[\psi(\tilde{\theta})]a\mathbf{1} + \mathbf{E}[\psi^2(\tilde{\theta})]G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + \mathbf{E}[\psi(\tilde{\theta})]G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] \\ &= \mu_\psi a\mathbf{1} + \mathbf{E}[\psi^2(\tilde{\theta})]G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + \mu_\psi G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})]. \end{aligned} \tag{9.31}$$

Substituting for $G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})]$ from (9.30), we have,

$$\begin{aligned} b\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] &= \mu_\psi a\mathbf{1} + \mathbf{E}[\psi^2(\tilde{\theta})]G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + \mu_\psi \left((bI - \mu_\psi G)\mathbf{E}[\mathbf{x}(\tilde{\theta})] - a\mathbf{1}\right) \\ &= \mu_\psi b\ \mathbf{E}[\mathbf{x}(\tilde{\theta})] + \left(\mathbf{E}[\psi^2(\tilde{\theta})] - \mu_\psi^2\right) G\mathbf{E}[\mathbf{x}(\tilde{\theta})] \\ &= \mu_\psi b\ \mathbf{E}[\mathbf{x}(\tilde{\theta})] + \sigma_\psi^2 G\mathbf{E}[\mathbf{x}(\tilde{\theta})] \\ &= (\mu_\psi bI + \sigma_\psi^2 G)\ \mathbf{E}[\mathbf{x}(\tilde{\theta})], \end{aligned}$$

where the second last line follows from the fact that $\sigma_\psi^2 = \mathbf{Var}[\psi(\tilde{\theta})] = \mathbf{E}[\psi^2(\tilde{\theta})] - \mu_\psi^2$. Therefore,

$$\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] = \left(\mu_\psi I + \frac{\sigma_\psi^2}{b}G\right)\ \mathbf{E}[\mathbf{x}(\tilde{\theta})]. \tag{9.32}$$

Substituting the expression for $\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})]$ in (9.29), we obtain

$$\begin{aligned} b\mathbf{E}[\mathbf{x}(\tilde{\theta})] &= a\mathbf{1} + \mu_\psi G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + G^T\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] \\ &= a\mathbf{1} + \mu_\psi G\mathbf{E}[\mathbf{x}(\tilde{\theta})] + G^T\left(\mu_\psi I + \frac{\sigma_\psi^2}{b}G\right)\mathbf{E}[\mathbf{x}(\tilde{\theta})] \\ &= a\mathbf{1} + \left[\mu_\psi(G + G^T) + \frac{\sigma_\psi^2}{b}G^TG\right]\mathbf{E}[\mathbf{x}(\tilde{\theta})]. \end{aligned} \tag{9.33}$$

Therefore, we obtain the following linear systems of equations:

$$\left[bI - \mu_\psi(G + G^T) - \frac{\sigma_\psi^2}{b}G^TG\right]\mathbf{E}[\mathbf{x}(\tilde{\theta})] = a\mathbf{1}.$$

From Assumption 4.4.1, we know that the coefficient matrix is invertible. Thus, we obtain,

$$\mathbf{E}[\mathbf{x}(\tilde{\theta})] = \left[bI - \mu_\psi(G + G^T) - \frac{\sigma_\psi^2}{b}G^TG\right]^{-1} a\mathbf{1} = aK_\psi\mathbf{1}, \tag{9.34}$$

where $K_\psi \triangleq \left[bI - \mu_\psi(G + G^T) - \frac{\sigma_\psi^2}{b}G^TG\right]^{-1}$.

[50]Note that since $\psi(\theta)$ and B are diagonal matrices we have, $\psi(\theta)B = B\psi(\theta)$.

Finally, from equations (9.28), (9.30), and (9.34), we obtain,

$$\begin{aligned}\mathbf{x}(\theta) &= b^{-1}\Big[a\mathbf{1} + \psi(\theta)GK_\psi\Lambda + ((bI - \mu_\psi G)K_\psi a\mathbf{1} - a\mathbf{1})\Big] \\ &= b^{-1}\Big[\psi(\theta)GK_\psi a\mathbf{1} + (bI - \mu_\psi G)\, K_\psi a\mathbf{1}\Big] \\ &= b^{-1}\Big[(\psi(\theta) - \mu_\psi)GK_\psi a\mathbf{1} + bK_\psi a\mathbf{1}\Big] \\ &= b^{-1}\Big(\psi(\theta) - \mu_\psi\Big)GK_\psi a\mathbf{1} + K_\psi a\mathbf{1}. \end{aligned} \tag{9.35}$$

Therefore,

$$x^{SB}(\theta_i) = \frac{a}{b}(\psi(\theta_i) - \mu_\psi)[GK_\psi\mathbf{1}]_i + a[K_\psi\mathbf{1}]_i.$$

To complete the proof, we next show $\underline{\mathbf{x}}^{SB} > \mathbf{0}$. Note that $\mu_\psi = \mathbf{E}[\psi(\theta_i)] = \underline{\theta}$. Thus,

$$\psi(\underline{\theta}) - \mu_\psi = -\frac{1 - F(\underline{\theta})}{f(\underline{\theta})} = -\frac{1}{f(\underline{\theta})}.$$

Define, $\alpha \triangleq \frac{1}{b}\ \frac{1}{f(\underline{\theta})}$. By Assumption 4.4.1, $bI - \mu_\psi(G + G^T) - \frac{\sigma_\psi^2}{b}G^TG$ is invertible. Thus,

$$\begin{aligned}K_\psi a\mathbf{1} &= \left[bI - \mu_\psi(G + G^T) - \frac{\sigma_\psi^2}{b}G^TG\right]^{-1} a\mathbf{1} \\ &= \left[I - \frac{\mu_\psi}{b}(G + G^T) - \frac{\sigma_\psi^2}{b^2}G^TG\right]^{-1} \frac{a}{b}\mathbf{1} \\ &= (I - A)^{-1}\frac{a}{b}\mathbf{1},\end{aligned}$$

where $A \triangleq \frac{\mu_\psi}{b}(G + G^T) + \frac{\sigma_\psi^2}{b^2}G^TG$. Then, we have for $\underline{\theta} = (\underline{\theta}_1, \underline{\theta}_2, \cdots, \underline{\theta}_n)$:

$$\begin{aligned}\mathbf{x}(\theta) &\geq \mathbf{x}(\underline{\theta}) \\ &= (I - \alpha G)K_\psi a\mathbf{1} \\ &= (I - \alpha G)(I - A)^{-1}b^{-1}a\mathbf{1} \\ &\overset{(a)}{=} (I - \alpha G)(I + A + A^2 + A^3 + \cdots)b^{-1}a\mathbf{1} \\ &= \Big(I - \alpha G + A - \alpha GA + A^2 - \alpha GA^2 + A^3 + \cdots\Big)b^{-1}a\mathbf{1} \\ &= \Big(I + (A - \alpha G) + (A - \alpha G)A + (A - \alpha G)A^2 + \cdots\Big)b^{-1}a\mathbf{1} \\ &> 0,\end{aligned}$$

where (a) follows since by Assumption 4.4.1, $\rho(A) < 1$, thus $(I - A)^{-1} = I + A + A^2 + \cdots$, and the last inequality is correct because $b^{-1}a > 0$ and $A - \alpha G \geq \mathbf{0}$ (since $\psi(\underline{\theta}) \geq 0$). Finally, monotonicity of $x^{SB}(\theta_i)$ is immediate since $\psi(\theta_i)$ is monotone in θ_i, completing the proof.

Edskip **E**dskip **Full information (first best)**: In this case there is no IC and all the PC constraint must bind, characterizing the payments. Thus:

$$t_i(\theta_i) = ax_i(\theta_i) - \frac{b}{2}x_i^2(\theta_i) + \theta_i x_i(\theta_i)\sum_{j=1}^{n} g_{ij}\mathbf{E}[x_j(\theta_j)].$$

Hence, the firms problem is written as

$$\begin{array}{cc} \max\limits_{x,t} & \sum\limits_{i=1}^{n}\mathbf{E}[t_i(\theta_i)] \\ \text{subject to} & \text{PC}. \end{array}$$

Plugging the payments in the objective function we have, and using FOC with respect to $x_i(\theta_i)$ give

$$0 = a - bx_i(\theta_i) + \theta_i\sum_{j=1}^{n} g_{ij}\mathbf{E}[x_j(\theta_j)] + \sum_{j=1}^{n} g_{ji}\mathbf{E}[\theta_j x_j(\theta_j)]. \tag{9.36}$$

Eqs. (9.36) can be re-written in the following compact form:

$$0 = a - bx_i(\theta_i) + \theta_i\sum_{j=1}^{n} g_{ij}\mathbf{E}[x_j(\theta_j)] + \sum_{j=1}^{n} g_{ji}\mathbf{E}[\theta_j x_j(\theta_j)], \tag{9.37}$$

where $\theta_i \in [\underline{\theta}, \bar{\theta}]$, for all i. The above equality is similar to (9.27), thus, by following the same argument as in the second best we obtain:

$$x_i^{FB}(\theta_i) = \frac{a}{b}(\theta_i - \mu)[GK\mathbf{1}]_i + a[K\mathbf{1}]_i,$$

where $K \triangleq \Big[bI - \mu(G+G^T) - \frac{\sigma^2}{b}G^T G\Big]^{-1}$. ■

Proof of Proposition 4.4.2. The first part is immediate from Proposition 4.4.1. For the second, consider the full information case. Using Proposition 4.4.1, $x_i^{FB}(\theta_i) = \frac{a}{b}(\theta_i - \mu)[GK\mathbf{1}]_i + a[K\mathbf{1}]_i$ where $K \triangleq \Big[bI - \mu(G+G^T) - \frac{1}{b}G^T\Sigma G\Big]^{-1}$, and σ is a diagonal matrix where $Var[\theta_i] = \sigma_i^2$ is in entry ii. Since $I = KK^{-1}$, then by chain rule we obtain,

$$\mathbf{0} = \frac{\partial I}{\partial \sigma_i^2} = \frac{\partial K\mathcal{K}^{-1}}{\partial \sigma_i^2} = \frac{\partial K}{\partial \sigma_i^2}\,K^{-1} + K\,\frac{\partial K^{-1}}{\partial \sigma_i^2}. \tag{9.38}$$

where $\mathbf{0}$ denotes a $n \times n$ zero matrix. Then

$$\frac{\partial K^{-1}}{\partial \sigma_i^2} = \frac{\partial}{\partial \sigma_i^2}\Big[bI - \mu(G+G^T) - \frac{1}{b}G^T\Sigma G\Big] = -\frac{1}{b}\frac{\partial}{\partial \sigma_i^2}\Big[G^T\Sigma G\Big] = -\frac{1}{b}G^T\,\frac{\partial \sigma}{\partial \sigma_i^2}\,G = -\frac{1}{b}G^T E_{ii}\,G,$$

where $E_{ii} = \frac{\partial \sigma}{\partial \sigma_i^2}$ is a matrix containing only one non-zero entry equal to one located at the intersection of the i'th row and i'th column. Finally, using the last equality in (9.38) we obtain

$$\frac{\partial K}{\partial \sigma_i^2} = -K \; \frac{\partial K^{-1}}{\partial \sigma_i^2} \; K = \frac{1}{b} K G^T E_{ii} G K.$$

From the last equality, since all the matrices are component-wise non-negative, we imply that $\frac{\partial K}{\partial \sigma_i^2}$ is a matrix with non-negative entries, that completes the proof. We note that one way to change σ_i while keeping its mean fixed is to change the support of $\Theta_i = [\underline{\theta} \; \bar{\theta}]$. The proof for the imperfect information case is similar. ∎

Proof of Proposition 4.4.3. Given (4.4), using $\mathbf{E}[\psi(\tilde{\theta}) - \mu_\psi] = 0$, we have

$$\eta \triangleq G\mathbf{E}[\mathbf{x}(\tilde{\theta})] = GK_\psi \Lambda, \tag{9.39}$$

where $\eta = (\eta_1, \eta_2, \cdots, \eta_n)^T$ and $\Lambda \triangleq a\mathbf{1}$.

In order to derive $\mathbf{E}\left[\sum_{i=1}^n t_i(\theta_i)\right]$, we decompose it into three terms as follows

$$\begin{aligned}\mathbf{E}\left[\sum_{i=1}^n t_i(\theta_i)\right] &= \mathbf{E}\sum_{i=1}^n \left[a x_i(\theta_i) - \frac{b}{2} x_i^2(\theta_i) + \psi(\theta_i) x_i(\theta_i) \sum_{j \neq i} g_{ij} \mathbf{E}[x_j(\tilde{\theta}_j)]\right] \\ &= \zeta_1 + \zeta_2 + \zeta_3,\end{aligned}$$

where

$$\begin{aligned}\zeta_1 &= \sum_{i=1}^n \mathbf{E}[\psi(\theta_i) x_i(\theta_i)] \left[\sum_{j \neq i} g_{ij} \mathbf{E}[x_j(\tilde{\theta}_j)]\right] = \sum_{i=1}^n \eta_i \mathbf{E}[\psi(\theta_i) x_i(\theta_i)], \\ \zeta_2 &= \frac{-b}{2} \sum_{i=1}^n \mathbf{E}[x_i^2(\theta_i)], \\ \zeta_3 &= a \sum_{i=1}^n \mathbf{E}[x_i(\theta_i)].\end{aligned}$$

It follows that

$$\begin{aligned}\zeta_1 &= \sum_{i=1}^n \eta_i \; \mathbf{E}[\psi_i(\tilde{\theta}_i) x_i(\tilde{\theta}_i)] = \eta^T \mathbf{E}[\psi(\tilde{\theta}) \mathbf{x}(\tilde{\theta})] \overset{(a)}{=} (GK_\psi \Lambda)^T \left(\mu_\psi I + \frac{\sigma_\psi^2}{b} G\right) K_\psi \Lambda \\ &= \frac{\sigma_\psi^2}{b} \Lambda^T K_\psi^T G^T G K_\psi \Lambda + \mu_\psi \Lambda^T K_\psi^T G^T K_\psi \Lambda,\end{aligned} \tag{9.40}$$

where (a) follows from (9.32), (that is, $\mathbf{E}[\psi(\tilde{\theta})\mathbf{x}(\tilde{\theta})] = \left(\mu_\psi + \frac{\sigma_\psi^2}{b} G\right) \mathbf{E}[\mathbf{x}(\tilde{\theta})] = \left(\mu_\psi + \frac{\sigma_\psi^2}{b} G\right) K_\psi \Lambda$).

Furthermore, we obtain

$$\begin{aligned}
\zeta_2 &= \frac{-b}{2}\mathbf{E}\left[\sum_{i=1}^{n} x_i^2(\tilde{\theta}_i)\right] = \frac{-b}{2}\mathbf{E}\left[\mathbf{x}(\tilde{\theta})^T\ \mathbf{x}(\tilde{\theta})\right] \\
&= \frac{-b}{2}\mathbf{E}\left[\left[\frac{1}{b}(\psi(\tilde{\theta}) - \mu_\psi)GK_\psi\Lambda + K_\psi\Lambda\right]^T \left[\frac{1}{b}(\psi(\tilde{\theta}) - \mu_\psi)GK_\psi\Lambda + K_\psi\Lambda\right]\right] \\
&= \frac{-b}{2}\mathbf{E}\left[\left[\Lambda^T K_\psi^T G^T(\psi(\tilde{\theta}) - \mu_\psi)\frac{1}{b} + \Lambda^T K_\psi^T\right] \left[\frac{1}{b}(\psi(\tilde{\theta}) - \mu_\psi)GK_\psi\Lambda + K_\psi\Lambda\right]\right] \\
&\overset{(a)}{=} \frac{-\sigma_\psi^2}{2b}\Lambda^T K_\psi^T G^T G K_\psi\Lambda - \frac{b}{2}\Lambda^T K_\psi^T K_\psi\Lambda, \qquad (9.41)
\end{aligned}$$

where (a) is followed by the fact that $\mathbf{E}[(\psi(\tilde{\theta}) - \mu_\psi)] = 0$ and $\mathbf{E}[(\psi(\tilde{\theta}) - \mu_\psi)^2] = \sigma_\psi^2$.

