Essays on Informational Frictions in Macroeconomics and Finance

by

Jennifer La'O

S.B., Massachusetts Institute of Technology (2005)

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 2010

Signature of Author ....
Department of Economics
15 May 2010

Certified by ....
George-Marios Angeletos
Professor of Economics
Thesis Supervisor

Certified by ....
Mikhail Golosov
Professor of Economics
Thesis Supervisor

Accepted by ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Chairman, Departmental Committee on Graduate Studies

Essays on Informational Frictions in Macroeconomics and Finance

by

Jennifer La'O

Submitted to the Department of Economics
on 15 May 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This dissertation consists of four chapters analyzing the effects of heterogeneous and asymmetric information in macroeconomic and financial settings, with an emphasis on short-run fluctuations. Within these chapters, I study the implications these informational frictions may have for the behavior of firms and financial institutions over the business cyle and during crises episodes.

The first chapter examines how collateral constraints on firm-level investment introduce a powerful two-way feedback between the financial market and the real economy. On one hand, real economic activity forms the basis for asset dividends. On the other hand, asset prices affect collateral value, which in turn determines the ability of firms to invest. In this chapter I show how this two-way feedback can generate significant expectations-driven fluctuations in asset prices and macroeconomic outcomes when information is dispersed. In particular, I study the implications of this two-way feedback within a micro-founded business-cycle economy in which agents are imperfectly, and heterogeneously, informed about the underlying economic fundamentals. I then show how tighter collateral constraints mitigate the impact of productivity shocks on equilibrium output and asset prices, but amplify the impact of "noise", by which I mean common errors in expectations. Noise can thus be an important source of asset-price volatility and business-cycle fluctuations when collateral constraints are tight.

The second chapter is based on joint work with George-Marios Angeletos. In this chapter we investigate a real-business-cycle economy that features dispersed information about underlying aggregate productivity shocks, taste shocks, and—potentially—shocks to monopoly power. We show how the dispersion of information can (i) contribute to significant inertia in the response of macroeconomic outcomes to such shocks; (ii) induce a negative short-run response of employment to productivity shocks; (iii) imply that productivity shocks explain only a small fraction of high-frequency fluctuations; (iv) contribute to significant noise in the business cycle; (v) formalize a certain type of demand shocks within an RBC economy; and (vi) generate cyclical variation in observed Solow residuals and labor wedges. Importantly, none of these properties requires significant uncertainty about the underlying fundamentals: they rest on the heterogeneity of information and the strength of trade linkages in the economy, not the level of uncertainty. Finally, none of these properties are symptoms of inefficiency: apart from undoing monopoly distortions or providing the agents with more information, no policy intervention can improve upon the equilibrium allocations.

The third chapter is also based on joint work with George-Marios Angeletos. This chapter investigates how incomplete information affects the response of prices to nominal shocks. Our baseline model is a variant of the Calvo model in which firms observe the underlying nominal shocks with noise. In this model, the response of prices is pinned down by three parameters: the precision of available information about the nominal shock; the frequency of price adjustment; and the degree of strategic complementarity in pricing decisions. This result synthesizes the broader lessons of the pertinent literature. However, this synthesis provides only a partial view of the role of incomplete information: once one allows for more general information structures than those used in previous work, one cannot quantify the degree of price inertia without additional information about the dynamics of higher-order beliefs, or of the agents' forecasts of inflation. We highlight this with three extensions of our baseline model, all of which break the tight connection between the precision of information and higher-order beliefs featured in previous work.

Finally, the fourth chapter studies how predatory trading affects the ability of banks and large trading institutions to raise capital in times of temporary financial distress in an environment in which traders are asymmetrically informed about each others' balance sheets. Predatory trading is a strategy in which a trader can profit by trading against another trader's position, driving an otherwise solvent but distressed trader into insolvency. The predator, however, must be sufficiently informed of the distressed trader's balance sheet in order to exploit this position. I find that when a distressed trader is more informed than other traders about his own balances, searching for extra capital from lenders can become a signal of financial need, thereby opening the door for predatory trading and possible insolvency. Thus, a trader who would otherwise seek to recapitalize is reluctant to search for extra capital in the presence of potential predators. Predatory trading may therefore make it exceedingly difficult for banks and financial institutions to raise credit in times of temporary financial distress.

Thesis Supervisor: George-Marios Angeletos
Title: Professor of Economics

Thesis Supervisor: Mikhail Golosov
Title: Professor of Economics

For my grandparents,

Virgilia, Estelita, Democrito, and Benjamin

Acknowledgements

I cannot overstate my debt to my advisor and co-author, George-Marios Angeletos. Marios has encouraged my development as an economist and incited my enthusiasm for macro ever since my days as an undergraduate. In graduate school, Marios has been an incredible mentor whose continual guidance and support has critically influenced my thinking and research. I cannot thank him enough for the countless discussions and insights he has passed on to me. I have truly enjoyed our time working together on our research; we have learned so much while sharing an excitement for our discoveries. Interacting with Marios throughout these years has been a true privilege.

I would like to also thank my thesis advisor Mike Golosov for his ongoing guidance. Mike's encouragement and advice have been invaluable to me throughout my years as a graduate student, and especially during the job market process. I am deeply indebted to Mike for all of his support. I also thank Guido Lorenzoni in particular for his insightful comments and useful discussions. Furthermore, I have benefitted greatly from classes and conversations with Daron Acemoglu, Olivier Blanchard, Ricardo Caballero, Veronica Guerrieri, Patrick Kehoe, Leonid Kogan, and Ivan Werning. Finally, I thank Ellen McGrattan for the substantial feedback she has given me on the first chapter in my thesis. My research has improved significantly thanks to the interactions I've had with all of these professors.

I would also like to thank my mentors during my undergraduate years at MIT. Through their teaching and advising, David Autor, Dora Costa, Sergei Izmalkov, Sendhil Mullainathan, Nancy Rose, as well as Marios Angeletos, all contributed to the formation of a strong interest in economics and encouraged my decision to pursue a PhD. Nancy in particular was not only a huge influence academically, but her unconditional support, especially through my last year of undergrad, gave me the peace of mind to start graduate school off on the right foot. Amazingly, by the end of graduate school, her guidance throughout the job market had again proven invaluable. I cannot thank her enough for everything she has done.

I am grateful for all of the classmates and officemates who have made my years in graduate school a great experience. I thank especially Ruchir Agarwal, Sergi Basco, Arthur Campbell, Dan Cao, Gabriel Carroll, David Cesarini, Sebastian Di Tella, Florian Ederer, Jesse Edgerton, Laura Feiveson, Stefano Giglio, Danny Gottlieb, Pablo Kurlat, Monica Martinez-Bravo, Sahar Parsa, Florian Scheuer, Alp Simsek, Johannes Spinnewijn, and Xiao Yu May Wang. I look forward to interacting and working with many of these friends and colleagues in the future.

This endeavor could not have been completed without the support of my family. I thank my father Edwin, my mother Debbie, and my brother Brian for their love and encouragement throughout my life.

Last but not least, I would like to extend my deepest gratitude to Luigi Iovino for countless discussions, constant support and encouragement.

Cambridge, MA

May 15, 2010

Contents

Chapter 1

Collateral Constraints and Noisy Fluctuations

1.1 Introduction

Standard business cycle models generally assume that the economy's financial structure is both indeterminate and irrelevant to real economic outcomes. However, many economists have argued that credit constraints and financial conditions may in fact be key factors in determining business cycle fluctuations. The idea of assigning a more central role to credit-market imperfections in explaining short-run macro fluctuations not only dates back to the work of Fisher and Keynes, but has furthermore gained renewed interest in light of the recent financial crisis. In accordance with this view, this chapter demonstrates an important and distinct role financial frictions may serve in generating significant *expectations*-driven business cycles.

An extensive literature on credit-market imperfections has incorporated financial frictions into standard macroeconomic models of the business cycle. In particular, many researchers have explored how collateral constraints interact with aggregate economic activity. In an economy in which firm-level debt must be fully secured by collateral, an endogenous two-way feedback arises between the financial market and the real economy. On one hand, firm output and activity forms the basis for pricing assets. On the other hand, asset prices affect collateral value, which in turn determines the ability of firms to invest. The seminal work of Kiyotaki-Moore (1997) demonstrated how under certain conditions this feedback may amplify the response of the economy to changes in technology or the income distribution. This has launched a wealth of research on collateral constraints; examples include Kiyotaki (1998), Krishnamurthy (2003), Caballero and Krishnamurthy (2001, 2006), Iacoviello (2005). Still others have questioned the quantitative significance of such an amplification mechanism, especially within less stylized environments and under more standard assumptions for preferences and technology (Kocherlakota, 2000; Cordoba and Ripoll, 2004). In any case, most of the existing literature focuses primarily on how collateral constraints affect the response of the economy to aggregate fundamental shocks, e.g. technology, endowments, and preferences, presuming that these are the main drivers of both business cycle and asset price fluctuations.

This comes as no surprise as the current dominant explanation of business cycle fluctuations are changes in fundamentals, the most important being technology shocks. At the same time, many economists find the idea of short-run fluctuations driven entirely by quarterly movements in the technological frontier somewhat implausible. For example, using structural VAR methods researchers have attempted to provide some empirical evidence that only a small fraction of output variation at business-cycle frequencies can be accounted for by technology shocks (Blanchard and Quah, 1989; Shapiro and Watson, 1988; Cochrane, 1994; Galí, 1999). Secondly, the finance literature has found it difficult to explain movements in the stock market on the basis of fundamentals alone; Shiller (1981) and LeRoy and Porter (1981) were the first to document the existence of "excess volatility", or residual variation in asset prices not accounted for by changes in technology or discount factors. Thus in general, while much research has been devoted to understanding the macroeconomic consequences of certain disturbances, the true *sources* of short-run business cycle and asset price fluctuations still remains very much an open question.

this chapter addresses these issues by providing a theory of why financial frictions may imply that an important and potentially significant driver of short-run fluctuations in aggregate output and asset prices may simply be "noise". By this I mean noise in available public information, or more generally, correlated errors in agents' expectations of underlying fundamentals. In this chapter I show how the feedback between asset prices and the real economy that results from collateral constraints may lead to significant expectations-driven business cycles when one breaks the assumption that information is commonly shared. Importantly, I find that when collateral constraints are tight and information is dispersed, the effects of noise can potentially account for a large portion of both business cycle and asset price fluctuations, while the relative contribution of technological shocks to these short-run fluctuations may be small. Thus, by allowing for heterogeneity in information, I identify a distinct and novel role for financial frictions.

Preview of the Model. The starting point of this chapter is a micro-founded one-sector real business cycle model in which firm debt must be fully secured by collateral.[1] I abstract from capital to simplify the analysis, but assume firms must borrow in order to pay their workers.[2] As in most models of collateral constraints, in this economy there exists a durable asset which plays a dual role: it serves both as collateral to secure firm debt, and its value depends on real economic activity.[3] The endogeneity of the collateral value is effectively what creates the two-way feedback between the financial market and the real economy: the tighter the constraint, the less any individually-constrained firm may produce; at the same time, the less the entire economy produces, the tighter the constraints. In this respect this model is very similar to the mechanism present in Kiyotaki-Moore (1997) as well as the rest of the literature on collateral constraints.

One key difference between this model and Kiyotaki-Moore (1997), however, is that Kiyotaki-Moore allows firms to be highly levered, in that they have borrowed heavily against the value of their collateral. For this reason they obtain amplification of productivity shocks, yet many have convincingly argued that this effect must be quantitatively small (Kocherlakota, 2000; Cordoba and Ripoll, 2004). In this chapter I abstract from the effects of leverage. I do so in order to isolate the contribution of the essential ingredients considered in this model. This model thus incorporates the salient feature that endogenous credit constraints create a feedback between the financial market and the real economy, but it is designed so that the Kiyotaki-Moore leveraged amplifier of productivity shocks is absent. I could argue that this doesn't seem to be such

[1]The underlying assumption that motivates this type of credit constraints is that lenders cannot force borrowers to repay their debts unless the debts are secured.

[2]I can show in an extension of the model that the inclusion of capital does not affect the paper's main qualitative results.

[3]Most collateral-constraint models allow the asset to be an input to production. Here I instead assume workers obtain utility from consuming its dividends. In either case, the main reason for this assumption is to make the asset price sensitive to aggregate economic activity.

a terrible assumption in its own right, as demonstrating the quantitative significance of the Kiyotaki-Moore amplifier is clearly a challenge. Nevertheless the main reason for making this modelling choice is that it allows me to focus solely on the contribution of the information structure. By construction, then, my benchmark economy is one in which collateral constraints have no effect on the business cycle properties of equilibrium output when information is complete.

As anticipated, a key ingredient in this chapter is the *heterogeneity* of information. I show how this type of informational constraint interacts with the aforementioned financial friction. Unlike all previous models of collateral constraints, I allow for production, employment, and asset demand decisions to take place under dispersed information about the underlying fundamentals. I then show how the introduction of dispersed information creates a distinct and important role for financial frictions with respect to the equilibrium behavior of the economy.

Finally, I assume that the durable asset is traded among firms and workers across all islands. Therefore the price of this asset is observed by all agents in the economy and in equilibrium partially aggregates the dispersed information.

Preview of the Results. First, as a benchmark I demonstrate that under complete information, by construction tighter collateral constraints have no effect on the business cycle properties of the equilibrium. The focus of this chapter, then, is to explore the implications of collateral constraints when information among agents is heterogeneous. The main results that obtain under dispersed information are the following:

(i) *Noise as a Source of the Business Cycle.* The combination of endogenous collateral constraints and dispersed information can lead to significant noise-driven fluctuations in aggregate output and asset prices. Specifically, I find that tighter collateral constraints mitigate the impact of productivity shocks on equilibrium output and asset prices, but amplify the impact of noise, and hence lead to a greater relative contribution of noise to business-cycle fluctuations. Noise can thus be a significant source of asset-price and output volatility when collateral constraints are tight.

(ii) *Positive Co-Movement.* Noise results in positive co-movement in employment, output, consumption, and asset prices. These fluctuations are independent of the underlying productivity shock; they are merely the product of common errors in agents' expectations.

(iii) *Labor Wedge.* Noise contributes to movements in the measured "labor wedge". In particular, the measured labor wedge is positively correlated with fluctuations in output when the main driver of the business cycle is noise.

The first result highlights the main contribution of this chapter. This is the first paper, to my knowledge, to provide a theory of how financial frictions may imply that movements in both the business cycle and asset prices may be driven primarily by noise, as opposed to fundamentals. In this sense, collateral constraints play a distinct role in comparison to the earlier literature. this chapter thus highlights the importance of considering heterogeneous information when firms are subject to collateral constraints.

The intuition for this result begins with the fact that asset prices, and hence the value of collateral, are based on average expectations of aggregate economic activity. When one relaxes the assumption of common knowledge and allows for heterogeneity of information, agents rely more on common sources of information regarding aggregate productivity, and less on private sources of information. This is because common sources of information are relatively better predictors of aggregate economic activity. As a result, asset prices become relatively more sensitive to this noise and less sensitive to fundamentals.

Consider then a small negative expectational error about technology, common to all agents. If agents believe common information will affect output, the small error depresses asset prices relatively more than what is dictated by its informational content about fundamentals alone. The decline in asset prices, however, lowers collateral values, which in turn leads to a fall in output. Thus, due to collateral constraints, the behavior of asset prices affects the behavior of real output. Moreover, the expected decline in output depresses asset prices even further. The two-way feedback just described amplifies the small expectational error, and in this way expectations about aggregate production become in many ways self-fullfilling.

Now consider a small positive innovation in aggregate productivity. Heterogeneity of information implies that this shock is not common knowledge. Hence even if agents on all islands are sufficiently well-informed of this innovation, because of lack of common knowledge asset prices are not extremely responsive. However, if asset prices move relatively little in response to the productivity innovation, constrained firms cannot react to the increase in productivity. But this inertia in expected output reinforces the dampened response of asset prices. Therefore, the two-way feedback generated by the collateral constraints mitigates the impact of the productivity shock.

In equilibrium, this behavior makes both asset prices and output more sensitive to noise and less sensitive to fundamentals. A robust prediction then of this model is that the tighter the collateral constraints, i.e. the more asset prices matter for firm production, the greater the relative contribution of noise to business cycle and asset price volatility. Thus, noise could very well be an important and dominant source of business cycle fluctuations.

The second and third main results then describe the implications of noise in this model in terms of observables.

The second result is notable as it shows how the noise-driven movements highlighted in this chapter are a possible interpretation of what many believe as a major source of business cycle fluctuations–namely, "demand shocks". Using structural VAR methods, many researchers claim to identify a certain type of shock, orthogonal to technology, which causes positive comovement in aggregate macroeconomic variables including employment, output and consumption. They moreover find that these shocks may account for the majority of short-run fluctuations. While New-Keynesians have designated these as "demand shocks", their structural interpretation remains a contentious issue. This model offers a possible explanation for these fluctuations. The noise-driven movements documented here resemble demand-driven fluctuations in the following sense. They have many of the same features often associated with such shocks: they contribute to positive co-movement in employment, output and consumption, yet at the same time are orthogonal to underlying movements in productivity.

The third result relates the implications of noise in this model to the large body of research documenting the observed variation in the "labor wedge" over the business cycle (Hall, 1997; Rotemberg and Woodford, 1999; Chari, Kehoe, and McGrattan, 2007; Shimer, 2009).[4] Chari, Kehoe, and McGrattan (2007), in particular, show that a large class of models are equivalent to a frictionless representative-agent model of the business cycle in which various types of time-varying wedges (labor, investment, efficiency) distort the equilibrium decisions of the agents. They furthermore document the empirical contribution of each wedge over certain business cycle episodes as well as over the entire postwar period and find that the efficiency and labor wedges

[4] The labor wedge is defined as the wedge between the marginal rate of substitution and the marginal product of labor that would obtain in a frictionless representative-agent model of the business cycle.

together account for the bulk of fluctuations while the investment wedge plays essentially no role. While their methodology is not meant to identify the source of the business cycle, it does suggest that the most promising models, in terms of relating to the data, are ones in which the underlying sources and frictions manifest themselves as labor wedges. Thus, this model may be useful for understanding certain business cycle periods, such as the recent crisis, in which the labor wedge plays an important role empirically.

Finally, in delivering these results this chapter makes two methodological contributions to the recent growing literature on dispersed information in macro. There are a number of papers, most notably Morris and Shin (2002), which highlight the potential implications of heterogeneous information in a more abstract class of games that feature strategic complementarity. In relation to this body of work, one of the contributions of this chapter is that it provides a novel microfoundation for these games. By this I mean the following: I show how the equilibrium allocation of the fully micro-founded business-cycle economy featured here can be represented as the Bayesian-Nash equilibrium of a game similar to those considered in the more abstract game-theoretic settings. This representation is useful, as the more abstract insights of this earlier work aids one in understanding how the two-way feedback between the financial market and the real economy results in noisy fluctuations when combined with dispersed information. In this respect, this chapter is complementary to Angeletos and La'O (2009), in that it considers a micro-founded real business cycle model in which agents have dispersed information about underlying fundamentals. However, the micro-foundations and hence the mechanisms considered in these two papers are substantially different. Angeletos-La'O focuses on the effects of specialization in a multisector economy, and how trade across sectors creates interlinkages among firms. this chapter instead considers a standard one-sector economy and highlights the two-way feedback that naturally arises between the financial and real sides of the economy when firms are constrained by collateral.

Secondly, the environment here features an asset whose price both aggregates and conveys information across islands. While most models with dispersed information abstract from this complication, this model calls for financial markets, and thus I demonstrate how one can allow for observable asset prices in a dispersed information environment yet preserve tractability. This is important, as one possible concern with heterogeneous information models is that their effects may disappear when financial market prices aggregate information. I show here that this concern is misguided; the effects of noise are robust to the inclusion of observable asset prices.

Layout. this chapter is organized as follows. The remainder of the introduction discusses the related literature. Section 1.2 introduces the model and describes the economic environment. In Section 1.3, I define and characterize the equilibrium of the economy. In Section 1.4 I consider the complete information benchmark. Section 1.5 explores the mechanism underlying the main results of this model in a simplified environment. In Section 1.6, I describe the method I use to obtain an approximate solution to the general equilibrium of the full model. I then present the results of this solution in Section 1.7 and highlight the main contributions of this chapter. Section 1.8 concludes.

Related Literature. this chapter is related to the literature on financial frictions in the business cycle. Within this literature, there are two general strands. One strand builds on the work by Townsend (1979) and Gale and Hellwig (1985) and focuses on the costly state verification problem between creditors and debtors. Papers such as Carlstrom and Fuerst (1997), Bernanke, Gertler, and Gilchrist (1999), Cooley, Marimon, and Quadrini (2004), Gertler, Gilchrist, and Natalucci (2007), and Christiano, Motto, and Rostagno (2008) build this problem into macroeconomic models, emphasizing the interaction between interest rate speads and aggregate quantities. The key mechanism here involves the link between "external financing premium" (the difference between the cost of funds raised externally and the opportunity cost of funds internal to the firm) and the net worth of borrowers.

this chapter, however, is more closely related to another strand of the financial frictions literature: one that emphasizes collateral constraints in models of the business cycle. Examples in this literature include Kiyotaki-Moore (1997), Kiyotaki (1998), Kocherlakota (2000), Krishnamurthy (2003), Caballero and Krishnamurthy (2001, 2006), Cordoba and Ripoll (2004), Iacoviello (2005), Liu, Wang, and Zha (2009). In this class of models, a firm's debt level is constrained by the value of its collateral. this chapter incorporates a similar collateral constraint friction into a business cycle model, but focuses on its interaction with heterogeneous information.

There is a wealth of empirical evidence that point towards the importance of financial frictions in aggregate fluctuations. First, there is considerable evidence of balance-sheet effects on investment, see Fazzari, Hubbard and Petersen (1988), Gertler, Hubbard, and Kashyap (1991), Gilchrist and Himmelberg (1995), Hubbard, Kashyap and Whited (1995), as well as on inventories, see Kashyap, Lamont and Stein (1994). Secondly, some authors have found that smaller firms are more sensitive to the business cycle; see Gertler and Gilchrist (1994), and Sharpe (1994).[5] Finally, Rajan and Zingales (1998) provide empirical evidence for a causal effect of financial development on growth.

this chapter furthermore contributes to the growing literature on informational frictions in macro.[6] Within this context, this chapter follows in methodology lines similar to Lucas (1972), Barro (1976), Townsend (1983) and others, by introducing a geographic segmentation of information among islands to formalize the dispersion of information. this chapter thus fits in with the recent revival of models of dispersed information, which has been duly influenced by the work of Morris and Shin (2002), Sims (2003), and Woodford (2003).

In considering expectations-driven business cycles, this chapter is related to the literature on news-driven business cycles; see, for example, Beaudry Portier (2004, 2006), Christiano et al.(2007), Jaimovich and Rebelo (2008), and Lorenzoni (2009). These papers also generate certain types of expectations-driven fluctuations. However, the mechanisms considered in these papers differs significantly from the one considered here. The news-shock literature focuses on how expectations of future productivity may drive current production within representative-agent environments. Thus, their results do not rest on the heterogeneity of information. Here, I show that the heterogeneity of information is key in bringing about a self-fulfilling feature of expectations: agents need not make large forecast errors when information is dispersed, they merely coordinate on the same information.

Finally, this chapter is also very related to an important contribution by Chen and Song (2010). Chen and Song (2010) study an economy in which a fraction of firms face financial

[5] However, recent work by Chari, Christiano, and Kehoe (2009), find evidence towards the contrary.

[6] Phelps (1970), Lucas (1972, 1975), Barro (1976), King (1982), Townsend (1983). More recent papers include Adam (2007), Amador and Weill (2007, 2008), Amato and Shin (2006), Angeletos and Pavan (2004, 2007, 2009), Bacchetta and Wincoop (2005), Collard and Dellas (2005), Hellwig (2002, 2005), Hellwig and Veldkamp (2008), Hellwig and Venkateswara (2008), Klenow and Willis (2007), Lorenzoni (2008, 2009), Luo (2008), Mackowiak and Wiederholt (2008, 2009), Mankiw and Reis (2002, 2006), Mertens (2009), Morris and Shin (2002, 2006), Moscarini (2004), Nimark (2008), Reis (2006, 2008), Rondina (2008), Sims (2003, 2006), Van Nieuwerburgh and Veldkamp (2006), Veldkamp (2006), Veldkamp and Woolfers (2007), and Woodford (2003a, 2008), Angeletos, Lorenzoni, Pavan (2009), Angeletos La'O (2009a, 2009b)

constraints. They then show how news of future productivity relaxes these financial constraints, allowing capital to flow to constrained projects. The reallocation of capital generates an increase in current aggregate TFP as well as positive co-movement in macro aggregates. Because there is no uncertainty in their model, news-driven fluctuations are in fact correlated with future productivity, despite the fact that they are independent of current productivity. In this sense, unlike this chapter, the positive-comovement generated in Chen and Song (2010) is remains correlated with fundamentals. If they were to add uncertainty under symmetric information to their model, errors in expectations of future productivity would drive current production. However, in this case, as in the news shocks literature, expectations would have no self-fulfilling role.

1.2 The Model

There is a (unit-measure) continuum of households, or "families", each consisting of a continuum of worker-members. There is a continuum of "islands", which define the boundaries of local labor markets as well as the "geography" of information: information is symmetric within an island, but asymmetric across islands. Each island is inhabited by a competitive representative firm which produces a consumption good; goods produced among islands are perfect substitutes.[7] Households are indexed by $k \in K = [0,1]$, its members by $j \in J = [0,1]$; and islands, as well as the representative firm on each island, are indexed by $i \in I = [0,1]$.

There are two stages. In stage 1, each household sends a worker to each of the islands. Once workers reach these islands, the representative firm on each island is endowed with a durable good, which I henceforth refer to as land, and learns of its own productivity for producing the commodity. This information is also revealed to all workers within the island. Thus, at this point workers and firms on each island have complete information regarding local productivity and land, but incomplete information regarding the productivities and land endowments of other islands.

Firms and workers then trade claims on land in a centralized (economy-wide) market; in particular, firms sell land claims, workers buy these claims, and the land price adjusts so as

[7] One may think of this as a continuum of identical competitive firms within an island. To supress notation I will refer to this continuum as a single competitive representative firm.

to clear the market. After the economy-wide land market has cleared, local (within-island) labor markets open: workers decide how much labor to supply, firms decide how much labor to demand, and local wages adjust so as to clear the local labor market. Finally, after employment and production choices are sunk, workers return home and the economy transits to stage 2.

In stage 2, all information that was previously dispersed becomes publicly known, and the commodity market opens. Quantities are now pre-determined by the exogenous productivities and the endogenous employment choices made during stage 1, but the price of the consumption good adjusts to clear the market. Finally, consumption of goods and housing takes place.

Households. The utility of household k is given by

$$\mathcal{U}_k = C_k^{1-\gamma} H_k^{\gamma} - N_k$$

where C_k is the household's consumption of the identical good produced by firms, and H_k and N_k are composites of the housing purchased and the labor supplied by the household's various workers in stage 1. The household has Cobb-Douglas utility in consumption and housing, where $\gamma \in (0,1)$ is the income share of land. The composites H_k and N_k are given by the following

$$H_k = \left[\int_J h_{jk}^{\frac{\nu-1}{\nu}} dj \right]^{\frac{\nu}{\nu-1}}, \qquad N_k = \int_J n_{jk} dj$$

where n_{jk} is the amount of labor supplied by family member j during stage 1 and h_{jk} is that family member's purchase of land claims.[8] Here $\nu > 0$ is the elasticity of substitution across housing purchased by different members of the household.

Households own equal shares of all firms in the economy. Accordingly, the budget constraint of household k is given by the following

$$PC_k + Q \int_J h_{jk} dj \leq \int_I \pi_i di + \int_J w_j n_{jk} dj$$

where P is the price of the consumption good, Q is the price of housing, π_i is the profit of firm i, and w_j is wage of family member j.

Finally the objective of the household is simply to maximize expected utility subject to

[8] Here utility is quasilinear in labor, but that assumption can easily be relaxed.

the budget and informational constraints faced by its members. Here, one should think of the members of each family as solving a team problem: they share the same objective (family utility) but have different information sets when making their labor-supply and asset-demand choices. Formally, during stage 1 the household sends off its workers to different islands with bidding instructions on how to supply labor and demand assets as a function of (i) the information that will be available to them on their island, (ii) the wage in their local labor market, and (iii) the price of the asset. In stage 2, the workers return home and the household collects all labor income and firm profits and consumes.

Firms. The output of the representative firm on island i is given by

$$y_i = A_i n_i^\theta \tag{1.1}$$

where A_i is the productivity on island i, n_i is the firm's employment, and θ parameterizes the degree of diminishing returns to production. Each firm is endowed with L_i units land, which they sell on the centralized land market at price Q. Thus, the firm's realized profits are given by

$$\pi_i = Py_i - w_i n_i + QL_i$$

The firm's objective is to maximize expected shareholder value.

I make the following critical assumption about firms: the firm cannot precommit to not run away with the revenues from its sales in stage 2. The goods, however, must be produced by workers in stage 1, before the product markets open. This poses a problem as workers, in light of the firm's commitment problem, refuse work in stage 1 unless paid up front.

To get around this problem, firms may use their land holdings as collateral to pay workers. This implies that the firm may hire as much employment as it likes, as long as the cost of doing so does not exceed the value of its land. The firm's problem then is given by the following. The competitive representative firm on island i, taking prices as given, maximizes expected profits weighted by the representative household's marginal value of wealth, λ,

$$\max \mathbb{E}_i \left[\lambda \; \pi_i\right]$$

subject to its production function (1.1) and subject to its collateral constraint given by

$$w_i n_i \leq \chi Q L_i \tag{1.2}$$

where $\chi > 0$ is the fraction of land which is collateralizable. The collateral constraint (1.2) dictates that the firm's cost of employment in stage 1 cannot exceed the value of its land, where land is valued mark-to-market at the price Q.

Labor, Land, and Product Markets. Labor markets operate in stage 1. Because labor cannot move across islands, the clearing conditions for the local labor markets are as follows

$$n_i = \int_K n_{ki} dk, \forall i$$

The land market also operates in stage 1, but unlike the labor market, land claims are traded in an economy-wide market. One can think of this as follows. There is a centralized market, not located on any island, to which workers submit demand schedules for land as functions of the price, and firms submit their land endowments. A walrasian auctioneer then sets the land price such that it clears the market. The clearing condition for the land market is as follows.

$$\int_I L_i di = \int_K \int_J h_{jk} dj dk$$

Finally, the consumption good is traded in a centralized market that operates in stage 2; the clearing condition for the product market is given by

$$\int_I y_i di = \int_K C_k dk$$

Fundamentals. Each island in this economy is subject to two types of shocks: technology and land endowment shocks. I assume the island-specific productivities $A_{k,t}$ are lognormally distributed in the cross-section of islands:

$$a_i \equiv \log A_i = \bar{a} + \xi_{a,i},$$

where $\bar{a}_t$ denotes the underlying aggregate productivity shock and $\xi_{a,i}$ is a purely idiosyncratic (i.e., island-specific) productivity shock. I assume $\xi_{a,i}$ is drawn from the Gaussian distribution with mean zero and variance $\sigma^2_{A,x} \equiv 1/\kappa_{A,x}$, is orthogonal to $\bar{a}$, and is i.i.d. across islands. Similarly, the island-specific land endowments are lognormally distributed in the cross-section of islands:

$$l_i \equiv \log L_i = \bar{l} + \xi_{l,i},$$

where $\bar{l}_t$ denotes the underlying aggregate land endowment and $\xi_{l,i}$ is a purely idiosyncratic (i.e., island-specific) land shock. I assume $\xi_{l,i}$ is drawn from the Gaussian distribution with mean zero and variance $\sigma^2_{L,x} \equiv 1/\kappa_{L,x}$, is orthogonal to both $\bar{l}$, and is i.i.d. across islands.

Finally, the aggregate shock $\bar{a}$ is drawn from a Gaussian distribution with mean zero and variance $\sigma^2_{A,0} \equiv 1/\kappa_{A,0}$, while the aggregate shock $\bar{l}$ is drawn from a Gaussian distribution with mean zero and variance $\sigma^2_{L,0} \equiv 1/\kappa_{L,0}$. The two aggregate shocks, $\bar{a}$ and $\bar{l}$ are orthogonal to each other.

Information. Aggregate productivity and the aggregate land endowment are assumed to be common knowledge at stage 2, when all production materializes and consumption takes place, but not in stage 1, when the key employment, production, and asset holding choices are made. Rather, at this stage workers and firms in any given island know perfectly the productivity of their own island but not the productivities of other islands. Local productivity and local land endowments thus serve also as valuable, but noisy, private signals of the distribution of productivities, land, and information in other islands.[9] In addition to these private signals, all islands observe an exogenous public signal of aggregate productivity:

$$z = \bar{a} + \varepsilon$$

where ε is drawn from the Gaussian distribution with mean zero and variance $\sigma^2_{A,z} \equiv 1/\kappa_{A,z}$, and is orthogonal to both $\bar{a}$ and $\bar{l}$. I assume the precisions (and equivalently the variances) of all fundamentals and signals, $\kappa_{A,x}$, $\kappa_{L,x}$, $\kappa_{A,0}$, $\kappa_{L,0}$, and $\kappa_{A,z}$ are common knowledge.

Here ε is what is referred to as "noise". In terms of this model, noise is literally the error in the public signal. However, this is just a convenient modelling device for common

[9] The assumption that firms and workers know perfectly their own productivities and endowments is not essential; all results go through if I allow for uncertainty about local as well as aggregate productivity.

information. More generally, one can think of noise merely as correlated errors in expectations about aggregate productivity.

Finally, firms and workers in each island observe the land asset price Q.

Remarks. The constraint imposed on firm-level expenditure in (1.2) is similar to those found in the larger literature on collateral constraints. In this class of models the borrowing capacity of a firm is constrained by the value of its collateral. Although this setup is described as though the firm literally takes its revenues from selling land claims and uses these revenues to pay its workers, one could allow for an intermediate step between these transactions in which a financial intermediary loans funds to the firm and secures this loan with collateral. The collateral, then, is the value of land held by the firm.

Most papers on collateral constraints provide a simple justification for this type of financial friction. Following Kiyotaki-Moore (1997), these papers interpret this type of constraint as reflecting the problem of control over assets (see Hart and Moore, 1989). The crucial feature is that the firm is needed to run or control the project but cannot precommit to repay its debt. If the firm defaults, the creditor can seize the borrower's land. Hence the firm's repayment is bounded from above by the value of its land.[10] Furthermore, the usual interpretation of the parameter χ stems from the consideration that it may be costly to liquidate seized land, hence the creditor can recoup only up to a fraction χ of the total value of the collateral assets. More generally, however, one can think of χ as parameterizing the tightness of collateral constraints in the economy. Finally, while it is true that few firms literally collateralize their loans using land, in reality there are many layers of collateralization that are ultimately grounded in a physical asset. In this sense, stylized collateral constraint environments do capture some element of reality.

The durable asset in this model plays a dual role: it serves both as an asset of value, and as collateral to secure firm debt. In this way the model preserves the key feature of all collateral constraint environments. However, as already mentioned in the introduction, this feature alone does not necessarily generate amplification. One key difference between this model

[10] This is not the same as exogenously ruling out the possibility of state-contingent contracts. As pointed out to me by Randy Wright, one could think that writing state-contingent contracts are possible in this framework so that loan markets are, in a sense, complete. However, the commitment problem of the firm makes it so that no worker nor creditor would ever want to write a state-contingent loan contract, as the firm would in every state have the incentive to run away with the money.

and Kiyotaki-Moore (1997), is that Kiyotaki-Moore allows firms to be highly levered in that they have borrowed heavily against the value of their collateral. In particular, firms enter each period already in debt, at which point these firms obtain new loans, subject to their collateral constraints, and use these loans to both repay old debt and to invest. In this environment, the response of the economy to changes in fundamentals such as productivity is amplified. In contrast, in this model I abstract from this effect and assume firms enter without any existing debt, yet must borrow in order to pay for their static operating cost of paying workers. Thus this model features a simple, static input-financing friction due to collateral constraints.[11] I show in Section 1.4 that in this case under complete information collateral constraints have no effect on the business cycle properties of the equilibrium. That is, the complete information benchmark is designed so that tighter constraints do not amplify the response of the economy to productivity shocks.

I assume that consumption choices in stage 2 take place in a centralized market, where information is homogeneous, and that households are "big families", with fully diversified sources of income. The big-family formulation allows for household members to be completely insured against idiosyncratic risk in consumption and hence guarantees that the economy admits a representative consumer in stage 2. This assumption is made merely to maintain high tractability in analysis despite the fact that some key economic decisions take place under heterogeneous information.

Shocks to the aggregate amount of land are introduced in this model for one reason and one reason only: so that in equilibrium the asset price is not fully-revealing of the underlying aggregate state. In other words, if I were to eliminate the shocks to the aggregate land endowment so that the only fundamental shock were that in aggregate productivity, the asset price would perfectly reveal aggregate productivity. In this case all agents would have complete information of the aggregate state of the economy. To guard against this, I introduce another disturbance in the asset price, so that as in any Noisy Rational Expectations equilibrium the mapping from the aggregate state to the asset price is not invertible. I do this by introducing shocks to the aggregate land endowment. This however could have been accomplished in a variety of ways;

[11]See Chari, Kehoe, McGrattan (2007) for an example of an input-financing friction due to differential interest rates.

for example, one could incorporate an aggregate preference shock. In any case, one may think of this as a substitute for the "noise traders" that are often used in the noisy rational expectations equilibrium models in finance. However, instead of bringing in agents from outside the model, shocks to the aggregate land endowment could be considered a more micro-founded modelling device that ensures a non fully-revealing asset price.