Following the same arguments we also obtain

$$\zeta_3 = a\mathbf{E}[\mathbf{x}(\tilde{\theta})] = a\mathbf{1}^T K_\psi\Lambda = \Lambda^T K_\psi\Lambda. \qquad (9.42)$$

Now, from (9.40), (9.41) and (9.42) we obtain

$$\begin{aligned}
\Pi_G = \mathbf{E}\left[\sum_{i=1}^{n} t_i(\theta_i)\right] &= \zeta_1 + \zeta_2 + \zeta_3 \\
&= \frac{\sigma_\psi^2}{2b}\Lambda^T K_\psi^T G^T G K_\psi\Lambda + \mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda - \frac{b}{2}\Lambda^T K_\psi^T K_\psi\Lambda + \Lambda^T K_\psi\Lambda. \qquad (9.43)
\end{aligned}$$

Note that, since $\mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda$ is a scalar, thus

$$\mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda = \left(\mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda\right)^T = \mu_\psi\Lambda^T K_\psi^T G K_\psi\Lambda.$$

Thus

$$\begin{aligned}
\mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda &= \frac{1}{2}\left[\mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda + \left(\mu_\psi\Lambda^T K_\psi^T G^T K_\psi\Lambda\right)^T\right] \\
&= \frac{1}{2}\left[\Lambda^T K_\psi^T\ \left[\mu_\psi\left(G^T + G\right)\right]\ K_\psi\Lambda\right].
\end{aligned}$$

Hence by the preceding equality, we can simplify (9.43) as

$$\begin{aligned}
\Pi_G &= \frac{\sigma_\psi^2}{2b}\Lambda^T K_\psi^T G^T G K_\psi \Lambda + \mu_\psi \Lambda^T K_\psi^T G^T K_\psi \Lambda - \frac{b}{2}\Lambda^T K_\psi^T K_\psi \Lambda + a K_\psi \Lambda \\
&= \frac{\sigma_\psi^2}{2b}\Lambda^T K_\psi^T G^T G K_\psi \Lambda + \frac{1}{2}\left[\Lambda^T K_\psi^T \left[\mu_\psi \left(G^T + G\right)\right] K_\psi \Lambda\right] - \frac{b}{2}\Lambda^T K_\psi^T K_\psi \Lambda + \Lambda^T K_\psi \Lambda \\
&= \frac{1}{2}\Lambda^T K_\psi^T \left[\frac{\sigma_\psi^2}{b} G^T G + \mu_\psi (G^T + G) - bI\right] K_\psi \Lambda + \Lambda^T K_\psi \Lambda \\
&\overset{a}{=} -\frac{1}{2}\Lambda^T K_\psi^T K_\psi^{-1} K_\psi \Lambda + \Lambda^T K_\psi \Lambda \\
&= -\frac{1}{2}\Lambda^T K_\psi^T \Lambda + \Lambda^T K_\psi \Lambda \\
&\overset{b}{=} \frac{1}{2}\Lambda^T K_\psi \Lambda \\
&\overset{c}{=} \frac{a^2}{2}\ \mathbf{1}^T K_\psi \mathbf{1},
\end{aligned}$$

where (a) follows from the definition of K_ψ and (b) follows from symmetry of K_ψ, i.e., $K_\psi = K_\psi^T$ and (c) follows because $\Lambda = a\mathbf{1}$ by its definition. The proof for the full information case follows similar steps and the final answer becomes $\Pi_G^{bi.} = \frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1}$.

■

Proof of Proposition 4.4.4. Recall that $b = 1$ and we are in the full information case. The same proof works for the incomplete information case as well.

Following the same steps as in the proof of Proposition 4.4.3, one can show

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1} = \frac{a^2}{2}\mathbf{1}^T \left[I - \mu(G + G^T) - \sigma^2 G^T G\right]^{-1} \mathbf{1}.$$

Since σ^2 is chosen to be small, thus Taylor expansion of $\Pi_G^{\text{bi.}}$ around $\sigma^2 = 0$ gives

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\mathbf{1}^T \left[S^{-1} + \sigma^2\ S^{-1}\ G^T G\ S^{-1}\right] \mathbf{1},$$

where $S^{-1} \triangleq [I - \mu_\delta (G + G^T)]^{-1}$.

Since G is symmetric (i.e. $G = G^T$) it has n district eigenvalues and is thus diagonalizable. Therefore, there exists an invertible matrix V so that $V^{-1} G V = \Lambda$, where Λ is a diagonal matrix with the eigenvalues of G on its diagonal, and columns of V are the eigenvectors of G. Therefore, $G = V \Lambda V^{-1}$. In addition, the eigenvectors can be chosen to form an orthonormal basis of $\mathbb{R}^n$, meaning that V is set to be orthonormal. That is $V^T V = V V^T = V^{-1} V = I$.

Moreover,

$$S^{-1} = \sum_{i=0}^{\infty} (2\xi \ G)^i = \sum_{i=0}^{\infty} \left(2\xi \ V\Lambda V^{-1}\right)^i = V\left(\sum_{i=0}^{\infty} (2\xi \ \Lambda)^i\right) V^{-1} = V\Lambda_1 V^{-1}, \tag{9.44}$$

where Λ_1 is a diagonal matrix where its $k-$th element is $\frac{1}{1-2\xi\lambda_k}$, (note that $2\xi\lambda_k \leq 2\xi\lambda_{max} < 1$, for all k). By following similar argument as in (9.44), one can also show:

$$S^{-1} \ G^T G \ S^{-1} = V\Lambda_1^2 \Lambda^2 V^{-1}. \tag{9.45}$$

Eq. (9.44) together with (9.45) yield:

$$\Xi \triangleq S^{-1} + \sigma^2 \ S^{-1} \ G^T G \ S^{-1} = V \left(\underbrace{\Lambda_1 + \sigma^2 (\Lambda_1 \Lambda)^2}_{\text{diagonal}} \right) V^{-1}. \tag{9.46}$$

The above equality characterizes eigenvalues of Ξ in terms of $\lambda_1, \lambda_2, \cdots, \lambda_n$. That is, $k-$th eigenvalue of Ξ is equal to $f(\lambda_k) = \frac{1-2\xi\lambda_k+\sigma^2\lambda_k^2}{(1-2\xi\lambda_k)^2}$. Further, it can be easily shown $f(\lambda)$ is increasing and convex in λ.

To wrap up the proof we make a use of the following Lemma.

Lemma 4.9.2. *Let ζ_{min} and ζ_{max} be the smallest and largest eigenvalues of the square matrix M (with n distinct eigenvalues). Then:*

$$\zeta_{min} x^T x \leq x^T M x \leq \zeta_{max} x^T x,$$

where $x \in \mathbb{R}^n$.

Proof. Let $v_1, v_2, \cdots, v_n$ denote the eigenvectors corresponding to the eigenvalues $\zeta_1, \zeta_2, \cdots, \zeta_n$ of M. Thus, there exists $\alpha_1, \cdots, \alpha_n$ such that $x = \sum_{i=1}^n \alpha_i v_i$. Using this fact, since

$$Mx = M\sum_{i=1}^n \alpha_i v_i = \sum_{i=1}^n \alpha_i \ Mv_i = \sum_{i=1}^n \alpha_i \ \zeta_i v_i \leq \zeta_{max} \sum_{i=1}^n \alpha_i v_i = \zeta_{max} x,$$

thus $x^T M x \leq \zeta_{max} x^T x$. And similarly the lower can be proved, completing the proof. ∎

Finally, the proof of the proposition is immediate by employing Lemma 4.9.2 and the fact that $\mathbf{1}^T\mathbf{1} = n$, i.e.,

$$0 < \frac{na^2}{2} f(\lambda_{min}) \leq \Pi_G^{bi.} \leq \frac{na^2}{2} \ f(\lambda_{\max})$$

where $f(\lambda) = \frac{1-2\xi\lambda+\sigma^2\lambda^2}{(1-2\xi\lambda)^2}$.

∎

Proof of Proposition 4.4.5. The proof follows similar steps as in the proof of Proposition 4.4.2 and is omitted. ■

Proof of Corollary 4.4.1. Given Proposition 4.4.3, the proof is immediate because

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1} = \frac{a^2}{2}\mathbf{1}^T \left[I - \mu(G+G^T) - \sigma^2 G^T G\right]^{-1} \mathbf{1}$$

thus $\frac{\partial \Pi_G^{bi.}}{\partial \sigma^2} = \frac{a^2}{2b}\ \mathbf{1}^T K\ G^T G\ K\mathbf{1} > 0$. The proof for the imperfect information case is similar.

To gain intution, in the following we consider a discrete distribution and change σ^2 while the mean is kept fixed and prove the corollary. Let us assume $\text{Prob}\{\theta_i = \underline{\theta}\} = 1 - \text{Prob}\{\theta_i = \bar{\theta}\} = \mathsf{v} > \frac{\Delta\theta}{\bar{\theta}}$, for all i. Without loss of generality, let $\bar{\theta} = \delta\ \underline{\theta} > 0$, for some δ. Further, since $\mathsf{v} > \frac{\Delta\theta}{\bar{\theta}}$, thus $\delta > \frac{1}{1-\mathsf{v}}$. Since $\mu = \mathbb{E}[\theta_i]$ is fixed and given, thus $\underline{\theta} = \frac{\mu}{1+(1-\mathsf{v})(\delta-1)}$. As a result,

$$\sigma^2 = \mathsf{v}(1-\mathsf{v})(\bar{\theta} - \underline{\theta})^2 = \mathsf{v}(1-\mathsf{v})\mu^2 \left(\frac{\delta - 1}{1 + (1-\mathsf{v})(\delta - 1)}\right)^2.$$

We note that since μ and v are fixed and given, thus σ^2 is only controlled by δ. It also observes that σ^2 is increasing in δ, i.e. $\frac{\partial \sigma^2}{\partial \delta} > 0$. Intuitively, increasing δ, increases $\Delta\theta = \bar{\theta} - \underline{\theta}$, while μ and v are both kept fixed, increasing σ^2. Thus, given Proposition 4.4.3, $\frac{\partial \Pi_G^{bi.}}{\partial \delta} = \frac{\partial}{\partial \sigma^2}[\frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1}]\frac{\partial \sigma^2}{\partial \delta} = [\frac{a^2}{2}\ \mathbf{1}^T \frac{\partial K}{\partial \sigma^2}\mathbf{1}]\frac{\partial \sigma^2}{\partial \delta} = [\frac{a^2}{2b}\ \mathbf{1}^T K\ G^T G\ K\mathbf{1}]\frac{\partial \sigma^2}{\partial \delta} > 0$. ■

Proof of Proposition 4.5.1. Recall that $b = 1$ and we are in the full information case. For ease of exposition define $S^{-1} \triangleq (I - \mu(G+G^T))^{-1}$. Following the same steps as in the proof of Proposition 4.4.3, one can show

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1} = \frac{a^2}{2}\mathbf{1}^T \left[I - \mu(G+G^T) - \sigma^2 G^T G\right]^{-1} \mathbf{1}.$$

Since σ^2 is chosen to be small, thus Taylor expansion of $\Pi_G^{\text{bi.}}$ around $\sigma^2 = 0$ implies

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\mathbf{1}^T \left[S^{-1} + \sigma^2\ S^{-1}\ G^T G\ S^{-1}\right] \mathbf{1}.$$

Since G is balanced (see Def. 4.5.1), thus $S^{-1}\mathbf{1} = \frac{1}{1-\xi\tau}\mathbf{1} = \zeta\mathbf{1}$ (note that τ is small enough so that all the matrix-invertibilities are preserved). Therefore, $\mathbf{1}^T S^{-1}\mathbf{1} = \zeta n$ and $\mathbf{1}^T S^{-1} G^T G S^{-1}\mathbf{1} = \zeta^2 \mathbf{1}^T G^T G \mathbf{1} = \sum_i \left(\sum_j g_{ij}\right)^2$, completing the proof.

The same argument holds for the second best contract, with the change that linearization (Taylor expansion) of $\Pi_G^{\text{bi.,SB}}$ is around $\sigma_\psi^2 = 0$.

■

Proof of Lemma 4.5.1. Recall that $b = 1$ and we are in the first best contract. The same proof works for the second best contract as well, note that there only θ changes to $\psi(\theta)$ and Taylor expansion will be around σ_ψ^2.

Following the same steps as in the proof of Proposition 4.4.3, one can show

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\ \mathbf{1}^T K \mathbf{1} = \frac{a^2}{2}\mathbf{1}^T \left[I - \mu(G+G^T) - \sigma^2 G^T G\right]^{-1} \mathbf{1}.$$

Since σ^2 is chosen to be small, thus Taylor expansion of $\Pi_G^{\text{bi.}}$ around $\sigma^2 = 0$ gives

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\mathbf{1}^T \left[S^{-1} + \sigma^2\ S^{-1}\ G^T G\ S^{-1}\right] \mathbf{1}.$$

By little algebra, it can be easily shown that

$$G^T G = \sum_i R_i^2 - \text{diag}\left(\sum_{j=1}^n g_{1j}^2, \sum_{j=1}^n g_{2j}^2, \cdots, \sum_{j=1}^n g_{nj}^2\right)$$

that implies:

$$\Pi_G^{\text{bi.}} = \frac{a^2}{2}\mathbf{1}^T \left[S^{-1} + \sigma^2\ S^{-1}\ \left(\sum_i R_i^2 - \text{diag}\left(\sum_{j=1}^n g_{1j}^2, \sum_{j=1}^n g_{2j}^2, \cdots, \sum_{j=1}^n g_{nj}^2\right)\right)\ S^{-1}\right] \mathbf{1}. \quad (9.47)$$

Next, we proceed with computing $\Pi_G^{\text{multi.}}$. As shown in Proposition 4.3.1, the optimal (first best) trade profile is $\mathbf{x}^{FB}(\theta) = a\ [bI - (M_\theta G + G^T M_\theta)]^{-1}\mathbf{1}$. Thus, for each type profile θ the (ex post) firm's profit is given by (let x_i denote agent i's allocation with respect to the type profile θ)

$$\begin{aligned}
\sum_i \left(a x_i - \frac{b}{2}x_i^2 + x_i\theta_i \sum_j g_{ij}x_j\right) &= a\mathbf{1}^T\mathbf{x} - \frac{b}{2}\mathbf{x}^T\mathbf{x} + (M_\theta \mathbf{x})^T G\mathbf{x} \\
&= a\mathbf{1}^T\mathbf{x} - \frac{b}{2}\mathbf{x}^T\mathbf{x} + \frac{1}{2}\mathbf{x}^T(M_\theta G + G^T M_\theta)\mathbf{x} \\
&= a\mathbf{1}^T\mathbf{x} - \frac{1}{2}\left[\mathbf{x}^T\ \ [bI - (M_\theta G + G^T M_\theta)]\ \ \mathbf{x}\right] \\
&= \frac{a^2}{2}\mathbf{1}^T[bI - (M_\theta G + G^T M_\theta)]^{-1}\mathbf{1},
\end{aligned}$$

where the last equality follows by substituting the optimal allocation trade in it. Therefore, (setting $b = 1$) the (ex-ante) firm's profit (using multilateral) contract becomes:

$$\Pi_G^{\text{multi.}} = \frac{a^2}{2}\mathbf{1}^T\ \mathbf{E}\left[I - (M_\theta G + G^T M_\theta)\right]^{-1}\ \mathbf{1}. \quad (9.48)$$