Finally, while the benchmark of this model is a one-sector economy, the model could be further enriched to allow for a multi-sector economy in which firms from different islands specialize in the production of differentiated goods. This could be introduced here by allowing C_k to be a CES composite of the consumption of differentiated commodities. Specifically, one could write $C_k = \left[\int_i c_{ik}^{\frac{\rho-1}{\rho}} di\right]^{\frac{\rho}{\rho-1}}$ where c_{ik} denotes household k's consumption of the differentiated good produced by the firm on island i, and where $\rho > 0$ is the elasticity of substitution across these commodities. For finite values of ρ, the model nests, as a special case in which there are no collateral constraints and no trading in the land market, the environment considered in Angeletos-La'O (2008, 2009a). Angeletos-La'O (2008, 2009a) show that the specialization of goods and the strength of trade linkages across islands leads to strategic complementarity in production decisions. Here, for the bulk of this chapter I omit this effect by assuming perfect substitution among goods, i.e. $\rho \to \infty$. When presenting numerical results for the full model in section 1.7, as an extension I return to allowing for imperfect substitution among goods and show how this additional effect interacts with the financial friction in general equilibrium.

1.3 Equilibrium definition and characterization

In this section I define and characterize the equilibrium of the economy.

1.3.1 Definition

Because of the symmetry of preferences across households and information within each island, each island admits a representative worker; that is, it is without any loss of generality to impose symmetry in the choices of workers within each island. Furthermore, because each family sends workers to every island and receives profits from every firm in the economy, each family's income is fully diversified during stage 2 and the economy admits a representative consumer in this stage.

I then use the following notation. First, let ω denote the "type" of an island in stage 1. This variable identifies the vector (a_i, l_i, z) and hence encodes all *exogenous* information available to an island about the local shocks as well as about the aggregate state. I thus let $A(\omega)$ and $L(\omega)$ denote the local productivity and land endowment in island ω, with $A(\omega) = e^{a_i}$, $L(\omega) = e^{l_i}$ for $\omega = (a_i, l_i, z)$, respectively. Next, let Ω denote the "aggregate state" of the economy. This variable denotes the realized distribution of ω in the cross-section of islands. Because Ω is a joint normal distribution with mean productivity $\bar{a}$, mean land endowment $\bar{l}$, public signal z, and a commonly known variance-covariance matrix, one may think of the vector $(\bar{a}, \bar{l}, z)$ as a sufficient statistic for the aggregate state Ω. The key informational friction in this model is that agents face uncertainty about the underlying aggregate state Ω. Finally, let $\mathcal{S}_\omega$ denote the set of possible types for each island, $\mathcal{S}_\Omega$ the set of probability distributions over $\mathcal{S}_\omega$, and $\mathcal{P}(\cdot)$ the prior probability measure for Ω over $\mathcal{S}_\Omega$.

I then formalize the information structure as follows. In stage 1, Nature draws a distribution $\Omega \in \mathcal{S}_\Omega$ using the measure $\mathcal{P}(\Omega)$. Nature then uses Ω to make independent draws of $\omega \in \mathcal{S}_\omega$, one for each island. The information that becomes available to an island in stage 1 is both ω and Q, where the asset price Q is an *endogenous* variable that in equilibrium depends on the aggregate state Ω. The island's local type ω and the asset price Q informs firms and workers on that island perfectly about their local shocks, but only imperfectly about the underlying aggregate state Ω.[12] As a result, for any given island, the labor supply and the housing demand of the local workers, the labor demand and the level of production of the local firms, and the wage that clears the local labor market, are functions of the island's ω and Q, but (conditional on Q) not on the aggregate state Ω. In stage 2, Ω becomes commonly known. Thus, the price that clears the commodity market in stage 2, and all aggregate outcomes, do depend on the current aggregate state Ω. I thus define an equilibrium as follows.

Definition 1 *An equilibrium consists of an employment strategy $n : \mathcal{S}_\omega \times \mathbb{R}_+ \to \mathbb{R}_+$ a production strategy $y : \mathcal{S}_\omega \times \mathbb{R}_+ \to \mathbb{R}_+$, a housing demand strategy $h : \mathcal{S}_\omega \times \mathbb{R}_+ \to \mathbb{R}_+$, a wage*

[12]Whether they agents uncertainty about their own local shocks is immaterial for the type of effects we analyze in this paper. Merely for convenience, then, I assume that the agents of an island learn their own local shocks in stage 1.

function $w : \mathcal{S}_\omega \times \mathbb{R}_+ \to \mathbb{R}_+$*, an aggregate output function* $Y : \mathcal{S}_\Omega \to \mathbb{R}_+$*, an aggregate employment function* $N : \mathcal{S}_\Omega \to \mathbb{R}_+$*, a commodity price function* $P : \mathcal{S}_\Omega \to \mathbb{R}_+$*, a land price function* $Q : \mathcal{S}_\Omega \to \mathbb{R}_+$*, and a consumption strategy* $C : \mathbb{R}_+ \to \mathbb{R}_+$*, such that the following are true:*

(i) The price that clears the commodity market in stage 2 is normalized so that $P(\Omega) = 1$ *for all* Ω*.*

(ii) The quantity $C(Y)$ *is the representative consumer's optimal demand for the consumption good when the aggregate output (income) is* Y*.*

(iii) The price that clears the land market in stage 1 is $Q(\Omega)$*.*

(iv) When the current asset price is Q*, the employment and output levels,* $n(\omega, Q)$ *and* $y(\omega, Q)$*, of the representative firm from island* ω *maximize the expected shareholder's value of this firm subject to its production function* $y(\omega, Q) = A(\omega)n(\omega, Q)^\theta$ *and collateral constraint* $w(\omega, Q)n(\omega, Q) \leq \chi QL(\omega)$*, taking into account that firms in other islands are behaving according to the same strategies, that the local wage is given by* $w(\omega, Q)$*, that the price of the commodity in stage 2 is given by* $P(\Omega) = 1$*, that the representative consumer will behave according to consumption strategy* C*, and that aggregate income will be given by* $Y(\Omega)$*.*

(v) When the current asset price is Q*, the quantity* $h(\omega, Q)$ *is the optimal housing demand of the typical worker in an island of type* ω*, and the local wage* $w(\omega, Q)$ *is such that the quantity* $n(\omega, Q)$ *is also this worker's optimal labor supply.*

(vi) When the aggregate state is Ω*, the aggregate output and employment indices are, respectively,*

$$Y(\Omega) = \int y(\omega, Q(\Omega))d\Omega(\omega) \quad \textit{and} \quad N(\Omega) = \int n(\omega, Q(\Omega))d\Omega(\omega).$$

This first condition simply states that, without any loss of generality, the numeraire for this economy is the consumption good. The rest of the conditions then represent a hybrid of the following: (i) for given production choices made in stage 1, a Walrasian equilibrium for the complete-information exchange economy that obtains in stage 2, (ii) for a given realization of the asset price, a Bayesian-Nash equilibrium for the incomplete-information game played among different islands in stage 1, and (iii) a Noisy Rational-Expectations equilibrium in the asset market in stage 1.

Allow me to expand on this. First, once production choices are sunk in stage 1, there is simply a competive exchange economy in stage 2 so that the product market clears as in any standard Walrasian equilibrium.

Second, in stage 1 for any given realization of the asset price, there is a competive equilibrium in the labor market on each island. However, note the following. Although firms within an island, in deciding how much labor to demand, face no uncertainty in terms of the price of their good, the asset price, nor the local wage, workers on the other hand, in deciding how much labor to supply and how much land to demand, do face uncertainty about the real income of their household in stage 2. In other words, in stage 1 workers face uncertainty about the marginal value of the wealth. But this then implies that the equilibrium local wage and the equilibrium asset price, and hence the amount of goods produced on each island, depends on the anticipated realized level of real aggregate income in stage 2. Aggregate income, however, is given by the level of production in *all* islands in the economy. Therefore, through general equilibrium price effects, the equilibrium employment and production decisions in the economy can be reduced to a certain incomplete-information game played among different islands, where the incentives of firms and workers on each island depend on their expectations of the choices of firms and workers in all other islands. In Section 1.5 I formalize this argument, and show that the equilibrium of this economy is isomorphic to the Bayesian-Nash equilibrium for the incomplete-information game just described.

Third, consider the asset market in stage 1. A worker on an island, in deciding how much land to demand, anticipates that the asset price affects his household's budget constraint and hence, as in any standard household optimization problem, the asset price affects the worker's demand for the asset. Given these demand functions, the equilibrium asset price clears the land market. But this implies that the asset price is a function mapping the distribution of agents' private information, and thereby the aggregate state Ω, to an asset price. Thus, by observing the price, a firm or worker on any island may use the price to extract information on the aggregate state of the economy. Thus, the Noisy Rational-Expectations equilibrium in the asset market imposes a fixed-point relation between the equilibrium allocations and the equilibrium information structure, and hence captures the notion that agents learn from prices. The reason it is considered "noisy" is due to the fact that agents are making inference of two aggregate shocks, the aggregate productivity shock and the aggregate land shock, from one price. Thus the price as a function of the underlying state is not invertible and therefore cannot be fully-revealing.

1.3.2 Characterization

In characterizing the equilibrium, I proceed as follows. First, I solve the team-problem of the representative household and characterize the optimal consumption, labor supply, and housing demand schedules for its members. I then solve for the firm's optimal level of production. Local labor-market clearing conditions allows me to solve for the partial equilibrium within each island. Finally, using economy-wide goods and land market clearing conditions, I characterize the general equilibrium for the economy.

Household optimality. Household members make choices both in stage 1 and stage 2.

Stage 2. In stage 2 the only choice the representative household faces is its consumption of goods. As goods are perfect substitutes, there is only one good the household buys at this stage. Thus this choice is quite simple: the household merely spends all of its income (net of wages, stock dividends, and payments for land) on consumption, and the price that clears this market is given by $P(\Omega) = 1$. In equilibrium, then, $C(\Omega) = Y(\Omega)$. It follows that the equilibrium consumption strategy is given by $C(Y) = Y$. Furthermore, the marginal value of wealth of the household in this stage is given by

$$\lambda(\Omega) = (1-\gamma)\left(\frac{H(\Omega)}{C(\Omega)}\right)^{\gamma} \tag{1.3}$$

and is thereby decreasing in aggregate consumption.

Stage 1. I now turn to the optimal choices of the household in stage 1. Again, one can think of this as the following: the household sends off its workers to different islands with instructions on how to supply labor and buy assets as a function of the wages, asset price, and information available to each member on his respective island. The optimal labor supply of the typical worker on island ω is then given by

$$w(\omega, Q) = \frac{1}{\mathbb{E}\left[\lambda(\Omega)\,|\,\omega, Q\right]}$$

which merely equates the wage with the marginal rate of substitution between consumption and leisure. Combining this equation with the marginal utility of wealth in (1.3) implies that the wage is increasing in the expected consumption of the household.

Secondly, the optimal housing demand of the typical worker on island ω is given by

$$\mathbb{E}\left[\left(\frac{\gamma}{1-\gamma}\right)\frac{C(\Omega)}{H(\Omega)}\frac{dH(\Omega)}{dh(\omega,Q)}\middle|\omega,Q\right]=Q(\Omega) \tag{1.4}$$

where

$$\frac{dH(\Omega)}{dh(\omega,Q)}=H(\Omega)^{\frac{1}{\nu}}\,h(\omega,Q)^{-\frac{1}{\nu}}$$

Condition (1.4) simply states that the price of land is equal to the expected marginal rate of substitution between land and consumption. Solving this for $h(\omega,Q)$ gives the following expression for the optimal demand for land of the typical worker on island ω when the asset price is Q, in terms of local expectations of aggregate output and land.

$$h(\omega,Q)=\left[\frac{1}{Q(\Omega)}\mathbb{E}\left[\left(\frac{\gamma}{1-\gamma}\right)\frac{C(\Omega)}{H(\Omega)}H(\Omega)^{\frac{1}{\nu}}\middle|\omega,Q\right]\right]^{\nu} \tag{1.5}$$

Firm Optimality. I now consider the firm's problem. In stage 1 the representative firm on island ω chooses production and employment to maximize the following objective,

$$\max\mathbb{E}\left[\lambda(\Omega)\left(y(\omega,Q)-w(\omega,Q)n(\omega,Q)\right)|\,\omega,Q\right], \tag{1.6}$$

subject to the production function $y(\omega,Q)=A(\omega)n(\omega,Q)^{\theta}$, and the collateral constraint given by

$$w(\omega,Q)\,n(\omega,Q)\leq\chi QL(\omega) \tag{1.7}$$

If credit constraints are not binding, the production choice of the typical firm on island ω is pinned down by the standard first-order condition

$$w(\omega,Q)=\theta\frac{y(\omega,Q)}{n(\omega,Q)} \tag{1.8}$$

which simply states that the unconstrained firm produces where the marginal cost of labor equals its marginal product. The representative firm on island ω would always like to produce the amount dictated in (1.8) whenever possible, as it is the firm's profit maximizing quantity. If the collateral constraint is binding however, the firm's output is pinned down by the constraint (1.7), satisfied at equality.

Partial Equilibrium. In equilibrium, the local labor market within each island must clear. By combining the optimal employment choices of the firms and workers, along with market clearing, I obtain the following result.

Lemma 2 *For any asset price Q, the equilibrium production of the typical firm on island ω is given by*

$$y(\omega, Q) = \min\left\{y^u(\omega, Q), y^c(\omega, Q)\right\} \tag{1.9}$$

where the functions $y^u : \mathcal{S}_\omega \times \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ and $y^c : \mathcal{S}_\omega \times \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ are defined as follows

$$y^u(\omega, Q) \equiv A(\omega)^{\frac{1}{1-\theta}}\left[\frac{\theta}{w(\omega, Q)}\right]^{\frac{\theta}{1-\theta}} \tag{1.10}$$

$$y^c(\omega, Q) \equiv A(\omega)\left[\chi Q \frac{L(\omega)}{w(\omega, Q)}\right]^{\theta} \tag{1.11}$$

where

$$w(\omega, Q) = \mathbb{E}\left[(1-\gamma)\left(\frac{L(\Omega)}{Y(\Omega)}\right)^{\gamma}\middle|\, \omega, Q\right]^{-1} \tag{1.12}$$

The functions y^u and y^c defined in (1.10) and (1.11), are equivalent to conditions (1.8) and (1.7), respectively, in which I have used the production function to solve for local equilibrium output. The function $y^c(\omega, Q)$ denotes the amount of output produced by the firm on island ω such that its collateral constraint is binding. The representative firm on island ω thus produces unconstrained output $y^u(\omega, Q)$ whenever possible unless this choice violates the collateral constraint, i.e. $y^u(\omega, Q) > y^c(\omega, Q)$, in which case the firm has no choice but to produce $y^c(\omega, Q)$. Therefore, the production of the typical firm on island ω is given by $y(\omega, Q) = \min\left\{y^u(\omega, Q), y^c(\omega, Q)\right\}$, when the asset price is Q.

Lemma 2 characterizes the equilibrium level of production on each island. One may think of this as the partial equilibrium within an island in the following sense. Given Q and local expectations of aggregate output and aggregate land endowments, the demand for labor by the firm is given by (1.9) while the supply of labor satisfies the wage equation given in (1.12).

Combined, these conditions fully describe the equilibrium levels of local employment and local output on an island in terms of the realized asset price and local expectations. These last objects–the asset price and local expectations of aggregate output–are, however, determined endogenously and depend on the general equilibrium properties of the economy.

One important lesson to take from Lemma 2 is that the asset price Q directly affects the production of the constrained firm. Because the value of collateral determines how much labor a constrained firm may hire, this firm's production is increasing in Q. This is precisely the feedback from the asset price to aggregate output. While obviously not a result in and of itself, this feature is important, and is highlighted in the following remark.

Remark 3 *The production of a constrained firm is an increasing function of the asset price. That is, there is a feedback from asset prices to real output.*

In contrast, for the unconstrained firms the asset price does not affect firm output, and hence this channel is absent.

Market Clearing and General Equilibrium. Finally, market clearing conditions then pin down the general equilibrium of this economy. First aggregate output is simply given by the total sum of output produced by all firms, $Y(\Omega) = \int y(\omega, Q) d\Omega(\omega)$.

I now consider market clearing in the land market. Land market clearing imposes that the sum of the individual housing demand equations of the workers, given in (1.5), aggregated over all islands, must be equal to the supply of land, i.e. $L(\Omega) = \int h(\omega) \, d\Omega(\omega)$. This market-clearing condition determines the equilibrium asset price, $Q(\Omega)$. The equations that determine equilibrium asset price are presented in the appendix. Here, for expositional purposes, I take a log-linear approximation (which will be justified later) and present the equilibrium price in the following terms.

Lemma 4 *In equilibrium, the asset price is given by*

$$\log Q(\Omega) = \left(\overline{\mathbb{E}} \log Y(\Omega) - \overline{\mathbb{E}} \log L(\Omega)\right) - \frac{1}{\nu}\left(\log L(\Omega) - \overline{\mathbb{E}} \log L(\Omega)\right) + const \tag{1.13}$$

where the operator $\overline{\mathbb{E}}$ denotes the average expectation in the population.

This condition has a simple interpretation. The first term is the average expectation of the marginal rate of substitution between land and consumption. Note that under complete information the land price is given by $Q(\Omega) = \frac{\gamma}{1-\gamma}\frac{Y(\Omega)}{L(\Omega)}$, so that the asset price is exactly equal to the marginal rate of substitution between land and consumption. The second term then is merely the average forecast error of the land endowment. This is due to the fact that under dispersed information, the average expectation will not be equal to the actual amount of land.

The main property I highlight here is the fact that the asset price is increasing in the average expectation of aggregate output. In fact, the elasticity of the asset price with respect to the average expectation of aggregate output is 1. This thereby represents the feedback from expectations of aggregate output to the asset price. This point is underscored in the following remark.

Remark 5 *The asset price is increasing in the average expectation of aggregate output. That is, there is a feedback from real output to asset prices.*

Remarks 3 and 5 form the basis in this model for what one would call the two-way feedback between the financial and real sides of the economy generated by collateral constraints. One one hand the asset price is sensitive to expectations of aggregate output, while at the same time output is sensitive to the asset price when firms are constrained by collateral.

In summary, the general equilibrium of the economy, as defined in Definition 1 is fully characterized by the following lemma.

Lemma 6 *The equilibrium levels of local and aggregate output and the equilibrium asset price are given by the joint solution to the following fixed point problem*

(i) The local production strategy $y(\omega, Q)$ satisfies all conditions in Lemma 2, $\forall(\omega, Q)$.

(ii) Aggregate output is given by $Y(\Omega) = \int y(\omega, Q(\Omega))d\Omega(\omega)$, $\forall\Omega$.

(ii) The asset price satisfies condition (1.13) in Lemma 4, $\forall\Omega$.

Condition (i) merely restates that the local production strategy must satisfy the optimality conditions of the firm and worker in the partial equilibrium within any island. Condition (ii) states that the product market must clear. Finally, condition (iii) is given by equilibrium in the asset market.

It is important to note that Q actually plays three roles in the equilibrium of this economy. First, as workers' demand functions for the asset are not completely inelastic, in equilibrium the asset price is such that it clears the land market. This is the standard function of a price in any Walrasian setting. The second role it plays is the value of collateral, and thereby directly affects the production of constrained firms. This is not a standard role of the asset price, as this effect would be absent in any frictionless business cycle setting. However, this effect would be present in any model which adds the specific friction of collateral constraints. Note that these first two roles just described form the basis for the main mechanism underlying any model with collateral constraints: it is specifically these two features of the asset price which create the two-way feedback between the financial and real sides of the economy. In the current model, too, these functions of the asset price are key to the results. Finally, the third role the asset price plays in this equilibrium is that it affects expectations by partially aggregating information in a dispersed information environment. This role would not be present in any standard complete-information environment, with or without financial frictions, but it is present in any model in which one introduces dispersed information with an observable price, i.e. a Noisy REE setting.

1.4 Complete Information Benchmark

In this section I solve for the general equilibrium of the economy under the assumption that information is commonly shared across islands. I therefore consider the following special case of the baseline environment

Case 7 *Firms and workers on every island have complete information about the aggregate state Ω in stage 1.*

When all information is commonly shared, aggregate output is also commonly known in equilibrium and hence all firms and workers have complete information about the aggregate shocks hitting the economy. I show that in this environment by design collateral constraints have only a level effect on the equilibrium value of aggregate output, but have no effect on the equilibrium sensitivity of output to underlying shocks.

Allow me to first state the result, and then show how it is true. The equilibrium level of aggregate output and the asset price may be characterized by the following lemma.

Lemma 8 *Suppose that information were complete as in Case 7. Then the equilibrium levels of aggregate output and the asset price are given by*

$$\begin{bmatrix} \log Y(\Omega) \\ \log Q(\Omega) \end{bmatrix} = m(\chi) + \begin{bmatrix} \phi_a \\ \phi_a \end{bmatrix} \bar{a}(\Omega) + \begin{bmatrix} \phi_{y,l} \\ \phi_{q,l} \end{bmatrix} \bar{l}(\Omega)$$

where

$$\phi_a = \frac{1}{1-\theta(1-\gamma)}, \qquad \phi_{y,l} = \frac{\theta\gamma}{1-\theta(1-\gamma)}, \qquad \phi_{q,l} = -\frac{1-\theta}{1-\theta(1-\gamma)}$$

and $m(\chi) \equiv \begin{bmatrix} m_y(\chi) & m_q(\chi) \end{bmatrix}'$ *is a* 2×1 *vector-valued function of* χ.

Consider the optimal production of each firm. Substituting the equilibrium behavior of aggregate output and the asset price into the conditions stated in Lemma 2, and using the fact that $a(\omega) = \bar{a}(\Omega) + \xi_a(\omega)$ and $l(\omega) = \bar{l}(\Omega) + \xi_l(\omega)$, I rewrite the equilibrium production on island ω as follows.

Lemma 9 *Suppose that information were complete as in Case 7. In equilibrium, the representative firm on island* ω *produces output*

$$\log y(\omega) = \min\{\log y^u(\omega), \log y^c(\omega)\}$$

where

$$\log y^u(\omega) = \phi_a \bar{a}(\Omega) + \phi_{y,l}\bar{l}(\Omega) + \frac{1}{1-\theta}\xi_a(\omega) - \frac{\theta}{1-\theta}\gamma m_y(\chi) + const^u \tag{1.14}$$

$$\log y^c(\omega) = \phi_a \bar{a}(\Omega) + \phi_{y,l}\bar{l}(\Omega) + \log\chi + \xi_a(\omega) + \theta\xi_l(\omega) + \theta(1-\gamma)m_y(\chi) + const^c \tag{1.15}$$

and $const^u$ *and* $const^c$ *are constants composed entirely of the parameters* θ *and* γ.

If there were no collateral constraints, the representative firm on island ω would always find it optimal to produce the amount dictated in (1.14), as it is its profit maximizing quantity. If the constraint is binding however, i.e. the amount the firm would produce if unconstrained is greater than that given by its constraint, then the firm's output is dictated by (1.15). The representative firm on island ω is thus constrained if and only if $\log y^u(\omega) > \log y^c(\omega)$. This implies that I may characterize the set of constrained firms in the following way.

Lemma 10 *Suppose that information were complete as in Case 7. In equilibrium, the representative firm on island ω is constrained if and only if ω is an element of the set $S_\chi \subset \mathcal{S}_\omega$, where the set S_χ is defined as*

$$S_\chi \equiv \{\omega : \xi_a(\omega) - (1-\theta)\,\xi_l(\omega) > \mu(\chi)\}$$

with

$$\mu(\chi) = \left(\frac{1-\theta}{\theta}\right)\log\chi + (1-\theta(1-\gamma))\,m(\chi) + const$$

Therefore, the constrained firms are those that have the greatest idiosyncratic shocks to their productivity and those that have the lowest idiosyncratic shocks to their land endowments. Furthermore, whether a firm is constrained or not is *independent* of the aggregate realizations of $\bar{a}(\Omega)$ and $\bar{l}(\Omega)$. This is due to the fact that in equilibrium, both constrained and unconstrained firms have the same sensitivity to aggregate shocks. Hence, an aggregate shock to productivity or land shifts all firms up or down in parallel, but does not affect the distribution of constrained and unconstrained firms relative to each other.

This property almost immediately implies that the behavior of aggregate output takes the one conjectured in Lemma 8. Let $\tilde{S}_\chi$ be the complement of S_χ, i.e. $S_\chi \cap \tilde{S}_\chi = \emptyset$ and $S_\chi \cup \tilde{S}_\chi = \mathcal{S}_\omega$; thus $\tilde{S}_\chi$ denotes the set of unconstrained firms. One can then express total output as the following

$$\log Y(\Omega) = \int_{\omega \in S_\chi} \log y^c(\omega)\, d\Omega(\omega) + \int_{\omega \in \tilde{S}_\chi} \log y^u(\omega)\, d\Omega(\omega) + disp(\chi, \sigma_{xa}, \sigma_{xl})$$

where the term *disp* captures the dispersion in production across islands, and hence is a function χ as well as the variance in idioysncratic productivity and land shocks, as well as χ. In general, the function *disp* could be either increasing or decreasing on χ.[13] From this expression it is immediate that the sensitivity of aggregate output to an innovation in aggregate productivity is ϕ_a. Thus, aggregate output indeed satisfies the equilibrium behavior posited in Lemma 8, where the functions $m_y(\chi)$ and $\mu(\chi)$ solve a fixed point problem given in the appendix.

[13] I find sufficient conditions so that *disp* is a decreasing function of χ.

Finally, note that the equilibrium asset price under complete information satisfies condition (1.13), that is,

$$\log Q(\Omega) = \log Y(\Omega) - \log L(\Omega) + const$$

thereby verifying the behavior of the asset price given in Lemma 8.

In Lemma 8, the coefficient ϕ_a summarizes the response of both aggregate output and the asset price to an aggregate productivity shock under complete information. Any variation in χ thereby moves the level of aggregate output and asset prices, but doesn't affect their overall sensitivity to aggregate fundamental shocks. Therefore, in this environment when information is commonly shared, tighter collateral constraints have no effect on the equilibrium behavior of the economy.

Proposition 11 *When information is complete, tighter collateral constraints have no effect on the business cycle properties of the equilibrium.*

This result provides an important benchmark for the business cycle properties of this economy under complete information. When information is commonly shared, both unconstrained and constrained firms exhibit the same sensitivity to aggregate movements in underlying fundamentals. Thus, changing the fraction of constrained firms in the economy has only a level effect on the equilibrium– constrained firms produce less than had they been unconstrained– however, it does not affect the business cycle properties of equilibrium aggregate output.

The natural question to ask in light of Proposition 11 is whether dispersed information significantly changes these results. Before moving to the general equilibrium of the full model with dispersed information, in the next section I provide intution for how and why introducing heterogeneity in information has a profound effect on the equilibrium properties of this economy.

1.5 Collateral Constraints and Dispersed Information

The aim of this section is to convey how the introduction of dispersed information creates a distinct role for financial frictions. In particular, I show how dispersed information interacts with the collateral constraints and delivers vastly different implications for the equilibrium of this economy than in the complete information benchmark. Towards this goal, I make a few important simplifying assumptions. These assumptions are made for purely pedagogical reasons, as the full model is solved numerically in the following sections. However, as I show here, these assumptions simplify the analysis so that the general equilibrium may be solved analytically and in closed-form. This serves to isolate the important mechanisms underlying the main results of the full model, illustrating the impact of dispersed information without any numerical complications.

Specifically, in this section I compare two economies: an economy in which all firms behave as if they were unconstrained and an economy in which all firms behave as if they were constrained. More precisely, I define the first of these economies as follows

Definition 12 *Let the "unconstrained economy" denote an economy in which the equilibrium levels of local and aggregate output and the equilibrium asset price are given by the joint solution to the following fixed point problem*

(i) The local production strategy is given by $y^u(\omega, Q)$ *in condition (1.10), Lemma 2,* $\forall (\omega, Q)$.

(ii) Aggregate output is given by $Y(\Omega) = \int y^u(\omega, Q(\Omega))d\Omega(\omega)$, $\forall\Omega$

(ii) The asset price satisfies condition (1.13) in Lemma 4, $\forall\Omega$.

Thus, the "unconstrained economy" equilibrium is essentially the same as the one described in Lemma 6, except for the following condition: in the "unconstrained economy" all firms follow a production strategy in which they produce their unconstrained optimal output. Conditions (ii) and (iii) are then identical to those in Lemma 6. Note that Definition 12 describes a special case of the full model in which χ approaches infinite. This definition thus describes the equilibrium of a basic one-sector static RBC economy with no collateral constraints, and in so doing, may also serve as an interesting benchmark when thinking about the role of dispersed information within the context of business cycles.

The "constrained economy" is defined in a similar manner.

Definition 13 *Let the "constrained economy" denote an economy in which the equilibrium levels of local and aggregate output and the equilibrium asset price are given by the joint solution to the following fixed point problem*

(i) The local production strategy is given by $y^c(\omega, Q)$ *in condition (1.11), Lemma 2,* $\forall (\omega, Q)$.

(ii) Aggregate output is given by $Y(\Omega) = \int y^c(\omega, Q(\Omega))d\Omega(\omega)$, $\forall\Omega$.

(ii) The asset price satisfies condition (1.13) in Lemma 4, $\forall\Omega$.

Again the equilibrium in the "constrained economy" is essentially the same as the one described in Lemma 6, except that all firms here behave as if they were constrained. That is, in the "constrained economy", all firms follow a production strategy in which they produce up to where their collateral constraint is binding.

Of course, neither equilibrium can be possible in the true general equilibrium of the full model. In the full model, for any finite χ, there will always be a non-zero measure of unconstrained firms which do not find it optimal to produce where their collateral constraint is binding, as well as a non-zero measure of constrained firms which cannot produce their unconstrained optimal amount.[14] Nevertheless, I consider these two economies here primarily for pedagogical reasons: comparing the two economies under dispersed information is instructive in understanding how dispersed information interacts with the collateral constraints. At the same time, assuming that all firms behave as if they were constrained or that all firms behave as if they were unconstrained greatly simplifies the analysis in such a way that the equilibrium fixed-point may be solved in closed form.

As for the information structure, in contrast to the complete information case presented in the last section, this section considers the impact of dispersed information. Instead of considering the full rich information structure outlined in the baseline model, however, for pedagogical reasons I again make another simplifying assumption. Recall that the asset price Q in the baseline environment serves three roles: it is at once the asset price, the value of collateral, as well as an endogenous source of information about the underlying aggregate state. The third role, in which the asset price partially aggregates and conveys information, thereby contributes to an indirect effect on firm output. This indirect effect, however, is subsidiary to the main mechanisms underlying this model, and hence becomes an unnecessary complication in terms of understanding the interaction between dispersed information and collateral constraints.

For this reason, in this section I completely abstract from the indirect effect of the asset price to aggregate information, in order to isolate the important aspects of this model. Towards

[14]Since i have modeled the idiosyncratic shocks to have a Gaussian structure, the support of these shocks is infinite. Thus for any χ there exists a non-zero measure of representative firms who will be unconstrained. This would not be true if the support of the idiosyncratic shocks were finite.

this goal I assume that all agents (in either economy) do not learn from the asset price. That is, although firms and workers on each island observe both their local information ω and the asset price Q, here I simply assume that agents only use their local information to update their priors about the aggregate state, but do *not* use the asset price. Put more simply, I assume no inference is made from observing Q. Moreover, as aggregate land endowment shocks were introduced in the baseline model solely to ensure that the asset price would not be fully-revealing, in this case they are unnecessary. Thus, I furthermore eliminate these shocks. This section therefore considers the following special information structure

Case 14 *Firms and workers on every island have incomplete and heterogenous information about the aggregate state Ω in stage 1:*

(i) The only information firms and workers use to update their priors about the aggregate state Ω in stage 1 is ω

(ii) Aggregate land is constant: $\sigma^2_{L,0} = 0$.

In later sections, I show that the exclusion of learning from the asset price in fact does not make much of a difference, i.e. the main results survive the partial aggregation of information through Q.

The Equilibrium under Dispersed Information. From Proposition 11, it is immediate that under complete information (Case 7) about the aggregate state, the equilibrium of the "unconstrained economy" and that of the "unconstrained economy" have identical business cycle properties. That is, the behavior of equilibrium aggregate output in either economy is identical up to a constant. In this sense, collateral constraints plays no role in affecting the equilibrium of the economy under complete information. I now compare this to the equilibrium in both economies under dispersed information.

Due to the Gaussian specification of shocks and information structure, one may conjecture a log-linear form for the equilibrium aggregate output and asset price. The behavior of these objects is characterized in the following lemma.

Lemma 15 *Suppose that information were dispersed as in Case 14.*

(i) In the unconstrained economy, the equilibrium aggregate output and asset price behave according to

$$\begin{bmatrix} \log Y(\Omega) \\ \log Q(\Omega) \end{bmatrix} = \phi_a \begin{bmatrix} 1 + d_{y,a} & -d_{y,\varepsilon} \\ 1 - d_{q,a} & d_{q,\varepsilon} \end{bmatrix} \begin{bmatrix} \bar{a}(\Omega) \\ \varepsilon(\Omega) \end{bmatrix} + const$$

where $d_{y,a}, d_{q,a}, d_{y,\varepsilon}, d_{q,\varepsilon}$ *are all positive.*

(ii) In the constrained economy, the equilibrium aggregate output and asset price behave according to

$$\begin{bmatrix} \log Y(\Omega) \\ \log Q(\Omega) \end{bmatrix} = \phi_a \begin{bmatrix} 1 - b_{y,a} & b_{y,\varepsilon} \\ 1 - b_{q,a} & b_{q,\varepsilon} \end{bmatrix} \begin{bmatrix} \bar{a}(\Omega) \\ \varepsilon(\Omega) \end{bmatrix} + const$$

where $b_{y,a}, b_{q,a}, b_{y,\varepsilon}, b_{q,\varepsilon}$ *are all positive and less than 1.*

This gives a closed-form solution of the equilibrium aggregate output and asset price in the two economies as a log-linear functions of the productivity shock $\bar{a}$ and the public error ε, which from now on I refer to as noise. The positive scalar b's and d's depend on the the precisions of the signals, as well as the parameters γ and θ.

First, note that aggregate output and asset prices in either economy respond disparately to an innovation in technology. Furthermore, dispersed information opens the door for possible effects of noise. While it is obvious that when one introduces incomplete information noise will have some effect, what is not obvious is how it has an effect. One stark distinction between these two economies is how equilibrium aggregate output reacts to noise: in the constrained economy output reacts positively to noise, while that in the unconstrained economy reacts negatively. I will elaborate on these distinctions shortly.

The differences in equilibrium behavior between the two economies implies that the introduction of dispersed information interacts with collateral constraints, creating a distinct role for the financial frictions. To understand how dispersed information interacts with the collateral constraints and implies results different from the complete info benchmark, it is instructive to understand the equilibrium of the economy at a game-theoretic level. In so doing, this translation explains how dispersed information is a key ingredient in this model. This translation is provided in the following subsection.

1.5.1 Strategic Complementarity and Positive Co-Movement

The equilibrium in either economy may be represented as the Bayesian-Nash equilibrium of a game, similar to those considered in more abstract game-theoretic settings. The literature on games with strategic complementarities (for example, Morris and Shin, 2002; Angeletos and Pavan, 2007; as well as others) considers games with a large number of players, each choosing an action in $\mathbb{R}_+$. In this game, as in the economic model, let ω identify the "types" of these players, and let Ω denote the aggregate state, i.e. the cross sectional distribution of ω. Furthermore, let the "action" of the type ω player be denoted by $\log y(\omega)$, and let $\log Y(\Omega) = \int \log y(\omega) d\Omega(\omega)$ denote the aggregate action; one can think of these respectively as the individual and aggregate levels of production. Finally, suppose the players in this game have payoffs which depend both on private fundamentals $a(\omega)$ as well as the aggregate level of production, so that players have the following "best-response functions"

$$\log y(\omega) = (1 - \alpha) \phi a(\omega) + \alpha \mathbb{E} [\log Y(\Omega) | \omega] \tag{1.16}$$

In this expression, the coefficient α identifies the slope of a player's best response to its expectation of the aggregate action—this is the standard definition of the degree of strategic complementarity (or strategic substitutability). If α is positive then economic decisions are said to be strategic complements; if α is negative then economic decisions are said to be strategic substitutes. Aggregating this condition over all players yields the following equilibrium condition.

$$\log Y(\Omega) = (1 - \alpha) \phi \bar{a}(\Omega) + \alpha \overline{\mathbb{E}} \log Y(\Omega) \tag{1.17}$$

Thus, the Bayesian-Nash equilibrium of this game reduces to a simple fixed-point relation between the aggregate action and the average expectation of aggregate action. In the aggregate equilibrium condition (1.17), α identifies the sensitivity of aggregate production to the average expectation of output.

I now translate the equilibrium of the economies defined in Definitions 12 and 13 in terms of the game just described. Consider the optimal production of each firm in partial equilibrium, as characterized in Lemma 2. Aggregation over individual firm production implies that the equilibrium in the unconstrained and constrained economies can be characterized as follows.