To simplify $\mathbf{E}\left[I - (M_\theta G + G^T M_\theta)\right]^{-1}$, we first write the Taylor expansion of

$\Lambda \triangleq \Big[I - (M_\theta G + G^T M_\theta)\Big]^{-1}$ around $\theta = \mu\mathbf{1}$. Thus we have:

$$\Big[I - (M_\theta G + G^T M_\theta)\Big]^{-1} = \Big[I - \mu(G + G^T)\Big]^{-1} + \sum_i (\theta_i - \mu)\left(\frac{\partial \Lambda}{\partial \theta_i}|_{\theta=\mu\mathbf{1}}\right)$$
$$+ \frac{1}{2!}\sum_i \sum_j (\theta_i - \mu)(\theta_j - \mu)\left(\frac{\partial^2 \Lambda}{\partial \theta_i \partial \theta_j}|_{\theta=\mu\mathbf{1}}\right)$$
$$+ \sum_{n=3}^{\infty} \frac{1}{n!} \sum_{\substack{\zeta_1, \zeta_2, \cdots, \zeta_n \geq 0 \\ \zeta_1 + \cdots + \zeta_n = n}} \left[\prod_{k=1}^{n} (\theta_k - \mu)^{\zeta_k}\right] \left(\frac{\partial^n \Lambda}{\partial \theta_1^{\zeta_1} \partial \theta_2^{\zeta_2} \cdots \partial \theta_n^{\zeta_n}}|_{\theta=\mu\mathbf{1}}\right). \tag{9.49}$$

Next, since $\mathbf{E}(\theta_i - \mu) = 0$, for all i, we have

$$\mathbf{E}\left[\Big[I - (M_\theta G + G^T M_\theta)\Big]^{-1}\right] = \Big[I - \mu(G + G^T)\Big]^{-1} + \frac{1}{2!}\sum_i \mathbf{E}[(\theta_i - \mu)^2]\left(\frac{\partial^2 \Lambda}{\partial \theta_i \partial \theta_j}|_{\theta=\mu\mathbf{1}}\right)$$
$$+ T_{>2}, \tag{9.50}$$

where

$$T_{>2} = \sum_{n=3}^{\infty} \frac{1}{n!} \sum_{\substack{\zeta_1, \zeta_2, \cdots, \zeta_n \\ \zeta_1 + \cdots + \zeta_n = n \\ \zeta_i \neq 1, \forall i}} \prod_{k=1}^{n} \mathbf{E}[(\theta_k - \mu)^{\zeta_k}] \left(\frac{\partial^n \Lambda}{\partial \theta_1^{\zeta_1} \partial \theta_2^{\zeta_2} \cdots \partial \theta_n^{\zeta_n}}|_{\theta=\mu\mathbf{1}}\right).$$

Notice that in $T_{>2}, \zeta_i \neq 1$ (for all i), since $\mathbf{E}(\theta_i - \mu) = 0$. Moreover, by the Assumption, since there exists $\hat{m} > 0$, such that $\mathbf{E}[(\theta_i - \mu)^k] < (\hat{m}\sigma)^k$, for all i and $k \geq 3$, thus for all $n \geq 3$,

$$\prod_{k=1}^{n} \mathbf{E}[(\theta_k - \mu)^{\zeta_k}] < (\hat{m}\sigma)^{\sum_k \zeta_k} = (\hat{m}\sigma)^n.$$

In addition, σ^2 is chosen to be small, implying $\hat{m}\sigma$ is small (since $\hat{m}$ is constant). Consequently,

$$T_{>2} < \sum_{n=3}^{\infty} \frac{1}{n!} \sum_{\substack{\zeta_1, \zeta_2, \cdots, \zeta_n \\ \zeta_1 + \cdots + \zeta_n = n \\ \zeta_i \neq 1, \forall i}} (\hat{m}\sigma)^n \left(\frac{\partial^n \Lambda}{\partial \theta_1^{\zeta_1} \partial \theta_2^{\zeta_2} \cdots \partial \theta_n^{\zeta_n}}|_{\theta=\mu\mathbf{1}}\right).$$

Above inequality along with the Taylor expansion in (9.49) and the fact that $\hat{m}$ is constant imply that

$$T_{>2} = O(\sigma^3).$$

Therefore, since σ^2 is chosen to be small, $T_{>2}$ is negligible and thus (9.50) becomes

$$\mathbf{E}\left[\left[I-(M_\theta G+G^T M_\theta)\right]^{-1}\right]=\left[I-\mu(G+G^T)\right]^{-1}+\frac{\sigma^2}{2!}\sum_i\left(\frac{\partial^2\Lambda}{\partial\theta_i^2}|_{\theta=\mu\mathbf{1}}\right) \tag{9.51}$$

Moreover, for all i:

$$\frac{\partial^2\Lambda}{\partial\theta_i^2}|_{\theta=\mu\mathbf{1}}=2S^{-1}\left(R_i\ S^{-1}\ R_i\right)S^{-1}.$$

Thus, (9.51) is simplified as

$$\mathbf{E}\left[\left[I-(M_\theta G+G^T M_\theta)\right]^{-1}\right]=\left[I-\mu(G+G^T)\right]^{-1}+\sigma^2\ S^{-1}\left(\sum_i R_i\ S^{-1}\ R_i\right)S^{-1}.$$

Plugging the above equality in (9.48) finally implies that

$$\begin{aligned}\Pi_G^{\text{multi.}}&=\frac{a^2}{2}\mathbf{1}^T\ \mathbf{E}\left[I-(M_\theta G+G^T M_\theta)\right]^{-1}\ \mathbf{1}\\&=\frac{a^2}{2}\mathbf{1}^T\left[S^{-1}+\sigma^2\ S^{-1}\left(\sum_i R_i\ S^{-1}\ R_i\right)S^{-1}\right]\mathbf{1},\end{aligned} \tag{9.52}$$

recall that $S^{-1}=\left[I-\mu(G+G^T)\right]^{-1}$.

The proof is complete by comparing (9.47) and (9.52), that is

$$\begin{aligned}\Pi_G^{\text{multi.}}-\Pi_G^{\text{bi.}}&=\frac{a^2}{2}\sigma^2\left\{\left[K^T\sum_{i=1}^n R_i(S^{-1}-I)R_i\ K\right]+\sum_{i=1}^n k_i^2\parallel \deg_{in}(i)\parallel_2^2\right\}\\&=\left(\frac{a^2}{2}\sigma^2\right)K^T\left\{\sum_{i=1}^n R_i(S^{-1}-I)R_i+\text{diag}[(G\circ G)\mathbf{1}]\right\}K\end{aligned}$$

since

$$K=S^{-1}\mathbf{1}=[k_1\ k_2\ \cdots\ k_n]$$

and

$$\begin{aligned}\sum_{i=1}^n k_i^2\parallel \deg_{in}(i)\parallel_2^2&=\mathbf{1}^T S^{-1}\text{diag}\left(\sum_{j=1}^n g_{1j}^2,\sum_{j=1}^n g_{2j}^2,\cdots,\sum_{j=1}^n g_{nj}^2\right)S^{-1}\mathbf{1}\\&=K^T\text{diag}[(G\circ G)\mathbf{1}]K.\end{aligned}$$

■

Proof of Proposition 4.5.2. We prove each part separately as follows.
Part (i) The proof of the first part is immediate by (5.1). Note that k_i, for all i, increases by adding g_{ij}, assuming all the invertibility assumptions are preserved.

Part (ii) Let $0 < \alpha < 1$. Thus, the wight reduction is captured by αG, introducing a new network denoted by $\tilde{G} = \alpha G$. Using Lemma 4.5.1, since $\tilde{G}^T\tilde{G} = \alpha^2 G^T G$ and $\tilde{R}_i = \alpha R_i$, for all i, thus we have:

$$\begin{aligned}
\Pi^{\text{multi.}}_{\tilde{G}} - \Pi^{\text{bi.}}_{\tilde{G}}\Big|_{\tilde{G}=\alpha G} &= \left(\frac{a^2}{2}\sigma^2\right)\ \alpha^2\ \mathbf{1}^T\tilde{S}^{-1}\left\{\sum_{i=1}^{n} R_i\ \tilde{S}^{-1}R_i\ - G^T G\right\}\tilde{S}^{-1}\mathbf{1} \\
&\leq \left(\frac{a^2}{2}\sigma^2\right)\ \alpha^2\ \mathbf{1}^T S^{-1}\left\{\sum_{i=1}^{n} R_i\ S^{-1}R_i\ - G^T G\right\} S^{-1}\mathbf{1} \\
&= \alpha^2\left(\Pi^{\text{multi.}}_{G} - \Pi^{\text{bi.}}_{G}\right),
\end{aligned}$$

where the inequality follows because $\tilde{S}^{-1} = (I - \mu(\alpha G + \alpha G^T))^{-1}$ is increasing in α, thus $\tilde{S}^{-1} \leq S^{-1}$ (component-wise), completing the proof.

Part (iii) Since G is symmetric (i.e. $G = G^T$) it has n distinct eigenvalues and is thus diagonalizable. Therefore, there exists an invertible matrix V so that $V^{-1}GV = \Lambda$, where Λ is a diagonal matrix with the eigenvalues of G on its diagonal, and columns of V are the eigenvectors of G. Therefore, $G = V\Lambda V^{-1}$. In addition, the eigenvectors can be chosen to form an orthonormal basis of $\mathbb{R}^n$, meaning that V is set to be orthonormal. That is $V^TV = VV^T = V^{-1}V = I$. Moreover,

$$S^{-1} = \sum_{i=0}^{\infty}(2\xi\ G)^i = \sum_{i=0}^{\infty}\left(2\xi\ V\Lambda V^{-1}\right)^i = V\left(\sum_{i=0}^{\infty}(2\xi\ \Lambda)^i\right)V^{-1} = V\Lambda_1 V^{-1}, \tag{9.53}$$

where Λ_1 is a diagonal matrix where its $k-$th element is $\frac{1}{1-2\xi\lambda_k}$, (note that $2\xi\lambda_k \leq 2\xi\lambda_{max} < 1$, for all k).

By following similar argument as in (9.53), one can also show:

$$\begin{aligned}
S^{-1}\ G^TG\ S^{-1} &= V\Lambda_1^2\Lambda^2V^{-1} \\
S^{-1}\left(\sum_{i=1}^{n}R_i\right)\ S^{-1}\left(\sum_{i=1}^{n}R_i\right)S^{-1} &= V\Lambda_1^3(2\Lambda)^2V^{-1}.
\end{aligned} \tag{9.54}$$

Note that $\sum_{i=1}^{n} R_i = G + G^T = 2G$, (since $G = G^T$).

Thus, (9.54) yields:

$$\Xi \triangleq S^{-1}(2G)\ \ S^{-1}(2G)\,S^{-1} -\ S^{-1}\ G^TG\ S^{-1} = V\left(\underbrace{\Lambda_1^3(2\Lambda)^2 - (\Lambda_1\Lambda)^2}_{\text{diagonal}}\right)V^{-1}. \tag{9.55}$$

The above equality characterizes eigenvalues of Ξ in terms of $\lambda_1, \lambda_2, \cdots, \lambda_n$. That is, $k-$th eigenvalue of Ξ is equal to $f(\lambda_k) = \lambda_k^2\ \frac{3+2\xi\lambda_k}{(1-2\xi\lambda_k)^3}$. Further, it can be easily shown $f(\lambda)$ is increasing and convex in λ.

Next, using Lemma 4.5.1, we have

$$\begin{aligned}
\Pi_G^{\text{multi.}} - \Pi_G^{\text{bi.}} &= \left(\frac{a^2}{2}\sigma^2\right) \ \mathbf{1}^T S^{-1} \left\{\sum_{i=1}^{n} R_i \ S^{-1} R_i \ - G^T G\right\} S^{-1}\mathbf{1} \\
&\leq \left(\frac{a^2}{2}\sigma^2\right) \ \mathbf{1}^T S^{-1} \left\{\left(\sum_{i=1}^{n} R_i\right) \ S^{-1} \left(\sum_{i=1}^{n} R_i\right) \ - G^T G\right\} S^{-1}\mathbf{1} \\
&= \left(\frac{a^2}{2}\sigma^2\right) \mathbf{1}^T \Xi \mathbf{1}.
\end{aligned}$$

Finally, the proof of the proposition is immediate by employing Lemma 4.9.2 and the fact that $\mathbf{1}^T\mathbf{1} = n$.

Part (iv) To prove this part, we only use the second term in (5.1). Let d_i denote agent i's in degree in G (thus, by assumption, $\sum_j g_{ij} = gd_i$). Suppose agent i_{max} has the maximum in degree in G. Thus: (i) by definition, $k_{i_{max}} > 1 + \xi(d_{i_{max}} g)$. (ii) $\| \deg_{in}(i_{max}) \|_2^2 = g^2 d_{i_{max}}$. Hence, (i) together with (ii) give:

$$\Pi_G^{\text{multi.}} - \Pi_G^{\text{bi.}} \geq \left(\frac{a^2}{2}\sigma^2\right) \ \sum_{i=1}^{n} k_i^2 \| \deg_{in}(i) \|_2^2 \geq \ \left(\frac{a^2}{2}\sigma^2\right) \ (\xi \ d_{i_{max}} g)^2 g^2 d_{i_{max}} = O(d_{i_{max}}^3 g^4),$$

completing the proof.

■

Proof of Proposition 4.8.1. Using Lemma 4.5.1, we have

$$\Pi_G^{\text{multi.}} - \Pi_G^{\text{bi.}} = \frac{a^2}{2}\sigma^2 \left\{ \underbrace{K^T \sum_{i=1}^{n} R_i (S^{-1} - I) R_i \ K}_{\triangleq T_1} \ + \ \underbrace{\sum_{i=1}^{n} k_i^2 \| \deg_{in}(i) \|_2^2}_{\triangleq T_2} \right\}. \tag{9.56}$$

We analyze T_1 and T_2, separately, as follows. Let G denote the adjacency matrix of a star network with n nodes, where agent 1 located at the center of it and the rest are at the periphery.

Given the definitions of R_i, for all i (that is, $R_i \triangleq E_i G + G^T E_i$, where E_i is the matrix with only ith diagonal set as 1 and other entries as zero) for any non-negative matrix A we clearly have:

$$\sum_{i=1}^{n} R_i \ A \ R_i \leq \left(\sum_{i=1}^{n} R_i\right) \ A \ \left(\sum_{i=1}^{n} R_i\right),$$

where here ($\leq$) is for component-wise comparison. Since S^{-1} is invertible and well defined, thus $S^{-1} - I$ is component-wise non-negative. Thus, using the above inequality we have:

$$\sum_{i=1}^{n} R_i \; (S^{-1} - I) \; R_i \leq \left(\sum_{i=1}^{n} R_i \right) \; (S^{-1} - I) \; \left(\sum_{i=1}^{n} R_i \right). \tag{9.57}$$

Focusing on directed star networks (with no parallel links) where node 1 is located at the center, the right-hand-side of the above inequality is achievable for star-inward networks in which

$$\sum_{i=1}^{n} R_i = R_1 \triangleq \kappa \begin{bmatrix} 0 & 1 & \cdots & 1 \\ 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \cdots & 0 \end{bmatrix}.$$

Recall that we have assumed $g_{ij} \in \{0, \kappa\}$ for all i, j, where $\kappa > 0$.

Moreover, since by definition $K \triangleq S^{-1}\mathbf{1} = [k_1 \; k_2 \; \cdots \; k_n]^T$ is the same for all of these star networks, thus, using (9.57), we also have

$$K^T \left(\sum_{i=1}^{n} R_i \; (S^{-1} - I) \; R_i \right) \; K \leq K^T \left(\sum_{i=1}^{n} R_i \right) \; (S^{-1} - I) \; \left(\sum_{i=1}^{n} R_i \right) K, \tag{9.58}$$

that means the right-hand-side is achievable with the star-inward network. Thus, in (9.56), maximum of T_1 is obtained star-inward networks.

We next complete the proof by considering T_2 in (9.56). Again, note that in all of these star networks $K = S^{-1}\mathbf{1} = [k_1 \; k_2 \; \cdots \; k_n]^T$ is the same, $(k_1 > k_2 = k_3 = \cdots = k_n)$. But in star inward network

$$T_2 = \kappa^2 (n-1) k_1^2. \tag{9.59}$$

We claim T_2 is maximum in the case of a star-inward network. To prove this, consider another star network in which agent 1 obtains externality from n_1 neighbors, and n_2 periphery nodes obtain externality from agent 1, thus $n_1 + n_2 = n - 1$. In this star network $T_2 = \kappa^2 n_1 k_1^2 + \kappa^2 n_2 k_2^2 = \kappa^2 (n_1 k_1^2 + n_2 k_2^2)$, that is lower than (9.59), because $k_1 > k_2$.

■

Proof of Proposition 4.8.2. To prove the proposition, we first show that if $\alpha_n = \frac{1}{\sqrt{(n-1)(\mu^2+\sigma^2+\epsilon)}}$, where ϵ is positive and constant, then $\left(I - \mu\left(G(n) + G(n)^T\right) + \sigma^2 G(n)^T G(n)\right)^{-1}$ is well-defined and non-negative. For ease of illustration denote $G(n) \triangleq \mu\left(G(n) + G(n)^T\right) - \sigma^2 G(n)^T G(n)$. To prove the invertibility, it is enough to show that $\rho\left(G(n)\right) < 1$, i.e., the maximum eigenvalue of $G(n)$, in absolute value, is less than 1. Note that $G(n)$ is symmetric, thus all of its eigenvalues are real. In the sequel, fix n. It is clear that:

$$G(n) + G(n)^T = \alpha_n \begin{bmatrix} 0 & 1 & \cdots & 1 \\ 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \cdots & 0 \end{bmatrix},$$

and

$$G(n)^T G(n) = \alpha_n^2 \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 1 & \cdots & 1 \end{bmatrix}.$$

Let λ be an eigenvalue of $G(n)$. By definition, $G(n)x = \lambda x$ does have a solution in x. Due to the symmetry, x has the following form, $x^T = [a\ b\ b\ \cdots\ b]$. Thus, due to the definition of $G(n)$ and the above two (matrix-from) equalities, $G(n)x = \lambda x$ yields

$$\begin{aligned} \lambda a &= \alpha_n \mu (n-1) b, \\ \lambda b &= \alpha_n \mu a + \sigma^2 \alpha_n^2 (n-1) b. \end{aligned}$$

Thus, by a little algebra, we have

$$\lambda^2 - \underbrace{\alpha_n^2 \sigma^2 (n-1)}_{\triangleq\, B} \lambda - \underbrace{\alpha_n^2 \xi^2 (n-1)}_{\triangleq\, C} = 0. \tag{9.60}$$

that is, $\lambda^2 - B\lambda - C = 0$. Now, note that $0 < C < B$ and $B + C < 1$, where the later is ensured by $\alpha_n = \frac{1}{\sqrt{(n-1)(\mu^2+\sigma^2+\epsilon)}}$, where $\epsilon > 0$ and is constant. Therefore, the roots of (9.60) are (in absolute value) less than 1, that is desired. This also implies $\left(I - \mu\left(G(n) + G(n)^T\right)\right)^{-1}$ is well-defined and non-negative. Next, we continue to complete the proof. By definition, we have:

$$\begin{aligned} k_1(n) = \left[\left(I - \mu\left(G(n) + G(n)^T\right)\right)^{-1} \mathbf{1}\right]_1 &> \left[\left(I + \mu\left(G(n) + G(n)^T\right)\right) \mathbf{1}\right]_1 \\ &= 1 + \mu \alpha_n (n-1). \end{aligned}$$

From the above inequality we obtain:

$$k_1(n) = O\left(\sqrt{n}\right).$$

Finally, using Lemma 4.5.1, for all n, we have:

$$\begin{aligned}
\Pi^{\text{multi.}}_{G(n)} - \Pi^{\text{bi.}}_{G(n)} &> \frac{a^2}{2}\sigma^2 \left\{ \sum_{i=1}^{n} k_i^2 \parallel \deg_{in}(i) \parallel_2^2 \right\} \\
&= \frac{a^2}{2}\sigma^2 \; k_1^2(n) \; \alpha_n^2 \; (n-1) \\
&= \frac{a^2}{2} \left(\frac{\sigma^2}{\xi^2 + \sigma^2 + \epsilon} \right) \; k_1^2(n).
\end{aligned}$$

The last equality implies that as n grows, $\Pi^{\text{multi.}}_{G(n)} - \Pi^{\text{bi.}}_{G(n)}$ tends to infinity, as $k_1(n)$ goes to infinity.

■

Proof of Proposition 4.8.3. Using Proposition 4.8.1, proving this result is equivalent to show the following.

Equivalent Proposition: Consider the sequence of Star-inward networks $\{G(n)\}_{n=2}^{\infty}$, where $\beta_n \triangleq g_{12} = g_{13} = \cdots = g_{1n} = \frac{1}{(n-1)\sqrt{\mu^2+\sigma^2+\epsilon}}$, where $\epsilon > 0$ is a constant, and the rest of entries are zero. Then we have:

$$\Pi^{\text{multi.}}_{G(n)} - \Pi^{\text{bi.}}_{G(n)} \;\longrightarrow\; 0.$$

We first note since $\beta_n = \frac{1}{(n-1)\sqrt{\mu^2+\sigma^2+\epsilon}}$, as shown in Proposition 4.8.2, $\left(I - \mu_\delta \left(G(n) + G(n)^T\right) - \sigma^2 G(n)^T G(n)\right)^{-1}$ is well-defined and non-negative, for all n.

Now, fix $n > 1$, and denote

$$R_1 \triangleq \beta_n \begin{bmatrix} 0 & 1 & \cdots & 1 \\ 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \cdots & 0 \end{bmatrix} = G + G^T,$$

thus, for any $t > 0$,

$$R_1^{2t+1} = \beta_n^{2t+1} \begin{bmatrix} 0 & (n-1)^t & \cdots & (n-1)^t \\ (n-1)^t & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ (n-1)^t & 0 & \cdots & 0 \end{bmatrix}, \tag{9.61}$$

and

$$R_1^{2t} = \beta_n^{2t} \begin{bmatrix} (n-1)^t & 0 & \cdots & 0 \\ 0 & (n-1)^{t-1} & \cdots & (n-1)^{t-1} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & (n-1)^{t-1} & \cdots & (n-1)^{t-1} \end{bmatrix}. \tag{9.62}$$

We first compute Bonacich centrality measures of agents in $G+G^T$. By definition, Bonacich centrality vector is given by: $[k_1\ k_2\ \cdots\ k_2]^T = \Big(I - \xi(G+G^T)\Big)^{-1}\mathbf{1}$. Due to the symmetry and the invertibility, $\Big(I - \xi(G+G^T)\Big)^{-1} = \sum_{t\geq 0}\Big(\xi(G+G^T)\Big)^t = \sum_{t\geq 0}(\xi R_1)^t$. Thus, by simple algebra, it follows that

$$k_1 = 1 + \beta_n\xi(n-1) + \sum_{t=2}^{\infty}(\beta_n\xi)^t\,(n-1)^{\lceil \frac{t}{2}\rceil} = O(1) = \text{constant}, \tag{9.63}$$

$$k_2 = 1 + \beta_n\xi + \sum_{t=2}^{\infty}(\beta_n\xi)^t\,(n-1)^{\lfloor \frac{t}{2}\rfloor} = O(1) = \text{constant}, \tag{9.64}$$

where the last equality in the above two equalities are followed since n is large and $\beta_n = \frac{1}{(n-1)\sqrt{\mu^2+\sigma^2+\epsilon}}$. Now, using Proposition 4.5.1, we have

$$\Pi_G^{\text{multi.}} - \Pi_G^{\text{bi.}} = \frac{a^2}{2}\sigma_\delta^2\left\{\underbrace{K^T\sum_{i=1}^{n}R_i(S^{-1}-I)R_i\ K}_{\triangleq T_1} \;+\; \underbrace{\sum_{i=1}^{n}k_i^2\parallel \deg_{in}(i)\parallel_2^2}_{\triangleq T_2}\right\}, \tag{9.65}$$

thus, in this set up, $T_2 = k_1^2\beta_n^2(n-1) \to 0$. This is because $\beta_n = \frac{1}{(n-1)\sqrt{\mu^2+\sigma^2+\epsilon}}$ and $k_1 = O(1)$. In the sequel of the proof, according to (9.65), we focus on T_1. According to the definitions of R_i (see proof of Proposition 4.5.1), we have:

$$\begin{aligned} T_1 &= K^T\sum_{i=1}^{n}R_i(S^{-1}-I)R_i\ K, \\ &= K^T R_1(S^{-1}-I)R_1\ K, \\ &= K^T R_1\left(\sum_{t=1}^{\infty}(\xi R_1)^t\right)R_1\ K, \\ &= K^T\left(\sum_{t=1}^{\infty}\xi^t\ R_1^{t+2}\right)K \\ &= \underbrace{K^T\left(\sum_{t=1}^{\infty}\xi^{2t-1}\ R_1^{2t+1}\right)K}_{\triangleq T_{odd}} + \underbrace{K^T\left(\sum_{t=2}^{\infty}\xi^{2t-2}\ R_1^{2t}\right)K}_{\triangleq T_{even}}. \end{aligned} \tag{9.66}$$

Recall that $K^T = [k_1 \ k_2 \ \cdots \ k_2]$. Further, using (9.61) and (9.62), for any t, we have:

$$K^T R_1^{2t+1} K = \beta_n^{2t+1} \ (2k_1 k_2 (n-1)^{t+1}),$$
$$K^T R_1^{2t} K = \beta_n^{2t} \ (k_1^2 + (n-1)k_2^2) \ (n-1)^t.$$

Plugging the above equalities into (9.66), and approaching n to infinity, we obtain:

$$T_{odd} = 2k_1 k_2 \sum_{t=1}^{\infty} \xi^{2t-1} (n-1)^{t+1} \beta_n^{2t+1} \to 0,$$
$$T_{even} = (k_1^2 + (n-1)k_2^2) \sum_{t=2}^{\infty} \xi^{2t-2} (n-1)^t \beta_n^{2t} \to 0,$$

That is because $O(k_1) = O(k_2) = 1$ and $\beta_n = O\left(\frac{1}{n}\right)$. Thus $T_1 \to 0$, that completes the proof. ∎

Proof of Proposition 4.5.3. To prove the result we show $\Pi_{G(n)}^{\text{multi.}} = \Pi_{G(n)}^{\text{bi.}} = O(n)$. Using (9.58) we have:

$$\begin{aligned}
\mathbf{1}^T S^{-1} \sum_i R_i S^{-1} R_i S^{-1} \mathbf{1} &= \zeta^2 \mathbf{1}^T \left(\sum_i R_i S^{-1} R_i \right) \mathbf{1} \\
&\leq \zeta^2 \mathbf{1}^T \left(\sum_i R_i \right) S^{-1} \left(\sum_i R_i \right) \mathbf{1} \\
&= \zeta^2 \mathbf{1}^T \left(G + G^T \right) S^{-1} \left(G + G^T \right) \mathbf{1} \\
&= \zeta^2 \tau^2 \mathbf{1}^T S^{-1} \mathbf{1} \\
&= \zeta^3 \tau^2 n \\
&= O(n).
\end{aligned}$$

Thus, the above inequality along with (9.52) imply that

$$\begin{aligned}
\Pi_{G(n)}^{\text{bi.}} \leq \Pi_{G(n)}^{\text{multi.}} &= \frac{a^2}{2} \mathbf{1}^T \left[S^{-1} + \sigma^2 \ S^{-1} \left(\sum_i R_i \ S^{-1} \ R_i \right) S^{-1} \right] \mathbf{1} \\
&\leq \frac{a^2}{2} \left[\zeta n + \sigma^2 \zeta^3 \tau^2 n \right] \\
&= O(n).
\end{aligned}$$

In addition, using Proposition 4.5.1, since the in degree at each node is the same, it clears that $\Pi_{G(n)}^{\text{bi.}} = O(n)$. Therefore, $\lim_{n\to\infty} \frac{\Pi_{G(n)}^{\text{multi.}}}{\Pi_{G(n)}^{\text{bi.}}} = O(1)$.

To complete the proof, we need to guaranty the invertibility of $T = \left(I - \xi(G + G^T) - \sigma^2 G^T G \right)$. Choosing small enough τ (independent of n) will ensure this, if we show $G^T G$ dose not blow up as n grows to infinity. Assuming $G+G^T$ is k-regular (k is finite), we obtain

$$\sum_j [G^T G]_{ij} = \text{number of walks of length 2 starting from } i = k^2,$$

which is independent of n. Meaning that by choosing small τ (independent of n) T can be diagonally dominant and of course invertible. ■

Proof of Proposition 4.5.4. Fix n. We first note that since

$$G(n) = \begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix},$$

thus,

$$G(n)^T G(n) = I = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}.$$

and

$$G(n) + G(n)^T = \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 1 \\ 1 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 1 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \ddots & \\ 0 & \cdots & 0 & 1 & 0 & 1 \\ 1 & 0 & \cdots & 0 & 1 & 0 \end{bmatrix}.$$

Therefore, when $2\xi + \sigma^2 < 1$, then $\left(I - \mu_\delta \left(G(n) + G(n)^T\right) - \sigma^2 G(n)^T G(n)\right)^{-1}$ is well-defined and non-negative[51]. Furthermore, Bonacich centrality measures of agents in $G(n) + G(n)^T$ (by definition) is given by: $[k_1 \ k_2 \ \cdots \ k_2]^T = S(n)^{-1}\mathbf{1} = \left(I - \xi(G(n) + G(n)^T)\right)^{-1}\mathbf{1}$. Due to the symmetry and the invertibility we have $k_1 = \cdots = k_n = k$, and since $\xi < \frac{1}{4}$, by simple algebra, we have $1 \leq k < 2$. That is, $k = O(1)$.

As shown in Proposition 4.5.1

$$\Pi^{\text{bi.}}_{G(n)} = \frac{a^2}{2}\mathbf{1}^T \left[S^{-1} + \sigma^2 \ S(n)^{-1} \ \underbrace{G(n)^T G(n)}_{=I} \ S(n)^{-1}\right]\mathbf{1} = \frac{a^2}{2}\left[\mathbf{1}^T S^{-1}\mathbf{1}^T + \sigma^2 \ \mathbf{1}^T S(n)^{-1} \ \ S(n)^{-1}\mathbf{1}\right]$$
$$= \frac{a^2}{2}\left[kn + \sigma^2 k^2 n\right] = O(n). \tag{9.67}$$

[51] This is because $I - \mu\left(G(n) + G(n)^T\right) - \sigma^2 G(n)^T G(n)$ becomes diagonally dominant (when $2\xi + \sigma^2 < 1$).

Next, we show $\Pi_{G(n)}^{\text{multi.}} = O(n)$.