Lemma 16 *Define the coefficients*

$$\alpha_u \equiv -\frac{\theta}{1-\theta}\gamma, \qquad \alpha_c \equiv \theta\,(1-\gamma)$$

so that $\alpha_u \in (-\infty, 0]$ *and* $\alpha_c \in (0,1)$

(i) In the unconstrained economy, the equilibrium level of aggregate output is given by the solution to the following fixed-point problem:

$$\log Y\,(\Omega) = (1-\alpha_u)\,\phi_a \bar{a}\,(\Omega) + \alpha_u \overline{\mathbb{E}} \log Y\,(\Omega) \tag{1.18}$$

(ii) In the constrained economy, the equilibrium level of aggregate output is given by the solution to the following fixed-point problem:

$$\log Y\,(\Omega) = (1-\alpha_c)\,\phi_a \bar{a}\,(\Omega) + \alpha_c \overline{\mathbb{E}} \log Y\,(\Omega) \tag{1.19}$$

This result establishes that the general equilibrium of either economy reduces to a simple fixed-point relation between aggregate output and the average expectation of aggregate output. It is then evident that the general equilibrium of either economy is isomorphic to the Bayesian-Nash equilibrium of the game described above. To see this, note that the conditions (1.18) and (1.19), are direct analogues of the aggregate best response in the abstract game given by (1.17). To further make the translation, think of an island in these economies as a "player" in the game, and the (partial) equilibrium level of output on that island as its corresponding "action". Finally, note that the "types" ω encodes the local shocks and local information sets of each island.

This game-theoretic representation for the equilibrium is useful in that facilitates a translation of some of the more abstract insights from the earlier literature. One of these abstract insights is the concept of strategic complementarity (or strategic substitutability). Note that the coefficients α_u and α_c identify the elasticity of output to variation in average expectations of aggregate output in the constrained and unconstrained economies, respectively. These coefficients are the analogues of the degree of strategic complementarity (or strategic substitutability) in the economic model. In the game described above strategic complementarity is merely an abstract idea or construct, yet here α_u and α_c are instead a result of the particular microfoundations of the economy. In what follows, I examine the micro underpinnings of these coefficients, and most importantly for the purposes of this chapter, show how collateral constraints translates into a novel source of strategic complementarity.

Strategic Complementarity and Strategic Substitutability. The coefficients α_u and α_c identify the sensitivity of output to expectations of aggregate output in the unconstrained and constrained economies, respectively. Understanding the specific micro-foundations for these coefficients are crucial for understanding the economics underlying the main results of this model. Thus, before considering how the information structure affects equilibrium outcomes, I first consider where these coefficients originate. These coefficients are composed of the underlying parameters governing preferences and technology in this economy: γ and θ. Allow me to rewrite these coefficients as follows.

$$\alpha_u = \frac{-\gamma}{\frac{1}{\theta} - 1}, \qquad a_c = \frac{1-\gamma}{\frac{1}{\theta}}$$

From this, one can see that there are similarities between the two coefficients. First, note that the numerators in both expressions contain the term $-\gamma$. This term controls the elasticity of the wage to expectations of aggregate output, and is driven by the income effect on labor supply.[15] Higher aggregate income causes households to desire more consumption of both goods and leisure. Thus, if workers are optimistic about aggregate income, they will work less for any given wage. In equilibrium the inward shift in labor supply leads to a reduction in the equilibrium output produces on an island. This negative income effect on labor, which is standard in any neoclassical setting, is thereby a source of strategic substitutability in both the unconstrained and constrained economy.

Secondly, the denominators in either expression merely reflect the fact that there are decreasing returns to scale in production: a one-percent increase in employment hired by the firm

[15]To see this clearly, note that labor is quasilinear in the consumption basket $C^{1-\gamma}L^{\gamma}$, where the share of income spent on land is γ. One can reexpress the wage as equal to the price of this consumption basket, which is given by $(1-\gamma)+\gamma Q$, where Q increases 1-for-1 with expectations of aggregate output. Thus, for a larger share of income spent on land, the price of this consumption basket increases, and in order to get workers to work, wages must increase.

leads to only a θ-percent increase in firm output. As the unconstrained firm's production is given by a condition involving the marginal cost of production (1.10), whereas the firm's collateral constrains the total cost of production (1.11), anything that independently shifts these two conditions will have an effect on output, but subject to a scale factor.

Finally, the primary difference between these two coefficients is that α_c contains a term in the numerator (specifically, the 1) that is unique to the constrained economy and originates directly from the collateral constraint. This term is given by the elasticity of the asset price to expectations of aggregate output. Optimism about aggregate output, as seen in condition 1.13, drives up the price of land. This is precisely the feedback from aggregate output to the asset price. Moreover, recall that the collateral constraint dictates that production is constrained by the price of land. An upward shift in asset prices thereby loosens the collateral constraint, allowing constrained firms to produce more. This is precisely the feedback from the asset price to real output. In total, optimism about aggregate output increases constrained firm production. This positive collateral constraint effect on output is not standard in a neoclassical setting, and as such is absent from the unconstrained economy's sensitivity α_u. Rather, this effect is a direct manifestation of the two-way feedback between the financial market and the real economy, and is hence a distinct *source of strategic complementarity* in the constrained economy.

In summary, note that for all parameter values, α_u is always negative whereas α_c is always positive. Therefore, in total, production in the unconstrained economy are strategic substitutes, whereas production in the constrained economy are strategic complements. As I show next, these total sensitivities are important determinants of the equilibrium when information is asymmetric across islands.

Implications of Dispersed Information. I now consider how the information structure affects the aggregate properties of the equilibrium. One of the most insightful and important results from the earlier more abstract literature on these games is that strategic complementarity is an important determinant of the equilibrium if and only if information is heterogeneous among players. To illustrate this, consider the abstract game presented above with the equilibrium described by the fixed point relation in (1.17). When information is commonly shared, strategic complementarity (or substitutability), in this game represented by α, plays no role. In fact, one can easily check that if $\log Y(\Omega) = \overline{\mathbb{E}} \log Y(\Omega)$, then both conditions (1.18) and (1.19) reduce to

$$\log Y(\Omega) = \phi \bar{a}(\Omega)$$

and hence confirms the result from the complete information benchmark presented in Lemma 8. Specifically, the sensitivities α_u and α_c do not affect equilibrium output in either economy under complete information.

Under dispersed information however, strategic complementarity and substitutability play a critical role in determining the response of aggregate output to underlying shocks. Under the information structure spelled out above, I obtain the following characterization of aggregate output.

Lemma 17 *The equilibrium level of aggregate output in this game is given by*

$$\log Y(\Omega) = \left(1 - \frac{\alpha \kappa_{0a}}{(1-\alpha)\kappa_{xa} + \kappa_{za} + \kappa_{0a}}\right)\phi \bar{a}(\Omega) + \frac{\alpha \kappa_{za}}{(1-\alpha)\kappa_{xa} + \kappa_{za} + \kappa_{0a}}\phi \varepsilon(\Omega) + const$$

and the average expectation of aggregate output is given by

$$\overline{\mathbb{E}} \log Y(\Omega) = \left(1 - \frac{\kappa_{0a}}{(1-\alpha)\kappa_{xa} + \kappa_{za} + \kappa_{0a}}\right)\phi \bar{a}(\Omega) + \frac{\kappa_{za}}{(1-\alpha)\kappa_{xa} + \kappa_{za} + \kappa_{0a}}\phi \varepsilon(\Omega) + const$$

Thus, the fixed point relation in (1.17) yields a closed form solution for the equilibrium level of aggregate output and the average expectation of output as log-linear functions of the fundamental and noise. This is the direct analogue of the results presented in Lemma 15, for the behavior of the aggregate output and the asset price, since $\log Q = \overline{\mathbb{E}} \log Y(\Omega)$. These expressions then confirm the results from the dispersed information presented in Lemma 15, with appropriately specified coefficients b's and d's.

From Lemma 17 we see the following: first, expectations of aggregate behavior are sensitive to the public signal. In particular, expectations of aggregate behavior in either economy respond positively to the noise in the public signal. While this would be true even if information were common but incomplete, dispersed information implies $\log Y(\Omega) \neq \overline{\mathbb{E}} \log Y(\Omega)$. Therefore, it is immediate from the fixed-point relation in (1.17) that noise affects equilibrium output, and, depending on the sign of α, the response of aggregate output to noise can be either positive or negative. When economic decisions are strategic substitutes ($\alpha < 0$), as in the unconstrained economy, the equilibrium level of output responds negatively to expectations of aggregate output. When instead economic decisions are strategic complements ($\alpha > 0$), as in the constrained economy, the equilibrium level of output responds positively to expectations of aggregate output. However, in either case, average expectations always respond positively to noise, and hence asset prices respond positively to expectations of aggregate output. Thus, a positive noise shock, or a correlated error in beliefs, can be thought of as collective optimism among agents about aggregate economic activity. Therefore, under dispersed information, the elasticity of production to expectations of aggregate economic activity then dictate how this collective optimism or pessimism manifests itself in equilibrium output.

Another observation can be made from these conditions about how much noise matters for equilibrium aggregate behavior and how this depends on the size of α. For now, I focus on the sign of the effect of noise, and how this could induce co-movement in variables, before moving to the relative contribution of noise shocks.

Positive Co-movement. In the unconstrained economy economic decisions are strategic substitutes ($\alpha_u < 0$), and therefore the equilibrium level of output responds negatively to expectations of aggregate output. While the mechanics for this result has just been described, the economics behind it also is straight-forward. In a standard one-sector neoclassical setting, greater income causes households to desire both more consumption and more leisure. Anticipation of greater wealth thus makes workers less willing to work, thereby reducing their labor supply, and eventually leading to a reduction in equilibrium output. Therefore, optimism about aggregate economic activity generates a recession.

It is relatively well known (though perhaps counter-intuitive) that in a standard one-sector neoclassical model, expectations of higher income or wealth cause a recession. This fact was first pointed out by Barro and King (1984), and later emphasized by Cochrane (1994). In fact, the difficulty of the neoclassical growth model to generate a boom in response to expectations of higher future TFP has been the main obstacle for the news-driven business cycle literature. See, for example the empirical work of Beaudry and Portier (2004, 2008), the theoretical papers of

Christiano et al.(2007), Jaimovich & Rebelo (2008), and Lorenzoni (2009), and finally the recent paper by Sims (2009), which offers an empirical challenge to the notion of booms being driven by news about future TFP. Although these papers focus primarily on expectations of future TFP as the driving force in dynamic representative-agent models, whereas the model considered here focuses on correlated errors in expectations among heterogeneously informed agents, the negative wealth effect on labor supply is a common denominator in any expectations-driven business cycle model.

Furthermore, it is also worth examining the cyclical behavior of aggregate consumption and employment. Aggregate consumption is given by

$$\log C(\Omega) = \log Y(\Omega) \tag{1.20}$$

while aggregate employment is given by

$$\log N(\Omega) = const + \frac{1}{\theta}\left(\log Y(\Omega) - \bar{a}(\Omega)\right) \tag{1.21}$$

It is then immediate that the response of aggregate consumption to noise is equal to that of output, while the response of employment to noise is proportional to that of aggregate output. Thus, in the unconstrained economy optimism about aggregate output not only implies that output falls, but that consumption and employment falls as well. At the same time, optimism about economic activity causes asset prices to increase, since these are directly tied to expectations of aggregate output. Therefore, not only does good news or optimism about output tend to cause a recession, but the implied negative comovement between aggregate output and asset prices seems difficult to reconcile with the procyclicality of stock prices found in the data.

On the other hand, in the constrained economy economic decisions are strategic complements ($\alpha_c > 0$), and hence the equilibrium level of output responds positively to expectations of aggregate output. Good news about aggregate output drives up the price of land, thereby increasing the value of collateral and allowing constrained firms to produce more. Therefore, in the constrained economy optimism about aggregate economic activity generates a boom in output–in this sense, expectations are partially self-fulfilling. Moreover, from conditions (1.20) and (1.21), it is apparent that in the constrained economy consumption and employment also respond positively to noise in the public signal. Good news about aggregate output thereby generates positive co-movement in output, employment and consumption and asset prices in the constrained environment. I summarize this finding in the following proposition.

Proposition 18 *Suppose that information were dispersed as in Case 14. In the constrained economy, aggregate output, employment, consumption, and asset prices exhibit positive co-movement in response to correlated errors in expectations of aggregate activity.*

Therefore, one of the main results in this chapter is that collateral constraints and dispersed information imply positive co-movement over the business cycle. Moreover, this positive co-movement is caused by variation in expectations orthogonal to underlying productivity.

1.5.2 The Relative Contribution of Noise

The previous subsection focused on the implications of noise in terms of observable aggregate macro variables. I now show how noise may be responsible for a larger fraction of business cycle and asset price fluctuations in the constrained economy than in the unconstrained economy. Towards this goal, I consider a certain decomposition of the variance of aggregate output and the asset price into the elements driven either by noise and or by technology. Using this decomposition, I define what I mean by the relative contribution of noise.

Definition 19 *Suppose a variable x is a linear combination of I independent random variables u_i with variances σ_i^2*

$$x = \sum_{i \in I} g_i u_i$$

where g_i are the loadings on each random variable u_i. The fraction of x explained by u_j is then given by

$$R \equiv \frac{g_j^2 \sigma_j^2}{\sum_{i \in I} g_i^2 \sigma_i^2}$$

This definition is relatively simple: the fraction of x explained by u_j is merely the variance of x due to u_j, divided by the total variance of x. Applying this definition to the log of aggregate output and the log of the asset price, I find the fraction of aggregate output volatility and the fraction of asset price volatility explained by noise. This leads to the following proposition.

Proposition 20 *Suppose that information were dispersed as in Case 14. Consider the two economies in Definitions 12 and 13, both with the same parameter values.*

(i) Noise contributes to a greater fraction of asset price volatility in the constrained economy than in the unconstrained economy.

(ii) For any set of parameters $(\theta, \kappa_{0a}, \kappa_{xa}, \kappa_{za})$, there exists a $\hat{\gamma}(\theta, \kappa_{0a}, \kappa_{xa}, \kappa_{za})$ such that for all $\gamma < \hat{\gamma}$, noise contributes to a greater fraction of output volatility in the constrained economy than in the unconstrained economy.

The first part of Proposition 20 compares the contribution of noise to asset price volatility across the two economies. From Lemma 17, note the following: $d_{q,a} < b_{q,a}$ and $d_{q,\varepsilon} < b_{q,\varepsilon}$. This immediately implies that the equilibrium asset price in the constrained economy has both a lower sensitivity to the aggregate productivity shock than in the unconstrained economy, and a greater sensitivity to noise than in the unconstrained economy. First, compared to the complete information benchmark, the asset price in both economies exhibit a dampened response to the aggregate productivity shock. However, as $d_{q,a} < b_{q,a}$, the asset price in the constrained economy has an even lower sensitivity to aggregate output than the unconstrained. Second, the asset price in both economies respond positively to noise. But, as $d_{q,\varepsilon} < b_{q,\varepsilon}$, the asset price in the constrained economy is in fact more sensitive to noise than in the unconstrained economy. Together, these conditions are sufficient to prove that the fraction of asset price volatility due to noise is always larger in the constrained economy than the unconstrained economy.

The second part of Proposition 20 compares the contribution of noise to aggregate output volatility across the two economies. While the fraction of asset price volatility due to noise is always larger in the constrained economy than the unconstrained economy, the same isn't always the case for output volatility. However, I show that for reasonable parameter values, noise contributes to a greater fraction of output variance in the constrained economy than in the unconstrained volatility.

Towards this goal, I can show that the fraction of output volatility accounted for by noise in the constrained economy is decreasing in the parameter γ, whereas the fraction of output volatility accounted for by noise in the unconstrained economy is increasing γ. Thus, for γ sufficiently low the constrained economy has a greater fraction of output variance accounted for by noise than the unconstrained economy.

The mechanics for why this is true is as follows. First, from Lemma 15, note that for all parameter values, aggregate output in the constrained economy has a lower sensitivity to aggregate productivity than in the unconstrained economy. This is because stronger complementarity, i.e. $\alpha_c > \alpha_u$, raises the anchoring effect on the common prior. In other words, in the constrained economy, since economic decisions are strategic complements, production is more sensitive to public sources of information and the prior than to idiosyncratic information. On the other hand, in the unconstrained economy, since economic decisions are strategic substitutes, local output relies more on idiosyncratic information than on public sources of information and the common prior. It follows that in the economy with higher strategic complementarity, that is, the constrained economy, equilibrium output is less responsive to underlying shifts in aggregate productivity. This is one force which makes productivity contribute relatively less to aggregate output volatility in the constrained economy relative to the unconstrained economy.

What then matters is the comparison across economies of the sensitivity of aggregate output to noise. From Lemma 15 one sees that equilibrium aggregate output in the constrained economy reacts positively to noise, while equilibrium aggregate output in the unconstrained economy reacts negatively noise. Since public information helps forecast the aggregate level of output, aggregate output in the economy with strategic complementarity (the constrained economy) reacts positively to public information whereas aggregate output in the economy with strategic substitutibility (the unconstrained economy) reacts negatively. In absolute terms, either sensitivity could be higher, and the coefficients that control these sensitivities are α_c and α_u. The greater α_c and α_u are in absolute terms, the greater the sensitivity of aggregate output to noise. Recall that γ controls the income effect on labor supply, and is hence the only source of strategic substitutibility in either economy. Thus, while a lower γ implies both a greater α_c and α_u, in terms of absolute values a lower γ implies a higher $|\alpha_c|$ but a lower $|\alpha_u|$. This then explains why the fraction of output volatility accounted for by noise in the constrained economy is decreasing in the parameter γ, whereas the fraction of output volatility accounted for by noise in the unconstrained economy is increasing γ. Therefore, in order for noise to contribute to a greater fraction of output variance in the constrained economy than in the unconstrained economy, a weak income effect on labor supply will suffice.

I illustrate this in the following figure. On the left-hand side I plot the fraction of output variance accounted for by noise in both the constrained and unconstrained economies. In this figure I let $\sigma_{0a} = 0.02$ for the standard deviation of the productivity innovation. I set the standard deviations of the private and public information at $\sigma_{xa} = \sigma_{za} = 4\sigma_0$. These values are arbitrary, but not implausible: when the period is interpreted as a quarter, the information about the current fundamentals and/or the current level of economic activity is likely to be very limited. Finally, I let $\theta = .99$ so that technology is near constant returns to scale– I soon show that this gives a lower bound in terms of $\hat{\gamma}$. Finally, the right hand side plots the fraction of asset volatility accounted for by noise in both the constrained and unconstrained economies for the same parameter values.

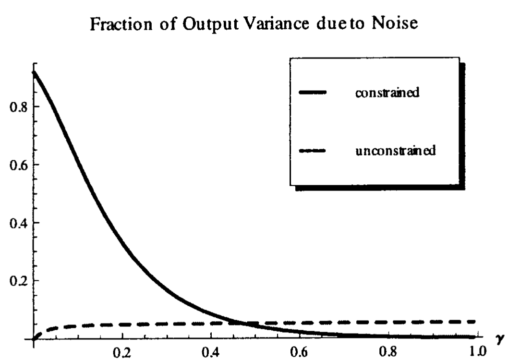

Figure 1.1a

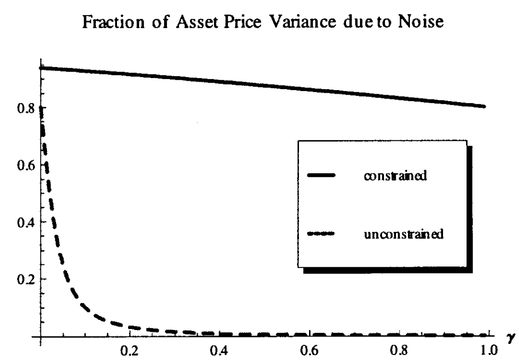

Figure 1.1b

As already stated, the asset price variance due to noise is always greater in the constrained economy than in the unconstrained economy. In terms of output volatility, one sees that the fraction of output volatility accounted for by noise in the constrained economy is decreasing in the parameter γ, and that the fraction of output volatility accounted for by noise in the unconstrained economy is increasing γ. Furthermore, for $\gamma = 0$ the fraction of output volatility acounted for by noise in the unconstrained economy is zero, whereas the fraction of output volatility accounted for by noise in the constrained economy is strictly positive. Together, this suggests that there exists a cutoff $\hat{\gamma}$, here $\hat{\gamma} \approx .45$, such that for any γ lower than $\hat{\gamma}$, the relative contribution of noise to aggregate output volatility is greater in the constrained economy than in the unconstrained economy.

In the following figure, I show that for given $\sigma_{xa}, \sigma_{za}, \sigma_{0a}$, the function $\hat{\gamma}$ is decreasing in θ. Thus for any set of shocks and information structure, the case where $\theta = .99$ gives the lowest possible $\hat{\gamma}$, in other words, the lowest upper bound for γ.

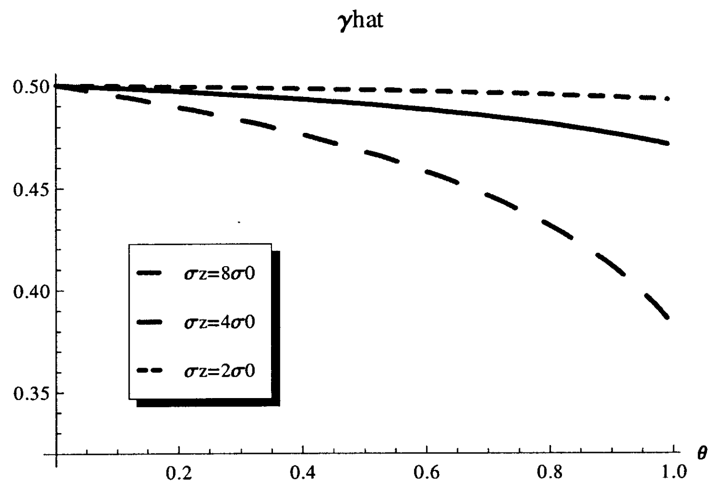

Figure 1.2

Thus, if I had assumed there had been a fixed factor in production such as capital, so that θ was equal to the income share of labor, say .6, the upper bound for γ would then be higher and hence easier to satisfy. Note from this figure that for high values of the standard deviation of noise in the public signal, the function $\hat{\gamma}$ is relatively flat in terms θ.

In any case, by considering γ sufficiently low, I am effectively assuming a small income effect on labor supply. For business-cycle frequencies, it seems quite plausible that the short-run wealth effect on labor supply is weak, and many papers calibrate their models as such. For example, Jaimovich and Rebelo (2009) use preferences such that the wealth effect on labor supply can be controlled and weakened; in their preferred calibrations, γ is between 0 and .25. Woodford (2003) also uses a similar number. Thus, γ between .1 and .3 seems to be a relevant range for this parameter in terms of matching empirically plausible short-run income effects on labor supply. Moreover, in terms of this model this number implies that about 75% of household income is spent on consumption goods, while the rest is spent on housing– a decomposition of wealth that seems fairly likely. Therefore, reasonable values for γ seem to always fall within the relevant range, i.e. below the $\hat{\gamma}$ given by plausible parameterizations of θ and σ_{za}/σ_{0a} shown above.

I therefore argue that the empirically plausible case is one in which γ is sufficiently small, and hence the relative contribution of noise to both asset price and output volatility is greater in the constrained economy than in the unconstrained economy.

1.5.3 Movements in the Labor Wedge

Finally, I study the cyclical properties of the labor wedge. Following the literature, I define the labor wedge $1 - \tau(\Omega)$ implicitly by

$$\frac{1}{(1-\gamma)\left(\frac{L(\Omega)}{Y(\Omega)}\right)^{\gamma}} = (1-\tau(\Omega))\,\theta\frac{Y(\Omega)}{N(\Omega)} \tag{1.22}$$

Namely, the labor wedge is the wedge between the marginal rate of substitution between consumption and leisure and the marginal product of labor.[16] First, note that under complete information, it is trivial to check that the labor wedge is always equal to 1 and does not vary in response to any shock. In contrast, the introduction of dispersed information changes this prediction. In this model, the equilibrium behavior of the measured labor wedge under dispersed information may be described in the following lemma.

Lemma 21 *Suppose that information were dispersed as in Case 14*

(i) In the unconstrained economy, the equilibrium labor wedge behaves according to

$$\log(1-\tau(\Omega)) = \begin{bmatrix} d_{y,a} & -d_{y,\varepsilon} \end{bmatrix} \begin{bmatrix} \bar{a}(\Omega) \\ \varepsilon(\Omega) \end{bmatrix} + const$$

[16] Note that in general, the labor wedge is defined by using standard CRRA preferences for consumption. Here, however, prefences are over both consumption and land, thus if I had defined the labor wedge in the usual way, there would be uninteresting effects coming from the misspecification of preferences by the econometrician. Instead, I opt for specifying preferences correctly and thus isolate the results produced solely from the dispersion of information.

(ii) In the constrained economy, the equilibrium labor wedge behaves according to

$$\log(1-\tau(\Omega)) = \begin{bmatrix} -b_{y,a} & b_{y,\varepsilon} \end{bmatrix} \begin{bmatrix} \bar{a}(\Omega) \\ \varepsilon(\Omega) \end{bmatrix} + const$$

The labor wedge thus exhibits very dissimilar responses to noise in the two economies. In the constrained economy, the labor wedge $1-\tau(\Omega)$ responds positively to noise, or correlated errors in expectations of aggregate output. The intution for why the labor wedge increases in response to noise is given by the following. When there is a positive noise shock, asset prices increase because agents think output will increase. As a result, firms become less constrained and increase their output, thus making the noise shock, or expectational error, somewhat self-fulfilling. In the process, though, firms hire more labor even if there has been no change in their productivity. Thus, the economy behaves as though there were a decrease in the tax on labor, distorting firms to increase employment, and thereby leading to an increase in the measured wedge between the MRS and the marginal product of labor.

In contrast, in the unconstrained economy the measured labor wedge is falls in response to noise. The intuition for this is the following. A positive noise shock causes workers to believe aggregate output will be high and thus workers would like to have more leisure and work less. This drives up wages, causing firms to reduce both employment and production. In other words, the inward shift of labor supply causes firms in equilibrium to reduce their employment even if there has been no change in their productivity. Thus, the economy behaves as though there were an increase in the labor tax, distorting firms to hire less workers, and thereby leading to a decrease in the measured labor wedge.

Proposition 22 *Suppose that information were dispersed as in Case 14. In the constrained economy, noise leads to a positive response in the measured labor wedge.*

Therefore, one of the main results of this model is that the financial friction considered here manifests itself as a time-varying labor wedge. Specifically, the interaction of collateral constraints and dispersed information implies that the labor wedge responds positively to common errors in expectations about aggregate activity. The empirical relevance of this statement will be discussed in the following sections.

The purpose of this section was to illustrate the feedback between asset prices and output, and show how dispersed information interacts with the financial frictions. Of course I have abstracted from many important features of the full model. In general equilibrium not all firms will be constrained and furthemore agents in the economy will learn from observing the asset price. In the following sections, I solve for the general equilibrium of the full economy and demonstrate how the theoretical insights highlighted in this section are not lost in the full model.

1.6 Equilibrium in the Full Model

In this section I describe the numerical method I use to solve for the equilibrium in the full model. Recall that in the previous section, equilibrium aggregate output could be solved for in closed-form in either economy. By imposing that all agents are either constrained or unconstrained, aggregate output in either case was log-linear in the underlying fundamentals. Combining this with the Gaussian information structure, firms and workers faced simple Bayesian inference problem, inducing log-linear production strategies and a fixed-point in equilibrium output.

In considering the full model, I now take into account that in equilibrium there are both constrained and unconstrained firms in the economy. As seen in the partial equilibrium characterization in Lemma 2, constrained and unconstrained firms react differently to their information sets. This creates a non-linearity in the aggregation of output across firms–a complication when agents must make optimal forecasts of aggregate output. Thus, it is mainly a problem of inference of aggregate behavior that prevents the equilibrium from admitting a closed-form analytical solution. In this sense, the numerical method I use is similar in spirit to that of Krusell-Smith (1998). Krusell-Smith (1998) examines an economy in which there is substantial heterogeneity in wealth, and show how the equilibrium can be approximated numerically, despite the fact that the state of the economy at any point in time is the distribution of wealth across agents–an infinite-dimensional object. Importantly, they show how a general equilibrium model with wealth heterogeneity may feature approximate aggregation, that is, aggregate endogenous variables can be described as a function of first moments, as opposed to the entire wealth distribution, without serious error.

Similarly, in this model there is heterogeneity in information, or types ω, across islands. The exogenous aggregate state variable in the economy is the realized distribution of types, Ω, in the cross-section of islands. The aggregate variables, $Y(\Omega)$ and $Q(\Omega)$, are endogenous functions of this state variable, and in theory could be highly complicated functions of the underlying distribution. I show how these aggregate endogenous variables may be approximated as simple log-linear functions of the first moments of the distribution of types. This greatly simplifies the analysis as firms and workers in stage 1 face manageable inference problems when making their employment, production, and asset demand decisions. In this sense, I follow a similar approximate aggregation technique as in Krusell-Smith (1998) in that I show how the endogenous aggregate variables can be simple functions of the distribution first moments, yet lead to small errors in the implied behavior of actual output.

I now outline my numerical strategy for computing the equilibrium. First, suppose that all agents in the economy perceive that aggregate output evolves according to a log-linear function of the first moments of Ω. Formally, I assume that agents perceive the law of motion for aggregate output Y to be given by a function F that belongs to a class $\mathcal{F}$ of log-linear functions

Definition 23 *Let $\mathcal{F}$ be the class of log-linear functions of the first moments of the aggregate state.*

$$\mathcal{F} \equiv \left\{ F : \mathcal{S}_\Omega \to \mathbb{R}_+ \left| \log F(\Omega) = \Phi_a \bar{a} + \Phi_l \bar{l} + \Phi_z z, \text{ and } \Phi_a, \Phi_l, \Phi_z \in \mathbb{R} \right. \right\}$$

This log-linear approximation of aggregate output greatly simplifies the inference problem of the firms and workers in Stage 1. I now show that combining this perceived law of motion with the Gaussian shocks and information structure generates a log-linear asset price and log-linear output strategies.

Information Structure and Asset Price First, for any F, there is a fixed-point relation between the equilibrium asset price and the equilibrium information structure. The asset price in this fixed point takes a log-linear form given in the following lemma.

Lemma 24 *For any F, the land price which clears the market is a log-linear function of the* *first moments of the aggregate state*

$$\log Q(\Omega; F) = \Psi_a \bar{a}(\Omega) + \Psi_l \bar{l}(\Omega) + \Psi_z z(\Omega)$$

where Ψ_a, Ψ_l, Ψ_z *are functions of* F.

For any asset price, there is a unique information structure generated by this price. This is because firms and workers in stage 1 observe the land price Q, and hence use the price to make inference about the underlying state. The log-linear nature of the asset price is important for tractability, as it preserves the Gaussian structure of the information. Since z is common knowledge, the asset price can be transformed into a simple Gaussian signal about the underlying aggregate productivity and land endowment:

$$\hat{q} \equiv \log Q(\Omega; F) - \Psi_{zz} z(\Omega) = \Psi_a \bar{a}(\Omega) + \Psi_l \bar{l}(\Omega)$$

I then solve the individual island's inference problem as follows. Let X_1 be the vector of aggregate fundamentals and X_2 the vector of signals observed by each island

$$X_1(\Omega) = \begin{bmatrix} \bar{a}(\Omega) \\ \bar{l}(\Omega) \end{bmatrix}, \qquad X_2(\omega) = \begin{bmatrix} a(\omega) \\ l(\omega) \\ z \\ \hat{q} \end{bmatrix}$$

and finally let $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$ be the unconditional variance-covariance matrix for these vectors. Note that while X_2 is different for each island, the matrix Σ will be the same for every island, since all signals are assumed to have the same stochastic properties. Using their signals, firms and workers update their priors and form their expectations of aggregate productivity and land holdings. Their conditional expectations are given by the following.

Lemma 25 *For any* F, *the expectations of firms and workers on island* ω, *conditional on their signals, are given by*

$$\mathbb{E}[X_1|\omega] = \Sigma_{12}\Sigma_{22}^{-1} X_2(\omega) \tag{1.23}$$

where Σ and $\hat{q}$ are functions of F

Thus the expectations of the aggregate state are linear functions of the signals, and are generated by the Bayesian inference made from the observed signals and the asset price. Note that this implies that the average expectation of the aggregate state is then given by

$$\mathbb{E}[X_1] = \Sigma_{12}\Sigma_{22}^{-1}\bar{X}_2$$

where $\bar{X}_2 = \left[\begin{array}{cccc} \bar{a} & \bar{l} & z & \hat{q} \end{array}\right]'$ is simply the mean of X_2. Therefore, the average expectation is furthermore linear in the first moments of Ω.

Furthermore, for any given information structure, there exists a unique asset price. Recall that the land price which clears the market is given by condition (1.13). Combining this with the linear average expectations that result from the agents' inference problem, the asset price that prevails in equilibrium can in fact be expressed in the log-linear form, as stated in Lemma 24.

Thus, for any F, there is a fixed-point between the equilibrium information structure and the equilibrium asset price, as standard in rational-expectations equilibria.[17]

Individual Firm Optimality Second, for any F, I may characterize the equilibrium output produced on each island. Lemma 2 describes the partial equilibrium output produced on an island, in terms of the equilibrium asset price and expectations of the aggregate state. Moreover, Lemma 25 chactererizes the expectations of the aggregate state of firms and workers on an island. Combining these results, one may find the equilibrium output produced on any island, as a function of its information set. Thus, given the law of motion F, each island's production can be represented by a decision rule $g(\omega, Q; F)$

Lemma 26 *The optimal production of typical firm on island ω is given by*

$$\log g(\omega, Q; F) = \min\left\{\log g^u(\omega, Q),\ \ \log g^c(\omega, Q)\right\}$$

[17]Note that for any given information structure, there exists a unique asset price function, and for any asset price function there exists a unique information structure generated by this price. Nevertheless, multiplicity can originate in the fixed-point relation between the two, as often is the case with rational-expectations equilibria. While this possibility is intriguing, but I will ignore because it is orthogonal to the goals of this paper.

where the functions g^u and g^c are given by

$$\begin{aligned}\log g^u(\omega, Q) &= \phi_a^u a_i + \phi_l^u l_i + \phi_z^u z + \phi_q^u \hat{q} \\ \log g^c(\omega, Q) &= \phi_a^c a_i + \phi_l^c l_i + \phi_z^c z + \phi_q^c \hat{q}\end{aligned}$$

where the coefficients $\{\phi_a^u, \phi_l^u, \phi_q^u, \phi_z^u; \phi_a^c, \phi_l^c, \phi_q^c, \phi_z^c\}$ are functions of F.

Here, the functions g^u and g^c merely represent the production of unconstrained and constrained firms, given in (1.10) and (1.11) of Lemma 2, in which I have used the agents' perceived law of motion for aggregate output, F, and their conditional expectations of the aggregate state (given by Lemma 25). Note that conditional on being either unconstrained or constrained, output is log-linear in the islands' signals.

Aggregation Finally, given such a production strategy g for individual islands and the distribution of types ω, it is then possible to derive the implied aggregate output by aggregating the output produced in the cross-section of islands. Aggregation over island production yields and aggregate output function which I call G,

$$G(\Omega; F) \equiv \int_{\omega} g(\omega, Q; F)\, d\Omega(\omega)$$

Thus, given the log-linear function F agents perceive for aggregate output, the optimal decisions of firms and workers induces an aggregate output function G which is not log-linear in the underlying state. Note that the only reason that the implied aggregate output is not log-linear in $\bar{a}, \bar{l}, z$ is due to aggregation over both constrained an unconstrained firms.

I may then compare the implied aggregate output function G to the perceived aggregate output function F on which agents base their behavior, by looking at the mean-squared error between the two. The mean squared error is calculated over all realizations of Ω, weighted by its pdf. I then minimize this mean-squared error over the the class of functions $\mathcal{F}$.

$$\min_{F \in \mathcal{F}} \int_{\Omega} (F(\Omega) - G(\Omega; F))^2 \, d\mathcal{P}(\Omega) \tag{1.24}$$

By this I mean that numerically I find coefficients Φ_a, Φ_l, Φ_z which minimize the mean squared error between what agents perceive and what is then induced in equilibrium. Therefore the approximate equilibrium is a function F such that, when taken as given by the agents, yields the best fit within the class $\mathcal{F}$ to the behavior of Y.

Note that had all firms been unconstrained, or all firms constrained, then the law of motion in equilibrium would in fact be log-linear. Thus, the approximation error approaches zero at either extreme: when almost all agents are constrained or when almost all agents are unconstrained. This is apparent from the previous section, in which I showed that the equilibrium could be solved analytically in either "unconstrained economy" and the "constrained economy", which implies that there is zero approximation error for the law of motion of aggregate output. I later show that it is only for the intermediate ranges where there is a non-zero, but small approximation error.

The results from this numerical approximate solution are presented in the following section.