As shown in Lemma 4.5.1 (see Eq. (9.52)) we have:

$$\begin{aligned}
\Pi_{G(n)}^{\text{multi.}} &= \frac{a^2}{2}\mathbf{1}^T\ \mathbf{E}\left[I - (M_\delta G(n) + G(n)^T M_\delta)\right]^{-1}\ \mathbf{1} \\
&= \frac{a^2}{2}\left[\underbrace{\mathbf{1}^T S(n)^{-1}\mathbf{1}}_{T_1} + \underbrace{\sigma^2\ \mathbf{1}^T\ S(n)^{-1}\left(\sum_i R_i\ S(n)^{-1}\ R_i\right) S(n)^{-1}\mathbf{1}}_{T_2}\right],
\end{aligned}$$

First we note that, by definition, $T_1 = kn = O(n)$. Next we show $T_2 = O(n)$. Notice that $S(n)^{-1}\mathbf{1} = k\mathbf{1}$ and $\mathbf{1}^T S(n)^{-1} = k\mathbf{1}^T$. Thus,

$$T_2 = \sigma^2 k^2\ \mathbf{1}^T\left(\sum_i R_i\ S(n)^{-1}\ R_i\right)\mathbf{1}.$$

Therefore,

$$T_2 \leq \sigma^2 k^2\ \ \mathbf{1}^T\left[\left(\sum_i R_i\right) S(n)^{-1}\left(\sum_i R_i\right)\right]\mathbf{1} = 2\sigma^2 k^2\ \ \mathbf{1}^T S(n)^{-1}\mathbf{1} = 2\sigma^2 k^3 n. \qquad (9.68)$$

Also since $S(n)^{-1} \geq I$, thus

$$T_2 \geq \sigma^2 k^2\ \mathbf{1}^T\left(\sum_i R_i\ R_i\right)\mathbf{1} \geq \sigma^2 k^2\ \ \mathbf{1}^T\mathbf{1} = \sigma^2 k^2 n.$$

Eq. (9.68) along with (9.69) imply that $T_2 = O(n)$. Thus, $\Pi_{G(n)}^{\text{multi.}} = O(n)$. This along with (9.67) complete the proof.

∎

Proof of Proposition 4.5.5. Without loss of generality, for the ratio analysis, one can set $\frac{a^2}{2} = 1$. Thus:

$$\begin{aligned}
\Pi_G^{multi.} &= \mathbf{1}^T K + \sigma^2\ K^T\left(\sum_i R_i\ S^{-1}\ R_i\right) K \\
&\leq \mathbf{1}^T K + \sigma^2\ K^T\left(\sum_i R_i\right)\ S^{-1}\ \left(\sum_i R_i\right) K \\
&= \mathbf{1}^T K + \sigma^2\ K^T\left(G + G^T\right)\ S^{-1}\ \left(G + G^T\right) K \\
&= \mathbf{1}^T K + \sigma^2\ \mathbf{1}^T\left(G + G^T\right)^2\ S^{-3}\mathbf{1}.
\end{aligned}$$

Let λ_{max} be the maximum eigenvalue of $G + G^T$. Note that invertibility ensures $\xi\lambda_{max} < 1$. Since $G + G^T$ is symmetric, it is diagonalizable. There exists orthonormal matrices P and

P^T such that $PP^{-1} = PP = PP^T = I$, and $P(G + G^T)P^T = \text{diag}(\lambda_1, \cdots, \lambda_n)$, where λ_i are eigenvalues of $G + G^T$.

Using the digitalization of $G + G^T$ yields

$$(G + G^T)^2 S^{-3} = P^T \text{diag}(f(\lambda_1), \cdots, f(\lambda_n))P$$

where $f(\lambda) \triangleq \frac{\lambda^2}{(1-\xi\lambda)^3}$, which is increasing and convex in λ.

Since $f(\lambda)$ is increasing in λ, thus the maximum eigenvalue of $(G + G^T)^2 S^{-3}$ becomes $f(\lambda_{max})$. This implies that $\mathbf{1}^T \left(G + G^T\right)^2 \ S^{-3}\mathbf{1} \leq nf(\lambda_{max})$. Therefore, we obtain:

$$\Pi_G^{multi.} \leq ||K||_1 + \sigma^2 \ n \ f(\lambda_{max}). \tag{9.69}$$

To wrap up the proof, recall that

$$\Pi_G^{bi.} = ||K||_1 + \sigma^2 K^T G^T G K \geq ||K||_1. \tag{9.70}$$

Since $||K||_1 \geq n$, the proof is complete. ∎

Proof of Proposition 4.5.6. Again, without loss of generality, for the ratio analysis, one can assume $\frac{a^2}{2} = 1$. Thus:

$$\begin{aligned}
\Pi_G^{multi.} &= \mathbf{1}^T K + \sigma^2 \ K^T \left(\sum_i R_i \ S^{-1} \ R_i\right) K \\
&\geq \mathbf{1}^T K + \sigma^2 \ K^T \left(\sum_i R_i^2\right) K \\
&= \mathbf{1}^T K + \sigma^2 \ K^T \left(G^T G + \tilde{M}\right) K,
\end{aligned}$$

where $\tilde{M} = \text{diag}(||d_1||_2^2, ||d_2||_2^2, \cdots, ||d_n||_2^2)$. Therefore,

$$\begin{aligned}
\frac{\Pi_G^{multi.}}{\Pi_G^{bi.}} &\geq \frac{\mathbf{1}^T K + \sigma^2 \ K^T \left(G^T G + \tilde{M}\right) K}{\mathbf{1}^T K + \sigma^2 \ K^T G^T G K} \\
&= 1 + \frac{\sigma^2 \ K^T \tilde{M} K}{\mathbf{1}^T K + \sigma^2 \ K^T G^T G K} \\
&\geq 1 + \frac{\sigma^2 ||\text{diag}(K) G||_F^2}{||K||_1 + ||K||_2^2}
\end{aligned}$$

where the last inequality follows by noting that the maximum eigenvalue of $\sigma^2 G^T G$ is less than 1, and therefore $\sigma^2 K^T G^T G K \leq K^T K = ||K||_2^2$.

∎

4.10 Appendix: Different centrality orders in $G+G^T$ and $\mathcal{G}+\mathcal{G}^T$

For ease of notation we first have the following definition.

Definition 4.10.1 (Modified Network: $\mathcal{G}$)**.** *Given the adjacency matrix G and the mean and variance of the (virtual) types, respectively μ_ψ and σ_ψ^2, the* modified network *$\mathcal{G}$ is defined as*

$$\mathcal{G} \triangleq \mu_\psi G + \frac{\sigma_\psi^2}{2b} G^T G. \tag{10.1}$$

The modified network for the first best allocation is similarly defined via using μ and σ^2 instead of μ_ψ and σ_ψ^2, respectively.

A quick comparison between Propositions 4.3.1 and 4.4.1 shows that multilateral and bilateral contracts may *not* necessarily induce the same order of allocations. To see this, Let $b=1$ and $\theta_i=\bar{\theta}$, for all i. Then, Proposition 4.3.1 implies $\mathbf{x}^{SB}_{\text{multilateral}} = a[I-\bar{\theta}(G+G^T)]^{-1}\mathbf{1}$, that is proportional to the Bonacich centrality measure in $G+G^T$. However, Proposition 4.4.1 implies $\mathbf{x}^{SB}_{\text{bilateral}} = a(\Delta\theta\ G+I)[I-(\mathcal{G}+\mathcal{G}^T)]^{-1}\mathbf{1}$, where $\mathcal{G}\triangleq \mu_\psi G + \frac{\sigma_\psi^2}{2b}G^TG$. That is proportional to the Bonacich centrality measure in $\mathcal{G}+\mathcal{G}$.

In this section we provide a condition under which centrality orders in the modified and the original networks are always diffident.

In general, the form of the modified network (see Def. 4.10.1) suggests that when G^TG is sufficiently different from G, even with a small amount of uncertainty (i.e. σ^2), a change on the central agents in $G+G^T$ in compare to $\mathcal{G}+\mathcal{G}^T$ may happen. In other words, in networks with high second order of connectivity, captured by G^TG, a small amount of uncertainty may make a change on the centrality orders of agents in the original network.

Following the above intuition, in the next theorem, we present a sufficient condition under which centrality orders in $G+G^T$ and $\mathcal{G}+\mathcal{G}^T$ are different. For ease of exposition, let us define $d_G^{out}(i)$ denoting the out-degree of agent i in G, i.e., $d_G^{out}(i) = |\{k : g_{ki}>0, k=1,2,\cdots,n\}|$.

Proposition 4.10.1. *Let agent i and j be both central in $G+G^T$. If $d_G^{out}(i) > d_G^{out}(j) = 0$, then Bonacich centrality of agent i is strictly higher than agent j in $\mathcal{G}+\mathcal{G}^T$.*

Proof of Proposition 4.10.1. By definition, we have

$$\frac{1}{b}(G^TG)_{rs} = \frac{1}{b}\sum_{k=1}^{n} \tilde{g}_{rk}\ g_{ks} = \frac{1}{b}\sum_{k=1}^{n} g_{kr}\ g_{ks}, \tag{10.2}$$

From (10.2) we obtain for any r, s:

$$(G^TG)_{rr} = \frac{1}{2}\sum_{k=1}^{n} g_{kr}^2 > 0, \qquad \text{if } d_G^{out}(r) > 0.$$
$$(G^TG)_{rs} = (G^TG)_{sr} = 0, \qquad \text{if } d_G^{out}(r) \times d_G^{out}(s) = 0.$$
$$(G^TG)_{rs} = (G^TG)_{sr} > 0, \qquad \text{if } \exists\, k \text{ s.t. } g_{kr} > 0 \text{ and } g_{ks} > 0. \tag{10.3}$$

Bonacich centrality measure of an agent in a given network counts the total number of (suitably weighted) walks of different length starting from the agent in the network. Next, consider agent i and j such that $d_G^{out}(i) > d_G^{out}(j) = 0$. In G^TG: since $d_G^{out}(j) = 0$, thus there is no walk starting from agent j, whereas, because $d_G^{out}(i) > 0$, there exits walks (at least self-loops) starting from agent i.

Therefore, because: (i) $d_G^{out}(i) > d_G^{out}(j) = 0$, (ii) agent i and j are both equivalently central in $G + G^T$, and (iii) $\mathcal{G} + \mathcal{G}^T = \mu(G+G^T) + \frac{\sigma^2}{b}G^TG$, it follows that Bonacich centrality of agent i is strictly higher than agent j in $\mathcal{G} + \mathcal{G}^T$. ■

The above theorem intuitively states that having a difference in the out-degrees of two central agents in $G + G^T$, might be crucial to make a change in their centralities with respect to the modified network. This theorem is substantiated by the following example.

Example 5 (In-ward-out-ward stars)**.** Consider the following network G with $2d+2$ agents[52] wherein $d_G^{out}(1) = d+1$, $d_G^{out}(i) = 1$, $d+3 \le i \le 2d+2$, and $d_G^{out}(2) = d_G^{out}(j) = 0$, $3 \le j \le d+2$, as it is depicted in Fig. 4-9. In the modified network (depicted in Fig. 4-9) the impact of G^TG is drawn by red lines. Since $d_G^{out}(1) = d+1$ and $d_G^{out}(i) = 1$, $d+3 \le i \le 2d+2$, thus due to G^TG node 1 and node i, $d+3 \le i \le 2d+2$, have, respectively, $d+1$ and 1 self-loops. Moreover, node 1 and i, $d+3 \le i \le 2d+2$, both feed agent 2, thus agents 1 and i are connected in $\mathcal{G}$ due to G^TG. Now, applying Theorem 4.10.1, it follows that agents 1 and 2 are both, equivalently, central in $G + G^T$. However, agent 1 has higher centrality than agent 2 in $\mathcal{G} + \mathcal{G}^T$.

[52]Suppose $g_{ij} \in \{0, k\}$, for any i, j, where $k > 0$ is sufficiently small enough such that Assumption 4.4.1 is satisfied.

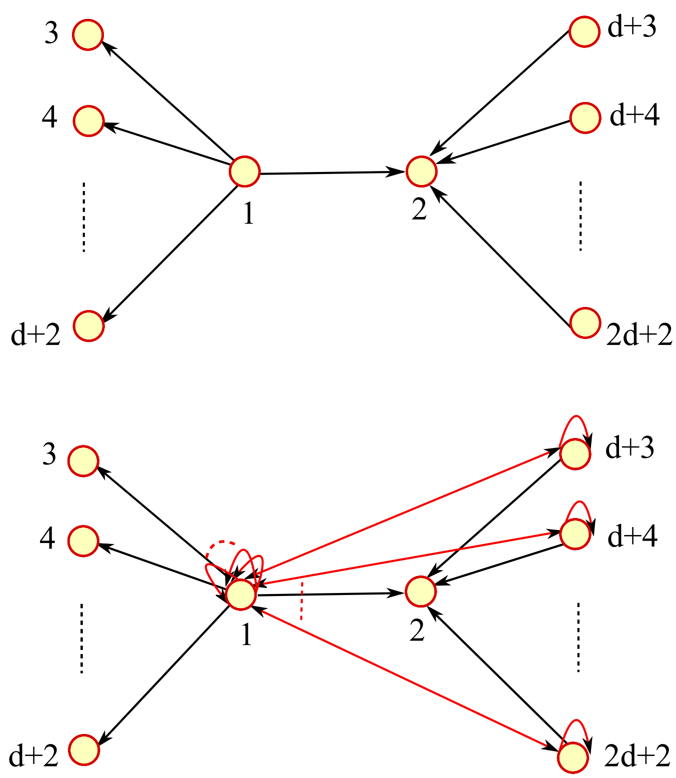

Figure 4-9: The top graph is G and the bottom is $\mathcal{G}$.

4.11 Appendix: Uncertainty in the direct utility

In section 4.3.2 we showed when uncertainty is in the network effect (i.e., incomplete information in the strength of interactions), the maximum distortion in the whole network is due to an agent that *she and her neighbors* **together** have high Bonacich centralities in the symmetrized network $G + G^T$, introducing the key agents. This is a novel feature for characterizing key agents in networks. In this appendix we show that when uncertainty is in the direct utility, then the structures of the first and second best allocations will become different than the main model where uncertainty is in the externality. Importantly, when uncertainty is in the direct utility, then the nature of consumptions and the overall distortion becomes *directly* related to Bonacich centrality in the symmetrized network, which is closely related to the previous studies (e.g., most notably Ballester, Calvo-Armengol and Zenou (2006*a*), Candogan, Bimpikis and Ozdaglar (2012*a*) and Bloch and Quérou (2013)).

Uncertainty in the direct utility In this case, the (ex post) utility of each agent is given by

$$u_i(\theta_i, x_i, x_{-i}, t_i) = \underbrace{\theta_i x_i - \frac{b}{2}x_i^2}_{\substack{\text{direct utility} \\ \text{(type dependent)}}} + \underbrace{x_i \sum_{j \neq i} g_{ij} x_j}_{\substack{\text{indirect utility} \\ \text{(network effect)}}} - \underbrace{t_i}_{\text{payment}}, \tag{11.1}$$

where x_i is the amount of the good she consumes, $x_{-i} \triangleq (x_1, x_2, \cdots, x_{i-1}, x_{i+1}, \cdots, x_n)$ is the consumption of other agents excluding agent i and t_i is the disutility charged for x_i by the firm. In addition, $\theta_i > 0$ measures the *intrinsic* marginal valuation for agent i. In the sequel, we call θ_i agent i's valuation factor, that is her type.

Assumption 4.11.1. *For each* $i = 1, 2, \cdots, n,\;\; b > \sum_{j \neq i}(g_{ij} + g_{ji})$.

We next characterize the first best and the second best solutions.

Proposition 4.11.1. *The first best and the second best allocations are as follows.*

(i) Under Assumption 4.11.1, the first best trade profile is given by:

$$\mathbf{x}^{FB}(\theta) = (bI - (G + G^T))^{-1} \boldsymbol{y}_{\boldsymbol{\theta}}, \tag{11.2}$$

for any $\theta \in [\underline{\theta}\ \bar{\theta}]^n$*, where* $\mathbf{y}_\theta \triangleq diag(\theta_1, \theta_2, \cdots, \theta_n)\mathbf{1}$.

(ii) Under Assumptions 4.11.1 and 4.3.2, the second best trade profile is given by:

$$\mathbf{x}^{SB}(\theta) = (bI - (G + G^T))^{-1} \mathbf{y}_{\psi,\theta}, \tag{11.3}$$

for any $\theta \in [\underline{\theta}\ \bar{\theta}]^n$*, where* $\mathbf{y}_{\psi,\theta} \triangleq diag(\psi(\theta_1), \psi(\theta_2), \cdots, \psi(\theta_n))\mathbf{1}$.

The proof of this result follows similar steps as in the proof of Proposition 4.3.1.

In contrast to Proposition 4.3.1, where uncertainty is in the externality, here θ and $\psi(\theta)$ appear *outside* of the matrix component of the answers (compare (11.2), (11.3) with Proposition 4.3.1). This is due to the location of uncertainty (i.e. which is in the direct utility or the externality), and, thus, is critical for the following result.

Given the definitions in 4.3.2 and 4.3.3, and the result in Proposition 4.11.1 we have the following result.

Proposition 4.11.2. *For any type profile* $\theta \in [\underline{\theta}\ \bar{\theta}]^n$*:*

(i) Distortion vector is given by:

$$\boldsymbol{d}(\theta) = (bI - (G + G^T))^{-1} \mathbf{y}_{\phi,\theta}, \tag{11.4}$$

where $\mathbf{y}_{\phi,\theta} \triangleq diag(\phi(\theta_1), \phi(\theta_2), \cdots, \phi(\theta_n))\mathbf{1}$.

(ii) Distortion is downward, *i.e.,* $\boldsymbol{d}(\theta) \geq \mathbf{0}$.

(iii) Let $[\boldsymbol{d}(\theta)]_i$ *denote the distortion on agent i's trade with regard to* θ. *Then*

$$[\boldsymbol{d}(\theta)]_i > 0.$$

if there exists at least one agent j *whose* $\theta_j \neq \bar{\theta}$.

(iv) Let $\theta_i \leq \bar{\theta}$ *and* $\theta_j = \bar{\theta}$, *for all* $j \neq i$. *Thus,* $\phi(\theta_j) = 0$, *for all* $j \neq i$. *Therefore,*

$$\begin{aligned}\mathcal{T}_i(\theta_i, \bar{\theta}_{-i}) &= \mathbf{1}^T \boldsymbol{d}(\theta_i, \bar{\theta}_{-i}) = (1/b)\left[\mathbf{1}^T \left(I - \frac{1}{b}(G + G^T)\right)^{-1}\right] \mathbf{y}_{\phi,\theta}\mathbf{1}, \\ &= (1/b)[k_1 \ k_2 \ \cdots \ k_n][0 \ 0 \ \cdots \phi(\theta_i) \ \cdots \ 0]^T \\ &= \left(\frac{\phi(\theta_i)}{b}\right) k_i,\end{aligned}$$

where k_i *is agent i's Bonacich centrality according to* $(I - \frac{1}{b}(G + G^T))^{-1}\mathbf{1}$.

The first three parts are intuitive. However, this last part of this result shows that the maximum distortion is due to an agent who has the highest Bonacich centrality, i.e., k_i, in $G+G^T$ which is closely related to the previously studies (e.g., Ballester, Calvo-Armengol and Zenou (2006*a*), Candogan, Bimpikis and Ozdaglar (2012*a*) and Bloch and Quérou (2013)). However, in section 4.3.2 we showed when uncertainty is in the externality, the maximum distortion in the whole network is due to an agent that *she and her neighbors* **together** have high Bonacich centralities in the symmetrized network $G + G^T$, introducing the key agents. This is a novel feature for characterizing key agents in networks which is due to the uncertainty in the externality (i.e., incomplete information in the strength of interactions).

Proof of Proposition 4.11.2. The first part of the proposition directly follows from the definition 4.3.2 and Proposition 4.11.1. To prove the second and third part, define $M \triangleq \left(I - \frac{1}{b}(G + G^T)\right)^{-1}$. Thus,

$$m_{ij} = \sum_{k=0}^{\infty} \frac{1}{b^k}(g_{ij} + g_{ji})^{[k]} \tag{11.5}$$

where m_{ij} is the ij entry of M and $(g_{ij} + g_{ji})^{[k]}$ is the ij enry of $(G_{ave})^k$. In other words, $(g_{ij} + g_{ji})^{[k]}$ is the total number of weighted walks of length k between agents i and j in G_{ave}, (see e.g., Jackson (2008)). Next, assume agent j's report is lower than the highest type, i.e., $\theta_j < \bar{\theta}$. Thus, $\phi(\theta_j) = \frac{1-F(\theta_j)}{f(\theta_j)} > 0$. Now, consider agent i, since G_{ave} is a connected network, thus agents i and j are connected through a path in G_{ave}. Further, as we shown in part (i), distortion vector is characterized as:

$$\mathbf{d}(\theta) = (1/b)\left(I - \frac{1}{b}(G + G^T)\right)^{-1} \ \mathbf{y}_{\phi,\theta}\mathbf{1},$$
$$= (1/b) \ M \ \mathbf{y}_{\phi,\theta}\mathbf{1},$$

where $\mathbf{y}_{\phi,\theta}\mathbf{1} = (\phi(\theta_1), \phi(\theta_2), \cdots, \phi(\theta_n))^T$ (a column vector). Since i and j are connected, $m_{ij} > 0$ (due to (11.5)), moreover, $\phi(\theta_j) > 0$ (due to $\theta_j < \bar{\theta}$), thus:

$$[\mathbf{d}(\theta)]_i \geq m_{ij}\phi(\theta_j) > 0. \tag{11.6}$$

The proof of the last part is immediate by the definition. ■

Finally, we wrap up this section with the following example.

Example 6 (Kite network)**.** Consider the following symmetric network, $g_{ij} = g_{ji} \in \{0, .1\}, i, j \in \{1, 2, \cdots, 5\}$, that captures the interrelations among the agents. Let $\theta_i \in \{\underline{\theta}, \bar{\theta}\}$ with $\text{Prob}\{\theta = \underline{\theta}\} = \frac{1}{2}$ (b is normalized to 1 and $\Delta\theta = 1$). The following table characterizes the distortion vector when only one agent is inefficient. As shown in the table, the distortion propagates throughout the network proportional to *Bonacich centrality*. To be precise, if distortion is due to say agent 1, agent 1 exerts the highest inefficacy in its allocation, and this affects others proportional to their Bonacich centrality as in (11.4) (notice that agent 2 is the central agent).

The third column of the table implies that overall distortion is proportional to Bonacich centrality of the agents in the network. Thus, distortion in agents 2's trade (the central agent) results in the highest overall distortion in the *whole* network[53]. Therefore, agent 2 is the key agent.

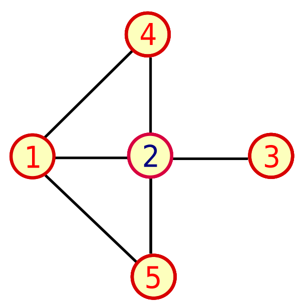

Figure 4-10: Kite network. Interconnection among the agents.

[53]Notice that $\frac{\phi(\theta)}{b} = 1$

Inefficient agent	Distortion vector	$\mathcal{T}(\theta_i, \bar{\theta}_{-i})$
1	(**1.2** **.38** .07 .31 .31)	2.27
2	(.38 **1.25** .25 .32 .32)	**2.52**
3	(.07 **.25** **1.05** .06 .06)	1.49
4	(.31 **.32** .06 **1.12** .12)	1.93
5	(.31 **.32** .06 .12 **1.12**)	1.93

Table 4.1: This table shows how distortion propagates throughout the network. Inefficiency in agent 2's trade (the central agent) creates the highest distortion in the whole network.

Bibliography

Abreu, Dilip, and Markus K. Brunnermeier. 2003. “Bubbles and Crashes.” *Econometrica*, 71(1): 173–204.

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. “Toward a theory of discounted repeated games with imperfect monitoring.” *Econometrica*, 1041–1063.

Abreu, Dilip, Paul Milgrom, and David Pearce. 1991. “Information and timing in repeated partnerships.” *Econometrica*, 1713–1733.

Acemoglu, Daron, Ali Kakhbod, and Asuman Ozdaglar. 2017. “Competition in Electricity Markets with Renewable Energy Sources.” *The Energy Journal*, 38(1).

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar. 2014. “Dynamics of Information Exchange in Endogenous Social Networks.” *Theoretical Economics*, 9(1): 41–97.

Acemoglu, D., M. Dahleh, I. Lobel, and A. Ozdaglar. 2011. “Bayesian learning in social networks.” *Review of Economic Studies*, 78: 1201–1236.

Acemoglu, D., V. Carvalho, A. Ozdaglar, and A Tahbaz-Salehi. 2012. “The Network Origins of Aggregate Fluctuations.” *Econometrica*, 80(5): 1977–2016.

Admati, Anant R., and Paul Pfleiderer. 1988. “A Theory of Intraday Patterns: Volume and Price Variability.” *Review of Financial Studies*, 1(1): 3–40.

Adrian, Tobias, Michael Fleming, Jonathan Goldberg, Morgan Lewis, Fabio Natalucci, and Jason Wu. 2013. “Dealer Balance Sheet Capacity and Market Liquidity during the 2013 Selloff in Fixed Income Markets.” *Working paper*.

Ahnert, Toni, and Ali Kakhbod. 2018. “Information Choice and Amplification of Financial Crises.” *Review of Financial Studies*, 30(6): 2130–2178.

Ajorlou, Amir, Ali Jadbabaei, and Ali Kakhbod. 2018. “Dynamic Pricing in Social Networks: The Word of Mouth Effect.” *Management Science*, 64: 971–979.

Allen, Franklin, and Douglas Gale. 1998. “Optimal Financial Crises.” *Journal of Finance*, 53(4): 1245–1284.

Amador, Manuel, and Pierre-Olivier Weill. 2010. "Learning from Prices: Public Communication and Welfare." *Journal of Political Economy*, 118: 866–907.

Ang, Andrew, Assaf A. Shtauber, and Paul C. Tetlock. 2013. "Asset Pricing in the Dark: The Cross-Section of OTC Stocks." *Review of Financial Studies*, 26(12): 2985–3028.

Angeletos, George-Marios, and Ivan Werning. 2006. "Crises and Prices: Information Aggregation, Multiplicity, and Volatility." *American Economic Review*, 96(5): 1720–1736.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2006. "Signaling in a Global Game: Coordination and Policy Traps." *Journal of Political Economy*, 114(3): 452–484.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2007. "Dynamic Global Games of Regime Change: Learning, Multiplicity and Timing of Attacks." *Econometrica*, 75(3): 711–756.

Ashcraft, Adam, and Darrell Duffie. 2007. "Systemic Illiquidity in the Federal Funds Market." *American Economic Review*, 97(2): 221–225.

Asquith, Paul, Thom Covert, and Parag Pathak. 2013. "The effects of mandatory transparency in financial market design: Evidence from the corporate bond market." *NBER Working paper 19417.*

Babus, Ana, and Cecilia Parlatore. 2017. "Strategic Fragmented Markets." *Working paper.*

Back, Kerry, and Hal Pedersen. 1998. "Continuous Auctions and Insider Trading." *Journal of Financial Markets*, 1(3): 385–402.

Back, Kerry, Ruomeng Liu, and Alberto Teguia. 2020. "Signaling in OTC Markets: Benefits and Costs of Transparency." *Journal of Financial and Quantitative Analysis*, 55(1): 47–75.

Ballester, C., A. Calvo-Armengol, and Y. Zenou. 2006*a*. "Who's who in networks. Wanted: The key player." *Econometrica*, 74: 1403–1417.

Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou. 2006*b*. "Who's who in networks. Wanted: The key player." *Econometrica*, 74(5): 1403–1417.

Banerjee, Abhijit, Arun Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science*, 341: DOI: 10.1126/science.1236498.

Banerjee, Snehal, Jesse Davis, and Naveen Gondhi. 2018. "When Transparency Improves, Must Prices Reflect Fundamentals Better?" *Review of Financial Studies*, 31(6): 2377–2414.

Bass, Frank. 1969. "A new product growth for model consumer durables." *Management Science*, 15(5): 215–227.

Bebchuk, Lucian A., and Itay Goldstein. 2011. "Self-fulfilling Credit Market Freezes." *Review of Financial Studies*, 24(11): 3519–3555.

Belhaj, M., S. Bervoets, and F. Deroian. 2016. "Efficient networks in games with local complementarities." *Theoretical Economics*, 11(1): 357–380.

Bergemann, Dirk, and Juuso Välimäki. 1997. "Market diffusion with two-sided learning." *RAND Journal of Economics*, 28(4): 773–795.

Bergemann, Dirk, and Juuso Välimäki. 2000. "Experimentation in markets." *Review of Economic Studies*, 67(2): 213–234.

Berger, Jonah, and Eric M. Schwartz. 2011. "What Drives Immediate and Ongoing Word of Mouth?" *Journal of Marketing Research*, 48(5): 869–880.

Bernstein, S., and E. Winter. 2012. "Contracting with Heterogeneous Externalities." *American Economic Journal: Microeconomics*, 4(2): 50–76.

Besbes, Omar, and Marco Scarsini. 2013. "On Information Distortions in Online Ratings." *Working paper.*

Bessembinder, Hendrik, and Kumar Venkataraman. 2004. "Does an electronic stock exchange need an upstairs market?" *Journal of Financial Economics*, 73(1): 3–36.

Bessembinder, Hendrik, and William Maxwell. 2008. "Markets: Transparency and the Corporate Bond Market." *Journal of Economic Perspectives*, 22(2): 217–234.

Billingsley, Patrick. 1995. *Probability and Measure.* Wiley Series in Probability and Mathematical Statistics.

Bimpikis, K., A. Ozdaglar, and E. Yildiz. 2016. "Competitive targeted advertising over networks." *Operation Research*, 64(3): 705–720.

Bloch, F. 2016. *Targeting and pricing in social networks. In: Y. Bramoullé, A. Galeotti and B. Rogers (Eds.).* Oxford: Oxford University Press:Oxford Handbook of the Economics of Networks.

Bloch, F., and N. Quérou. 2013. "Pricing in social networks." *Games and Economic Behavior*, 80: 263–281.

Board, Simon. 2008. "Durable-Goods Monopoly with Varying Demand." *Review of Economic Studies*, 75(2): 391–413.

Bollobás, Béla. 2001. *Random Graphs.* Academic Press.

Bosshardt, Joshua, and Ali Kakhbod. 2020*a*. "The Effect of Liquidity Regulation on U.S. Bank Risk-Taking: Evidence from Non-Performing Loans and CDS Spreads." *Working paper.*

Bosshardt, Joshua, and Ali Kakhbod. 2020*b*. "Liquidity Regulation, Bank Complexity and Policy." *Working paper.*

Bosshardt, Joshua, and Ali Kakhbod. 2020*c*. "Managerial Protections, Capital Requirements, and Bank Lending." *Working paper.*

Bosshardt, Joshua, and Ali Kakhbod. 2020*d*. "Why did Firms Draw Down their Credit Lines during the COVID-19 Shutdown?" *Working paper.*

Bramoulle, Y., and R. Kranton. 2007*a*. "Public goods in networks." *Journal of Economic Theory*, 135: 478–494.

Bramoullé, Yann, and Rachel Kranton. 2007*b*. "Public goods in networks." *Journal of Economic Theory*, 135(1): 478–494.