1.7 Numerical Results

In this section I present the numerical results to the full model. For these simulations, I use values similar to those used in Section 1.5. In particular, I let $\sigma_{0a} = 0.02$ for the standard deviation of the productivity innovation and set the standard deviation of land shocks to the same value $\sigma_{0l} = \sigma_{0a}$. Furthermore, $\sigma_{za} = \sigma_0$ so that the standard deviation of public information is equal to that of productivity. I set the standard deviations of the private information at $\sigma_{xa} = \sigma_{xl} = 4\sigma_0$. I let $\theta = .9$ so that technology is near constant returns to scale, and $\gamma = .1$ to ensure an empirically plausible income effect on labor supply. I set $\nu = 2$, an arbitrary number but robustness checks show that the quantitative results are not extremely sensitive to this value, as it merely controls how sensitive the asset price is to forecast errors about land shocks. Finally, I do not pick any specific value for χ, the fraction of land which is collateralizable. Rather, I focus on comparative statics with respect to χ, and study how the sensitivitiy of equilibrium output to productivity, land, and noise varies as I vary the tightness of the collateral constraints.

In the following figure I graph the fraction of constrained firms as a function of χ. This merely illustrates that for lower values of χ, i.e. tighter collateral constraints, there are more constrained firms in equilibrium.

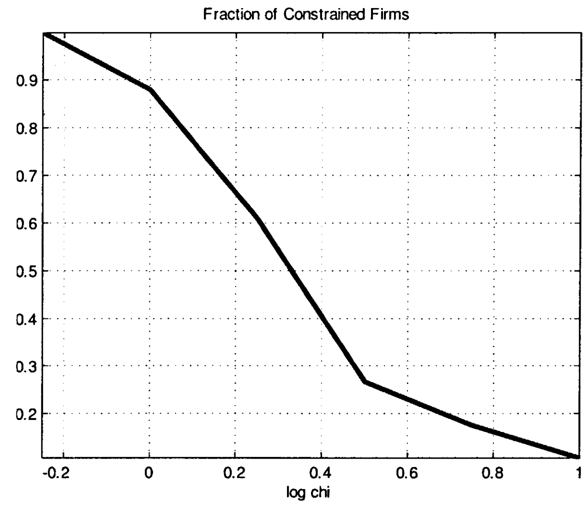

Figure 1.3

Thus, tighter collateral constraints imply that in equilibrium a greater fraction of firms are constrained. I now examine how the tightness of collateral constraints affects the response of equilibrium aggregate output and asset prices to shocks in underlying productivity and noise. The following figure graphs the sensitivity of equilibrium output to productivity and to noise, for different values of χ.

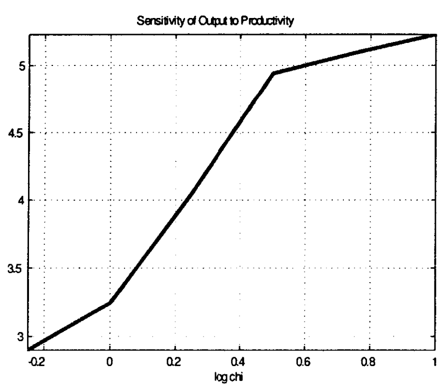

Figure 1.4a

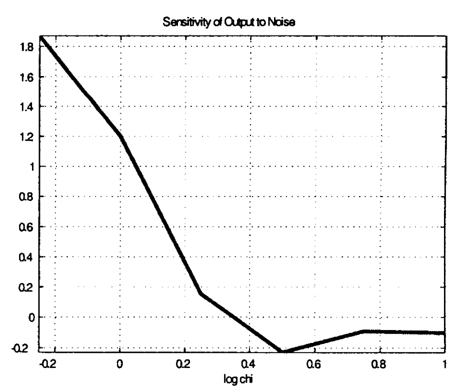

Figure 1.4b

I find that tighter collateral constraints, i.e. lower values of χ, lead to a lower sensitivity of equilibrium output to productivity and a greater sensitivity to noise. Thus, when information is dispersed, tighter collateral constraints mitigate the impact of productivity yet amplify the impact of noise.

1.7.1 The Relative Contribution of Noise

Here I show how noise may be an important source of business cycle fluctuations. Towards this goal, I define the variance decomposition of aggregate output and the asset price as in Definition 19. In the following figure I graph the fraction of output variance due to noise on the left, and the fraction of asset price variance due to noise on the right.

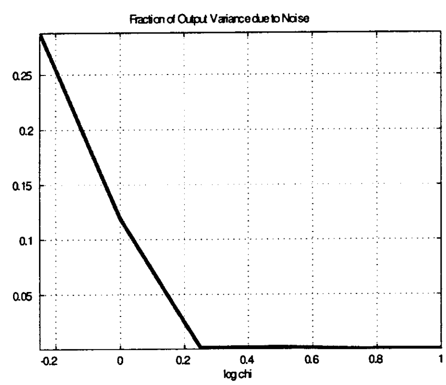

Figure 1.5a

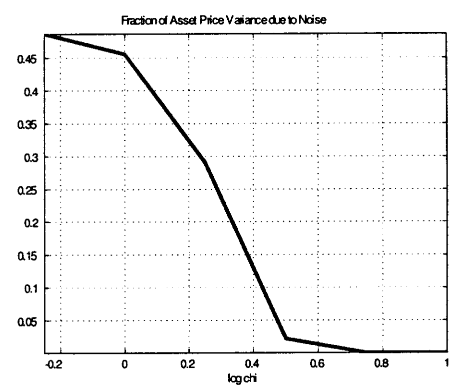

Figure 1.5b

I find that tighter collateral constraints, i.e. lower values of χ, lead to a greater relative contribution of noise to both aggregate output and asset price fluctuations. This follows from the previous discussion on how tighter collateral constraints dampen the impact of productivity shocks, but amplify the impact of noise. Noise may therefore be an important source of asset-price and output volatility when collateral constraints are tight and information is dispersed.

1.7.2 Positive Co-Movement

I now examine how noise affects other observable aggregate variables; in particular I show how noise produces positive co-movement in aggregate output, employment, consumption, and asset prices. In the following figure, I graph the sensitivity of output, consumption, employment, and asset prices to noise.

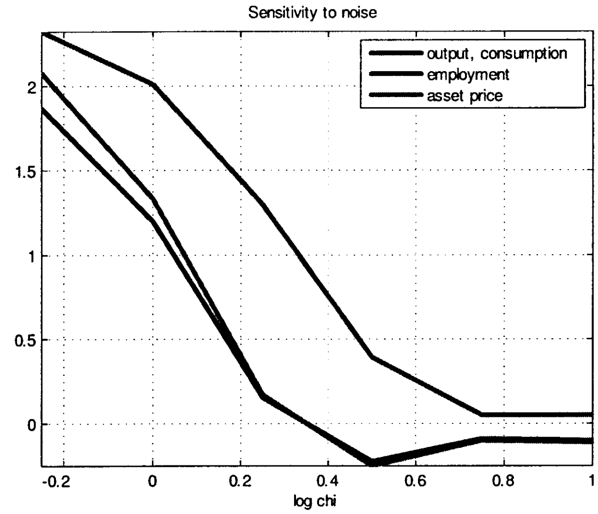

Figure 1.6

Note that when collateral constraints are relatively loose, i.e. for high values of χ, positive noise or, in other words, optimism about aggregate output leads to a small increase in asset prices, since these are directly tied to expectations of aggregate output. At the same time, however, optimism causes aggregate output, consumption and employment fall. Therefore, when collateral constraints are relatively loose, good news or optimism about output tends to cause a recession.

In contrast, when collateral constraints are relatively tight, i.e. for low values of χ, both the equilibrium level of output and the equilibrium asset price respond positively to common errors in expectations of aggregate output. Moreover, consumption and employment also respond positively to noise. Thus, good news about aggregate output thereby generates positive co-movement in output, consumption, employment, and asset prices as in Proposition 18.

This is important, as these co-movement properties coincide with many of the stylized facts about business cycles one observes in the data; see, for example Stock and Watson (1998). Cyclicality of the stock market and employment are standard features of business cycles. Moreover, many often think of booms as periods in which agents are optimistic about aggregate production (and likewise recessions are periods in which agents are pessimistic). In this model, when collateral constraints are tight, good news about aggregate output drives up the price of land, thereby increasing the value of collateral and allowing constrained firms to produce more. Therefore, optimism about aggregate economic activity generates a boom in output–in this sense, expectations are partially self-fulfilling.

Finally, these findings further demonstrate how the noise-driven movements found here are a possible interpretation of what many call "demand shocks". The noise-driven movements documented here resemble demand-driven fluctuations in the following sense. They have many of the same features often associated with such shocks: they contribute to positive co-movement in employment, output and consumption, yet at the same time are orthogonal to underlying movements in productivity, and are merely the product of noise in agents' information.

1.7.3 Movements in the Labor Wedge

I now examine the labor wedge, as defined in (1.22). In the following figure, I graph the sensitivity of the labor wedge to productivity and noise.

Figure 1.7

First, note that the labor wedge is sensitive to noise. Therefore, the underlying sources and financial frictions considered in this model manifest themselves as labor wedges when information is dispersed. This finding relates the noise-driven fluctuations in this model to the large body of research documenting the variation in the "labor wedge" over the business cycle (Hall, 1997; Rotemberg and Woodford, 1999; Chari, Kehoe, and McGrattan, 2007; Shimer, 2009). These researchers have documented that variation in the lavor wedge plays an important role in accounting for business-cycle fluctuations. Thus, this model may be useful for understanding certain business cycle periods, in which the labor wedge plays an important role empirically.

Second, one sees that when collateral constraints are loose, the labor wedge falls in response to noise, but when collateral constraints are tight, the labor wedge increases in response to noise. Thus, as in Proposition 22, when information is dispersed and firms are constrained, noise leads to a positive response in the measured labor wedge. Multiple authors have documented that movements in the labor wedge (as defined as the wedge, as opposed to the implicit tax) is highly cyclical and exhibits sharp decreases during recessions. For example, while productivity has stayed relatively flat during the current recession, the labor wedge has decreased significantly. The labor wedge produced by this model is thus consistent with these empirical properties: when information is dispersed and collateral constraints are tight, the observed cyclical fluctuation in the labor wedge could well be the result of small common errors in expectations.

Finally, there is evidence that a large part of business cycle fluctuations in the data can be accounted for by labor wedges as opposed to intertemporal wedges. Chari, Kehoe, and McGrattan (2007), in particular, document the empirical contribution of various types of time-varying wedges (labor, investment, efficiency) over certain business cycle episodes as well as over the entire postwar period and find that the efficiency and labor wedges together account for the bulk of fluctuations while the investment (or intertemporal) wedge plays essentially no role. Thus, they argue that models in which financial frictions show up primarily as intertemporal wedges are not promising in terms of relating to the data, while models in which financial frictions show up as efficiency or labor wedges may well be. This model is one in which the financial frictions manifest themselves as movements in the labor wedge, and hence may provide guidance in how models which feature financial frictions models could be more in line with these empirical findings.

1.7.4 Further Results

Mean squared error. To show that the approximation method is reasonable, I present the mean squared error of the approximation method as in (1.24). The following figure graphs the mean-squared error for different values of χ.

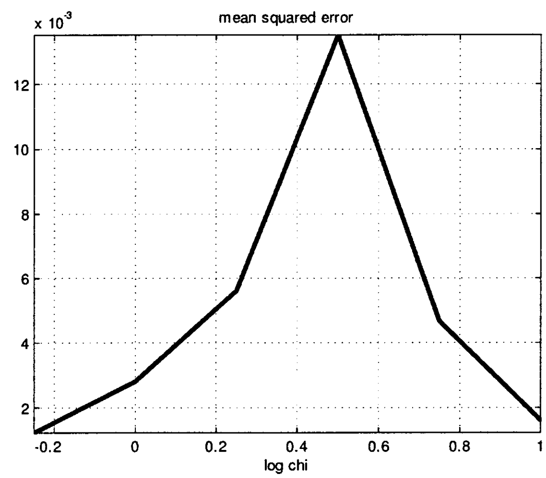

Figure 1.8

Recall that had all firms been unconstrained, or all firms constrained, then the law of motion in equilibrium would in fact be log-linear, and therefore the approximation error approaches zero at either extreme. Thus it is only for the intermediate ranges where there is a non-zero, but small approximation error.

Cross sectional distribution of firms. I may also describe which types are constrained in equilibrium. I find that $\phi_a^u > \phi_a^c$, which implies that the unconstrained are more senstivite to their private idiosyncractic productivity shock than the constrained. In the following figure, I graph the optimal production of a firm who faces no constraint, and the constraint on production, both as a function of their productivity holding all other variables constant.

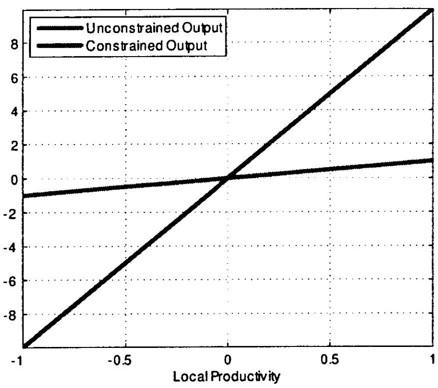

Figure 1.9a

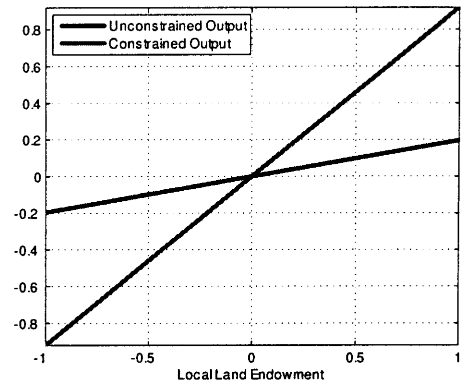

Figure 1.9b

In terms of idiosyncratic productivity, since the unconstrained are more senstivite to their private idiosyncractic productivity shock than the constrained, the unconstrained cross the constrained from below. Thus, the constrained will be those with the highest productivity shocks. This occurs for two reasons. First, having higher productivity implies that a firm would like to produce more. Secondly, the constrained will be those who received the highest private signals, i.e. the ones with the highest expectations of output. This is due to the fact that the constraint depends on the asset price which is an average of all expectations of future output. Thus, firms with the greatest expectations will be the most constrained by their collateral values.

Trade Linkages and Demand Externalities. Next, I consider the effects of trade linkages across islands on the equilibrium sensitivity to the shocks.

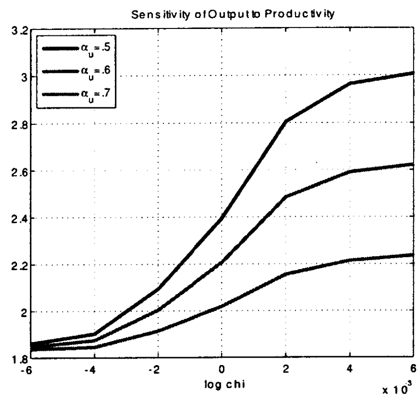

Figure 1.10a

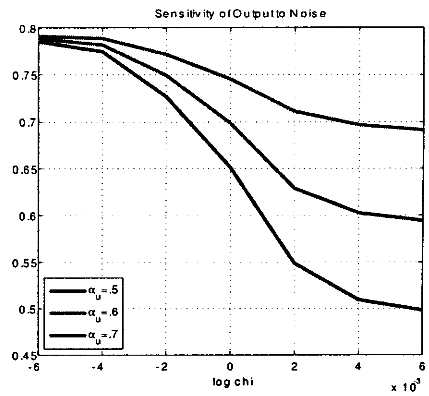

Figure 1.10b

Thus, for a given χ, lower values of ρ, the parameter that controls the substitutabiliy across goods, imply a greater sensitivity to noise and lower sensitivity to productivity.

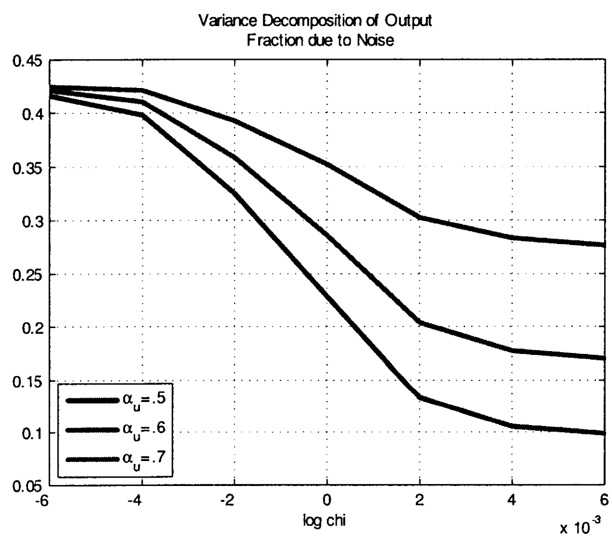

Figure 1.11

Thus demand externalities matters only for how much unconstrained firms care about forecasting the level of economic activity relatively to forecasting the underlying economic fundamentals.

This introduces an additional affect of financial frictions, as the behavior of constrained firms now affects the production choices of the unconstrained.

1.8 Concluding Remarks

This chapter provides a theory of how financial frictions may imply that movements in both the business cycle and asset prices may be driven primarily by noise, as opposed to fundamentals. Collateral constraints on firm-level investment introduce a powerful two-way feedback between the financial market and the real economy. In this chapter I show how this two-way feedback can generate significant expectations-driven fluctuations in asset prices and macroeconomic outcomes when information is dispersed. Noise can thus be an important source of asset-price volatility and business-cycle fluctuations when collateral constraints are tight.

Chapter 2

Noisy Business Cycles

2.1 Introduction

There is a long tradition in macroeconomics, going back to Phelps (1970), Lucas (1972, 1975), Barro (1976), and King (1982), that breaks monetary neutrality by adding informational frictions. This tradition has recently been revived by Mankiw and Reis (2002), Sims (2003), Woodford (2003a, 2008), Mackowiak and Wiederholt (2009a), and others. While this work proposes new formalizations of the origins of informational frictions, most of it remains focused on the same old theme: how imperfect information about the underlying *monetary* shocks can break monetary neutrality.[1]

In this chapter, we depart from the pertinent literature in one fundamental way: we abstract from monetary factors and, instead, focus on the dispersion of information about the *real* shocks hitting the economy. We do so by introducing heterogeneous information in an otherwise canonical micro-founded real-business-cycle model, where nominal prices are fully flexible. We then show how this heterogeneity can have profound implications for the business cycle and can indeed accommodate a somewhat "Keynesian" view of the business cycle *without* any rigidity in nominal prices. In our framework, the bulk of short-run fluctuations is driven, not by technology shocks, but rather by a certain type of "noise". This noise generates positive co-movement in all key macroeconomic variables. Furthermore, the resulting fluctuations may

[1]Exemptions to this statement inlude Sims (2003) and Graham and Wright (2008), who study how informational frictions impact the response of consumption to income or productivity shocks.

look to an outside observer much alike Keynesian demand shocks, even though their origin and their policy implications are very different.

Motivation. Our departure from the pertinent literature is motivated by the following considerations. First, the empirical relevance of theories that require significant lack of information, or some type of unawareness, about the current monetary policy is debatable. Indeed, the older generation of the aforementioned literature succumbed to the criticism that such information is widely, readily, and cheaply available.[2] Second, we contend that the dispersion of information about the *real* shocks hitting the economy is more severe than the one about the conduct of monetary policy. In the recent crisis, for example, there appears to be far more uncertainty, and disagreement, about non-monetary factors such as the value of certain assets, the health of the financial system, or the broader economic fundamentals. And yet, the pertinent literature has little to say about how the heterogeneity of information about the real underlying economic fundamentals matters for business cycles. Finally, even if one is ultimately interested in a monetary model, understanding the positive and normative properties of its underlying real backbone is an essential first step.

Motivated by these considerations, this chapter introduces dispersed information in an otherwise canonical RBC model, where nominal prices are flexible and monetary factors are irrelevant. We first show that the dispersion of information can significantly alter certain positive properties of the RBC paradigm—indeed in ways that might imply that technology shocks explain only a small fraction of high-frequency business cycles, while at the same time helping overcome certain criticisms that New-Keynesians have raised against the RBC paradigm. We next show that this significant change in the positive properties of the RBC paradigm happens without affecting one important normative lesson: as long as there are no monopoly distortions, the equilibrium allocations coincide with the solution to a certain planning problem, leaving no room for stabilization policies.

These results should not be interpreted narrowly as an attack against the New-Keynesian paradigm. Our primary goal is to provide a clean theoretical benchmark for the positive and

[2]The new generation attempts to escape this criticism by postulating that, even if such information is readily available, it may still be hard to update one's information sufficiently frequently (Mankiw and Reis, 2002) or to process and absorb such information sufficiently well (Woodford, 2003a, 2008; Mackowiak and Wiederholt, 2009a).

normative implications of dispersed information. Abstracting from nominal frictions best serves this purpose. And yet, our framework is rich enough to nest the real backbone of New-Keynesian models. Our framework and results may thus prove equally useful for RBC and New-Keynesian analysts alike. In this regard, we believe that our paper makes not only a specific contribution into business-cycle theory but also a broader methodological contribution.

Preview of model. The backbone of our model is a canonical RBC economy. We abstract from capital to simplify the analysis, but allow for a continuum of differentiated commodities. This multi-good (or multi-sector) specification serves two purposes. First and foremost, it introduces a certain type of general-equilibrium, or trading, interactions that, as further highlighted in Angeletos and La'O (2009b), play a crucial role for aggregate fluctuations when, and only when, information is dispersed; this is true whether each of the goods is produced in a competitive or monopolistic fashion. Second, when combined with monopoly power, this specification permits us to nest the real backbone of New-Keynesian models, facilitating a translation of our results to such models.[3] Accordingly, while the core of our analysis focuses on shocks to technology (TFP), in principle we also allow for two other types of shocks to the fundamentals of the economy: taste shocks (shocks to the disutility of labor), and mark-up shocks (shocks to the elasticity of demand). However, none of our results rests on the presence of either monopoly power or these additional shocks.

The only friction featured in our model is that certain economic decisions have to be made under heterogeneous information about the aggregate shocks hitting the economy. The challenge is to incorporate this informational friction without an undue sacrifice in either the micro-foundations or the tractability of the analysis. Towards this goal, we formalize this friction with a certain geographical segmentation, following similar lines as Lucas (1972), Barro (1976), Townsend (1983), and Angeletos and La'O (2008, 2009b). In particular, we assume that each period firms and workers meet in different "islands" and have to make their employment and production decisions while facing uncertainty about the shocks hitting other islands. At the same time, we assume that consumption choices take place in a centralized market, where

[3]Indeed, all the results we document in this paper directly extend to a New-Keynesian variant as long as monetary policy replicates flexible-price allocations, which in certain cases is the optimal thing to do (Angeletos and Lao, 2008).

information is homogenous, and that households are "big families", with fully diversified sources of income. This guarantees that our economy admits a representative consumer and maintains high tractability in analysis despite the fact that some key economic decisions take place under heterogeneous information.

Preview of results. As mentioned, the core of our analysis focuses on the special case where firms are competitive and the only shocks hitting the fundamentals of the economy are technology (TFP) shocks which makes the analysis directly comparable to the RBC paradigm.

(i) In standard RBC models (e.g., Hansen, 1985; Prescott, 1986), macroeconomic outcomes respond fast and strongly to technology shocks. We show that the dispersion of information induces inertia in the response of macroeconomic outcomes. Perhaps paradoxically, this inertia can be significant even if the agents face little uncertainty about the underlying shocks.

(ii) Some researchers have argued that employment responds negatively to productivity shocks in the data; have pointed out that that this fact is inconsistent with standard RBC models; and have used this fact to argue in favor of New-Keynesian models (e.g., Galí, 1999; Basu, Fernald and Kimball, 2006; Galí and Rabanal, 2004). Although this fact remains debatable (e.g., Christiano, Eichenbaum and Vigfusson, 2003; McGrattan, 2004), we show that the dispersion of information can accommodate it within the RBC paradigm.

(iii) In the RBC paradigm, technology shocks account for the bulk of short-run fluctuations. Many economists have argued that this is empirically implausible and have favored New-Keynesian alternatives. We show that the dispersion of information can induce technology shocks to explain only a small fraction of the high-frequency variation in the business cycle. And yet, the entire business cycle remains neoclassical in its nature: monetary factors play no role whatsoever.

(iv) What drives the residual business-cycle fluctuations in our model is a certain type of noise, namely correlated errors in expectations of the underlying technology shocks. Most interestingly, we show that the fraction of short-run volatility that is due to such noise can be arbitrarily high even if the agents are nearly perfectly informed about the underlying technology shocks.

(v) These noise-driven fluctuations feature positive co-movement between employment, output, and consumption. In so doing, they help formalize a certain type of "demand shocks" within an RBC setting. The associated errors in forecasting economic activity can be interpreted as variation in expectations of "aggregate demand ". They help increase the relative volatility of employment while decreasing its correlation with output. An identification strategy as in Blanchard and Quah (1989) or Galí (1999) would likely identify these shocks as "demand" shocks.

(vi) These noise-driven fluctuations involve countercyclical variation in measured labor wedges, and procyclical variation in Solow residuals, consistent with what observed in the data. Once again, these cyclical variations can be significant even if the agents are nearly perfectly informed about the underlying technology shocks.

While we stop short of quantifying these results, we hope that they at least highlight how the heterogeneity of information has a very different mark on macroeconomics outcomes than the uncertainty about fundamentals—a point that we further elaborate on in Angeletos and La'O (2009b). Indeed, what drives our results is not per se the level of uncertainty about the underlying technology or other shocks, but rather the lack of common knowledge about them: our effects are consistent with an arbitrarily small level of uncertainty about the underlying fundamentals.

At the same time, the lack of common knowledge does not *alone* explain the magnitude of our effects. Rather, this depends crucially on the strength of trade linkages among the firms and workers our economy. This idea is formalized by our game-theoretic representation. A measure of the trade linkages in our economy, namely the elasticity of substitution across different goods, maps one-to-one to the degree of strategic complementarity in the game that represents our economy. One can then extrapolate from earlier more abstract work on games of strategic complementarity (Morris and Shin, 2002, Angeletos and Pavan, 2007a) that the strength of trade linkages in our economy may play a crucial role in determining the equilibrium effects of heterogeneous information. We conclude that our findings hinge on the combination of heterogeneous information with strong trade linkages—but they do not hinge on the level of uncertainty about the underlying fundamentals.

We finally seek to understand the normative content of the aforementioned findings. Clearly, a planner could improve welfare by aggregating the information that is dispersed in the economy, or otherwise providing the agents with more information. But this provides no guidence on whether the government should stabilize the fluctuations that originate in noise, or otherwise interfere with the way the economy responds to available information. To address this issue, one has to ask whether a planner can improve upon the equilibrium allocations *without* changing the information structure.

We show that the answer to this question is essentially negative. In particular, in the special case of our model where firms are competitive, there is indeed no way in which the planner can raise welfare without changing the information that is available to the economy. As for the more general case where firms have monopoly power, the best the planner can do is merely to undo the monopoly distortion, much alike what he is supposed to do when information is commonly shared. We conclude that, insofar the information is taken as exogenous, the key normative lessons of the pertinent business-cycle theory survive the introduction of dispersed information, no matter how severely the positive lessons might be affected.

Layout. The remainder of the introduction discusses the related literature. Section 2.2 introduces the model. Section 2.3 characterizes the general equilibrium. Sections 2.4 and 2.5 explore the implications for business cycles. Section 2.6 studies efficiency. Section 2.7 concludes.

Related literature. The macroeconomics literature on informational frictions has a long history, a revived present, and—hopefully—a promising future.[4] Among this literature, most influential in our approach have been Morris and Shin (2002), Woodford (2003a), and Angeletos and Pavan (2007, 2009). Morris and Shin (2002) were the first to highlight the potential implications of asymmetric information, and higher-order beliefs, for settings that feature strategic complementarity. Woodford (2003a) exploited the inertia of higher-order beliefs to generate inertia in the response of prices to nominal shocks in a stylized model of price setting. Finally, Angeletos and Pavan (2007a, 2009) provided a methodology for studying the positive and normative properties of a more general class of games with strategic complementarity and dispersed information.

[4]Recent contributions include Adam (2007), Amador and Weill (2007, 2008), Amato and Shin (2006), Angeletos, Lorenzoni and Pavan (2009), Angeletos and Pavan (2004, 2007a, 2007b, 2009), Bacchetta and Wincoop (2005), Coibion and Gorodnichenko (2008), Collard and Dellas (2005b), Graham and Wright (2008), Hellwig (2002, 2005), Hellwig and Veldkamp (2008), Hellwig and Venkateswara (2009), Klenow and Willis (2007), Lorenzoni (2008, 2009), Luo (2008), Mackowiak and Wiederholt (2009a, 2009b), Mankiw and Reis (2002, 2006), Morris and Shin (2002, 2006), Moscarini (2004), Nimark (2008), Porapakkarm and Young (2008), Reis (2006, 2009), Rodina (2008), Sims (2003, 2006), Van Nieuwerburgh and Veldkamp (2006), Veldkamp (2006), Veldkamp and Woolfers (2007), and Woodford (2003a, 2008).

Part of our contribution in this chapter, and in two companion papers (Angeletos and La'O, 2008, 2009), is to show how the equilibrium and efficient allocations of fully micro-founded business-cycle economies can be represented as the Perfect Bayesian equilibria of a certain class of games with strategic complementarity, similar to those considered in Morris and Shin (2002) and Angeletos and Pavan (2007a, 2009). This representation is useful, as it facilitates a translation of some of the more abstract insights of this earlier work within a macroeconomic context. At the same time, the specific micro-foundations are crucial for understanding both the positive and the normative implications of the particular form of complementarity that we identify in this chapter. Indeed, it is only these micro-foundations that explain either why this complementarity turns out to be irrelevant for the business cycle when information is commonly shared, or why it has none of the welfare implications conjectured in Morris and Shin (2002).

Our main contribution, however, is with regard to business-cycle theory. In this chapter, we show how dispersed information can significantly alter the positive properties of the RBC paradigm. In Angeletos and La'O (2008), we extend the analysis by introducing nominal frictions and by allowing information to get aggregated through certain price and quantity indicators; we then explore a number of novel implications for optimal fiscal and monetary policy. Finally, in Angeletos and La'O (2009b), we show how the heterogeneity of information opens the door to a certain type of sentiment-driven (or sunspot-like) fluctuations despite the uniqueness of equilibrium. Combined, this work highlights how the heterogeneity of information has very distinct implications for the business cycle than the uncertainty about the underlying economic fundamentals.

This also explains how our approach differentiates from the recent literature on "news shocks" (Barsky and Sims, 2009; Beaudry and Portier, 2004, 2006; Christiano, Ilut, Motto, and Rostagno (2008); Gilchrist and Leahy, 2002; Jaimovich and Rebelo, 2009; Lorenzoni, 2008). These papers also feature noise-driven fluctuations. However, these fluctuations obtain within representative-agent models, do not rest on the heterogeneity of information, and are bound to vanish when the uncertainty about the fundamentals is small enough. Furthermore, these papers generate positive co-movement in the key macroeconomic variables *only* by introducing exotic preferences (e.g., Jaimovich and Rebelo, 2009) or sticky prices and suboptimal monetary policy (e.g., Lorenzoni, 2008). In contrast, our paper generates positive co-movement without either of these features.

Interestingly, Kydland and Prescott (1982) had also allowed for noise shocks, only to be discarded in subsequent work; but they, too, did not allow heterogeneous information and hence could not have considered the effects we identify here. Finally, there are numerous papers that consider geographical and trading structures similar to the one in our model (e.g., Lucas and Prescott, 1974; Rios-Rull and Prescott, 1992; Alvarez and Shimer, 2008), but also rule out heterogeneous information about the aggregate economic fundamentals. To recap, it is the heterogeneity of information that is both the distinctive feature of our approach and the key to the results of this chapter.[5]

2.2 The Model

There is a (unit-measure) continuum of households, or "families" , each consisting of a consumer and a continuum of workers. There is a continuum of "islands" , which define the boundaries of local labor markets as well as the "geography" of information: information is symmetric within an island, but asymmetric across islands. Each island is inhabited by a continuum of firms, which specialize in the production of differentiated commodities. Households are indexed by $h \in H = [0, 1]$; islands by $i \in I = [0, 1]$; firms and commodities by $(i, j) \in I \times J$; and periods by $t \in \{0, 1, 2, ...\}$.

Each period has two stages. In stage 1, each household sends a worker to each of the islands. Local labor markets then open, workers decide how much labor to supply, firms decide how much labor to demand, and local wages adjust so as to clear the local labor market. At this point, workers and firms in each island have perfect information regarding local productivity, but imperfect information regarding the productivities in other islands. After employment and production choices are sunk, workers return home and the economy transits to stage 2. At this point, all information that was previously dispersed becomes publicly known, and commodity markets open. Quantities are now pre-determined by the exogenous productivities and the endogenous employment choices made during stage 1, but prices adjust so as to clear product markets.

[5]It is worth noting that our approach is also different from the Mirrless literature, which allows for private information about idiosyncratic shocks but rules out private (heterogeneous) information about aggregate shocks.

Households. The utility of household h is given by

$$u_i = \sum_{t=0}^{\infty} \beta^t \left[U(C_{h,t}) - \int_I S_{i,t} V(n_{hi,t}) di, \right]$$

with

$$U(C) = \frac{C^{1-\gamma}}{1-\gamma} \qquad \text{and} \qquad V(n) = \frac{n^{1+\epsilon}}{1+\epsilon}.$$

Here, $\gamma \geq 0$ parametrizes the income elasticity of labor supply,[6] $\epsilon \geq 0$ parameterizes the Frisch elasticity of labor supply, $n_{hi,t}$ is the labor of the worker who gets located on island i during stage 1 of period t, $S_{h,t}$ is an island-specific shock to the disutility of labor, and $C_{h,t}$ is a composite of all the commodities that the household purchases and consumes during stage 2.

This composite, which also defines the numeraire used for wages and commodity prices, is given by the following nested CES structure:

$$C_{h,t} = \left[\int_I c_{hi,t}^{\frac{\rho-1}{\rho}} di \right]^{\frac{\rho}{\rho-1}}$$

where

$$c_{hi,t} = \left[\int_J c_{hij,t}^{\frac{\eta_{it}-1}{\eta_{it}}} dj \right]^{\frac{\eta_{it}}{\eta_{it}-1}}$$

and where $c_{hij,t}$ is the quantity household h consumes in period t of the commodity produced by firm j on island i. Here, η_{it} is a random variable that determines the period-t elasticity of demand faced by any individual firm within a given island i, while ρ is the elasticity of substitution across different islands. Letting the within-island elasticity η differ from the across-islands elasticity ρ permits us to distinguish the degree of monopoly power (which will be determined by the former) from the strength of trade linkages and the associated degree of strategic complementarity (which will be determined by the latter). In fact, a case of special interest that we will concentrate on for much of our analysis is the limit where monopoly power vanishes ($\eta \to \infty$) while the strategic complementarity remains non-trivial ($\rho < \infty$); this case nests a canonical, competitive RBC economy. At the same time, letting the within-island

[6] Note that risk aversion and intertemporal substitution play no role in our setting because all idiosyncratic risk is insurable and there is no capital. Therefore, γ only controls the sensitivity of labor supply to income for given wage.

elasticity to be finite and random permits us to introduce monopoly power and mark-up shocks, thus facilitating a translation/extension of our results to the New-Keynesian framework.

Households own equal shares of all firms in the economy. The budget constraint of household h is thus given by the following:

$$\int_{I\times J} p_{ij,t}c_{hij,t}d(j,k) + B_{h,t+1} \leq \int_{J\times I} \pi_{ij,t}d(i,j) + \int_I w_{it}n_{hi,t}dk + R_t B_{h,t},$$

where $p_{ij,t}$ is the period-t price of the commodity produced by firm j on island i, $\pi_{ij,t}$ is the period-t profit of that firm, w_{it} is the period-t wage on island i, R_t is the period-t nominal gross rate of return on the riskless bond, and $B_{h,t}$ is the amount of bonds held in period t.

The objective of each household is simply to maximize expected utility subject to the budget and informational constraints faced by its members. Here, one should think of the worker-members of each family as solving a team problem: they share the same objective (family utility) but have different information sets when making their labor-supply choices. Formally, the household sends off during stage 1 its workers to different islands with bidding instructions on how to supply labor as a function of (i) the information that will be available to them at that stage and (ii) the wage that will prevail in their local labor market. In stage 2, the consumer-member collects all the income that the worker-member has collected and decides how much to consume in each of the commodities and how much to save (or borrow) in the riskless bond.

Asset markets. Asset markets operate in stage 2, along with commodity markets, when all information is commonly shared. This guarantees that asset prices do not convey any information. The sole role of the bond market in the model is then to price the risk-free rate. Moreover, because our economy admits a representative consumer, allowing households to trade risky assets in stage 2 would not affect any of the results.

Firms. The output of firm j on island i during period t is given by

$$q_{ij,t} = A_{i,t}(n_{ij,t})^{\theta}$$

where $A_{i,t}$ is the productivity in island i, $n_{ij,t}$ is the firm's employment, and $\theta \in (0,1)$ parameterizes the degree of diminishing returns in production. The firm's realized profit is given by

$$\pi_{ij,t} = p_{ij,t}q_{ij,t} - w_{i,t}n_{ij,t}$$

Finally, the objective of the firm is to maximize its expectation of the representative consumer's valuation of its profit, namely, its expectation of $U'(C_t)\pi_{ij,t}$.

Labor and product markets. Labor markets operate in stage 1, while product markets operate in stage 2. Because labor cannot move across islands, the clearing conditions for labor markets are as follows:

$$\int_J n_{ij,t}dj = \int_H n_{hi,t}dh \ \forall i$$

On the other hand, because commodities are traded beyond the geographical boundaries of islands, the clearing conditions for the product markets are as follows:

$$\int_H c_{hij,t}dh = q_{ij,t} \ \forall (i,j)$$

Fundamentals and information. Each island in our economy is subject to three types of shocks: shocks to the technology used by local firms (TFP shocks); shocks to the disutility of labor faced by local workers (taste shocks); and shocks to the elasticity of demand faced by local firms, causing variation in their monopoly power (mark-up shocks). We allow for both aggregate and idiosyncratic components to these shocks.

The aggregate fundamentals of the economy in period t are identified by the joint distribution of the shocks $(A_{it}, S_{it}, \eta_{it})$ in the cross-section of islands.[7] Let Ψ_t denote this distribution. The standard practice in macroeconomics is to assume that Ψ_t is commonly known in the beginning of period t. In contrast, we consider situations where information about Ψ_t is imperfect and, most importantly, heterogeneous. We thus assume that different islands observe only noisy private (local) signals about Ψ_t in stage 1, when they have to make their decentralized employment and production choices. On the other hand, we assume that Ψ_t becomes common known in stage 2, when agents meet in the centralized commodity and financial markets.

For our main theoretical results we do not need to make any special assumptions about the information that is available to each island. For example, we can impose a Gaussian structure as

[7] In special cases (as with Assumption 1 later on), this distribution might be conveniently parameterized by the mean values of the shocks; but in general the aggregate fundamentals are identified by the entire distribution.

in Morris and Shin (2002). Alternatively, we could allow some islands to be perfectly informed and others to be imperfectly informed, mimicking the idea in Mankiw and Reis (2002) that only a fraction of the agents update their information sets in any given point of time. To some extent, we could even interpret the noise in these signals as the product of rational inattention, as in Sims (2003) and Woodford (2003a). More generally, we do not expect the details of the origins of noise to be crucial for our positive results.

We thus start by allowing a rather arbitrary information structure, as in the more abstract work of Angeletos and Pavan (2009). First, we let ω_t denote the "type" of an island during period t. This variable encodes all the information available to an island about the local shocks as well as about the cross-sectional distribution of shocks and information in the economy. Next, we let Ω_t denote the distribution of ω_t in the cross-section of islands. This variable identifies the aggregate state of the economy during period t; note that the aggregate state now includes not only the cross-sectional distribution Ψ_t of the shocks but also the cross-sectional distributions of the information (signals). Finally, we let $\mathcal{S}_\omega$ denote the set of possible types for each island, $\mathcal{S}_\Omega$ the set of probability distributions over $\mathcal{S}_\omega$, and $\mathcal{P}(\cdot|\cdot)$ a probability measure over $\mathcal{S}_\Omega^2$.[8]

We then formalize the information structure as follows. In the beginning of period t, and conditional on Ω_{t-1}, Nature draws a distribution $\Omega_t \in \mathcal{S}_\Omega$ using the measure $\mathcal{P}(\Omega_t|\Omega_{t-1})$.[9] Nature then uses Ω_t to make independent draws of $\omega_t \in \mathcal{S}_\omega$, one for each island. In the beginning of period t, before they make their current-period employment and production choices, agents in any given island get to see only their own ω_t; in general, this informs them perfectly about their local shocks, but only imperfectly about the underlying aggregate state Ω_t. In the end of the period, however, Ω_t becomes commonly known (ensuring that Ψ_t also becomes commonly known).

To recap, the key informational friction in our model is that agents face uncertainty about the underlying aggregate state Ω_t. Whether they face uncertainty about their own local shocks is immaterial for the type of effects we analyze in this chapter. Merely for convenience, then, we assume that the agents of an island learn their own local shocks in stage 1. We can thus

[8] To avoid getting distracted by purely technical issues, our proofs treat $\mathcal{S}_\omega$ and $\mathcal{S}_\Omega$ as if they were finite sets. However, none of our results hinges on this restriction.

[9] Note that we have imposed that the aggregate state Ω_t follows a Markov process; apart from complicating the notation, nothing changes if we let the aforementioned probability measure depend on all past aggregate states.

express the shocks as functions of ω_t: we denote with $A(\omega_t)$ the local productivity shock, with $S(\omega_t)$ the local taste shock, and with $\eta(\omega_t)$ the local mark-up shock.

2.3 Equilibrium

In this section we characterize the equilibrium by providing a game-theoretic representation that turns out to be instrumental for our subsequent analysis.

2.3.1 Definition

Because each family sends workers to every island and receives profits from every firm in the economy, each family's income is fully diversified during stage 2. This guarantees that our model admits a representative consumer and that no trading takes place in the financial market. To simplify the exposition, we thus set $B_t = 0$ and abstract from the financial market. Furthermore, because of the symmetry of preferences, technologies and information within each island, it is without any loss of generality to impose symmetry in the choices of workers and firms within each island. Finally, because of the absence of capital and the Markov restriction on the aggregate state, Ω_{t-1} summarizes all the payoff-relevant public information as of the beginning of period t. Recall then that the additional information that becomes available to an island in stage 1 is only ω_t. As a result, the local levels of labor supply, labor demand, wage, and output can all depend on Ω_{t-1} and ω_t, but not the current aggregate state Ω_t. On the other hand, the commodity prices in stage 2, and all aggregate outcomes, do depend on Ω_t. We thus define an equilibrium as follows.

Definition 27 *An equilibrium consists of an employment strategy $n : \mathcal{S}_\omega \times \mathcal{S}_\Omega \to \mathbb{R}_+$ a production strategy $q : \mathcal{S}_\omega \times \mathcal{S}_\Omega \to \mathbb{R}_+$, a wage function $w : \mathcal{S}_\omega \times \mathcal{S}_\Omega \to \mathbb{R}_+$, an aggregate output function $Q : \mathcal{S}_\Omega^2 \to \mathbb{R}_+$, an aggregate employment function $N : \mathcal{S}_\Omega^2 \to \mathbb{R}_+$, a price function $p : \mathcal{S}_\omega \times \mathcal{S}_\Omega^2 \to \mathbb{R}_+$, and a consumption strategy $c : \mathbb{R}_+^3 \to \mathbb{R}_+$, such that the following are true:*

(i) The price function is normalized so that

$$P(\Omega_t, \Omega_{t-1}) \equiv \left[\int p(\omega, \Omega_t, \Omega_{t-1})^{1-\rho} d\Omega_t(\omega)\right]^{\frac{1}{1-\rho}} = 1$$

for all (Ω_t, Ω_{t-1}).

(ii) The quantity $c(p, p', Q)$ *is the representative consumer's optimal demand for any commodity whose price is* p *when the price of all other commodities from the same island is* p' *and the aggregate output (income) is* Q.

(iii) When the current aggregate state is Ω_t *and the past aggregate state is* Ω_{t-1}*, the price that clears the market for the product of the typical firm from island* ω_t *is* $p(\omega_t, \Omega_t, \Omega_{t-1})$*; the employment and output levels of that firm are, respectively,* $n(\omega_t, \Omega_{t-1})$ *and* $q(\omega_t, \Omega_{t-1})$*, with* $q(\omega_t, \Omega_{t-1}) = A(\omega_t)n(\omega_t, \Omega_{t-1})^\theta$ *; and the aggregate output and employment indices are, respectively,*

$$Q(\Omega_t, \Omega_{t-1}) = \left\{ \int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right\}^{\frac{\rho}{\rho-1}} \quad \textit{and} \quad N(\Omega_t, \Omega_{t-1}) = \int n(\omega, \Omega_{t-1}) d\Omega_t(\omega).$$

(iv) The quantities $n(\omega_t, \Omega_{t-1})$ *and* $q(\omega_t, \Omega_{t-1})$ *are optimal from the perspective of the typical firm in island* ω_t*, taking into account that firms in other islands are behaving according to the same strategies, that the local wage is given by* $w(\omega_t, \Omega_{t-1})$*, that prices will be determined in stage 2 so as to clear all product markets, that the representative consumer will behave according to consumption strategy* c*, and that aggregate income will be given by* $Q(\Omega_t, \Omega_{t-1})$.

(v) The local wage $w(\omega_t, \Omega_{t-1})$ *is such that the quantity* $n(\omega_t, \Omega_{t-1})$ *is also the optimal labor supply of the typical worker in an island of type* ω_t.

Note that condition (i) simply means that the numeraire for our economy is the CES composite defined when we introduced preferences. The rest of the conditions then represent a hybrid of a Walrasian equilibrium for the complete-information exchange economy that obtains in stage 2, once production choices are fixed, and a subgame-perfect equilibrium for the incomplete-information game played among different islands in stage 1.

Let us expand on what we mean by this. When firms in an island decide how much labor to employ and how much to produce during stage 1, they face uncertainty about the prices at which they will sell their product during stage 2 and hence they face uncertainty about the marginal return to labor. Similarly, when workers in an island decide how much labor to supply, they face uncertainty about the real income their household will have in stage 2 and hence face uncertainty about the marginal value of the wealth that they can generate by working more.

But then note that firms and workers in each island can anticipate that the prices that clear the commodity markets and the realized level of real income are, in equilibrium, determined by the level of employment and production in other islands. This suggests that we can solve for the general equilibrium of the economy by reducing it to a certain game, where the incentives of firms and workers in an island depend on their expectations of the choices of firms and workers in other islands. We implement this solution strategy in the following.

Remark. To simplify notation, we often use q_{it} as a short-cut for $q(\omega_t, \Omega_{t-1})$, Q_t as a short-cut for $Q(\Omega_t, \Omega_{t-1})$, $\mathbb{E}_{it}$ as a short-cut for $\mathbb{E}[\cdot|\omega_t, \Omega_{t-1}]$, and so on; also, we drop the indices h and j, because we know that allocations are identical across households, or across firms within an island.

2.3.2 Characterization

Towards solving for the equilibrium, consider first how the economy behaves in stage 2. The optimal demand of the representative consumer for a commodity from island i whose price is p_{it} when the price of other commodities in the same island is p'_{it} is given by the following:

$$c_{it} = \left(\frac{p_{it}}{p'_{it}}\right)^{-\eta_{it}} \left(\frac{p'_{it}}{P_t}\right)^{-\rho} C_t,$$

where $P_t = 1$ by our choice of numeraire.[10] In equilibrium, $C_t = Q_t$. It follows that the equilibrium consumption strategy is given by $c(p, p', Q) = p^{-\eta} (p')^{\eta - \rho} Q$. Equivalently, the inverse demand function faced by a firm during period t is

$$p_{it} = (p'_{it})^{1-\frac{\rho}{\eta_{it}}} q_{it}^{-\frac{1}{\eta_{it}}} Q_t^{\frac{1}{\eta_t}} \tag{2.1}$$

Consider now stage 1. Given that the marginal value of nominal income for the representative household is $U'(C_t)$ and that $C_t = Q_t$ in equilibrium, the objective of the firm is simply

$$\mathbb{E}_{it} \left[U'(Q_t) \left(p_{it} q_{it} - w_{it} n_{it} \right) \right].$$

[10]To understand this condition, note that $c'_{it} = \left(\frac{p'_{it}}{P_t}\right)^{-\rho} C_t$ is the demand for the busket of commodities produced by a particular island; the demand for the commodity of a particular firm in that islands is then $c_{it} = \left(\frac{p_{it}}{p'_{it}}\right)^{-\eta_i} c'_{it}$.

Using (2.1), we conclude the typical firm on island ω_t maximizes the following objective:

$$\mathbb{E}_{it}\left[U'\left(Q_t\right)\left(\left(p'_{it}\right)^{1-\frac{\rho}{\eta_{it}}}Q_t^{\frac{1}{\eta_{it}}}q_{it}^{1-\frac{1}{\eta_{it}}}-w_{it}n_{it}\right)\right],$$

where $q_{it}=A_{it}n_{it}^{\theta}$. As long as $1>(1-\frac{1}{\eta_t})\theta>0$ (which we assume to be always the case), the above objective is a strictly concave function of n_t, which guarantees that the solution to the firm's problem is unique and that the corresponding first-order condition is both necessary and sufficient. This condition is simply given by equating the expected marginal cost and revenue of labor, evaluated under local expectation of the equilibrium pricing kernel:

$$\mathbb{E}_{it}\left[U'(Q_{it})\right]w_{it}=\left(\frac{\eta_{it}-1}{\eta_{it}}\right)\mathbb{E}_{it}\left[U'\left(Q_t\right)\left(p'_{it}\right)^{1-\frac{\rho}{\eta_{it}}}\left(\frac{Q_t}{q_{it}}\right)^{\frac{1}{\eta_{it}}}\right]\left(\theta A_{it}n_{it}^{\theta-1}\right). \tag{2.2}$$

Next, note that, since all firms within an island set the same price in equilibrium, it must be that $p'_{it}=p_{it}$. Along with (2.1), this gives

$$p'_{it}=p_{it}=\left(\frac{q_{it}}{Q_t}\right)^{-\frac{1}{\rho}}. \tag{2.3}$$

This simply states that the equilibrium price of the typical commodity of an island relative to the numeraire is equal to the MRS between that commodity and the numeraire. Finally, note that the optimal labor supply of the typical worker on island i is given by equating the local wage with the MRS between the numeraire and leisure:

$$w_{it}=\frac{S_{it}n_{it}^{\epsilon}}{\mathbb{E}_{it}\left[U'(Q_t)\right]} \tag{2.4}$$

Conditions (2.3) and (2.4) give the equilibrium prices and wages as functions of the equilibrium allocation. Using these conditions into condition (2.2), we conclude that the equilibrium allocation is pinned down by the following condition:

$$S_{it}n_{it}^{\epsilon}=\left(\frac{\eta_{it}-1}{\eta_{it}}\right)\mathbb{E}_{it}\left[U'\left(Q_t\right)\left(\frac{q_{it}}{Q_t}\right)^{-\frac{1}{\rho}}\right]\left(\theta A_{it}n_{it}^{\theta-1}\right). \tag{2.5}$$

This condition has a simple interpretation: it equates the private cost and benefit of effort in each island. To see this, note that the left-hand side is simply the marginal disutility of an extra unit of labor in island i; as for the right-hand side, $\frac{\eta_{it}-1}{\eta_{it}}$ is the reciprocal of the local monopolistic mark-up, $U'\left(Q_t\right)\left(\frac{q_{it}}{Q_t}\right)^{-\frac{1}{\rho}}$ is the marginal utility of an extra unit of the typical local commodity, and $\theta A_{it} n_{it}^{\theta-1}$ is the corresponding marginal product of labor.

Note that condition (2.5) expresses the equilibrium levels of local employment n_{it} and local output q_{it} in relation to the local shocks and the local expectations of aggregate output Q_t. Using the production function, $q_{it} = A_{it} n_{it}^{\theta}$, to eliminate n_{it} in this condition, and reverting to the more precise notation of Definition 1 (i.e., replacing q_{it} with $q(\omega_t, \Omega_{t-1})$, Q_t with $Q(\Omega_t, \Omega_{t-1})$, A_{it} with $A(\omega_t)$, and so on), we reach the following result.

Proposition 28 *Let*

$$f(\omega) \equiv \log\left\{\theta^{\frac{\theta}{1-\theta+\epsilon+\gamma\theta}}\left(\frac{\eta(\omega)-1}{\eta(\omega)}\right)^{\frac{\theta}{1-\theta+\epsilon+\gamma\theta}} S(\omega)^{-\frac{\theta}{1-\theta+\epsilon+\gamma\theta}} A(\omega)^{\frac{1+\epsilon}{1-\theta+\epsilon+\gamma\theta}}\right\}$$

be a composite of all the local shocks hitting an island of type ω and define the coefficient

$$\alpha \equiv \frac{\frac{1}{\rho}-\gamma}{\frac{1}{\rho}+\frac{1-\theta+\epsilon}{\theta}} < 1$$

The equilibrium levels of local and aggregate output are the solution to the following fixed-point problem:

$$\log q\left(\omega_t, \Omega_{t-1}\right) = (1-\alpha)\, f(\omega_t) + \alpha \log\left\{\mathbb{E}\left[Q(\Omega_t, \Omega_{t-1})^{\frac{1}{\rho}-\gamma}\middle|\,\omega_t, \Omega_{t-1}\right]^{\frac{1}{\frac{1}{\rho}-\gamma}}\right\} \quad \forall(\omega_t, \Omega_{t-1}) \tag{2.6}$$

$$Q(\Omega_t, \Omega_{t-1}) = \left[\int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega)\right]^{\frac{\rho}{\rho-1}} \quad \forall(\Omega_t, \Omega_{t-1}). \tag{2.7}$$

This result establishes that the general equilibrium of our economy reduces to a simple fixed-point relation between local and aggregate output. In so doing, it offers a game-theoretic representation of our economy, similar to the one established in Angeletos and La'O (2009b) for a variant economy with capital. To see this, consider a game with a large number of players, each choosing an action in $\mathbb{R}_+$. Identify a "player" in this game with an island in our economy and interpret the level of output of that island as the "action" of the corresponding player.

Next, identify the "types" of these players with ω_t, which encodes the local shocks and local information sets in our economy. Finally, let their "best responses" be given by condition (2.6). It is then evident that the Perfect Bayesian equilibrium of this game identifies the general equilibrium of our economy.

Note then that the variable $f(\omega_t)$ conveniently summarizes all the local economic fundamentals, while the coefficient α identifies the degree of strategic complementarity in our economy. To see this more clearly, consider a log-linear approximation to conditions (2.6) and (2.7):

$$\log q_{it} = const + (1-\alpha) f_{it} + \alpha \mathbb{E}_{it} [\log Q_t], \tag{2.8}$$

$$\log Q_t = const + \int \log q_{it} di, \tag{2.9}$$

where *const* capture second- and higher-order terms.[11] It is then evident that the coefficient α identifies the slope of an island's best response to the activity of other islands—which is the standard definition of the degree of strategic complementarity.

Finally, note that Proposition 28 holds no matter the information structure. This is important. While much of the recent literature has focused on specific formalizations of the information structure (e.g. Mankiw and Reis, 2002; Sims, 2003; Woodford, 2003a), our result indicates that the information structure typically matters *only* by pinning down the agents' forecasts of economic activity. We would thus invite future researchers to pay more attention on the theoretical and empirical properties of these forecasts as opposed to the details of the information structure.

2.3.3 Specialization, trade and strategic complementarity

As evident from Proposition 28, the degree of complementarity, α, is a monotone function of the elasticity of substitution across the commodities of different islands, ρ. In what follows, we adopt the convention that variation in α represents variation in ρ for given other parameters. We also interpret α as a measure of the strength of trade linkages in our economy. These

[11] In general, these second- and higher-order terms may depend on the underlying state and the above is only an approximation. However, when the underlying shocks and signals are jointly log-normal with fixed second moments (as imposed by Assumption 1 in the next section), these terms are invariant, the approximation error vanishes, and conditions (2.8) and (2.9) are exact.

choices are motivated by the following observations. First, if we consider a variant of our model where each household lives and works only in one island and consumes only the products of that island, then Proposition 1 holds with $\alpha = 0$; in this sense, it is precisely the trade linkages across different islands that introduces strategic interdependence ($\alpha \neq 0$). In this sense, there is indeed a close relation between our model and models of international trade where different countries specialize in the production of certain goods but consume goods from all over the world. Second, while α depends, not only on ρ, but also on ϵ, γ, and θ, these other parameters affect the composite shock f and matter for equilibrium allocations whether islands (agents) are linked or not; in contrast, ρ affects only α. For these reasons, we henceforth use the notions of strategic complementarity, elasticity of substitution across islands, and strength of trade linkages, as synonymous to one another. However, we also note that strong complementarity in our model does not strictly require low ρ: if the wealth effect of labor supply is small ($\gamma \longrightarrow 0$), the Frisch elasticity is high ($\epsilon \longrightarrow 0$), and production is nearly linear ($\theta \longrightarrow 1$), then the degree of complementarity is high ($\alpha \longrightarrow 1$) no matter what ρ is.

The insight that trade introduces a form of strategic complementarity even in neoclassical, perfectly-competively settings is likely to extend well beyond the boundaries of the model we have considered here or the variant in Angeletos and La'O (2009b). We believe that this insight has been under-appreciated in prior work on business cycles for two reasons. First, the two welfare theorems have thought us that it rarely helps, and it can often be misleading, to think of Walrasian settings as games. And second, the type of strategic complementarity we highlight here is simply irrelevant for the business cycle when information is commonly shared.

To understand what we mean by the last point, consider the response of the economy to a *symmetric* aggregate shock (i.e., a shock that keeps the level of heterogeneity invariant). Formally, let $\bar{f}_t$ denote the cross-sectional average of the composite fundamental f_{it} and consider any shock that varies the average fundamental, $\bar{f}_t$, without varying the cross-sectional distribution of the idiosyncratic components of the fundamentals, $\xi_{it} \equiv f_{it} - \bar{f}_t$. When all information is commonly shared, aggregate output is also commonly known in equilibrium. Condition (2.6) then reduces to

$$\log q_{it} = (1-\alpha)(\bar{f}_t + \xi_{it}) + \alpha \log Q_t. \tag{2.10}$$

It is then immediate that the entire cross-sectional distribution of $\log q_{it}$ moves one-to-one with

$\bar{f}_t$, which establishes the following.

Proposition 29 *Suppose that information is commonly shared and that the level of heterogeneity is invariant. Then the equilibrium levels of aggregate output is given by*

$$\log Q_t = const + \bar{f}_t.$$

Recall that, by its definition, the composite shock depends on ϵ and γ but not on ρ. It is then evident that the response of the economy to the underlying aggregate productivity, taste, or mark-up shocks is independent of ρ. In this sense, the business cycle is indeed independent of the degree of strategic complementarity that is induced by trade.

The intuition behind this result is further explained in Angeletos and La'O (2009b). The key is that the strength of trade linkages matters only for how much agents care about forecasting the level of economic activity *relatively* to forecasting the underlying economic fundamentals. But when information is symmetric (commonly shared), any uncertainty the agents face about the level of economic activity reduces to the one that they face about the underlying economic fundamentals, which renders the degree of strategic complementarity irrelevant. In contrast, when information is asymmetric (dispersed), agents can face additional uncertainty about the level of economic activity, beyond the one they face about the fundamentals. The strength of trade linkages then dictates precisely the impact on equilibrium outcome of this residual uncertainty about economic activity.

This is important. It is precisely the aforementioned property that makes dispersed information distinct from uncertainty about the fundamentals—for it is only the *heterogeneity* of information that breaks the coincidence of forecasts of economic activity with the forecasts of the underlying fundamentals when the equilibrium is unique. We further elaborate on this point in Angeletos and La'O (2009b), showing how dispersed information can open the door to a certain type of sunspot-like fluctuations. We refer the reader to that paper for a more thorough discussion of this important, broader insight. In what follows, we concentrate on how this broader insight helps understand why the combination of dispersed information with the aforementioned type of complementarity can have a significant impact on the positive properties of the RBC paradigm.

2.3.4 Relation to complementarity in New-Keynesian models

The familiar condition that characterizes optimal target prices in the New-Keynesian paradigm (e.g., Woodford, 2003b) looks like the following:

$$p_{i,t} = (1 - \xi)\mathcal{Y}_t + \xi p_t + z_{i,t}, \tag{2.11}$$

where $p_{i,t}$ is the target price of a firm (in logs), $\mathcal{Y}_t$ is nominal GDP, p_t is the aggregate price level, $z_{i,t}$ captures idiosyncratic productivity or demand shocks, and ξ is a coefficient that is interpreted as the degree of strategic complementarity in pricing decisions. If we compare the above condition with condition (2.8) in our model, the resemblance is striking. The only noticeable difference seems to be that the relevant choice variable is a price in the New-Keynesian model, while it is a quantity in our model. However, there are some crucial differences behind this resemblance.

First, condition (2.11) does not *alone* pin down the equilibrium. Rather, it must be combined with other conditions regarding the determination of $\mathcal{Y}_t$, the nominal GDP level. In contrast, condition (2.8) offers a complete, self-contained, representation of the equilibrium in our model.

Second, the endogeneity of $\mathcal{Y}_t$ undermines the meaning of condition (2.11). For example, letting y_t denote real GDP and using $\mathcal{Y}_t = p_t + y_t$, condition (2.11) can also be restated as $p_{i,t} = p_t + (1 - \xi)y_t + z_{i,t}$; but then the degree of complementarity appears to be 1, not ξ. In fact, this alternative representation is more informative when money is neutral, because y_t is then exogenous to nominal factors and this condition determines only relative prices. But even when money is non-neutral, ξ fails to identify the degree of complementarity in pricing decisions simply because nominal GDP is far from exogenous—at the very least because monetary policy responds to variation in p_t and y_t. Once this endogeneity is incorporated, the complementarity in pricing decisions is different from ξ and becomes sensitive to policy parameters. In contrast, in our model the degree of strategic complementarity is pinned down only by preferences and technologies, and is completely invariant to monetary policy.

Third, the comparative statics of the complementarity in our model (α) with respect to deeper preference and technology parameters are different from those of its New-Keynesian counterpart (ξ). In particular, note that α decreases with ρ (the elasticity of substitution across different goods), decreases with ϵ (the inverse of the Firsch elasticity of labor supply), and increases with θ (the degree of diminishing returns to labor). Hence, what contributes to strong complementarity in our model is low substitutability in the commodity side, so that trade is crucial, along with high substitutability in the labor and production side, as in Hansen (1985) and King and Rebelo (2000). As one of our discussants highlighted, the opposite comparative statics hold for ξ in the New-Keynesian paradigm. That's interesting. Nevertheless, it is important to bear in mind that our notion of complementarity may have little to do with either the degree of monopoly power or the price elasticities of individual demands. In our model, that latter are pinned down by η (the within-island elasticity of substitution), while the degree of strategic complementarity is pined down by ρ (the across-island elasticity).

Last, but not least, the complementarity highlighted in the New-Keyenesian framework would vanish if firms were setting real (indexed) prices. In this sense, the New-Keyenesian complementarity in is a nominal phenomenon, whereas ours is a real phenomenon.

2.4 Dispersed information and the business cycle

In this section we seek to illustrate how the introduction of dispersed information can impact the positive properties of the RBC paradigm. To facilitate this task, we impose a Gaussian specification on the shocks and the information structure, similar to the one in Morris and Shin (2002), Woodford (2003a), Angeletos and Pavan (2007), and many others.

The shocks and the available information satisfy the following properties:

(i) The aggregate shock $\bar{f}_t$ follows a Gaussian AR(1) or random walk process:

$$\bar{f}_t = \psi \bar{f}_{t-1} + \nu_t,$$

where ψ parameterizes the persistence of the composite shock and ν_t is a Normal innovation, with mean 0 and variance $\sigma_\nu^2 \equiv 1/\kappa_f$, i.i.d. over time.

(ii) The local shock f_t is given by

$$f_t = \bar{f}_{it} + \xi_{it},$$

where ξ_{it} is a purely idiosyncratic shock, Normally distributed with mean zero and variance σ_ξ^2, orthogonal to $\bar{f}_t$, and i.i.d. across islands.

(iii) The private information of an island about the aggregate shock $\bar{f}_t$ is summarized in a Gaussian sufficient statistic x_{it} such that

$$x_{it} = \bar{f}_t + \varsigma_{it},$$

where ς_{it} is noise, Normally distributed with mean zero and variance $\sigma_x^2 \equiv 1/\kappa_x$, orthogonal to both $\bar{f}_t$ and ξ_{it}, and i.i.d. across islands.[12]

(iv) The public information about the aggregate shock $\bar{f}_t$ is summarized in a Gaussian sufficient statistic y_t such that

$$y_t = \bar{f}_t + \varepsilon_t,$$

where ε_t is noise, Normally distributed with mean zero and variance $\sigma_\varepsilon^2 \equiv 1/\kappa_y$, and orthogonal to all other variables.

This specification imposes a certain correlation in the underlying productivity, taste and mark-up shocks: for the composite shock f_{it} to follow a univariate process as above, it must be that all the three type of shocks are moved by a single underlying factor. However, this is only for expositional simplicity. We can easily extend our results to a situation where each of the shocks follows an independent Gaussian process, or consider a more general correlation structure among the shocks.

2.4.1 Closed-form solution

Under Assumption 1, we can identify ω_t with the vector (f_t, x_t, y_t). Because Ω_t is then a joint normal distribution with mean $(\bar{f}_t, \bar{f}_t, y_t)$ and an invariant variance-autocovariance matrix, we can also reduce the aggregate state variable from Ω_t to the more convenient vector $(\bar{f}_t, y_t)$. Next, we can guess and verify that there is always an equilibrium in which $\log q_{it}$ is linear in $(\bar{f}_{t-1}, f_{it}, x_{it}, y_t)$ and $\log Q_t$ is linear in $(\bar{f}_{t-1}, \bar{f}_t, y_t)$. We then find the coefficients of these

[12] Note that the local fundamental f_{it} is itself a private signal of $\bar{f}_t$. However, by the fact that we define x_{it} as a sufficient statistic of *all* the local private information, the informational content of f_{it} is already included in x_{it}.

linear functions by the familiar method of undetermined coefficients. Finally, we can use an independent argument to rule out any other equilibrium. We thereby reach the following result.

Proposition 30 *Under Assumption 1, the equilibrium level of local output is given by*

$$\log q_{it} = const + \varphi_{-1}\bar{f}_{t-1} + \varphi_f f_{it} + \varphi_x x_{it} + \varphi_y y_t, \tag{2.12}$$

where the coefficients $(\varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$ *are given by*

$$\begin{aligned} \varphi_{-1} &= \left\{ \frac{\kappa_f}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \right\} \alpha\psi & \varphi_f &= (1-\alpha) \\ \varphi_x &= \left\{ \frac{(1-\alpha)\kappa_x}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \right\} \alpha & \varphi_y &= \left\{ \frac{\kappa_y}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \right\} \alpha \end{aligned} \tag{2.13}$$

This result gives a closed-form solution of the equilibrium level of output in each island as a log-linear function of the past aggregate fundamental $\bar{f}_{t-1}$, the current local fundamental f_{it}, the local (private) signal x_{it}, and the public signal y_t. Note then that the equilibrium level of output is necessarily an increasing function of the local fundamental f_{it}: $\varphi_f > 0$ necessarily. To interpret this sign, note that higher f means a higher productivity, a lower disutility of labor, or a lower monopolistic distortion. But whether and how local output depends on $\bar{f}_{t-1}$, x_{it} and y_t is determined by the degree of strategic complementarity α.

To understand this, note that local output depends on these variables only because these variables contain information about the current aggregate shocks and, in so doing, help agents forecast the aggregate level of output. But when $\alpha = 0$, the demand- and supply side effects that we discussed earlier perfectly offset each other, so that at the end economic decisions are not interdependent: local incentives depend only the local fundamentals and not on expectations of aggregate activity. It follows that the dependence of local output to $\bar{f}_{t-1}$, x_{it} and y_t vanishes when $\alpha = 0$. On the other hand, if $\alpha \neq 0$, local output depends on $\bar{f}_{t-1}$, x_{it} and y_t because, and only because, these variables help predict aggregate output. In particular, when economic decisions are strategic complements ($\alpha > 0$), the equilibrium level of output in each island responds positively to expectations of aggregate output; in this case, the coefficients φ_{-1}, φ_x, and φ_y are all positive. When instead economic decisions are strategic substitutes ($\alpha > 0$), the equilibrium level of output in each island responds negatively to expectations of aggregate output; in this case, the coefficients φ_{-1}, φ_x, and φ_y are all negative. As mentioned earlier, we view the case in which $\alpha > 0$, and hence in which economic activity responds positively to good news about aggregate fundamentals, as the empirically most relevant scenario. For this reason, our subsequent discussion will focus on this case; however, our results apply more generally.

2.4.2 Remark on interpretation of noise and comparative statics

Before we proceed, we would like to emphasize that one should not give a narrow interpretation to the signal y_t, or its noise ε_t. This signal is not meant to capture only purely public information; rather, it is a convenient modeling device for introducing correlated errors in beliefs of aggregate fundamentals. Indeed, the results we document below can easily be re-casted with a more general information structure, one that allows agents to observe multiple private signals and introduce imperfect cross-sectional correlation in the errors of these private signals; the origin of noise, then, is not only the public signal, but also the correlated errors in the private signals of the agents. We invite the reader to keep this more general interpretation of what "noise" stands for in our model: it is a acronym for all sources of correlated errors in expectations of the fundamentals.[13]

Similarly, we would like to warn the reader not to focus on the comparative statics of the equilibrium with respect to the precisions of private and public information, κ_x and κ_y. These comparative statics fail to isolate the distinct impact of the *heterogeneity* of information, simply because they confound a change in the heterogeneity of information with a change in the overall precision of information.[14] Furthermore, if we had allowed for multiple private signals with correlated errors, it would be unclear whether an increase in the precision of a certain signal raises or reduces the heterogeneity of information. With this in mind, in what follows we focus on the comparative statics with respect to α. These comparative statics best isolate the distinct role of dispersed information, simply because the degree of complementarity matters

[13] In fact, one could go further and interpret "noise" as a certain type of sentiment shocks, namely shocks that do not move at all the agents' beliefs about the fundamentals and nevertheless move equilibrium outcomes. With a unique-equilibrium model as ours, such shocks cannot exist when information is commonly shared; but emerge robustly once information is dispersed. See Angeletos and La'O (2009b).

[14] For example, an increase in κ_x would increase the heterogeneity of information, but would also increase the overall precision of information; and while the former effect would tend to amplify the volatility effects we have documented here, the latter effect would work in the opposite direction.

for aggregate fluctuations in our model *only* by regulating the impact of the heterogeneity of information.[15]

2.4.3 Macroeconomic responses to fundamentals and noise

We now study how the dispersion of information and the strength of trade linkages affect aggregate fluctuations. Towards this goal, we aggregate condition (2.12) and use the fact that $\bar{f}_t = \psi \bar{f}_{t-1} + \nu_t$ to obtain the following characterization of aggregate output.

Corollary 31 *Under Assumption 1, the equilibrium level of aggregate output is given by*

$$\log Q_t = const + \psi \bar{f}_{t-1} + \varphi_\nu \nu_t + \varphi_\varepsilon \varepsilon_t, \tag{2.14}$$

where

$$\varphi_\nu \equiv \varphi_f + \varphi_x + \varphi_y = 1 - \frac{\alpha \kappa_f}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f} \quad \text{and} \quad \varphi_\varepsilon \equiv \varphi_y = \frac{\alpha \kappa_y}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f}, \tag{2.15}$$

and where $\nu_t = \bar{f}_t - \psi \bar{f}_{t-1}$ is the innovation in the fundamentals, ψ is the persistence in the fundamentals, $\varepsilon_t = y_t - \bar{f}_t$ is the aggregate noise.

Condition (2.14) gives the equilibrium level of aggregate output as a log-linear function of the past aggregate fundamentals, $\bar{f}_{t-1}$, the current innovation in the fundamentals, ν_t, and the current noise, ε_t. Consider the impact effect of an innovation in fundamentals. This effect is measured by the coefficient φ_ν. Because the latter is a decreasing function of the precisions κ_x and κ_y, we have that the impact effect of an innovation in fundamentals decreases with the level of noise. This is essentially the same insight as the one that drives the real effects of monetary shocks in both the older macro models with informational frictions (e.g., Lucas, 1972; Barro, 1976) and their recent descendants (e.g., Mankiw and Reis, 2002): the less informed economic agents are about the underlying shocks, the less they respond to these shocks. Clearly, this

[15] Angeletos and Pavan (2007a) propose that a good measure of the "commonality "of information (an inverse measure of the heterogeneity of information) is the cross-sectional correlation of the errors in the agents' forecasts of the fundamentals: holding constant the variance of these forecast errors, an increase in the correlation implies that agents can better forecast one another's actions, even though they cannot better forecast the fundamentals. Following this alternative route would deliver similar insights as the ones we document here.

is true no matter whether agents interact with one another—it is true even in a single-agent decision problem.

More interestingly, we find that φ_ν is a decreasing function of α. That is, the more economic agents care about aggregate economic activity, the weaker the response of the economy to innovations in the underlying fundamentals. At the same time, we find that φ_ε is an increasing function of α. That is, the more economic agents care about aggregate economic activity, the stronger the equilibrium impact of noise. These properties originate from the interaction of strategic complementarity with dispersed information. Indeed, if the underlying shock was common knowledge (which here can be nested by taking the limit as the public signal becomes infinitely precise, $\kappa_y \to \infty$), then both φ_ν and φ_ε would cease to depend on α. But as long as information is dispersed, a higher α reduces φ_ν and raises φ_ε. This highlights how strategic complementarity becomes crucial for the business cycle once information is dispersed.

Corollary 32 *When information is dispersed, and only then, stronger complementarity dampens the impact of fundamentals on output and employment, while amplifying the impact of noise.*

The key intuition behind this result is the same as the one in the more abstract work of Morris and Shin (2002) and Angeletos and Pavan (2007a). Public information and past fundamentals (which here determine the prior about the current fundamentals) help forecast the aggregate level of output relatively better than private information. The higher α is, the more the equilibrium level of output in any given island depends on the local forecasts of aggregate output and the less it depends on the local current fundamentals. It follows that a higher α induces the equilibrium output of each island to be more anchored to the past aggregate fundamentals, more sensitive to public information, and less sensitive to private information. The anchoring effect of past aggregate fundamentals explains why aggregate output responds less to any innovation in the fundamentals, while the heightened sensitivity to noisy public information explains why aggregate output responds more to noise. A similar anchoring effect of the common prior underlies the inertia effects in Woodford (2003a), Morris and Shin (2006), and Angeletos and Pavan (2007a), while the heightened sensitivity to public information is the same as the one in Morris and Shin (2002). However, as mentioned before, we favor a more general interpretation of the signal y_t, not as a public signal, but rather as a source of correlated noise in forecasts of economic fundamentals.