Bramoullé, Y., R. Kranton, and M. D'Amours. 2014. "Strategic Interaction and Networks." *American Economic Review*, 104(3): 898–930.

Bromiley, Paul A. 2003. "Products and convolutions of Gaussian distributions." *TINA Vision Memo 2003-003, Manchester, United Kingdom.*

Brunnermeier, Markus K, and Lasse H Pedersen. 2005. "Predatory Trading." *The Journal of Finance*, 60(4): 1825–1863.

Bryant, John. 1980. "A model of reserves, bank runs, and deposit insurance." *Journal of Banking and Finance*, 4(4): 335–344.

Caballero, Ricardo J. 2010. "Macroeconomics after the Crisis: Time to Deal with the Pretense-of-Knowledge Syndrome." *Journal of Economic Perspectives*, 24(4): 85–102.

Caballero, Ricardo J, and Alp Simsek. 2013. "Fire Sales in a Model of Complexity." *The Journal of Finance*, 68(6): 2549–2587.

Caballero, Ricardo J, and Arvind Krishnamurthy. 2008. "Collective Risk Management in a Flight to Quality Episode." *The Journal of Finance*, 63(5): 2195–2230.

Cabral, Luis M.B., David J. Salant, and Glenn A. Woroch. 1999. "Monopoly pricing with network externalities." *International Journal of Industrial Organization*, 17(2): 199–214.

Campbell, Arthur. 2012. "Word of mouth model of sales." *Working paper.*

Campbell, Arthur, Dina Mayzlin, and Jiwoong Shin. 2013. "Managing Buzz." *Working paper.*

Candogan, O., K. Bimpikis, and A. Ozdaglar. 2012*a*. "Optimal Pricing in Networks with Externalities." *Operation Research*, 60(4): 883–905.

Candogan, Ozan, Kostas Bimpikis, and Asuman Ozdaglar. 2012*b*. "Optimal Pricing in Networks with Externalities." *Operations Research*, 60(4): 883–905.

Carapella, Francesca, and Stephen Williamson. 2015. "Credit Markets, Limited Commitment, and Government Debt." *Review of Economic Studies*, 82(3): 963–990.

Carlsson, Hans, and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica*, 61(5): 989–1018.

Çelebi, Oğuzhan, Ali Kakhbod, and Kevin K. Li. 2020. "Speed Competition and Segmentation in Illiquid Markets." *Working paper.*

Cetorelli, Nicola, and Linda S. Goldberg. 2014. "Measures of Global Bank Complexity." *FRBNY Economic Policy Review / December 2014.*

Chari, V V, and Ravi Jagannathan. 1988. " Banking Panics, Information, and Rational Expectations Equilibrium." *Journal of Finance*, 43(3): 749–61.

Chau, Minh, and Dimitri Vayanos. 2008. "Strong-Form Efficiency with Monopolistic Insiders." *Review of Financial Studies*, 21: 2275–2306.

Chen, Y., Y. Zenou, and J. Zhou. 2018. "Competitive Pricing Strategies in Social Networks." *RAND Journal of Economics.*

Chevalier, Judith, and Dina Mayzlin. 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research*, 43(3): 345–354.

Coase, Ronald. 1972. "Durability and monopoly." *Journal of Law and Economics*, 15: 143–149.

Colombo, Luca, Gianluca Femminis, and Alessandro Pavan. 2014. "Information Acquisition and Welfare." *The Review of Economic Studies*, forthcoming.

Conlisk, John, Eitan Gerstner, and Joel Sobel. 1984. "Cyclic Pricing by a Durable Goods Monopolist." *Quarterly Journal of Economics*, 99(3): 489–505.

Corbo, J., A. Calvo-Armengol, and D. Parkes. 2007. "The importance of network topology in local contribution games." *Proceedings of the 3rd international Workshop on Internet and Network Economics.*

Corless, Robert M., Gaston H. Gonnet, David E. G. Hare, David J. Jeffrey, and Donald E. Knuth. 1996. "On the Lambert W function." *Advances in Computational Mathematics*, 5(1): 329–359.

Corsetti, Giancarlo, Amil Dasgupta, Stephen Morris, and Hyun Song Shin. 2004. "Does One Soros Make a Difference? A Theory of Currency Crisis with Large and Small Traders." *Review of Economic Studies*, 71(1): 87–113.

Corsetti, Giancarlo, Bernardo Guimaraes, and Nouriel Roubini. 2006. "International lending of last resort and moral hazard: A model of IMF's catalytic finance." *Journal of Monetary Economics*, 53(3): 441–471.

Csorba, G. 2008. "Screening contracts in the presence of positive network effects." *International Journal of Industrial Organization*, 26(1): 213–226.

Dang, Tri Vi, Gary Gorton, and Bengt Holmström. 2012. "Ignorance, Debt and Financial Crises." *Mimeo.*

Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam. 1998. "Investor psychology and security market under- and over-reactions." *Journal of Finance*, 53(6): 1839–85.

Dasgupta, Amil. 2007. "Coordination and Delay in Global Games." *Journal of Economic Theory*, 134: 195–225.

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. "In Search of Attention." *The Journal of Finance*, 66(5): 1461–1499.

DeGroot, Morris H. 1970. *Optimal Statistical Decisions.* John Wiley and Sons (Wiley Classics Library Edition).

Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank runs, deposit insurance and liquidity." *Journal of Political Economy*, 91: 401–419.

Dichter, Ernest. 1966. "How Word-of-Mouth Advertising Works." *Harvard Business Review*, 44(6): 147–166.

Duffie, Darrell. 2012. *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets.* Princeton University Press.

Duffie, Darrell. 2017. "Financial Regulatory Reform After the Crisis: An Assessment." *Management Science.*

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen. 2005. "Over-the-Counter Markets." *Econometrica*, 73(6): 1815–1847.

Easley, David, and Maureen O'Hara. 1987. "Price, trade size, and information in securities markets." *Journal of Financial Economics*, 19(1): 69–90.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar. 2007. "Corporate Bond Market Transaction Costs and Transparency." *Review of Financial Studies*, 62(3): 1421–1451.

Ellison, Glenn. 1994. "Cooperation in the prisoner's dilemma with anonymous random matching." *The Review of Economic Studies*, 61(3): 567–588.

Erdős, Paul, and Alfred Rényi. 1959. "On Random Graphs." *Publicationes Mathematicae*, 6: 290–297.

Erdős, Paul, and Alfred Rényi. 1960. "On the Evolution of Random Graphs." *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5: 17–61.

Erdős, Paul, and Alfred Rényi. 1961. "On the Strength of Connectedness of a Random Graph." *Acta Mathematica Scienta Hungary*, 12(1-2): 261–267.

Fainmesser, I.P., and A. Galeotti. 2016*a*. "Pricing Network Effects." *Review of Economic Studies*, 83(1): 165–198.

Fainmesser, I.P., and A. Galeotti. 2016*b*. "Pricing Network Effects: Competition." *Working paper.*

Fama, Eugene F. 1998. "Market efficiency, long-term returns and behavioral finance." *Journal of Financial Economics*, 49(3): 283–306.

Farrell, J., and G. Saloner. 1985. "Standardization, compatibility, and innovation." *RAND Journal of Economics*, 16(1): 70–83.

Financial Crisis Inquiry Report. 2011. "National Commission on the Causes of the Financial and Economic Crisis in the United States." Final Report.

Frankel, David, Stephen Morris, and Ady Pauzner. 2003. "Equilibrium selection in global games with strategic complementarities." *Journal of Economic Theory*, 108(1): 1–44.

Fudenberg, D, D Levine, and Eric Maskin. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica*, 62(5).

Fudenberg, Drew, and David K Levine. 1994. "Efficiency and observability with long-run and short-run players." *Journal of Economic Theory*, 62(1): 103–135.

Galeotti, A., and S. Goyal. 2009*a*. "Influencing the influencers: a theory of strategic diffusion." *RAND Journal of Economics*, 40(3): 509–532.

Galeotti, Andrea, and Andrea Mattozzi. 2011. "Personal Influence: Social Context and Political Competition." *American Economic Journal: Microeconomics*, 3(1): 307–327.

Galeotti, Andrea, and Sanjeev Goyal. 2009*b*. "Influencing the influencers: a theory of strategic diffusion." *RAND Journal of Economics*, 40(3): 509–532.

Galeotti, Andrea, Sanjeev Goyal, Matthew O. Jackson, Fernando Vega-Redondo, and Leeat Yariv. 2010. "Network games." *Review of Economic Studies*, 77(1): 218–244.

Gallego, Guillermo, and Garrett van Ryzin. 1994. "Optimal dynamic pricing of inventories with stochastic demand over finite horizons." *Management Science*, 40(8): 999–1020.

Garrett, Daniel F. 2013. "Incoming demand with private uncertainty." *Working paper*.

Garriott, Corey, and Adrian Walton. 2016. "Retail Order Flow Segmentation." *Bank of Canada Staff Working Paper 2016-20*.

Gershkov, Alex, and Benny Moldovanu. 2009. "Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach." *American Economic Journal: Microeconomics*, 1(2): 168–198.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders." *Journal of Financial Economics*, 14(1): 71–100.

Godes, David, and Dina Mayzlin. 2004. "Using Online Conversations to Study Word of Mouth Communication." *Marketing Science*, 23(4): 545–560.

Goldberg, Linda S. 2016. "Cross-Border Banking Flows and Organizational Complexity in Financial Conglomerates." *Mimeo*.

Goldenberg, Jacob, Sangman Han, Donald R. Lehmann, and Jae Weon Hong. 2009. "The Role of Hubs in the Adoption Process." *Journal of Marketing*, 73: 1–13.

Goldstein, Itay. 2012. "Empirical Literature on Financial Crises: Fundamentals vs. Panic." In *The Evidence and Impact of Financial Globalization.* , ed. Gerard Caprio. Amsterdam:Elsevier.

Goldstein, Itay, and Ady Pauzner. 2005. "Demand Deposit Contracts and The Probability of Bank Runs." *Journal of Finance*, 60(3): 1293–1327.

Goldstein, Itay, and Haresh Sapra. 2013. "Should Banks' Stress Test Results be Disclosed? An Analysis of the Costs and Benefis." *Foundations and Trends in Finance*, 8: 1–54.

Goldstein, Itay, and Liyan Yang. 2017. "Information Disclosure in Financial Markets." *Annual Review of Financial Economics*, 9(101): 101–125.

Goldstein, Itay, and Liyan Yang. 2018. "Good Disclosure, Bad Disclosure." *Journal of Financial Economics*, forthcoming.

Goldstein, Itay, and Philip Bond. 2015. "Government Intervention and Information Aggregation by Prices." *Journal of Finance*, 70(6): 2777–2811.

Goldstein, Michael A., Edith S. Hotchkiss, and Erik R. Sirri. 2007. "Transparency and Liquidity: A Controlled Experiment on Corporate Bonds." *Review of Financial Studies*, 20(2): 235–273.

Golub, B., and M.O. Jackson. 2010. “Naive learning in social networks: Convergence, influence, and the wisdom of crowds.” *American Economic Journal: Microeconomics*, 2: 112–149.

Golub, B., and M.O. Jackson. 2012. “How homophily affect the speed of learning and best response dynamics.” *Quarterly Journal of Economics*, 127: 1287–1338.

Gorton, Gary, and George Pennacchi. 1990. “Financial Intermediaries and Liquidity Creation.” *Journal of Finance*, 45(1): 49–71.

Gorton, Gary, and Guillermo Ordonez. 2014. “Collateral Crises.” *American Economic Review*, 104(2): 343–378.

Green, Richard, Burton Hollifield, and Norman Schürhoff. 2007*a*. “Dealer intermediation and price behavior in the aftermarket for new bond issues.” *Journal of Financial Economics*, 86(3): 643–682.

Green, Richard, Burton Hollifield, and Norman Schürhoff. 2007*b*. “Financial Intermediation and the Costs of Trading in an Opaque Market.” *Review of Financial Studies*, 20(2): 275–314.

Green, Richard, Dan Li, and Norman Schürhoff. 2010. “Price Discovery in Illiquid Markets: Do Financial Asset Prices Rise Faster Than They Fall?” *Journal of Finance*, 65(5): 1669–1702.

Grossman, Sanford J. 1992. “The Informational Role of Upstairs and Downstairs Trading.” *The Journal of Business*, 65(4).

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. “On the Impossibility of Informationally Efficient Markets.” *American Economic Review*, 70(3): 393–408.

Grossman, Sanford J., and Merton H. Miller. 1988. “Liquidity and Market Structure.” *Journal of Finance*, 43(3): 735–777.

Guerrieri, Veronica, and Guido Lorenzoni. 2009. “Liquidity and Trading Dynamics.” *Econometrica*, 77(6): 1751–1790.

Guerrieri, Veronica, and Robert Shimer. 2014. “Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales, and Flight to Quality.” *American Economic Review*, 104(7): 1875–1908.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson. 1981. “Foundations of Dynamic Monopoly and the Coase Conjecture.” *Bell Journal of Economics*, 12: 112–128.

Hanson, Ward, and Kipp Martin. 1990. “Optimal Bundle Pricing.” *Management Science*, 36(2): 155–174.

Hartline, Jason, Vahab Mirrokni, and Mukund Sundararajan. 2008. "Optimal marketing strategies over social networks." *Proc. 17th Internat. Conf. World Wide Web (WWW '08) (ACM, New York)*, 189–198.

Hartman-Glaser, Barney. 2017. "Reputation and signaling in asset sales." *Journal of Financial Economics*, 125(2): 245–265.

Hellwig, Christian, and Laura Veldkamp. 2009. "Knowing what others know: Coordination motives in information acquisition." *Review of Economic Studies*, 76: 223–251.

Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski. 2006. "Self-Fulfilling Currency Crises: The Role of Interest Rates." *American Economic Review*, 96: 1769–1787.

Herr, Paul M., Frank R. Kardes, and John Kim. 1991. "Effects of Word-of-Mouth and Product-Attribute Information on Persuasion: An Accessibility-Diagnosticity Perspective." *Journal of Marketing Research*, 17(4): 454–462.

Hervas-Drane, Andres. 2015. "Recommended for you: The effect of word of mouth on sales concentration." *Intern. J. of Research in Marketing*, 32(15): 207–218.

He, Zhiguo, and Asaf Manela. 2016. "Information Acquisition in Rumor-Based Bank Runs." *The Journal of Finance.*

Hong, Harrison, Terence Lim, and Jeremy Stein. 2000. "Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies." *Journal of Finance*, 55(1): 265–295.

Hörner, Johannes, and Larry Samuelson. 2011. "Managing Strategic Buyers." *Journal of Political Economy*, 119(3): 379–425.

Horn, R., and C. Johnson. 2012. *Matrix Analysis.* Cambridge University Press.

Iachan, Felipe S., and Plamen T. Nenov. 2015. "Information Quality and Crises in Regime-Change Games." *Journal of Economic Theory*, 158, Part B: 739–768.

Ifrach, Bar, Costis Maglaras, and Marco Scarsini. 2011. "Monopoly Pricing in the Presence of Social Learning." *Working paper.*

Ifrach, Bar, Costis Maglaras, and Marco Scarsini. 2013. "Bayesian Social Learning from Consumer Reviews." *Working paper.*

Jacklin, Charles J, and Sudipto Bhattacharya. 1988. "Distinguishing Panics and Information-Based Bank Runs: Welfare and Policy Implications." *Journal of Political Economy*, 96(3): 568–592.

Jackson, M. 2008. *Social and Economic Networks.* Princeton, NJ.:Princeton University Press.