As another way to appreciate the aforementioned result, consider following variance-decomposition exercise. Let $\log \hat{Q}_t$ be the projection of $\log Q_t$ on *past* fundamentals. The residual, which is given by $\log \tilde{Q}_t \equiv \log Q_t - \log \hat{Q}_t = \varphi_\nu \nu_t + \varphi_\varepsilon \varepsilon_t$, can be interpreted as the "high-frequency component" of aggregate output. Its total variance is $Var(\log \tilde{Q}_t) = \varphi_\nu^2 \sigma_\nu^2 + \varphi_\varepsilon^2 \sigma_\varepsilon^2$, where σ_ν^2 ($\equiv 1/\kappa_f$) is the variance of the innovation in the fundamentals and σ_ε^2 ($\equiv 1/\kappa_y$) is the variance of the noise. The fraction of the high-frequency variation in output that originates in noise is thus given by the following ratio:[16]

$$R_{noise} \equiv \frac{Var(\log \tilde{Q}_t | \nu_t)}{Var(\log \tilde{Q}_t)} = \frac{\varphi_\varepsilon^2 \sigma_\varepsilon^2}{\varphi_\nu^2 \sigma_\nu^2 + \varphi_\varepsilon^2 \sigma_\varepsilon^2}.$$

Since a higher α raises φ_ε and reduces φ_ν, it necessarily raises this fraction: the more agents care about the aggregate level of economic activity, the more the high-frequency volatility in output that is driven by noise.

We can then further highlight the distinct nature of dispersed information by showing that, as long as α is high enough, the contribution of noise to short-run fluctuations can be large even if the level of noise is small. Note that the overall precision of an agent's posterior about the underlying fundamentals is given by $\kappa = \kappa_0 + \kappa_x + \kappa_y$. We can then show the following.

Proposition 33 *When information is dispersed and α is sufficiently high, agents can be arbitrarily well informed about the fundamentals ($\kappa \approx \infty$) and, yet, the high-frequency variation in aggregate output can be driven almost exclusively by noise ($R_{noise} \approx 1$).*

Clearly, this is not possible when information is commonly shared. In that case, the contribution of noise on the business cycle is tightly connected to the precision of information and vanishes as this precision becomes infinite. In contrast, when information is dispersed, the contribution of noise in the business cycle can be high even when the precision of information is arbitrarily high. What makes this possible is the combination of heterogeneous information with a sufficiently strong degree of strategic complementarity induced by trade linkages. In particular, a sufficiently strong complementarity induces agents to disregard any valuable private

[16] This fraction equals 1 minus the R-square of the regression of $\log \tilde{Q}_t$ on the innovation ν_t.

information they may have about the underlying shocks and instead focus almost exclusively on noisy public information. As a result, even if this private information happens to be arbitrarily precise, it is not utilized in equilibrium. Note then how this result also contrasts with our earlier observation that this particular type of strategic complementarily would have been irrelevant for the business cycle had information been commonly shared.

Finally, it is worth noting how the dispersion of information and trade linkages affect the cyclical behavior of aggregate employment. The latter is given by

$$\log N_t = const + \tfrac{1}{\theta}(\log Q_t - \bar{a}_t),$$

where $\bar{a}_t$ is the aggregate productivity shock (i.e., the cross-sectional average of $\log A_{i,t}$). It is then immediate that the response of employment to an aggregate shock in either tastes or monopoly power is proportional to that of output. The same is true for the response to noise. More interestingly, the response of employment to an aggregate productivity shock may now turn from a positive sign under common information to a negative sign under dispersed information. To see this, let $\beta \equiv \frac{\partial \bar{f}_t}{\partial \bar{a}_t} = \frac{1+\epsilon}{1-\theta+\epsilon+\theta\gamma} > 0$. When information is commonly shared, the sensitivity of output to an innovation to aggregate productivity is simply β, and that of employment is $\frac{1}{\theta}(\beta-1)$. When, instead, information is dispersed, the corresponding sensitivities are $\varphi_\nu\beta$ for output and $\frac{1}{\theta}(\varphi_\nu\beta-1)$ for employment, with φ_ν as in (2.15). Suppose $\beta > 1$, which means that employment responds positively to a productivity shock under common information, as in any plausible calibration of the RBC framework. As noted earlier, φ_ν is necessarily lower than 1 and is decreasing in α. It follows that, when information is dispersed, stronger trade linkages dampen the response of employment and may actually turn it negative.

2.5 Slow learning and numerical illustration

The preceding has focused on a setting where the underlying shocks become common knowledge within a period. Although this permitted a sharp theoretical analysis of the distinct implications of dispersed information, and of its interaction with trade linkages, it makes it hard to map our results to either empirical business cycles or calibrated RBC models. We now seek to illustrate how incorporating slower learning can facilitate a better mapping between our analysis and the data.

Towards this goal, we need to relax the assumption that the aggregate state, Ω_t, becomes publicly revealed at the end of each period. Accommodating this possibility in a fully micro-founded way would require that there is no centralized commodity trading: with centralized trading, equilibrium prices are likely to reveal the state. However, allowing for decentralized trading would complicate the analysis by introducing informational externalities and/or by letting the relevant state space explode as in Townsend (1983). We are currently exploring some possibilities along these lines. However, for the current purposes, we opt for tractability and expositional simplicity.

In particular, we assume that firms and workers do not ever learn Ω_t, either directly or indirectly from prices and past outcomes. Rather, they only keep receiving exogenous signals about the current fundamentals, of the same type as in Assumption 1, and they use these signals to update each period their beliefs about the underlying state. Think of this as follows. Each firm has two managers: one who decides the level of employment and production; and another who sells the product, receives the revenue, and sends the realized profits to the firm's shareholders. The two managers share the same objective—maximize firm valuation—but do not communicate with one another. Moreover, the first manager never receives any signals on economic activity. He only observes the exogenous local private and public signals. Similarly, the consumers, who observe all the prices in the economy, fail to communicate this information to the workers in their respective families. The workers also base their decisions solely on the exogenous signals.

Needless to say, this specification of the learning process is not particularly elegant. However, it would also be naive to take it too literally: the exogenous signals that we allow firms and workers to receive each period are meant to capture more generally the multiple sources of information that these agents may have. To the extent that the underlying shocks do not become common knowledge too fast, more plausible formalizations of the learning process, albeit highly desirable, need not impact the qualitative properties we wish highlight here.[17]

[17] The learning process we assume here is similar to the one in Woodford (2003a). We refer the reader to Amador and Weill (2008), Angeletos and La'O (2008), Angeletos and Pavan (2009), Hellwig (2002), and Lorenzoni (2008) for some alternative formalizations of the learning process. None of these alternative formalizations would crucially affect the positive results we document in this section; the key here is only that learning is slow, not

Under the aforementioned specification, equilibrium behavior continues to be characterized by the same best-response-like condition as in the baseline model:

$$\log q_{i,t} = (1-\alpha) f_{i,t} + \alpha \mathbb{E}_{i,t}\left[\log Q_t\right], \tag{2.16}$$

where we have normalized the constant to zero. The only difference is in the information that underlies the expectation operator in this condition. Finally, for concreteness, we henceforth focus on productivity shocks as the only shock to fundamentals: $f_{i,t} = \beta \log A_{i,t}$, with $\beta \equiv \frac{1+\epsilon}{1-\theta+\epsilon+\theta\gamma}$.

The procedure we follow to solve for the equilibrium dynamics is based on Kalman filtering and is similar to the one in Woodford (2003a). We guess and verify that the aggregate state can be summarized in a vector X_t comprised of the aggregate fundamental and aggregate output:

$$X_t \equiv \begin{bmatrix} \bar{f}_t \\ \log Q_t \end{bmatrix}, \tag{2.17}$$

Firms and workers in any given island never observe the state, but instead receive the following vector of signals each period:

$$z_{it} \equiv \begin{bmatrix} x_{it} \\ y_t \end{bmatrix} = \begin{bmatrix} \bar{f}_t + \varsigma_{it} \\ \bar{f}_t + \varepsilon_t \end{bmatrix} \tag{2.18}$$

As emphasized before, y_t should not be taken too literally—it is a convenient modeling device for introducing common noise in the agents' forecasts of the state of the economy. Finally, we guess and verify that the state vector X_t follows a simple law of motion:

$$X_t = M X_{t-1} + m_\nu \nu_t + m_\varepsilon \varepsilon_t \tag{2.19}$$

where M is a 2×2 matrix, while m_ν and m_ε are 2×1 vectors. We then seek to characterize the equilibrium values of M, m_ν, and m_ε.

In each period t, firms and workers start with some prior about X_t and use the new signals

the details of how this learning takes place. However, the endogeneity of learning may have distinct normative implications; see Angeletos and La'O (2008) and Angeletos and Pavan (2009) on this issue.

that they receive in the beginning of period t to update their beliefs about X_t. Local output is then determined Condition (2.16) then givens local output as a function of the local belief about X_t. Aggregating across islands, we obtain the aggregate level of output. In equilibrium, the law of motion that aggregate output follows must match the one believed by the firms. Therefore the equilibrium is a fixed point between the law of motion believed by agents and used to form their forecasts of the aggregate state, and the law of motion induced by the optimal output and employment decisions that firms and workers are making following their signal extraction problem. We characterize the fixed point of this problem in the Appendix and use its solution to numerically simulate the impulse responses of output and employment to positive innovations in v_t and ε_t.

For our numerical simulations, we interpret a period as a quarter. Accordingly, we let $\sigma_\nu = 0.02$ for the standard deviation of the productivity innovation and $\psi = 0.99$ for its persistence. Next, we set $\theta = .60$ and $\epsilon = .5$, which correspond to an income share of labor equal to 60% and a Frisch elasticity of labor supply equal to 2. These parameter values are broadly consistent with the literature. Less standard is our choice of γ. Recall that in our setting there is no capital, implying that labor income is the only source of wealth, the elasticity of intertemporal substitution is irrelevant, and γ only controls the income elasticity of labor supply. We accordingly set $\gamma = .2$ to ensure an empirically plausible income effect on labor supply. Next, we set the standard deviations of the noises as $\sigma_x = \sigma_y = 5\sigma_v$. These values are arbitrary, but they are not implausible: when the period is interpreted as a quarter, the information about the current innovations to fundamentals and/or the current level of economic activity is likely to be very limited. Finally, we do not pick any specific value for α (equivalently, ρ). Rather, we study how the variance decomposition of the high-frequency components of output and employment varies as we vary α from 0 to 1 (keeping in mind that a higher α means stronger trade linkagess or, equivalently, a lower ρ).

2.5.1 Impulse responses to productivity and noise shocks

Figure 1 plots the impulse responses of aggregate output and employment to a positive innovation of productivity, for various degrees of α. (The size of the innovation here, and in all other impulse responses we report, is equal to one standard deviation.) Clearly, if aggregate productivity were common knowledge, then output would follow the same AR(1) process as aggregate productivity itself. This is simply because there is no capital in our model. The same thing happens when information is dispersed but there is no strategic complementarity in output decisions ($\alpha = 0$). This is simply because when $\alpha = 0$ islands are effectively isolated from one another; but as each island knows perfectly its own productivity, the entire economy responds to the aggregate shock as if the aggregate shock had been common knowledge.

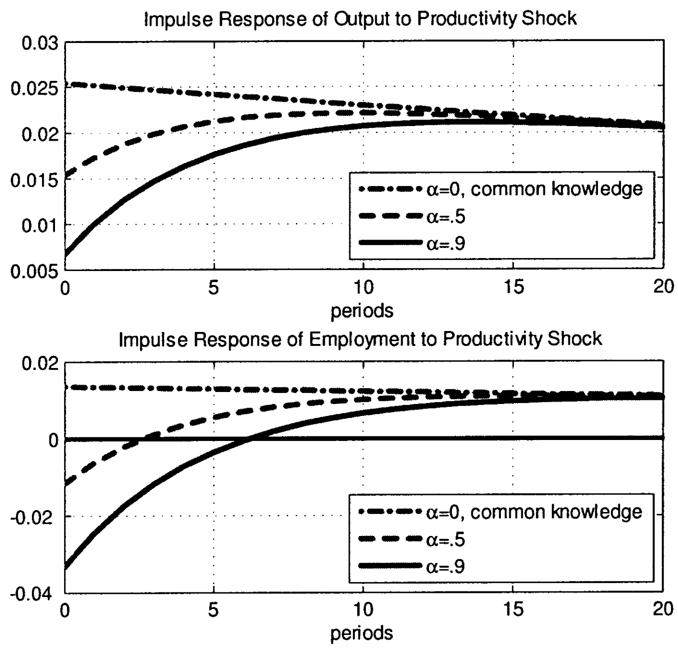

Figure 1

In contrast, when information is dispersed but islands are interconnected ($\alpha \neq 0$), employment and output in one island depends crucially on expectations of employment and output in other islands. As a result, even though each island remains perfectly informed about their local fundamentals, each island responds less to the shock than what it would have done had the shock been common knowledge, precisely because each island expects output in other islands to respond less. Note then that the key for the response of each island is not per se whether the island can disentangle an aggregate shock from an idiosyncratic shock. Even if a particular island was perfectly informed about the aggregate shock, as long as $\alpha > 0$ the island will respond less to this shock than under common knowledge if it expects the other island to respond less, presumably because the other island has imperfect information about the shock. Thus, the key for the inertia in the response of aggregate outcomes is the uncertainty islands face about one another's response, not necessarily the uncertainty they themselves face about the aggregate shock.

As evident in Figure 1, the equilibrium inertia is higher the higher the degree of strategic complementarity. This is because of two reasons. First, there is a direct effect: the higher α is, the less the incentive of each island to respond to the underlying shock for any given expectation of the response of other islands. But then there is also an indirect, multiplier-like, effect: as all other islands are expected to respond less to the underlying shock, each individual island finds it optimal to respond even less.

At the same time, the inertia vanishes in the long-run: the long-run response of the economy to the shock is the same as with common knowledge. This seems intuitive: as time passes, agents become better informed about the underlying aggregate shock. However, that's only part of the story. First, note that agents are always perfectly informed about their own fundamentals, so there is no learning in this dimension. Second, recall that agents do not care per se about the aggregate fundamentals, so the fact that they are learning more about them is per se inconsequential. Rather, the key is that agents in each island are revising their forecasts of the output of other islands. What then drives the result that inertia vanishes in the long-run is merely that forecasts of aggregate output eventually converge their common-knowledge counterpart.[18]

[18] It may be hard to fully appreciate this point, because how fast output forecasts converge to their common-knowledge counterpart is itself pinned down by the speed of learning about the underlying aggregate productivity shock. However, with richer information structures, one can disentangle the speed of adjustment in output forecasts from the speed of learning about the fundamentals. It is then only the former that matters for the result. See Angeletos and La'O (2009a) for a related example within the context of a Calvo-like monetary model.

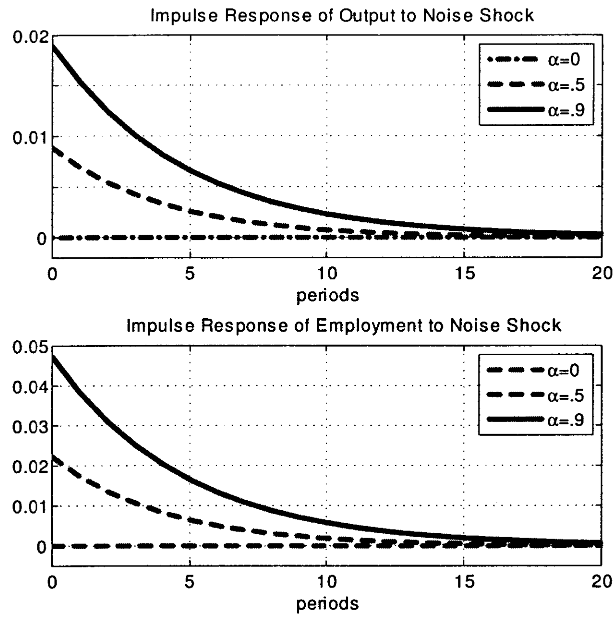

Figure 2

Finally, a salient property of the response of employment is that, for high α, the short-run impact of a productivity shock on employment turns from positive to negative; this happens for parameters values for which the model would have generate a strong positive response had information been symmetric. We find this striking. The baseline RBC paradigm has long been criticized for generating a near perfect correlation between employment and labor productivity, whereas in the data this correlation is near zero. In our setting, this correlation could be close to zero or even turn negative if α is sufficiently high. Of course, correlations may confound the effects of multiple shocks. Some authors in the structural VAR literature have thus sought to show that identified technology shocks lead to a reduction in employment and have then argue that this as a clear rejection of the RBC paradigm (e.g., Galí, 1999; Galí and Rabanal, 2004). Here, we have shown that the dispersion or information may accommodate this fact without invoking sticky prices.

It is worth noting that there are few variants of the baseline RBC model that can also accommodate a negative response of employment to technology shocks, through very different mechanisms than ours. See Collard and Dellas (2005a), Francis and Ramey (2003a), Rotemberg (2003), Wen (2001), and the discussion is Section 4.2 of Galí and Rabanal (2004). Most interestingly for our purposes, as Collard and Dellas (2005a) emphasize, the RBC paradigm faces a tension between, on the one hand, accounting for the negative response of employment to technology shocks and, on the other hand, maintaining the proposition that business cycles are driven by technology shocks. In our framework, this tension is still present, but it is only complementary to our own view about the business cycle: the central position of our approach is that it is the uncertainty agents face about one another's beliefs and responses, not the underlying technology shocks, that explain the bulk of short-run fluctuations.

At the same time, note that it is the dispersion of information, not the uncertainty about the technology shock, that causes employment to fall. If agents had been imperfectly informed about the productivity shock but information had been common, then they could fail to increase their employment as much as they would have done with perfect information, but they would not have reduced their employment—for how could they respond to the shock by reducing employment if they were not aware of the shock in the first place? Thus, employment falls in our model precisely because each agent is well informed about the shock but the shock is not common knowledge.

Turning to the effects of noise, in Figure 2 we consider the impulse responses of output and employment in response to a positive innovation in ε_t. As emphasized before, this should be interpreted as a positive error in expectations of aggregate output, rather than as an error in expectations of aggregate fundamentals. When $\alpha = 0$, such forecast errors are irrelevant, simply because individual incentives do not depend on forecasts of aggregate activity. But when $\alpha = 0$, they generate a positive response in output and employment, thus becoming partly self-fulfilling. Furthermore, the stronger the complementarity, the more pronounced the impact of these errors on aggregate employment and output.

The figure considers a positive noise shock, which means a positive shift in expectations about economic activity. The impact of a negative shift in expectations is symmetric. Note that when these shocks occur, output, employment and consumption move in the same direction, without any movement in TFP. The resulting booms and recessions could thus be (mis)interpreted as a certain type of demand shocks. We will return to this point in a moment.

Finally, note that the impact of these noise shocks on output and employment can be quite persistent, even though the noise itself is not. This is simply because the associated forecast errors are themselves persistent.

2.5.2 Variance decomposition and forecast errors

Comparing the responses of employment with those of output to the two shocks, we see that the former is smaller than the latter in the case of productivity shocks but quite larger in the case of noise. This is simply because productivity shocks have a double effect on output, both directly and indirectly through employment, while the noise impacts output only through employment. But then the response of employment to noise is bound to be stronger than that of output as long as there are diminishing returns to labor ($\theta < 1$), and the more show the lower θ. It follows that noise contributes to a higher relative volatility for employment, while productivity shocks contribute in the opposite direction. In the standard RBC framework, employment may exhibit a higher volatility than output to the extent that there are powerful intertemporal substitution effects (which here we have ruled out since we have also ruled out capital). However, the RBC framework is known to lack in this dimension. Our results here indicate how noise could help improve the performance of the RBC framework in this dimension.

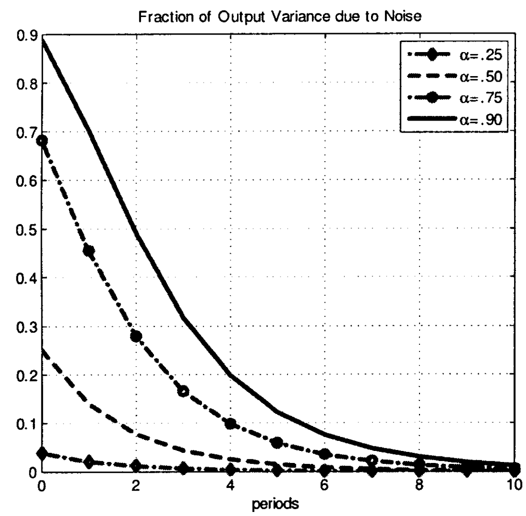

Figure 3a

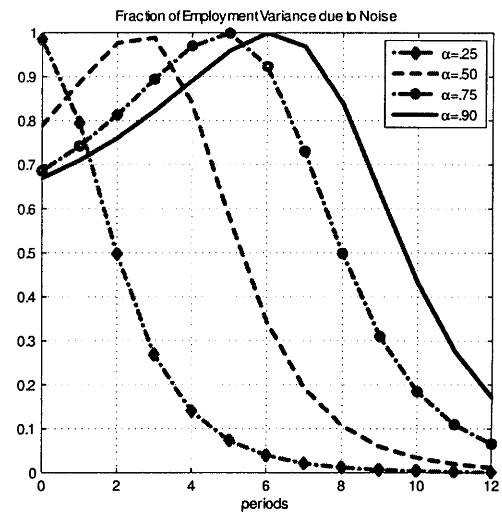

Figure 3b

Comparing Figures 1 and 2, it is evident that low-frequency movements in employment and output are dominated by the productivity shocks, while noise contributes relatively more to high-frequency movements. To further illustrate this property, in Figure 3 we plot the variance decomposition of output and employment at different time horizons. For sufficiently strong strategic complementarity, productivity shocks explain only a small fraction of the high-frequency variation in output—short-run fluctuations are driven mostly by noise. As for employment, the contribution of noise is quite dramatic.

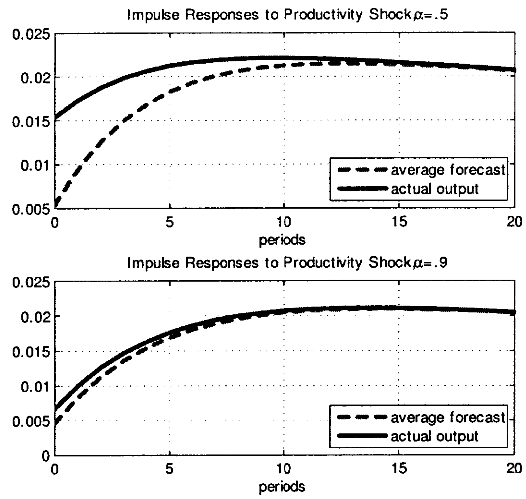

Figure 4a

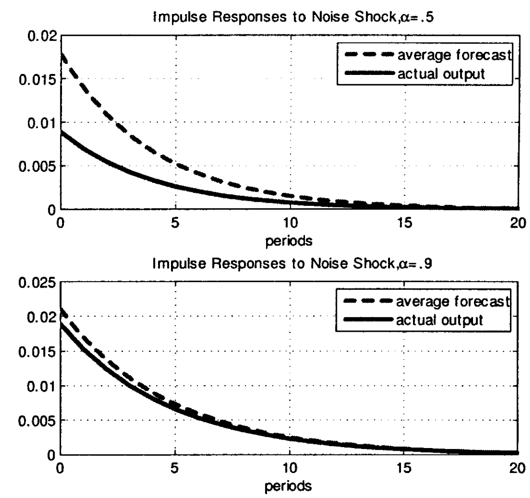

Figure 5b

Finally, Figure 4 plots the dynamics of the average forecast of aggregate output and the true level of aggregate output in response to a productivity or noise shock. The average forecast error is the distance between the two aforementioned variables. A salient feature of this figure is that forecast errors are *smallest* when the degree of strategic complementarity is highest.

This is crucial. We earlier showed that a higher degree of strategic complementarity, α, leads to both more inertia in the response of output and employment to productivity shock, and to a bigger impact of noise. In this sense, the deviation from the common-knowledge benchmark is highest when α is highest. However, one should not expect that these large deviations will show up in large forecast errors. To the contrary, a higher α implies that actual economic activity is more driven by forecasts of economic activity, so that at the end a higher α guarantees that the forecast errors are smaller. It follows that, as we vary α, the magnitude of the deviations of actual outcomes from their common-knowledge counterparts is *inversely* related to the magnitude of the associated forecast errors. Indeed, both the inertia and the impact of noise become nearly self-fulfilling as α gets closer to 1.

Combined, these results illustrate the distinct mark that dispersed information can have on macroeconomic outcomes once combined with strategic complementarity. Not only can the effects we have documented be significant, but they are also consistent with small errors in the agents' forecasts of either the underlying economic fundamentals or the level of economic activity.

2.5.3 Demand shocks, new-Keynesian models, and structural VARs

Many economists have found the idea that short-run fluctuations are driven primarily by technology shocks implausible either on a priori grounds or on the basis of certain structural VARs. Blanchard and Quah (1989) were the first to attempt to provide some evidence that short-run fluctuations are driven by "demand" rather than "supply" shocks, albeit with the caveat that one cannot know what the shocks they identify really capture. Subsequent contributions by Galí (1999), Basu, Fernald and Kimball (2006), Galí and Rabanal (2004), and others have tried to improve in that dimension. One way or another, though, this basic view that business cycles are not driven by technology shocks appears to underly the entire New-Keynesian literature.

Our findings here are consistent with this view. In our environment, technology shocks may explain only a small fraction of the high-frequency volatility in macroeconomic outcomes. However, the residual fluctuations have nothing to do with monetary shocks. Rather, they are the product of the noise in the agents' information. Importantly, to the extent that information is dispersed and trade linkages are important, this noise might be quite small and nevertheless explain a big fraction of the high-frequency volatility in macroeconomic outcomes.

Furthermore, the noise-driven fluctuations we have documented here, albeit being purely neoclassical in their nature, they could well be interpreted as some kind of "demand" or "monetary" shocks in the following sense. This is because they share many of the features often associated with such shocks: they contribute to positive co-movement in employment, output and consumption; they are orthogonal to the underlying productivity shocks; they are closely related to shifts in expectations of aggregate demand; and they explain a large portion of the high-frequency variation in employment and output while vanishing at low frequencies.[19]

To better appreciate this, suppose that we generate data from our model using a random-walk specification for the productivity shock and let an applied macroeconomist—preferably

[19] Of course, further exploring under what conditions our noise-driven fluctuations can be associated also with procyclical nominal prices requires a monetary extension of the model.

of the new-keynesian type—to run a structural VAR as in Blanchard and Quah (1989) or Galí (1999). One would then correctly identify the underlying innovations to productivity by the shock that is allowed to have a long-run effect on output or labor productivity, and the underlying noise shocks by the residual.[20] In the language of Blanchard and Quah, the productivity shocks would be interpreted as "supply shocks" and the noise shocks as "demand shocks ". However, the latter would have no relation to sticky prices and the like. To the contrary, both type of shocks emerge from a purely supply-side mechanism. In the language of Galí (1999) and others, on the other hand, the productivity shocks would be interpreted as "technology shocks ". Furthermore, as already noted, the short-run response of employment to these identified shocks would be negative for high enough α; but this would no favor a sticky-price interpretation.

As mentioned in the introduction, a growing literature explores, within the context of either RBC or New-Keynesian models, the complementary idea that noisy news about future productivity contribute to short-run fluctuations (Barsky and Sims, 2009; Beaudry and Portier, 2004; 2006; Christiano et al., 2008; Gilchrist and Leahy, 2002; Jaimovich and Rebelo, 2009; and Lorenzoni, 2008). Furthermore, Lorenzoni (2008) interprets the resulting fluctuations as "demand shocks" and discusses how they help match related facts. However, there are some crucial differences between this line of research and our work. First and foremost, all these papers focus on fluctuations that originate from uncertainty about a certain type of fundamentals (namely future productivity), not on the distinct type of uncertainty that emerges when information is heterogeneous and that we highlight in our work.[21] Second, the "demand shocks" in Lorenzoni (2008) confound real shocks with monetary shocks. By this we mean the following. Since there is no capital in his model (as in ours), expectations of future productivity would have been irrelevant for current macroeconomic outcomes had nominal prices been flexible; the only reason then that news about future productivity cause demand-like fluctuations is that they cause an expansion in monetary policy away from the one that would replicate flexible-price allocations. A similar comment applies to all the New-Keynesian representatives of this line of research: by

[20] Incidentally, note that it is unclear whether the econometric issues studied in Blanchard, L'Huillier, and Lorenzoni (2009) apply to our model.

[21] In his baseline model, Lorenzoni considers a representative-agent model with symmetric information. In an extension, he allows for dispersed information, but only to facilitate a more plausible calibration of the model.

focusing on monetary policies that fail to replicate the flexible-price allocations, they confuse noise shocks with monetary surprises. In contrast, our "demand shocks" obtain in an RBC setting and are completely unrelated to monetary policy.

Finally, note that a positive productivity shock in our model induces a small impact on output at high frequencies, followed by a large persistence response at lower frequencies.[22] Again these properties are consistent with the estimated dynamics of "technology" shocks.

More generally, note that in many New-Keynesian models sticky prices dampen the response of output to productivity shocks relative to the RBC framework and help get a negative response for employment. As noted earlier, some researchers argue that these properties seem to be more consistent with the data than their RBC counterparts. However, what is a success for these models appears to be only a failure for monetary policy: the only reason that the response of the economy to productivity shocks in the baseline New-Keynesian model differs from that in the baseline RBC model is that monetary policy fails to replicate flexible-price allocations, which is typically the optimal thing to do. Here, instead, we obtain the same empirical properties without introducing sticky properties and without presuming any suboptimality for policy.

Galí, López-Salido and Vallés (2003) argue that the negative empirical response of employment to technology shocks has vanished in the VolckerÐGreenspan era, while it was prevalent earlier on. Within the context of New-Keynesian models, this finding is consistent with the idea that, by shifting focus to price stability, monetary policy has come closer to being optimal during this later period of the data. However, this finding is also consistent within the context of our model with the possibility that advances in information and communication technologies, as well as improved policy transparency, may have contributed to a reduction in the heterogeneity of information. Thus, neither the empirical findings of Galí, López-Salido and Vallés (2003) help discriminate New-Keynesian models from our theory.

Finally, our approach may also have intriguing implications for the identification of monetary shocks. One of the standard identification strategies is based on the idea that monetary policy often reacts to measurement error in the level of aggregate economic activity (Bernanke and Mihov, 1995; Christiano, Eichenbaum and Evans, 1999). In particular, consider the idea

[22] In our numerical exercises, the impact of the productivity shock vanishes asymptotically, only because we have assumed that $\bar{a}_t$ is (slowly) mean-reverting. If instead we assume that $\bar{a}_t$ is a random walk, then the long-run impact of a productivity shocks becomes positive, while the rest of the results remain unaffected.

that measurement error justifies the existence of random shocks to monetary policy, which are orthogonal to the true underlying state of the economy. If one then traces the impact of these particular shocks on subsequent aggregate outcomes, one can escape the endogeneity problem and identify the impact of monetary shocks. However, these measurement errors, or more generally any forecast errors that the central bank makes about current and future economic activity, are likely to be correlated with the corresponding forecast errors of the private sector. But then the so-identified monetary shocks may actually be proxying for the real effects of the forecast errors of the private sector, which unfortunately are not observed by the econometrician.

2.5.4 Labor wedges and Solow residuals

Many authors have argued that a good theory of the business cycle must explain the observed variation in the labor wedge and the Solow residual (e.g., Hall, 1997; Rotemberg and Woodford, 1999; Chari, Kehoe, and McGrattan, 2007; Shimer, 2009). We now consider the implications of our model for these two key characteristics of the business cycle.

Following the literature, we define the labor wedge $\tau_{n,t}$ implicitly by

$$\frac{N_t^{\varepsilon-1}}{C_t^{-\gamma}} = (1 - \tau_{n,t})\,\theta\frac{Q_t}{N_t}.$$

The left panel of Figure 5 plots the impulse response of the labor wedge to a positive productivity and a positive noise shock. The labor wedge follows very different dynamics in response to the two types of shocks. In particular, a positive productivity shock induces a positive response in the labor wedge, implying positive comovement of the labor wedge with output. On the other hand, a positive noise shock produces a negative response in the observed labor wedge, implying a negative comovement with output.

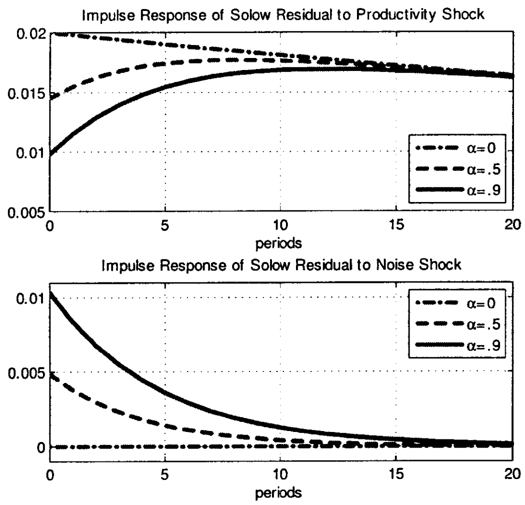

Figure 5a

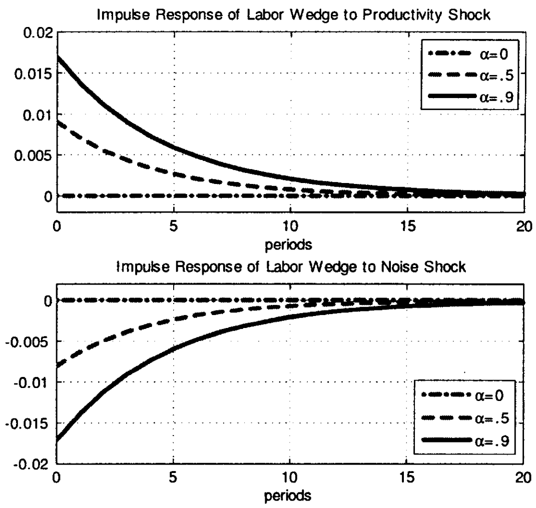

Figure 5b

Multiple authors have documented that variation in the labor wedge plays a large role in accounting for business-cycle fluctuations during the post-war period. Importantly, the labor wedge is highly countercyclical, exhibiting sharp increases during recessions. Shimer (2009) surveys the facts and the multiple explanations that have been proposed for the observed countercyclicality of the labor. These include taxes, shocks to the disutility of labor, mark-up shocks, fluctuations in wage-setting power, and Shimer's preferred explanation, search frictions in the labor market. Here, we have found that noise offers another possible explanation for the same fact.

We finally consider the potential implications of our results for observed Solow residuals. Towards this goal, we now introduce a variable input in the production function; the optimal use of this input responds to shocks, but is unobserved by the econometrician and is thus absorbed in the Solow residual. As in King and Rebelo (2000), our preferred interpretation of this input is capital utilization. The only caveat is that in our model capital exogenously fixed. However, we could introduce capital following the same approach as Angeletos and La'O (2009b), without affecting the qualitative points we seek to make here.

We denote the unobserved input by χ_{it}; we let the gross product of a firm be $\tilde{q}_{it} = \tilde{A}_{it}\chi_{it}^{1-\tilde{\theta}}n_{it}^{\tilde{\theta}}$;

and we specify the cost of this input in terms of final product as $\delta\chi_{it}^{1+\xi}$, where $\xi, \delta > 0$. The net product of a firm is then $q_{it} = \tilde{q}_{it} - \delta\chi_{ut}^{1+\xi}$. Solving out for the optimal level of this input, The optimal level of this input is given by equating its marginal product with its marginal cost: $(1-\tilde{\theta})\frac{q_{it}}{\chi_{it}} = \delta\left(1+\xi\right)\chi_{it}^{\xi}$. We thus obtain obtain the following reduced-form production function:

$$q_{it} = A_{it}n_{it}^{\theta} \tag{2.20}$$

where $\theta \equiv \left(\frac{1+\xi}{\tilde{\theta}+\xi}\right)\tilde{\theta}$ and $A_{it} \equiv \left(\frac{1+\xi}{\tilde{\theta}+\xi}\right)\tilde{A}_{it}^{\frac{1+\xi}{\tilde{\theta}+\xi}}$. Our analysis then remains intact, provided we reinterpret the production function in the above way. Accordingly, we set $\tilde{\theta} = .6$ and $\xi = .1$ (a preferred value in King and Rebelo, 2000), which implies $\theta = .88$. We also re-calibrate the underlying aggregate productivity shocks so that the observed Solow residual ($SR_t \equiv \log Q_t - \theta \log N_t$) implied by the common-knowledge version of the model continues to have a standard deviation of 0.02 and a persistence of 0.99.

The right panel of Figure 5 plots the dynamic response of the Solow residual to a productivity or a noise shock. Both shocks raise the measured Solow residual, but only the innovation in productivity has a persistent effect. Moreover, these responses of the Solow residual mirror those of output. It follows that the Solow residual and output move tightly together, much alike in a standard RBC model, although employment has the more distinct behavior we mentioned earlier.

Finally, it is worth noting that additional variation in measured Solow residuals could obtain from variation in the dispersion of information, simply because the dispersion of information affects the cross-sectional allocations to resources. Note in particular that the observed heterogeneity in forecast surveys is highly countercyclical, suggesting that the dispersion of information may also be countercyclical. Exploring how such variation in the dispersion of information affects the business cycle is left for future work.

2.5.5 Discussion

While the characterization of equilibrium in Section 3 allowed for arbitrary information structures, the more concrete positive results that we documented thereafter presumed a specific, Gaussian information structure (Assumption 1). However, we do not expect any of the predictions we have emphasized to be unduly sensitive to the details of the information structure.

We build this expectation on the following observations. Proposition 1 permits us to map our economy to a class of games with linear best responses, like those studied in Morris and Shin (2002) and Angeletos and Pavan (2007, 2009). In this class of games, one can show under arbitrary information structures that a stronger strategic complementarity makes equilibrium outcomes less sensitive to first-order beliefs (the forecasts of the fundamentals) and more sensitive to higher-order beliefs (the forecasts of the forecasts of others). One can then proceed to show quite generally that higher-order beliefs are more sensitive to the initial common prior, to public signals, and to signals with strongly correlated errors, than lower-order beliefs, simply because these pieces of information are relatively better predictors of the forecasts of others. It follows that higher-order beliefs are less sensitive to innovations in the fundamentals and more sensitive to common sources of noise than lower-order beliefs. Combined, these observations explain why stronger complementarity dampens the response of the economy to innovations in fundamentals while amplifying the impact of noise—which are the key properties that drive the results we documented in Sections 4 and 5. We conclude that these results are not unduly sensitive to the details of the underlying information structure; rather, they obtain from robust properties of higher-order beliefs and the very nature of the general-equilibrium interactions in our economy.

Our analysis has implications, not only for aggregate fluctuations, but also for the cross-sectional dispersion of prices and quantities. As evident from condition (2.23), a higher α necessarily reduces the sensitivity of local output to local fundamentals, while increasing the sensitivity to expectations of aggregate output. When information is commonly shared, all agents share the same expectation of aggregate output, and hence heterogeneity in output (and thereby in prices) can originate *only* from heterogeneity in fundamentals (productivities, tastes, etc). It then follows that a higher α necessarily reduces cross-sectional dispersion in output and prices, simply because it dampens the only source of heterogeneity. However, once information is dispersed, there is an additional source of heterogeneity: different firms have different expectations of aggregate economic activity. It then follows that a higher α dampens the former source of heterogeneity while amplifying the latter. We conclude that, once information is dispersed, the impact of complementarity on cross-sectional dispersion is ambiguous—which also implies that evidence on the cross-sectional dispersion of prices and quantities may provide little guidance for a quantitative assessment of our results.[23]

Similarly, evidence on the size of monopolistic mark-ups, or the elasticity of demands faced by individual firms, do not necessarily discipline the magnitude of our results. This is for two reasons. First, in our model, the mark-up and the elasticity of individual demands identify only η, whereas it is ρ that matters for complementarity. And second, as evident from the definition of α, a high complementarity in our model is consistent with *any* value of ρ, provided that there is a sufficiently small wealth effect on labor supply in the short run, a sufficiently high Frisch elasticity (as in Hansen, 1985), and nearly linear returns to labor in the short run (as in King and Rebelo, 2000).

Finally, it is worth noting that our results need not be subject to the critique that Hellwig and Venkateswaran (2009) raise against Woodford (2003a). That paper considers a New-Keynesian model in which firms cannot tell apart aggregate monetary shocks from idiosyncratic productivity or demand shocks; this is essentially the same as in Lucas (1972), except that firms are monopolistic, and can be viewed as a micro-foundation of Woodford (2003a). For a particular calibration of that model, the aforementioned confusion induces firms to adjust their prices a lot in response to monetary shocks even when these shocks are unobserved. In effect, nominal prices adjust a lot to monetary shocks, albeit for the "wrong reasons ". These findings are interesting on their own right—and may also complement our motivation for focusing on *real* rather than monetary shocks. However, one cannot possibly extrapolate from that paper to the likely quantitative importance of our results. First, the core mechanism of that paper does not apply to our context: if firms were to confuse aggregate shocks for local ones in our model, this confusion would only reinforce our results.[24] And second, the quantitative findings of that paper are based on a number of heroic assumptions, which might serve certain purposes but are out of place in our own context.[25]

[23] One of our discussants made the opposite argument. But his argument was based on the premise that a higher α necessarily reduces cross-sectional dispersion. This happens to be true under the specific signal structure we introduced in Assumption 1 but, as just explained, is not true in general.

[24] To see this, recall from Proposition 3 and Corollary 3 that the response of equilibrium output to an idiosyncratic shock in fundamentals is given by $\varphi_f = 1 - \alpha$, while its response to an aggregate shock is given by $\varphi_\nu = 1 - \alpha \frac{\kappa_f}{(1-\alpha)\kappa_x + \kappa_y + \kappa_f}$. As long as $\alpha > 0$, φ_f is smaller than φ_ν, which means that mistaking an aggregate shock for an idiosyncratic shock only helps dampen the response of the economy to the aggregate shock.

[25] In particular, Hellwig and Venkateswaran (2009) assume that workers are perfectly informed about the monetary shocks, so that nominal wages adjust one-to-one with them. When firms face constant real marginal costs and iso-elastic demands, this assumption can *alone* guarantee that prices will move one-to-one with monetary

With these observations we are not trying to escape the need for a serious quantitative exercise, nor are we ready to speculate on the outcome of such an exercise. We are only trying to provide some guidance for any future quantitative exploration of our results. The key effects we have documented in this chapter hinge only on (i) the sensitivity of individual output to forecasts of aggregate output and (ii) the sensitivity of these forecasts to the underlying shocks. We are thus skeptical that micro evidence on prices or quantities can alone provide enough guidance on the quantitative importance of our results. We instead propose that a quantitative assessment of our results should rely more heavily on survey evidence about the agents' forecasts of economic activity. Indeed, these forecasts concisely summarize all the informational effects in our model, and their joint stochastic behavior with actual outcomes speaks to the heart of our results.

In this regard, we find the approach taken in Coibion and Gorodnichenko (2008) particularly promising. this chapter uses survey evidence to study how the agents' forecasts of certain macroeconomic outcomes respond to certain structural shocks (with the latter being identified by specific structural VARs). In effect, the exercises conducted in that paper are empirical analogues of the theoretical exercise we conducted in Figure 4 for the case of productivity shocks. Combined, the empirical investigation of that paper and the theoretical one of our paper indicate how the focus in recent research could be shifted away from the details of the underlying informational frictions to the joint stochastic properties of the agents' forecasts and the actual macroeconomic outcomes.

2.6 Efficiency

The positive properties we have documented are intriguing. However, their normative content is unclear. Is the potentially high contribution of noise to business-cycle fluctuations, or the

shocks even if firms cannot tell whether their nominal wages have moved because of nominal or idiosyncratic reasons. Clearly, the empirical relevance of this assumption may be questionable even within the context of that paper. As for our own context, we see no good reason for assuming a priori that workers are perfectly informed about the aggregate *real* shocks hitting the economy. Furthermore, Hellwig and Venkateswaran (2009) assume that firms are free to adjust their action at no cost and at a daily or weekly frequency. When that action is interpreted as a nominal price (as in that paper), this assumption serves a useful pedagogical purpose: it helps isolate information frictions from sticky prices. But once that action is interpreted as a real employment or investment choice (as in our model), this assumption makes no sense: the "stickiness" of real employment and investment decisions is a matter of technology, not a matter of contracts.

potentially high inertia in the response of the economy to innovations in productivity, a symptom of inefficiency?

It is obvious that a planner could improve welfare if he could centralize all the information that is dispersed in society and then dictate allocations on the basis of all this information. But this would endow the planner with a power that seems far remote from the powers that policy makers have in reality. Furthermore, the resulting superiority of centralized allocations over their decentralized equilibrium counterparts would not be particularly insightful, since it would be driven mostly by the assumption that the planner has the superior power to overcome the information frictions imposed on the market. Thus, following Angeletos and Pavan (2007a, 2009) and Angeletos and La'O (2008), we contend that a more interesting question—on both practical and conceptual grounds—is to understand whether a planner could improve upon the equilibrium while being subject to the same informational frictions as the equilibrium.

This motivates us to consider a constrained efficiency concept that permits the planner to choose any resource-feasible allocation that respects the geographical segmentation of information in the economy—by which we simply mean that the planner cannot make the production and employment choices of firms and workers in one island contingent on the private information of another island. A formal definition of this efficiency concept and a detailed analysis of efficient allocations can be found, for a variant model, in Angeletos and La'O (2008). Here we focus on the essence.

Because of the concavity of preferences and technologies, efficiency dictates symmetry in consumption across households, as well as symmetry across firms and workers within any given island. Using these facts, we can represent the planning problem we are interested in as follows.

Planner's problem. *Choose a pair of local production and employment strategies, $q : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$ and $n : \mathcal{S}_\omega \times \mathcal{S}_\Omega \rightarrow \mathbb{R}_+$, and an aggregate output function, $Q : \mathcal{S}_\Omega^2 \rightarrow \mathbb{R}_+$, so as to maximize*

$$\int_{\mathcal{S}_\Omega} \left[U(Q(\Omega_t, \Omega_{t-1})) - \int_{\mathcal{S}_\omega} \tfrac{1}{1+\epsilon} S(\omega) n(\omega, \Omega_{t-1})^{1+\epsilon} d\Omega_t(\omega) \right] d\mathcal{P}(\Omega_t|\Omega_{t-1}) \tag{2.21}$$

subject to

$$q(\omega, \Omega_{t-1}) = A(\omega) n(\omega, \Omega_{t-1})^\theta \; \forall (\omega, \Omega_{t-1}) \tag{2.22}$$

$$Q(\Omega_t, \Omega_{t-1}) = \left[\int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right]^{\frac{\rho}{\rho-1}} \quad \forall (\Omega_t, \Omega_{t-1}) \tag{2.23}$$

where $\mathcal{P}(\Omega_t|\Omega_{t-1})$ *denotes the probability distribution of* Ω_t *conditional on* Ω_{t-1}.

This problem has a simple interpretation. $U(Q(\Omega_t, \Omega_{t-1})$ is the utility of consumption for the representative household; $\frac{1}{\epsilon} S(\omega) n(\omega, \Omega_{t-1})^{\epsilon}$ is the marginal disutility of labor for the typical worker in a given island; and the corresponding integral is the overall disutility of labor for the representative household. Furthermore, note that, once the planner picks the production strategy q, the employment strategy n is pinned down by (2.22) and the aggregate output function Q is pinned down by (2.22). The reduced-form objective in (2.21) is thus a functional that gives the level of welfare implied by any arbitrary production strategy that the planner dictates to the economy.

Because this problem is strictly concave, it has a unique solution and this solution is pinned down by the following first-order condition:[26]

$$S_{it} n_{it}^{\epsilon} = \mathbb{E}_{it} \left[U'(Q_t) \left(\frac{q_{it}}{Q_t} \right)^{-\frac{1}{\rho}} \right] \left(\theta A_{it} n_{it}^{\theta-1} \right). \tag{2.24}$$

This condition simply states that the planner dictates the agents to equate the social cost of employment in their island with the *local* expectation of the social value of the marginal product of that employment. Essentially the same condition characterizes (first-best) efficiency in the standard, symmetric-information paradigm. The only difference is that there expectations are conditional on the commonly-available information set, while here they are conditional on the locally-available information sets.

As with equilibrium, we can use $q_{it} = A_{it} n_{it}^{\theta}$ to eliminate n_{it} in the above condition, thereby reaching the following result.

Proposition 34 *Let*

$$f^*(\omega) \equiv \log \left\{ \theta^{\frac{1}{\frac{\epsilon}{\theta}+\gamma-1}} \left(\frac{A(\omega)}{S(\omega)} \right)^{\frac{\frac{\epsilon}{\theta}}{\frac{\epsilon}{\theta}+\gamma-1}} \left(\frac{A(\omega)}{S(\omega)} \right)^{\frac{\frac{\epsilon}{\theta}}{\frac{\epsilon}{\theta}+\gamma-1}} \right\}$$

[26] Because of the continuum, the efficient allocation is determined only for *almost* every ω. For expositional simplicity, we bypass the *almost* qualification throughout the paper.

be a composite of the local productivity and taste shocks. The efficient strategy $q : S_\omega \times S_\Omega \to \mathbb{R}_+$ *is the fixed point to the following:*

$$\log q(\omega_t, \Omega_{t-1}) = (1-\alpha) f^*(\omega_t) + \alpha \log \left\{ \mathbb{E}\left[Q(\Omega_t, \Omega_{t-1})^{\frac{1}{\rho}-\gamma} \middle| \omega_t, \Omega_{t-1} \right]^{\frac{1}{\frac{1}{\rho}-\gamma}} \right\} \ \forall (\omega_t, \Omega_{t-1}), \quad (2.25)$$

$$Q(\Omega_t, \Omega_{t-1}) = \left[\int q(\omega, \Omega_{t-1})^{\frac{\rho-1}{\rho}} d\Omega_t(\omega) \right]^{\frac{\rho}{\rho-1}} \ \forall (\Omega_t, \Omega_{t-1}). \quad (2.26)$$

A number of remarks are worth making. First, note that the composite shock f_t^* plays a similar role for the efficient allocation as the composite shock f_t played for the equilibrium: it identifies the fundamentals that are relevant from the planner's point of view. This is evident, not only from the above result, but also directly from the planner's problem: using $q_t = A_t n_t^\theta$ to eliminate n_t in the expression for welfare given in the planner's problem, we can express welfare as a simple function of the production strategy and the composite shock f_t^* alone.

Second, note that Proposition 34 permits a game-theoretic interpretation of the efficient allocation, much alike what Proposition 28 did for equilibrium: the efficient allocation of the economy coincides with the Bayes-Nash equilibrium of a game in which the different players are the different islands of the economy and their best responses are given by (2.25).

Third, note that, apart from the different composite shock, the structure of the fixed point that characterizes the efficient and the equilibrium allocation is the same: once we replace $f^*(\omega_t)$ with $f(\omega_t)$, condition (2.25) coincides with its equilibrium counterpart, condition (2.6). And because $f^*(\omega_t) = f(\omega_t)$ for every ω_t if and only if there is no monopoly power, the following is immediate.

Corollary 35 *In the absence of monopoly distortions, the equilibrium is efficient, no matter the information structure.*

This result establishes that neither the presence of noise nor the dispersion of information are per se sources of inefficiency. This result might sound bizarre in light of our earlier results that the economy can feature extreme amplification effects, with a tiny amount of noise contributing to large aggregate fluctuations. However, it should be ex post obvious. What causes these large positive effects is the combination of dispersed information and strong complementarity. But neither one introduces a wedge between the equilibrium and the planner. Indeed, the geographical segmentation of information is similar to a technological constraint that impacts equilibrium and efficient allocations in a completely symmetric way. As for the complementarity, it's origin is preferences and technologies, not any type of market inefficiency, guaranteeing that private motives in coordinating economic activity are perfectly aligned with social motives. It follows that, when stronger complementarity amplifies the impact of noise, it does so without causing any inefficiency.[27]

We can generalize this result for situations where firms have monopoly power, to the extent that there are no aggregate shocks to monopoly power, as follows.

Corollary 36 *Suppose that information is Gaussian (Assumption 1 holds) and there are no aggregate mark-up shocks ($\bar{f}_t^* - \bar{f}_t$ is fixed). Then, the the business cycle is efficient in the sense the gap $\log Q_t - \log Q_t^*$ between the equilibrium and the efficient level of output is invariant.*

If we allow for mark-up shocks, then clearly the equilibrium business cycle ceases to be efficient. But this is true irrespectively of whether information is dispersed or commonly shared. We conclude that the dispersion of information per se is not a source of inefficiency, whether one considers a competitive RBC or a monopolistic New-Keynesian model. We further discuss the implications of this result for optimal policy and the social value of information in Angeletos and La'O (2008).

We conclude this section with an important qualification. While our efficiency results allowed for an arbitrary information structure, they restricted the information structure to be exogenous to the underlying allocations. This ignores the possibility that information gets endogenously aggregated through prices, macro indicators, and other channels of social learning—which is clearly an important omission. We address this issue, too, in Angeletos and La'O (2008), by allowing information to get partly aggregated through certain price and quantity indicators. We first show that a planner who internalizes the endogeneity of the information contained in these indicators will choose a different allocation than the equilibrium. This typically means that the planner likes to increase the sensitivity of allocations to private information, so as to increase the precision of the information that gets revealed by the available macroeconomic indicators. We then explore policies that could help in this direction.

[27] As mentioned earlier, this is the opposite of what happens in Morris and Shin (2002).

2.7 Concluding remarks

The pertinent macroeconomics literature has used informational frictions to motivate why economic agents may happen, or choose, to be partly unaware about the shocks hitting the economy. Sometimes the informational friction is exogenous, sometimes it is endogenized. Invariably, though, the main modeling role of informational frictions seems to remain a simple and basic one: to limit the knowledge that agents have about the underlying shocks to economic fundamentals.

Our approach, instead, seeks to highlight that the *heterogeneity* of information may have a very distinct mark on macroeconomic outcomes than the uncertainty about fundamentals. We highlighted this in this chapter by showing how the heterogeneity of information can induce significant inertia in the response of the economy to productivity shocks, and can also generate significant noise-driven fluctuations, even when the agents are well informed about the underlying fundamentals. In Angeletos and La'O (2009b), we further show that the heterogeneity of information can open the door to a novel type of sentiment shocks—namely shocks that are independent of either the underlying fundamentals or the agents' expectations of the fundamentals and nevertheless cause variation in the agents' forecasts of economic activity and thereby in actual economic activity, despite the uniqueness of equilibrium. This in turn permits a broader interpretation of what noise stood for in the present paper: noise could be interpreted more generally as any variation in the forecasts of economic activity that is orthogonal by fundamentals.

In this chapter, we focused on the dispersion of information about the *real* shocks hitting the economy, ruling out sticky prices and dismissing any lack of common knowledge about innovations to monetary policy. This, however, does not mean that we see no interesting interaction between dispersed information and nominal frictions. It only means that we find it a good modeling benchmark to assume common knowledge of the current monetary policy. Where we instead see an intriguing interaction between our approach and monetary policy is the following dimension: when there is dispersed information about the underlying *real* shocks hitting the economy and nominal prices are rigid, the response of monetary policy to any information that becomes available about these shocks may be crucial for how the economy responds to these shocks in the first place. This point was first emphasized at a more abstract level by Angeletos and Pavan (2007b, 2009) and is further explored by Angeletos and La'O (2008) and Lorenzoni (2009) within new-Keynesian variants of the economy we have studied in this chapter.

We conclude with a comment on the alternative formalizations of informational frictions. For certain questions, one formalization might be preferable to another; for example, if one wishes to understand which particular pieces of information agents are likely to pay more attention to, Sims (2003) offers an elegant, intriguing, and micro-founded methodology. However, for certain other questions, the specifics of any particular formalization may prove unnecessary, or even distracting. The results we have emphasized in this chapter appear to hinge only on the heterogeneity of information, not on the specific details of the information structure. To highlight this, we showed that the information structure matters for economic outcomes *only* through its impact on the agents' forecasts of aggregate economic activity. We would thus invite other researchers not to commit to any particular formalization of the information structure (including ours), but rather to take a more flexible approach to the modeling of informational frictions. After all, the data cannot possibly inform us about the details of the information structure. What, instead, the data can do is to inform us about the stochastic properties of the agents' forecasts of economic activity—which, as mentioned, is the only channel through which the dispersion of information matters of economic behavior. Thus, in our view, it is only this evidence that should help discipline the theory.

Appendix

Proof of Proposition 28. The characterization of the equilibrium follows directly from the discussion in the main text. Its existence and uniqueness can be obtained by showing that the equilibrium coincides with the solution to a concave planning problem. For the case that there is no monopoly power ($\eta = \infty$), this follows directly from our analysis in Section 6 and in Proposition 34. A similar result can be obtained for the case with monopoly power.

Proof of Proposition 29. This follows from the discussion in the main text.

Proof of Proposition 30. Suppose that, conditional on ω_t and Ω_{t-1}, $Q(\Omega_t, \Omega_{t-1})$ is log-normal, with variance independent of ω_t; that this is true under the log-normal structure for the underlying shocks and signals we will prove shortly. Using log-normality of Q in condition (2.6), we infer that the equilibrium production strategy must satisfy condition (2.8) with

$$const = \frac{\alpha}{2}\left(\frac{1}{\rho} - \gamma\right) \mathrm{Var}\left[\log Q(\Omega_t, \Omega_{t-1})|\omega_t, \Omega_{t-1}\right]$$

and $\mathrm{Var}\left[\log Q(\Omega_t, \Omega_{t-1})|\omega_t, \Omega_{t-1}\right] = \mathrm{Var}\left[\log Q(\Omega_t, \Omega_{t-1})|\Omega_{t-1}\right]$.

We now guess and verify a log-linear equilibrium under the log-normal specification for the shock and information structure. Suppose the equilibrium production strategy takes a log-linear form: $\log q_t = \varphi_0 + \varphi_{-1}\bar{f}_{t-1} + \varphi_f f_t + \varphi_x x_t + \varphi_y y_t$, for some coefficients $(\varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$. Aggregate output is then given by

$$\log Q(\Omega_t, \Omega_{t-1}) = \varphi_0' + \varphi_{-1}\bar{f}_{t-1} + (\varphi_f + \varphi_x)\bar{f}_t + \varphi_y y_t \tag{2.27}$$

where $\varphi_0' \equiv \varphi_0 + \frac{1}{2}\left(\frac{\rho-1}{\rho}\right)\left[\frac{\varphi_f^2}{\kappa_\xi} + \frac{\varphi_x^2}{\kappa_x} + 2\frac{\varphi_f\varphi_x}{\kappa_x}\right]$. It follows that $Q(\Omega_t, \Omega_{t-1})$ is indeed log-normal, with

$$\begin{aligned} \mathbb{E}\left[\log Q(\Omega_t, \Omega_{t-1})|\omega_t, \Omega_{t-1}\right] &= \varphi_0' + \varphi_{-1}\bar{f}_{t-1} + (\varphi_f + \varphi_x)\mathbb{E}\left[\bar{f}_t|\omega_t, \Omega_{t-1}\right] + \varphi_y y_t \quad (2.28) \\ Var\left[\log Q(\Omega_t, \Omega_{t-1})|\omega_t, \Omega_{t-1}\right] &= (\varphi_f + \varphi_x)^2\left(\frac{1}{\kappa_f + \kappa_x + \kappa_y}\right) \quad (2.29) \end{aligned}$$

where $\mathbb{E}\left[\bar{f}_t|\omega_t, \Omega_{t-1}\right] = \frac{\kappa_f}{\kappa_f+\kappa_x+\kappa_y}\psi f_{t-1} + \frac{\kappa_x}{\kappa_f+\kappa_x+\kappa_y}x_t + \frac{\kappa_y}{\kappa_f+\kappa_x+\kappa_y}y_t$. Substituting these expressions into (2.8) gives us

$$\begin{aligned}\log q\left(\omega_t, \Omega_{t-1}\right) &= const + (1-\alpha)\, f\left(\omega\right) + \alpha\left({\varphi_0}' + \varphi_{-1}\bar{f}_{t-1} + \varphi_y y_t\right) \\ &\quad + \alpha(\varphi_f + \varphi_x)\left(\frac{\kappa_f}{\kappa_f+\kappa_x+\kappa_y}\psi f_{t-1} + \frac{\kappa_x}{\kappa_f+\kappa_x+\kappa_y}x_t + \frac{\kappa_y}{\kappa_f+\kappa_x+\kappa_y}y_t\right)\end{aligned}$$

For this to coincide with $\log q\left(\omega\right) = \varphi_0 + \varphi_{-1}\bar{f}_{t-1} + \varphi_f f + \varphi_x x + \varphi_y y$ for every (f, x, y), it is necessary and sufficient that the coefficients $(\varphi_0, \varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$ solve the following system:

$$\begin{aligned}\varphi_0 &= const + \alpha\varphi_0' \\ \varphi_f &= 1-\alpha \\ \varphi_x &= \alpha(\varphi_f + \varphi_x)\left(\frac{\kappa_x}{\kappa_f+\kappa_x+\kappa_y}\right) \\ \varphi_{-1} &= \alpha\varphi_{-1} + \alpha(\varphi_f + \varphi_x)\left(\frac{\kappa_f}{\kappa_f+\kappa_x+\kappa_y}\right)\psi \\ \varphi_y &= \alpha\varphi_y + \alpha(\varphi_f + \varphi_x)\left(\frac{\kappa_y}{\kappa_f+\kappa_x+\kappa_y}\right)\end{aligned}$$

The unique solution to this system for $(\varphi_{-1}, \varphi_f, \varphi_x, \varphi_y)$ is the one given in the proposition; φ_0 is then uniquely determined from the first equation of this system along with the definition of *const* and φ_0'.

Proof of Proposition 33. The result follows by a triple limit. First, take $\alpha \to 1$; next, take $\kappa_y \to 0$; and finally, take $\kappa_x \to \infty$. It is easy to check that this triple limit implies $\kappa \to \infty$ and $R \to 1$. That is, the precision of the agents posterior about the fundamentals (the mean squared forecast error) converges to zero, while the fraction of the high-frequency variation in output that is due to noise converges to 100%.

Kalman filtering for dynamic extension. The method we use in solving this equilibrium is similar to that found in Woodford (2003b).

State Vector and Law of Motion. We guess and verify that the relevant aggregate state variables of the economy at time t are $\bar{f}_t$ and $\log Q_t$ and thus define state vector X_t in (2.17) accordingly.

Claim. *The dynamics of the economy are given by the following law of motion*

$$X_t = MX_{t-1} + m_v v_t + m_\varepsilon \varepsilon_t \tag{2.30}$$

with

$$M \equiv \begin{bmatrix} \psi & 0 \\ M_{21} & M_{22} \end{bmatrix}, m_v \equiv \begin{bmatrix} 1 \\ m_{v2} \end{bmatrix}, m_\varepsilon \equiv \begin{bmatrix} 0 \\ m_{\varepsilon 2} \end{bmatrix}. \tag{2.31}$$

The coefficients $(M_{21}, M_{22}, m_{v2}, m_{\varepsilon 2})$ *are given by*

$$\begin{aligned} M_{21} &= \psi\left(K_{21} + K_{22}\right) && (2.32) \\ M_{22} &= \psi\left(1 - K_{21} - K_{22}\right) && (2.33) \\ m_{u2} &= 1 - \alpha\left(1 - K_{21} - K_{22}\right) && (2.34) \\ m_{\eta 2} &= \alpha K_{22} && (2.35) \end{aligned}$$

and

$$K \equiv \begin{bmatrix} K_{11} & K_{21} \\ K_{21} & K_{22} \end{bmatrix}$$

is the matrix of kalman gains, defined by

$$K \equiv \mathbb{E}\left[\left(X_t - \mathbb{E}_{i,t-1}\left[X_t\right]\right)\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)'\right] \mathbb{E}\left[\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)'\right]^{-1} \tag{2.36}$$

We verify this claim in the following and describe the procedure for finding the fixed point.

Observation Equation. In each period t, firms and workers on island i observe vector $z_{i,t}$, as in (2.18), of private and public signals. In terms of the aggregate state and error terms, island i's observation equation takes the form

$$z_{i,t} \equiv \begin{bmatrix} e_1' \\ e_1' \end{bmatrix} X_t + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \varsigma_{it} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \varepsilon_t \tag{2.37}$$

where e_j is defined as a column vector of length two where the j-th entry is 1 and all other entries are 0.

Forecasting and Inference. Island i's $t-1$ forecast of z_t^i is given by

$$\mathbb{E}_{i,t-1}\left[z_{i,t}\right] = \begin{bmatrix} e_1' \\ e_1' \end{bmatrix} \mathbb{E}_{i,t-1}\left[X_t\right]$$

where $\mathbb{E}_{i,t-1}\left[X_t\right]$ is island's i's $t-1$ forecast of X_t. Combining this with the law of motion (2.30), it follows that $\mathbb{E}_{i,t-1}\left[X_t\right] = M\mathbb{E}_{i,t-1}\left[X_{t-1}\right]$.

To form minimum mean-squared-error estimates of the current state, firms and workers on each island use the kalman filter to update their forecasts. Updating is done via

$$\mathbb{E}_{i,t}\left[X_t\right] = \mathbb{E}_{i,t-1}\left[X_t\right] + K\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right), \tag{2.38}$$

where K is the 2×2 matrix of Kalman gains, defined in (2.36). Substitution of island i's $t-1$ forecast of z_t^i into (2.38) gives us

$$\mathbb{E}_{i,t}\left[X_t\right] = \left(I - K\begin{bmatrix} e_1' \\ e_1' \end{bmatrix}\right) M\mathbb{E}_{i,t-1}\left[X_{t-1}\right] + Kz_{i,t} \tag{2.39}$$

Let $\bar{\mathbb{E}}_t\left[X_t\right] \equiv \int_I \mathbb{E}_{i,t}\left[X_t\right] di$ be the time t average expectation of the current state. Aggregation over (2.39) implies

$$\bar{\mathbb{E}}_t\left[X_t\right] = \left(I - K\begin{bmatrix} e_1' \\ e_1' \end{bmatrix}\right) M\bar{\mathbb{E}}_{t-1}\left[X_{t-1}\right] + K\int z_{i,t}di$$

Finally, using the fact that aggregration over signals yields $\int z_{i,t}di = \begin{bmatrix} e_1' \\ e_1' \end{bmatrix} X_t + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \varepsilon_t$, it follows that the average expectation evolves according to

$$\begin{aligned} \bar{\mathbb{E}}_t\left[X_t\right] &= K\begin{bmatrix} e_1' \\ e_1' \end{bmatrix} MX_{t-1} + \left(I - K\begin{bmatrix} e_1' \\ e_1' \end{bmatrix}\right) M\bar{\mathbb{E}}_{t-1}\left[X_{t-1}\right] \\ &+K\begin{bmatrix} e_1' \\ e_1' \end{bmatrix} m_v v_t + K\left(\begin{bmatrix} e_1' \\ e_1' \end{bmatrix} m_\varepsilon + \begin{bmatrix} 0 \\ 1 \end{bmatrix}\right)\varepsilon_t \end{aligned} \tag{2.40}$$

where M,m_v, m_ε are given by (2.31).

Characterizing Aggregate Output. Local output in each island is determined by the best-response-like condition in (2.16), which may be rewritten as $\log q_{i,t} = (1-\alpha) f_t + \alpha e_2' \mathbb{E}_{i,t}[X_t]$. Aggregating over this condition, we find that aggregate output must satisfy

$$\log Q_t = (1-\alpha)\bar{f}_t + \alpha e_2' \bar{\mathbb{E}}_t [X_t] \tag{2.41}$$

Substituting our expression for $\bar{\mathbb{E}}_t [X_t]$ from (2.40) into (2.41), gives us

$$\begin{aligned} \log Q_t &= [(1-\alpha)\psi + \alpha\psi(K_{21}+K_{22})]\bar{f}_{t-1} + [\alpha M_{21} - \alpha\psi(K_{21}+K_{22})]\bar{\mathbb{E}}_{t-1}\left[\bar{f}_{t-1}\right] \\ &\quad + \alpha M_{22}\bar{\mathbb{E}}_{t-1}[\log Q_{t-1}] + [(1-\alpha)+\alpha(K_{21}+K_{22})]v_t + \alpha K_{22}\varepsilon_t \end{aligned}$$

Moreover, rearranging condition (2.41), we find that $\bar{\mathbb{E}}_t[\log Q_t] = \frac{1}{\alpha}\left(\log Q_t - (1-\alpha)\bar{f}_t\right)$. Finally, using this condition in the above equation gives us

$$\begin{aligned} \log Q_t &= [(1-\alpha)\psi + \alpha\psi(K_{21}+K_{22}) - M_{22}(1-\alpha)]\bar{f}_{t-1} + M_{22}\log Q_{t-1} \\ &\quad + [\alpha M_{21} - \alpha\psi(K_{21}+K_{22})]\bar{\mathbb{E}}_{t-1}\left[\bar{f}_{t-1}\right] + [1-\alpha+\alpha(K_{21}+K_{22})]v_t + \alpha K_{22}\varepsilon_t \end{aligned}$$

For this to coincide with the law of motion conjectured in (2.30) and (2.31) for every $(\bar{f}_{t-1}, \log Q_{t-1}, v_t, \varepsilon_t)$, it is necessary and sufficient that the coefficients $(M_{21}, M_{22}, m_{v2}, m_{\varepsilon 2})$ solve the following system:

$$\begin{aligned} M_{21} &= (1-\alpha)\psi + \alpha\psi(K_{21}+K_{22}) - M_{22}(1-\alpha) \\ m_{v2} &= 1-\alpha+\alpha(K_{21}+K_{22}) \\ m_{\varepsilon 2} &= \alpha K_{22} \\ 0 &= \alpha M_{21} - \alpha\psi(K_{21}+K_{22}) \end{aligned}$$

The unique solution to this system for $(M_{21}, M_{22}, m_{v2}, m_{\varepsilon 2})$ is the one given in the proposition. Therefore, given the kalman gains matrix K, we can uniquely identify the coefficients of the law of motion of X_t.

Kalman Filtering. Let us define the variance-covariance matrices of forecast errors as

$$
\begin{aligned}
\Sigma &\equiv \mathbb{E}\left[\left(X_t - \mathbb{E}_{i,t-1}\left[X_t\right]\right)\left(X_t - \mathbb{E}_{i,t-1}\left[X_t\right]\right)'\right] \\
V &\equiv \mathbb{E}\left[\left(X_t - \mathbb{E}_{i,t}\left[X_t\right]\right)\left(X_t - \mathbb{E}_{i,t}\left[X_t\right]\right)'\right]
\end{aligned}
$$

These matrices will be the same for all islands i, since their observation errors are assumed to have the same stochastic properties. Using these matrices, we may write K as the product of two components:

$$
\mathbb{E}_i\left[\left(X_t - \mathbb{E}_{i,t-1}\left[X_t\right]\right)\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)'\right] = \Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix}
$$

and

$$
\begin{aligned}
\mathbb{E}_i\left[\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)'\right] &= \begin{bmatrix} e_1' \\ e_1' \end{bmatrix} \Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_v^2 \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \\
&+\sigma_\varepsilon^2 \left(\begin{bmatrix} e_1' \\ e_1' \end{bmatrix} m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} m_\varepsilon' \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \right)
\end{aligned} \tag{2.42}
$$

Therefore, K is given by

$$
K = \left(\Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix}\right)\left(\sigma_z^2\right)^{-1} \tag{2.43}
$$

where $\sigma_z^2 \equiv \mathbb{E}_i\left[\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)\left(z_{i,t} - \mathbb{E}_{i,t-1}\left[z_{i,t}\right]\right)'\right]$ is given by (2.42).

Finally, what remains to determine is the matrix Σ. The law of motion implies that matrices Σ and V satisfy

$$
\Sigma = MVM' + \sigma_v^2 m_v m_v' + \sigma_\varepsilon^2 m_\varepsilon m_\varepsilon',
$$

In addition, the forecasting equation (2.39) imply these matrices must further satisfy

$$
V = \Sigma - \left(\Sigma \begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix}\right)\left(\sigma_z^2\right)^{-1} \left(\begin{bmatrix} e_1' \\ e_1' \end{bmatrix} \Sigma + \sigma_\varepsilon^2 \begin{bmatrix} 0 \\ 1 \end{bmatrix} m_\varepsilon' \right)
$$

Combining the above two equations, we obtain the stationary *Ricatti Equation* for Σ:

$$\begin{aligned}\Sigma \;=\;& M\Sigma M' - M\left(\Sigma\begin{bmatrix} e_1 & e_1 \end{bmatrix} + \sigma_\varepsilon^2 m_\varepsilon \begin{bmatrix} 0 & 1 \end{bmatrix}\right)\left(\sigma_z^2\right)^{-1}\left(\begin{bmatrix} e_1' \\ e_1' \end{bmatrix}\Sigma + \sigma_\varepsilon^2 \begin{bmatrix} 0 \\ 1 \end{bmatrix} m_\varepsilon'\right)M' \\ & + \sigma_v^2 m_v m_v' + \sigma_\varepsilon^2 m_\varepsilon m_\varepsilon' \end{aligned} \tag{2.44}$$

where M, m_v, m_ε are functions of the kalman gains matrix K, and K is itself a function of Σ and m_ε. The variance-covariance matrix Σ, the kalman gains matrix K, and the law of motion matrices M, m_v, m_ε are thus obtained by solving the large non-linear system of equations described by (2.32)-(2.35), (2.43), and (2.44). This system is too complicated to allow further analytical results; we thus solve for the fixed point numerically.