Jackson, M., and Y. Zenou. 2015. *Games on networks. In P. Young and S. Zamir (Eds.).* Amsterdam: Elsevier Publisher:Handbook of Game Theory, Vol. 4.

Jackson, Matthew O., and Brian W. Rogers. 2007*a*. "Meeting Strangers and Friends of Friends: How Random Are Social Networks?" *American Economic Review*, 97(3): 890 – 915.

Jackson, Matthew O., and Brian W. Rogers. 2007*b*. "Relating Network Structure to Diffusion Properties through Stochastic Dominance." *B. E. Journal of Theoretical Economics: Advances in Theoretical Economics*, 7(1): 1 – 13.

Jackson, Matthew O., and Leeat Yariv. 2007. "Diffusion of Behavior and Equilibrium Properties in Network Games." *American Economic Review*, 97(2): 92 – 98.

Jadbabaei, Ali, and Ali Kakhbod. 2019. "Optimal Contracting in Networks." *Journal of Economic Theory*, 183(9): 1094–1153.

JJ, Laffont., and J. Tirole. 1990. "Optimal Bypass and Cream Skimming." *The American Economic Review*, 80: 1040–1061.

Kacperczyk, Marcin, and Philipp Schnabl. 2010. "When Safe Proved Risky: Commercial Paper during the Financial Crisis of 2007-2009." *Journal of Economic Perspectives*, 24(1): 29–50.

Kakhbod, Ali. 2013*a*. "Pay-to-bid auctions: To bid or not to bid." *Operations Research Letters*, 41(5): 462–467.

Kakhbod, Ali. 2013*b*. *Resource Allocation in Decentralized Systems with Strategic Agents: An Implementation Theory Approach.* Springer.

Kakhbod, Ali, and Demosthenis Teneketzis. 2012*a*. "An Efficient Game Form for Multi-Rate Multicast Service Provisioning." *IEEE Journal on Selected Areas in Communications*, 30(11): 2093–2104.

Kakhbod, Ali, and Demosthenis Teneketzis. 2012*b*. "An Efficient Game Form for Unicast Service Provisioning." *IEEE Transactions on Automatic Control*, 57(2): 392–404.

Kakhbod, Ali, and Demosthenis Teneketzis. 2012*c*. "Power Allocation and Spectrum Sharing in Multi-User, Multi-Channel Systems With Strategic Users." *IEEE Transactions on Automatic Control*, 57(9): 2338–2342.

Kakhbod, Ali, and Fei Song. 2020*a*. "Dynamic price discovery: Transparency vs. information design." *Games and Economic Behavior*, 122(7): 203–232.

Kakhbod, Ali, and Fei Song. 2020*b*. "Liquidity, Welfare and Transparency in Dynamic OTC markets." *Working paper.*

Kakhbod, Ali, and Giacomo Lanzani. 2020. "Market Power and Asymmetric Learning in Product Markets." *Working paper.*

Kakhbod, Ali, and Morteza Zadimoghaddam. 2011. "On the construction of prefix-free and fix-free codes with specified codeword compositions." *Discrete Applied Mathematics*, 159: 2269–2275.

Kakhbod, Ali, and S. M. Sadegh Tabatabaei Yazdi. 2010. "On describing the routing capacity regions of networks." *Mathematical Methods of Operations Research*, 72: 95–106.

Kakhbod, Ali, and Uliana Loginova. 2018. "When does introducing verifiable communication choices improve welfare?" *Working paper.*

Kakhbod, Ali, Ashutosh Nayyar, Shruti Sharma, and Demosthenis Teneketzis. 2014. "Power allocation and spectrum sharing in wireless networks: an implementation theory approach." In *Mechanisms and Games for Dynamic Spectrum Allocation.* , ed. M. Honig H. Vincent Poor T. Alpcan, H. Boche, 108–144. Cambridge University Press.

Kakhbod, Ali, Asuman Ozdaglar, and Ian Schneider. 2021. "Selling wind." *The Energy Journal*, 42(1).

Kakhbod, Ali, Uliana Loginova, Andrey Malenko, and Nadya Malenko. 2018. "Advising the Management." *Working paper.*

Kandori, Michihiro. 1992. "The use of information in repeated games with imperfect monitoring." *The Review of Economic Studies*, 59(3): 581–593.

Katona, Zsolt, Peter Zubcsek, and Miklos Sarvary. 2011. "Network Effects and Personal Influences: The Diffusion of an Online Social Network." *Journal of Marketing Research*, 48(3): 425–443.

Katz, Elihu, and Paul Lazarsfeld. 1955. *Personal Influence: The Part Played by People in the Flow of Mass Communications.* Free Press, Glencoe, IL.

Katz, M., and C. Shapiro. 1986. "Technology adoption in the presence of network externalities." *Journal of Political Economy*, 94(4): 822–841.

Kendall, Chad. 2015. "The Time Cost of Information in Financial Markets." *Mimeo.*

Kisgen, Darren J. 2006. "Credit Ratings and Capital Structure." *Journal of Finance*, 61(3): 1035–1072.

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy*, 105: 211–248.

Kothari, S. P., and Jerold B. Warner. 1997. "Measuring long-horizon security price performance." *Journal of Financial Economics*, 43(3): 301–309.

Krugman, Paul. 1979. "A model of balance-of-payments crises." *Journal of Money, Credit and Banking*, 11(3): 311–25.

Kyle, Albert S. 1985*a*. "Continuous auctions and insider trading." *Econometrica*, 1315–1335.

Kyle, Albert S. 1985*b*. "Long-lived information and intraday patterns." *Econometrica*, 53(6): 1315–1335.

Laffont, J., and D. Martimort. 2002. *The Theory of Incentives: The Principal-Agent Model.* Princeton University Press.

Lagos, Ricardo, and Guillaume Rocheteau. 2009. "Liquidity in Asset Markets with Search Frictions." *Econometrica*, 77(2): 403–26.

Lang, Mark H., and Russell J. Lundholm. 1996. "Corporate Disclosure Policy and Analyst Behavior." *Accounting Review*, 71(4): 467–492.

Larson, Nathan. 2013. "Dynamic Viral Marketing on a Social Network." *Working paper.*

Leduc, M. V., M. O. Jackson, and R. Johari. 2017. "Pricing and referrals in diffusion on networks." *Games and Economic Behavior*, 104: 568–594.

Lester, Benjamin, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones. 2018. "Market-making with Search and Information Frictions." *Working paper.*

Linnainmaa, Juhani T., and Gideon Saar. 2012. "Lack of anonymity and the inference from order flow." *Review of Financial Studies*, 25(5): 1414–1456.

López-Pintado, Dunia. 2008. "Diffusion in Complex Social Networks." *Games and Economic Behavior*, 62(2): 573 – 590.

Lovász, L. 2007. "Eigenvalues of graphs."

Mailath, George, Masahiro Okuno-Fujiwara, and Andrew Postlewaite. 1993. "Belief-Based Refinements in Signaling Games." *Journal of Economic Theory*, 60: 241–276.

Malamud, Semyon, and Marzena Rostek. 2017. "Decentralized Exchange." *American Economic Review*, 107(11): 3320–3362.

Maskin, E., and J. Riley. 1984. "Monopoly with Incomplete Information." *Rand Journal of Economics*, 15: 171–196.

Massa, Massimo, and Andrei Simonov. 2003. "Reputation and interdealer trading: A microstructure analysis of the treasury bond market." *Journal of Financial Markets*, 6(2).

McCloughan, Kaitlin. 2013. "New report on app developers' attitudes towards mobile advertising." *Available at: http://appflood.com/blog/app-marketing-report-2013.html.*

Metz, Christina E. 2002. "Private and Public Information in Self-fulfilling Currency Crises." *Journal of Economics*, 76(1): 65–85.

Mohajer, Soheil, and Ali Kakhbod. 2011. "Anti-uniform Huffman codes." *IET Communications*, 5(9): 1213–1219.

Mohajer, Soheil, Payam Pakzad, and Ali Kakhbod. 2012. "Tight Bounds on the Redundancy of Huffman Codes." *IEEE Transactions on Information Theory*, 58(11): 6737–6746.

Morris, Stephen, and Hyun S. Shin. 1998. "Unique equilibrium in a model of self-fulfilling currency attacks." *American Economic Review*, 88(3): 587–597.

Morris, Stephen, and Hyun S. Shin. 2003. "Global games: theory and applications." Cambridge University Press.

Morris, Stephen, and Hyun S. Shin. 2004. "Coordination Risk and the Price of Debt." *European Economic Review*, 48: 133–153.

Mussa, M., and C. Rosen. 1978. "Monopoly and Product Quality." *Journal of Economic Theory*, 18: 301–317.

Myatt, David, and Chris Wallace. 2012. "Endogenous Information Acquisition in Coordination Games." *Review of Economic Studies*, 79(1): 340–374.

Nagel, Joachim. 2016. "Markets Committee Electronic trading in fixed income markets." *Technical Report*.

Newman, Mark E., Steven H. Strogatz, and Duncan J. Watts. 2001. "Random Graphs with Arbitrary Degree Distributions and their Applications." *Physical Review E*, 64(2): 026118.

Nikitin, Maxim, and R. Todd Smith. 2008. "Information acquisition, coordination, and fundamentals in a financial crisis." *Journal of Banking and Finance*, 32(6): 907–914.

Nocke, Volker, and Martin Peitz. 2007. "A Theory of Clearance Sales." *Economic Journal*, 117(522): 964–990.

Oberfield, Ezra. 2012. "Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture." *Working paper*.

Obstfeld, Maurice. 1996. "Models of currency crises with self-fulfilling features." *European Economic Review*, 40(3-5): 1037–47.

Papanastasiou, Yiangos, Nitin Bakhshi, and Nicos Savva. 2013. "Social Learning from Early Buyer Reviews: Implications for New Product Launch." *Working paper*.

Parker, Geoffrey G., and Marshall W. Van Alstyne. 2005. "Two-Sided Network Effects: A Theory of Information Product Design." *Management Science*, 51(10): 1494–1504.

Phillips, Robert. 2005. *Pricing and Revenue Optimization.* Stanford Business Books.

Probst, F., L. Grosswiele, and R. Pfleger. 2013. "Who will lead and who will follow: Identifying Influential Users in Online Social Networks." *Business and Information Systems Engineering*, 5(3): 179–193.

Radner, Roy, Ami Radunskaya, and Arun Sundararajan. 2013. "Dynamic pricing of network goods with boundedly rational consumers." *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 111(1): 99–104.

Rapoport, Anatol. 1957. "A Contribution to the Theory of Random and Biased Nets." *Bulletin of Mathematical Biophysics*, 19(4): 257–277.

Reuters. 2011. "U.S. money funds diverge on European bank paper." *July 5, 2011.*

Rochet, Jean-Charles, and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the European Economic Association*, 1(4): 990–1029.

Rochet, Jean-Charles, and Xavier Vives. 2004. "Coordination failures and the lender of last resort: was Bagehot right after all?" *Journal of the European Economic Association*, 2(6): 1116–1147.

Satariano, Adam. 2014. "WhatsApp's Founder Goes From Food Stamps to Billionaire." *Bloomberg, available at: http://www.bloomberg.com/news/2014-02-20/whatsapp-s-founder-goes-from-food-stamps-to-billionaire.html.*

Segal, I. 1999. "Contracting with Externalities." *Quarterly Journal of Economics*, 114: 337–388.

Segal, I. 2003. "Coordination and Discrimination in Contracting with Externalities: Divide and Conquer." *Journal of Economic Theory*, 113: 147–181.

Seppi, Duane J. 1990. "Equilibrium Block Trading and Asymmetric Information." *Journal of Finance*, 45(1): 73–94.

Shimer, Robert, and Iván Werning. 2019. "Efficiency and information transmission in bilateral trading." *Review of Economic Dynamics*, 33: 154–176.

Shin, E. 2017. "Monopoly pricing and diffusion of social network goods." *Games and Economic Behavior*, 102: 162–178.

Shleifer, Andrei, and Robert W. Vishny. 1992. "Liquidation Values and Debt Capacity: A Market Equilibrium Approach." *The Journal of Finance*, 47(4): 1343–1366.

Shleifer, Andrei, and Robert W. Vishny. 1997. "The Limits of Arbitrage." *The Journal of Finance*, 52(1): 35–55.

Sobel, Joel. 1991. "Durable Goods Monopoly with Entry of New Consumers." *Econometrica*, 59(5): 1455–1485.

Stephen, Andrew T., and Olvier Toubiaz. 2010. "Deriving Value from Social Commerce Networks." *Journal of Marketing Research*, 47(2): 215–228.

Stokey, Nancy L. 1979. "Intertemporal Price Discrimination." *Quarterly Journal of Economics*, 93(3): 355–371.

Stokey, Nancy L. 1981. "Rational expectations and durable goods pricing." *Bell Journal of Economics*, 12: 112–128.

Sundararajan, A. 2008*a*. "Local network effects and complex network structure." *The B.E. Journal of Theoretical Economics*, 7(1).

Sundararajan, Arun. 2008*b*. "Local network effects and complex network structure." *The B.E. Journal of Theoretical Economics*, 7(1): 1935–1704.

Su, Xuanming. 2007. "Intertemporal Pricing with Strategic Customer Behavior." *Management Science*, 53(5): 726–741.

Szkup, Michal, and Isabel Trevino. 2015. "Information Acquisition and Transparency in Global Games." *Journal of Economic Theory*, 160: 387–428.

Talluri, Kalyan T., and Garrett van Ryzin. 2004. *The Theory and Practice of Revenue Management (International Series in Operations Research & Management Science).* Springer.

Trusov, Michael, Randolph E. Bucklin, and Koen Pauwels. 2008. "Effects of Word-of-Mouth Versus Traditional Marketing: Findings from an Internet Social Networking Site." *Journal of Marketing*, 73: 90–102.

Tuttle, Laura. 2014. "OTC Trading: Description of Non-ATS OTC Trading in National Market System Stocks." *Technical Report*.

Verrecchia, R. E. 2001. "Essays on Disclosure." *Journal of Accounting and Economics*, 32: 97–180.

Vives, Xavier. 2005. "Complementarities and Games: New Developments." *Journal of Economic Literature*, 43: 437–479.

Vives, Xavier. 2014. "Strategic Complementarity, Fragility, and Regulation." *Review of Financial Studies*, 27(12): 3547–3592.

Wang, Jiang. 1993. "A Model of Intertemporal Asset Prices under Asymmetric Information." *Review of Economic Studies*, 102: 249–282.

Wang, Jiang. 1994. "A Model of Competitive Stock Trading Volume." *Journal of Political Economy*, 102: 127–168.

Watts, Duncan J. 2002. "A Simple Model of Cascades on Random Networks." *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 99(9): 5766–5771.

Watts, Duncan J., and Peter S. Dodds. 2007. "Influentials, Networks, and Public Opinion Formation." *Journal of Consumer Research*, 34(2): 441–458.

Watts, Duncan J., and Steven H. Strogatz. 1998. "Collective Dynamics of Small World Networks." *Nature*, 393: 440–442.

Yang, Ming. 2015. "Coordination with Flexible Information Acquisition." *Journal of Economic Theory*, 158: 721–738.

Young, Peyton H. 2009. "Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning." *American Economic Review*, 99(5): 1899 – 924.

Yu, Man, Laurens Debo, and Roman Kapuscinski. 2013. "Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods." *Working paper.*

Zhou, J., and Y. Chen. 2015. "Key leaders in social networks." *Journal of Economic Theory*, 157: 212–235.

Zhu, Haoxiang. 2012. "Finding a Good Price in Opaque Over-the-Counter Markets." *Review of Financial Studies*, 25(4): 1255–1285.

Zhu, Haoxiang. 2014. "Do Dark Pools Harm Price Discovery?" *Review of Financial Studies*, 27(3): 747–789.