Proof of Proposition 34. The planner's problem is strictly convex, guaranteeing that its solution is unique and is pinned down by its first-order conditions. The Lagrangian of this problem can be written as

$$\begin{aligned}\Lambda \;=\;& \int_{S_\Omega}\left[U(Q(\Omega_t,\Omega_{t-1})) - \int_{S_\omega}\tfrac{1}{1+\epsilon}S(\omega)e^{-\frac{1+\epsilon}{\theta}a}q(\omega,\Omega_{t-1})^{\frac{1+\epsilon}{\theta}}d\Omega_t(\omega)\right]d\mathcal{F}(\Omega_t|\Omega_{t-1}) \\ & + \int_{S_\Omega}\lambda(\Omega_t)\left[Q(\Omega_t,\Omega_{t-1})^{\frac{\rho-1}{\rho}} - \int_{S_\omega}q(\omega,\Omega_{t-1})^{\frac{\rho-1}{\rho}}d\Omega_t(\omega)\right]d\mathcal{F}(\Omega_t|\Omega_{t-1})\end{aligned}$$

The first-order conditions with respect to $Q(\Omega)$ and $q(\omega)$ are given by the following:

$$U'\left(Q(\Omega_t,\Omega_{t-1})\right) + \lambda(\Omega_t)\left(\frac{\rho-1}{\rho}\right)Q(\Omega_t,\Omega_{t-1})^{-\frac{1}{\rho}} = \tag{2.45}$$

$$\int_{S_\Omega}\left[-\frac{1}{\theta}S(\omega)e^{-\frac{1+\epsilon}{\theta}a}q(\omega,\Omega_{t-1})^{\frac{1+\epsilon}{\theta}-1} - \lambda(\Omega_t)\left(\frac{\rho-1}{\rho}\right)q\left(\omega,\Omega_{t-1}\right)^{-\frac{1}{\rho}}\right]\mathcal{F}\left(\Omega_t|\omega,\Omega_{t-1}\right) = \tag{2.46}$$

where $\mathcal{F}\left(\Omega_t|\omega,\Omega_{t-1}\right)$ denotes the posterior about Ω_t (or, equivalently, about $\bar{f}_t$ and y_t) given ω_t. Restating condition (2.45) as $\lambda(\Omega_t)\left(\frac{\rho-1}{\rho}\right) = -U'\left(Q(\Omega_t,\Omega_{t-1})\right)Q(\Omega_t,\Omega_{t-1})^{\frac{1}{\rho}}$ and substituting this into condition (2.46), gives condition (2.25), which concludes the proof.

Chapter 3

Incomplete Information, Higher-Order Beliefs and Price Inertia

3.1 Introduction

How much, and how quickly, do prices respond to nominal shocks? This is one of the most fundamental questions in macroeconomics: it is key to understanding the sources and the propagation of the business cycle, as well as the power of monetary policy to control real economic activity.

To address this question, one strand of the literature has focused on menu costs and other frictions in adjusting prices; this includes both convenient time-dependent models and state-dependent adjustment models. Price rigidities are then identified as the key force behind price inertia. Another strand of the literature has focused on informational frictions; this strand highlights that firms may fail to adjust their price to nominal shocks, not because it is costly or impossible to do so, but rather because they have imperfect information about these shocks. The older literature formalized this imperfection as a geographical dispersion of the available information (Lucas, 1972, Barro, 1976); more recent contributions have proposed infrequent updating of information (Mankiw and Reis, 2002; Reis, 2006) or rational inattention (Sims, 2003; Woodford, 2003, 2008; Machowiak and Wiederholt, 2008). One way or another, though, the key driving force behind price inertia is that firms happen, or choose, to have imperfect knowledge of the underlying nominal shocks.

The starting point of this chapter is a bridge between these two approaches. In particular, our baseline model is a hybrid of Calvo (1983), Morris and Shin (2002), and Woodford (2003): on the one hand, firms can adjust prices only infrequently, as in Calvo; on the other hand, firms observe the underlying nominal shocks only with noise, similarly to Morris-Shin and Woodford.

Within this baseline model, the response of prices to nominal shocks—and hence also the real impact of these shocks—is characterized by the interaction of three key parameters: the precision of available information about the underlying nominal shocks (equivalently, the level of noise in the firms' signals of these shocks); the frequency of price adjustment; and the degree of strategic complementarity in pricing decisions. That all three parameters should matter is obvious; but their interaction is also interesting. The combination of sticky prices and strategic complementarity implies that the incompleteness of information can have lasting effects on inflation and real output even if the shocks become commonly known very quickly. This is because firms that have full information about the shock at the time they set prices will find it optimal to adjust only partly to the extent that other firms had only incomplete information at the time they had set their prices. Moreover, incomplete information can help make inflation peak after real output, which seems consistent with available evidence based on structural VARs.

These findings synthesize, and marginally extend, various lessons from the pertinent literature with regard to how frictions in either price adjustment or information about the underlying nominal shocks impact the response of prices to these shocks. This synthesis has its own value, as it provides a simple and tractable incomplete-information version of the Calvo model that could readily be taken to the data. Nevertheless, this synthesis is not the main contribution of the paper. Rather, the main contribution of the paper is, first, to highlight that the aforementioned lessons miss the distinct role that higher-order beliefs play in the dynamics of price adjustment and, second, to show how this distinct role can be parsimoniously accommodated within our Calvo-like framework.

The basic idea behind our contribution is simple. The precision of the firms' information about the underlying nominal shock identifies how fast the firms' forecasts of the shock respond to the true shock: the more precise their information, the faster their forecasts converge to the truth. However, this need not also identify how fast the forecasts of the forecasts of others may adjust. In other words, the precision of available information pins down the response of first-order beliefs, but not necessarily the response of higher-order beliefs. But when prices are strategic complements, the response of the price level to the underlying shock depends heavily on the response of higher-order beliefs. It follows that neither the precision of available information nor the degree of price rigidity suffice for calibrating the degree of price inertia at the macro level.

To better understand this point, it is useful to abstract for a moment from sticky prices. Assume, in particular, that all firms can adjust their prices in any given period and that the prices they set are given by the following simple pricing rule:

$$p_i = (1-\alpha)\mathbb{E}_i\theta + \alpha\mathbb{E}_i p$$

where p_i is the price set by firm i, θ is nominal demand, p is the aggregate price level, $\alpha \in (0,1)$ is the degree of strategic complementarity in pricing decisions, and $\mathbb{E}_i$ denotes the expectation conditional on the information of firm i. Aggregating this condition and iterating over the expectations of the price level imply that the aggregate price level must satisfy the following condition:

$$p = (1-\alpha)\left(\bar{E}^1 + \alpha\bar{E}^2 + \alpha^2\bar{E}^3 + \ldots\right),$$

where $\bar{E}^k$ denotes the k^{th}-order average forecast of θ.[1] It then follows that the response of price level p to an innovation in θ depends on the response of the entire sequence of different orders of beliefs, $\{\bar{E}^k\}_{k=1}^{\infty}$, to that shock.

When information is perfect or at least commonly shared, then $\bar{E}^k = \bar{E}^1$ for all k. It then follows that the response of prices to an innovation in θ depends merely on the response of first-order beliefs, which in turn is pinned down by the precision of the available information about θ. When, instead, information is dispersed, higher-order beliefs need not coincide with

[1] The k^{th}-order average forecasts are defined recursively as follows: $\bar{E}^1$ is the cross-sectional mean of the firms' forecasts of the underlying nominal shock; $\bar{E}^2$ is the cross-sectional mean of the firms' forecasts of $\bar{E}^1$; and so on.

first-order beliefs.

The potential role of higher-order beliefs has been noted before by Morris and Shin (2002), Woodford (2003), and others. However, in the standard Gaussian example used in the pertinent literature, the sensitivity of higher-order beliefs to the shock is tightly connected to that of first-order beliefs: higher-order beliefs can be less sensitive to the underlying shock only if the precision of information about the shock is lower, in which case first-order beliefs are also less sensitive. It follows that in the standard Gaussian example the precision of information about the underlying nominal shock remains the key determinant of the response of the price level to the shock. However, once one goes away from the standard Gaussian example, this tight connection between first- and higher-order beliefs can break—and the break can be quite significant.

We highlight the crucial and distinct role of higher-order beliefs with three variants of our baseline model. All three variants retain the combination of infrequent price adjustment (as in Calvo) and noisy information about the underlying shocks (as in Morris-Shin and Woodford), but differentiate in the specification of higher-order beliefs.

In the first variant, firms face uncertainty, not only about the size of the aggregate nominal shock, but also about the precision of the signals that *other* firms receive about this shock. This extension helps isolate the role of higher-order beliefs or, equivalently, the role of strategic uncertainty: it shows how this additional source of uncertainty about the distribution of precisions can impact the response of higher-order beliefs to the underlying shocks, and thereby the response of prices, without necessarily affecting the response of first-order beliefs.

In the second variant, firms hold heterogeneous priors about the stochastic properties of the signals that other firms receive. In this economy, firms expect the beliefs of others to adjust more slowly to the underlying shocks than their own beliefs. They thus behave in equilibrium as if they lived in a economy where all other firms had less precise information than what they themselves have. This in turn helps rationalize why equilibrium prices may adjust very slowly to the underlying nominal shocks even if the frequency of price adjustment is arbitrarily high and each firm has arbitrarily precise information about the underlying nominal shock. Once again, the key is the inertia of higher-order beliefs; heterogeneous priors is a convenient modeling device.

The aforementioned two variants focus on how higher-order beliefs impact the propagation of nominal shocks in the economy. The third and final variant shows how higher-order beliefs can be the source of fluctuations in the economy—how they can themselves be one of the "structural" shocks. In particular, variation in higher-order beliefs that is orthogonal to either the underlying nominal shocks or the firms information about these shocks can generate fluctuations in inflation and real output that resemble those generated by "cost-push" shocks.

Combined, these findings point out that a macroeconomist who wishes to quantify the response of the economy to its underlying structural shocks, or even to identify what are these structural shocks in the first place, may need appropriate information, not only about the degree of price rigidity and the firms' information (beliefs) about these shocks, but also about the stochastic properties of their forecasts of the forecasts of others (higher-order beliefs).

Because the hierarchy of beliefs is an infinitely dimensional object, incorporating the distinct role of higher-order beliefs in macroeconomic models may appear to be a challenging task. Part of the contribution of the paper is to show that this is not the case. All the models we present here are highly parsimonious and nevertheless allow for rich dynamics in higher-order beliefs.

The rest of the paper is organized as follows. Section 3.2 discusses the relation of our paper to the literature. Section 3.3 studies our baseline model, which introduces incomplete information in the Calvo model. Section 3.4 studies the variant with uncertainty about the precisions of one another. Section 3.5 studies the variant that with heterogeneous priors. Section 3.6 turns to cost-push shocks. Section 3.7 concludes with suggestions for future research. The details of the proofs of all the formal results can be found in the appendix.

3.2 Related literature

The macroeconomics literature on informational frictions has a long history, going back to Phelps (1970), Lucas (1972, 1975), Barro (1976), King (1983) and Townsend (1983). Recently, this literature has been revived by Mankiw and Reis (2002), Morris and Shin (2002), Sims (2003), Woodford (2003), and subsequent work.[2] this chapter contributes to this literature

[2]See, e.g., Amato and Shin (2006), Angeletos and La'O (2008, 2009a, 2009b), Angeletos and Pavan (2007), Bacchetta and Wincoop (2005), Collard and Dellas (2005), Hellwig (2005), Lorenzoni (2008, 2009), Mackowiak and Wiederholt (2008, 2009), Nimark (2007, 2008), and Reis (2006).

in two ways: first, by studying the interaction of incomplete information with price rigidities within the Calvo model; second, and most importantly, by furthering our understanding of the distinct role of higher-order beliefs. Closely related in this respect are Angeletos and La'O (2009a, 2009b), which emphasize other dimensions in which dispersed information has very distinct implications for the business cycle than uncertainty about the fundamentals.

Our paper is highly complementary to the papers by Woodford (2003) and Morris and Shin (2002, 2006). These papers document how higher-order beliefs may respond less to information about the underlying shocks than first-order beliefs, simply because they are more anchored to the common prior.[3] However, by adopting the convenience of a popular but very specific Gaussian information structure, they have also restricted attention to settings where the response of higher-order beliefs is tightly tied to the response of first-order beliefs: in their settings, the response of higher-order beliefs is a monotone transformation of the response of first-order beliefs, thus precluding any *independent* role for higher-order beliefs.

Our paper, instead, highlights that, whereas the response of first-order beliefs to the underlying nominals shocks is pinned down solely by the level of uncertainty about these shocks, the response of higher-order beliefs depends also on other sources of uncertainty, such as uncertainty about the precision of others' information. Furthermore, it shows how with heterogeneous priors it is possible that higher-order beliefs respond very little, or even not at all, to the underlying shocks even if all firms are nearly perfectly informed about these shocks (in which case first-order beliefs respond nearly one-to-one with the shock).

Closely related are also the papers by Nimark (2008) and Klenow and Willis (2007). The former paper studies a similar framework as ours, namely a Calvo model with incomplete information, along with a more complex learning dynamics: unlike our paper and rather as in Woodford (2003), the underlying shocks do not become common knowledge after one-period delay. The latter paper studies a menu-cost model with sticky information. Much alike our baseline model, both papers study the interaction of price rigidities and informational frictions. However, these papers do disentangle the role of higher-order beliefs as this chapter does. In particular, the quantitative importance of incomplete information in these papers is tied to the

[3]A similar role of higher-order beliefs has been highlighted by Allen, Morris and Shin (2005), Angeletos, Lorenzoni and Pavan (2007) and Bacchetta and Wincoop (2005) within the context of financial markets.

precision of information about the underlying shocks at the time of price changes. In contrast, our paper shows how one can disentangle the dynamics of higher-oder beliefs from the speed of learning, and uses this to argue that significant price inertia at the macro level can be consistent with both significant price flexibility at the micro level and fast learning about the underlying nominal shocks.

Finally, the heterogeneous-priors variant of this chapter builds on Angeletos and La'O (2009b). That paper considers a real-business-cycle model in which firms have dispersed information about the underlying productivity shocks. It is then shown that dispersed information opens the door to a certain type of sunspot-like fluctuations—i.e., fluctuations that cannot be explained by variation in *either* the underlying economic fundamentals or the firms' beliefs about these fundamentals. These fluctuations obtain also under a common prior, but are easier to model with heterogeneous priors. The present paper complements this other work by illustrating how these type of fluctuations can take the form of cost-push shocks in a new-keynesian model, and how heterogeneous priors can also help rationalize significant inertia in the response of prices to nominal shocks.

3.3 The Calvo Model with Incomplete Information

This section considers a variant of the Calvo model that allows firms to have dispersed private information about aggregate nominal demand.

Households and firms. The economy is populated by a representative household and a continuum of firms that produce differentiated commodities. Firms are indexed by $i \in [0,1]$. Time is discrete, indexed by $t \in \{0,1,2,...\}$. There is no capital, so that there is no saving in equilibrium. Along with the fact that there is a representative household, this implies that it is without serious loss of generality to abstract from asset trading, provided that one abstracts from the role of such markets in aggregating information.[4]

The preferences of the household are given by $\sum_t \beta^t U(C_t, N_t)$, with

$$U(C_t, N_t) = \log C_t - N_t,$$

[4]For endogenous aggregation of information within a business-cycle context, see Angeletos and La'O (2008).

where $\beta \in (0,1)$ is the discount rate, N_t is the labor supplied by the household,

$$C_t = \left[\int C_{i,t}^{\frac{\eta-1}{\rho}} di \right]^{\frac{\eta}{\eta-1}}$$

is the familiar CES aggregator, $C_{i,t}$ is the consumption of the commodity produced by firm i, and $\eta > 0$ is the elasticity of substitution across commodities. As usual, this specification implies that the demand for the commodity of firm i is given by

$$C_{i,t} = \left(\frac{P_{i,t}}{P_t} \right)^{-\eta} C_t,$$

where $P_t \equiv \left[\int P_{i,t}^{\eta-1} di \right]^{\frac{1}{\eta-1}}$ is the aggregate price index.

The output of firm i, on the other hand, is given by

$$Y_{i,t} = A_{i,t} L_{i,t}^{\epsilon},$$

where $A_{i,t}$ is the idiosyncratic productivity shock and $\epsilon \in (0,1)$ parameterizes the degree of diminishing returns.

By the resource constraint for each commodity, $Y_{i,t} = C_{i,t}$ for all i and therefore aggregate output is given by $Y_t = C_t$. Finally, aggregate nominal demand is given by the following quantity-theory or cash-in-advance constraint:

$$P_t C_t = \Theta_t.$$

Here, Θ_t denotes the level of aggregate nominal demand (aggregate nominal GDP), is assumed to be exogenous, and defines the "monetary shock" of our model.

In what follows, we use lower-case variables for the logarithms of the corresponding upper-case variables: $\theta_t \equiv \log \Theta_t$, $y_t \equiv \log Y_t$, $p_t \equiv \log P_t$, and so on. We also assume that all exogenous shocks are log-normally distributed, which guarantees that the equilibrium admits an exact log-linear solution.

Shocks and information. Aggregate nominal demand is assumed to follow a random walk:

$$\theta_t = \theta_{t-1} + v_t$$

where $v_t \sim \mathcal{N}(0, \sigma_\theta^2)$ is white noise. Each period has two stages, a morning and an evening. Let $\mathcal{I}_{i,t}^1$ and $\mathcal{I}_{i,t}^2$ denote the information set of firm i during, respectively, the morning and the evening of period t. Information about the current level of nominal demand is imperfect during the morning but perfect during the evening. The information that a firm has about θ_t during the morning is summarized in a Gaussian private signal of the following form:

$$x_{i,t} = \theta_t + \varepsilon_{i,t},$$

where $\varepsilon_{i,t} \sim \mathcal{N}(0, \sigma_x^2)$ is purely idiosyncratic noise (i.i.d. across firms). Pricing choices (for the firms that have the option to set prices) are made in the morning, while information about θ_t is imperfect; employment and consumption choices are made in the evening, once θ_t has been publicly revealed. Finally, the idiosyncratic productivity shock $a_{i,t}$ follows a random walk and it is known to the firm from the beginning of the period.

The information of firm i in the morning of period t is therefore given by $\mathcal{I}_{i,t}^1 = \mathcal{I}_{i,t-1}^2 \cup \{x_{i,t}, a_{i,t}\}$, while her information in the evening of the same period is given by $\mathcal{I}_{i,t}^2 = \mathcal{I}_{i,t-1}^1 \cup \{\theta_t\}$. We will see in a moment that, in equilibrium, p_{t-1} and y_{t-1} are functions of $\{\theta_{t-1}, \theta_{t-2}, \ldots\}$; it would thus make no difference if past values of the price level and real GDP had been included in the information set of the firm. Similarly, $y_{i,t}$ is a function of $\mathcal{I}_{i,t}$; it would thus make no difference if the realized level of a firm's demand had been included into its evening information set.

The assumption that θ_t becomes common knowledge at the end of each period is neither random nor inconsequential. If one wishes information about θ_t to remain dispersed after the end of period t, one needs somehow to limit the aggregation of information that takes place through commodity markets. That requires a more decentralized trading structure and complicates the necessary micro-foundations, which seems an undue cost given that the contribution of the paper is not quantitative. Furthermore, the dynamics become much less tractable: as in Townsend (1983), Nimark (2008), and others, firms must now keep tract of the entire history of their information in order to forecast the forecasts of others, and the equilibrium dynamics would cease to have any simple recursive structure. Here, we avoid all these complications, and keep the analysis highly tractable, only by assuming, as in Lucas (1972), that θ_t becomes common knowledge after a short delay. However, it is important to recognize that, in so doing, we impose a fast convergence of beliefs about the past shocks and also rule out any heterogeneity in the agents' expectations of future shocks beyond the one in their beliefs about the current shock. One may expect that relaxing these properties would add to even more inertia, both because firms would learn more slowly (Woodford, 2003) and because expectations of future shocks could be more anchored to public information (Morris and Shin, 2006).

Price-setting behavior. Consider for a moment the case where prices are flexible and the current nominal shock is common knowledge at the moment firms set prices. The optimal price set by firm i in period t is then given by

$$p_{i,t}^{*} = \mu + mc_{i,t}$$

where $\mu \equiv \frac{\eta}{\eta-1}$ is the monopolistic mark-up and $mc_{i,t}$ is the nominal marginal cost the firm faces in the evening of period t. The latter is given by

$$mc_{i,t} = w_t + \frac{1-\epsilon}{\epsilon} y_{i,t} - \frac{1}{\epsilon} a_{i,t}$$

where w_t is the nominal wage rate in period t. From the representative household's optimality condition for work,

$$w_t - p_t = c_t.$$

From the consumer's demand,

$$c_{i,t} - c_t = -\eta(p_{i,t} - p_t).$$

From market clearing, $c_{i,t} = y_{i,t}$ and $c_t = y_t$. Finally, from the cash-in-advance constraint, aggregate real output is given by

$$y_t = \theta_t - p_t.$$

Combining the aforementioned conditions, we conclude that the "target" price (i.e., the flexible-price full-information optimal price) of firm i in period t is given by

$$p^*_{i,t} = (1-\alpha)\theta_t + \alpha p_t + \xi_{i,t} \tag{3.1}$$

where

$$\alpha \equiv 1 - \frac{1}{\epsilon + (1-\epsilon)\eta} \in (0,1)$$

defines the degree of strategic complementarity in pricing decisions, and where

$$\xi_{i,t} \equiv \frac{\epsilon}{\epsilon + (1-\epsilon)\eta}\mu - \frac{1}{\epsilon + (1-\epsilon)\eta} a_{i,t}$$

is simply a linear transformation of the idiosyncratic productivity shock. We henceforth normalize the mean of the idiosyncratic productivity shock so that the cross-sectional mean of $\xi_{i,t}$ is zero.

If prices had been flexible and θ_t had been common knowledge in the begging of period t, the firm would set $p_{i,t} = p^*_{i,t}$ in all periods and states. However, we have assumed that firms have only imperfect information. Moreover, following Calvo (1983), we assume that a firm may change its price only with probability $1-\lambda$ during any given period, where $\lambda \in (0,1)$. It then follows that, in the event that a firm changes its price, the price it chooses is equal (up to a constant) to its current expectation of a weighted average of the current and all future target prices:

$$p_{i,t} = \mathbb{E}_{i,t}\left[(1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\, p^*_{t+j}\right] \tag{3.2}$$

where $\beta \in (0,1)$ is the discount factor, λ is the probability that the firm won't have the option to adjust its price (a measure of how sticky prices are), and $\mathbb{E}_{i,t}$ is the expectation conditional on the information set of firm i in period t.[5]

Combining conditions (3.1) and (3.2) implies that the price set by any firm that gets the

[5]To be precise, condition (3.2) should have been written as $p_{i,t} = const + \mathbb{E}_{i,t}\left[(1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\, p^*_{t+j}\right]$, where *const* is an endogenous quantity that involves second-order moments and that emerges due to risk aversion. However, under the log-normal structure of shocks and signals that has been imposed, these second-order moments are invariant with either the shock or the information of the firms, and it is thus without any loss of generality to ignore the aforementioned constant.

chance to adjust its price in period t is given by

$$p_{i,t} = (1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_{t+j} + \alpha\mathbb{E}_{i,t}p_{t+j} + \mathbb{E}_{i,t}\xi_{i,t+j}\right] \tag{3.3}$$

In the remainder of the paper, we treat condition (3.3) as if it were an exogenous behavioral rule, with the understanding though that this rule is actually fully micro-founded in equilibrium.

Equilibrium dynamics. The economy effectively reduces to a dynamic game of incomplete information, with condition (3.3) representing the best response of the typical firm. The equilibrium notion we adopt is standard Perfect-Bayesian Equilibrium.[6] Because of the linearity of the best-response condition (3.3) and the Gaussian specification of the information structure, it is a safe guess that the equilibrium strategy will have a linear form. It is thus safe to conjecture the existence of equilibria in which the price set by a firm in period t is a linear function of $(p_{t-1}, \theta_{t-1}, x_{i,t}, \xi_{i,t})$:

$$p_{i,t} = P(p_{t-1}, \theta_{t-1}, x_{i,t}, \xi_{i,t}) = b_1 p_{t-1} + b_2 x_{i,t} + b_3\theta_{t-1} + \xi_{i,t} \tag{3.4}$$

for some coefficients b_1, b_2, b_3. This particular guess is justified by the following reasoning: p_{t-1} is bound to matter because of a fraction of firms cannot adjust prices; $x_{i,t}$ because it conveys information about the current nominal shock θ_t; θ_{t-1} because it is the prior about θ_t; and $\xi_{i,t}$ for obvious reasons.

Given this guess, and given the fact that only a randomly selected fraction $1-\lambda$ of firms can adjust prices in any given period, we infer that the aggregate price level must satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\int\int P(p_{t-1}, \theta_{t-1}; x, \xi)dF_t(x)dG_t(\xi)$$

where F_t denotes the cross-sectional distribution of the private signals (conditional on the current shock θ_t) and G_t denotes the cross-sectional distribution of the idiosyncratic shocks. Given that P is linear, that the cross-sectional average of x_t is θ_t, and that cross-sectional average

[6] We can safely ignore out-of-equilibrium paths by assuming that firms observe (at most) the cross-section distribution of prices, which guarantees that no individual deviation is detectable.

of ξ_t is 0, the above condition can be re-written as $p_t = \lambda p_{t-1} + (1-\lambda) P(p_{t-1}, \theta_{t-1}; \theta_t, 0)$, or equivalently as

$$p_t = c_1 p_{t-1} + c_2 \theta_t + c_3 \theta_{t-1} \tag{3.5}$$

where

$$c_1 = \lambda + (1-\lambda) b_1, \quad c_2 = (1-\lambda) b_2 \quad c_3 = (1-\lambda) b_3. \tag{3.6}$$

Next, note that condition (3.3) can be stated in recursive form as

$$p_{i,t} = (1-\beta\lambda) \left[(1-\alpha) \mathbb{E}_{i,t}\theta_t + \alpha \mathbb{E}_{i,t} p_t + \xi_{i,t} \right] + (\beta\lambda) \mathbb{E}_{i,t} p_{i,t+1}.$$

Using (3.4) and (3.5) into the right-hand side of the above condition, we infer that the price must satisfy

$$p_{i,t} = (1-\beta\lambda) \left[(1-\alpha) \mathbb{E}_{i,t}\theta_t + \alpha \mathbb{E}_{i,t} p_t + \xi_{i,t} \right] + (\beta\lambda) \left[b_1 \mathbb{E}_{i,t} p_t + b_2 \mathbb{E}_{i,t}\theta_t + b_3 \mathbb{E}_{i,t}\theta_{t+1} + \mathbb{E}_{i,t}\xi_{i,t+1} \right]. \tag{3.7}$$

Next, note that

$$\mathbb{E}_{i,t}\theta_{t+1} = \mathbb{E}_{i,t}\theta_t = \frac{\kappa_x}{\kappa_x + \kappa_\theta} x_{it} + \frac{\kappa_x}{\kappa_x + \kappa_\theta} \theta_{t-1} \qquad \text{and} \qquad \mathbb{E}_{i,t}\xi_{i,t+1} = \xi_{i,t},$$

where $\kappa_x \equiv \sigma_x^{-2}$ is the precision of the firms' signals and $\kappa_\theta \equiv \sigma_\theta^{-2}$ is the precision of the common prior about the innovation in θ. Using these facts, and substituting p_t from (3.5) into (3.7), the left-hand side of (3.7) can be rewritten as a linear function of p_{t-1}, $x_{i,t}$, θ_{t-1}, and $\xi_{i,t}$. For this to coincide with our conjecture in (3.4), it is necessary and sufficient that the coefficients (b_1, b_2, b_3) solve the following system:

$$\begin{aligned}
b_1 &= (1-\beta\lambda)\alpha c_1 + (\beta\lambda) b_1 c_1 \\
b_2 &= \left[(1-\beta\lambda)(1-\alpha+\alpha c_2) + (\beta\lambda)(b_1 c_2 + b_2 + b_3) \right] \frac{\kappa_x}{\kappa_x + \kappa_\theta} \\
b_3 &= \left[(1-\beta\lambda)(1-\alpha+\alpha c_2) + (\beta\lambda)(b_1 c_2 + b_2 + b_3) \right] \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \\
&\quad + (1-\beta\lambda)\alpha c_3 + (\beta\lambda) b_1 c_3
\end{aligned} \tag{3.8}$$

We conclude that an equilibrium is pinned down by the joint solution of (3.6) and (3.8).

Combining the conditions for c_1 and b_1 implies that c_1 must solve the following equation:

$$c_1 = \lambda + (1 - \lambda) \left(\frac{1 - \beta\lambda}{1 - \beta\lambda c_1} \right) \alpha c_1.$$

This equation admits two solutions: one with $c_1 > 1$ and another with $c_1 \in (\lambda, 1)$. We ignore the former solution because it leads to explosive price paths and henceforth limit attention to the latter solution. Note that this solution is independent of the information structure; indeed, the coefficient c_1, which identifies the endogenous persistence in the price level, coincides with the one in the standard (complete-information) Calvo model.

Given this solution for c_1, the remaining conditions define a linear system that admits a unique solution for the remainder of the coefficients. It is straightforward to check that the solution satisfies

$$c_1 + c_2 + c_3 = 1,$$

which simply means that the price process is homogenous of degree one in the level of nominal demand. Furthermore,

$$c_2 = \frac{\lambda(1 - c_1)}{\lambda + c_1 \frac{\kappa_\theta}{\kappa_x}}$$

which identifies the sensitivity of the price level to the current innovation in nominal demand as an increasing function of the precision of available information. We therefore reach the following characterization of the equilibrium.

Proposition 37 *(i) There exists an equilibrium in which the pricing strategy of a firm is given by*

$$p_{i,t} = b_1 p_{t-1} + b_2 x_{i,t} + b_3 \theta_{t-1}$$

and the aggregate price level is given by

$$p_t = c_1 p_{t-1} + c_2 \theta_t + c_3 \theta_{t-1}$$

for some positive coefficients (b_1, b_2, b_3) *and* (c_1, c_2, c_3).

(ii) The equilibrium values of the coefficients (c_1, c_2, c_3) *satisfy the following properties:* c_1 *is increasing in* λ*, increasing in* α*, and invariant to* κ_x/κ_θ*;* c_2 *is non-monotone in* λ*, decreasing*

in α, and increasing in κ_x/κ_θ; c_3 is non-monotone in α, and decreasing in κ_x/κ_θ.

The comparative statics described in part (ii) are illustrated in Figures 1, 2 and 3. The baseline parameterization used for these figures, as well as for the impulse responses reported later on, is as follows. We identify the length of a period with one year; this seems a good benchmark for how long it takes for macro data to become widely available.[7] We accordingly set $\beta = .95$ (which means a discount rate of about 1% per quarter), $\lambda = .20$ (which means a probability of price change equal to 1/3 per quarter), and $\alpha = .85$ (which means a quite strong complementarity in pricing decisions); these values are consistent with standard calibrations of the Calvo model. Lacking any obvious estimate of the precision of information about the underlying shocks, we set $\kappa_x/\kappa_\theta = 1$; this means that the variance of the forecast error of the typical firm about the current innovation in nominal demand is one half the variance of the innovation itself.[8]

Figure 1 plots the coefficients c_1, c_2, and c_3 as functions of the Calvo parameter λ, the probability the firm does *not* revise its price in a given period. Note that c_1 is increasing in λ, c_3 is decreasing in λ, and c_2 is non-monotonic in λ (it increases for low values but decreases for high values). Figure 2 plots these coefficients as functions of α, the degree of strategic complementarity in pricing decisions. Note that c_1 is an increasing function in α, c_2 is a decreasing function in α, and c_3 is non-monotonic in α. Finally, Figure 3 plots these coefficients as functions of κ_x/κ_θ, the ratio of the precision of private signals to the precision of the prior. Clearly, c_2 is increasing in this ratio, while c_3 is decreasing and c_1 is invariant.

[7] A qualification is due here. The fact that these data are widely available suggests that most agents are likely to be well informed about them. This in turn implies that their first-order beliefs are likely to converge to the truth very fast. However, to the extent that this fact is not common knowledge, it is possible that higher-order beliefs do not converge as fast, which could contribute to further inertia in the response of prices.

[8] Note that only the ratio κ_x/κ_θ, and not the absolute values of κ_x and κ_θ, matter for the equilibrium coefficients c_1, c_2, and c_3. In other words, once we fix κ_x/κ_θ, κ_θ is only a scaling parameter. By implication, the impulse responses of the economy to a one-standard-deviation change in v are invariant to κ_θ.

Figure1

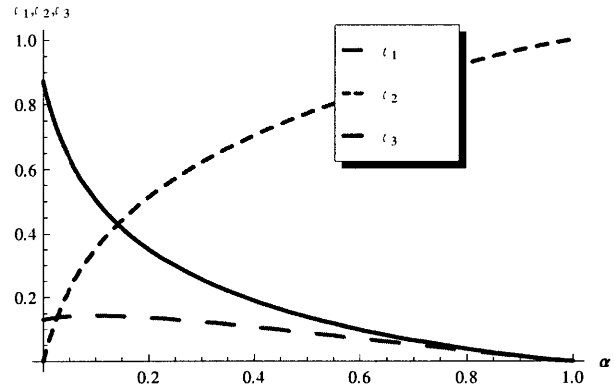

Figure2

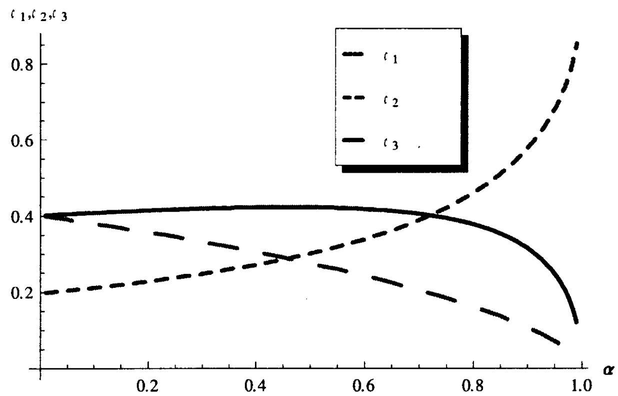

Figure3

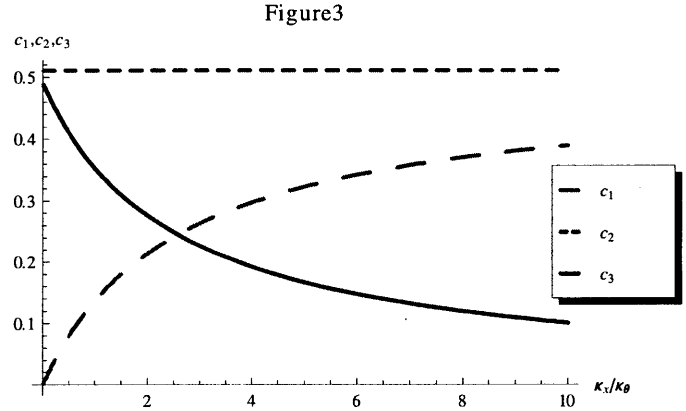

The comparative statics described above are a hybrid of the results found in sticky-price Calvo models and in the incomplete-information literature. As in the standard Calvo model, the aggregate price level is persistent due to the fact that some firms cannot adjust prices. In our model, the coefficient which characterizes the persistence of the aggregate price process is c_1. We find that this coefficient is unaffected by the incompleteness of information. In this sense, the persistence of prices in our baseline model is the same as in the standard Calvo model. This property, however, hinges on our assumption that the nominal shock becomes common knowledge only with a delay of one period. If we increase the length of this delay, then we can obtain more persistence, similarly to Woodford (2003) or Nimark (2008), but only till the shock becomes common knowledge; after that point, any subsequence persistence is driven solely by the Calvo rigidity. The property, then, that c_1 is increasing in the Calvo parameter λ should be familiar: it is almost the mechanical implication of the fact that a fraction λ of firms do not adjust prices. The impact of α on c_1 is also familiar: even under complete information, firms who can adjust their price following a monetary shock will find it optimal to stay closer to the past price level the higher the degree of strategic complementarity between them and the firms that cannot adjust (and that are thus stuck to the past price level).

Where the incompleteness of information has a bite in our baseline model is on the coefficients c_2 and c_3, which characterize, respectively, the price impact of the current and the past shock for any given past price level. To understand how the precision of information affects these coefficients, consider the choice of the price-setting firm. The price chosen by a firm is a linear combination of past prices and past nominal shocks (which are common knowledge among all firms) and the firm's own expectation of current nominal demand (which is unknown in the current period). Aggregating across firms gives the aggregate price level as a linear combination of past price levels and past nominal shocks and of the *average* expectation of θ_t. As in any static incomplete information model with Gaussian signals, the firm's own expectation of the fundamental is merely a weighted combination of his private signal and the common prior, which here coincides with θ_{t-1}. If firms have less precise private information relative to the prior, i.e., lower κ_x/κ_θ, they place less weight on their private signals than on their prior when forming their expectations of θ_t. As a result, the average expectation is less sensitive to the current shock θ_{t-1} and more anchored to the past shock θ_{t-1}. This explains why less precise information (a lower κ_x/κ_θ) implies a lower c_2 and a higher c_3.

Impulse responses. The above analysis highlights how introducing incompleteness of information into the Calvo model dampens the response of prices to the underlying nominal shocks—the precision of information becomes a key parameter for the dynamics of inflation along with the Calvo parameter and the degree of strategic complementarity. To further appreciate this, we now study how the precision of information affects the impulse responses of the inflation rate and real output to an innovation in nominal demand.

Figures 4 and 5 plot these impulse responses. (Inflation in period t is given by $p_t - p_{t-1}$, while real output is $y_t = \theta_t - p_t$.) As before, we identify the period with a year and set $\beta = .95$, $\lambda = .20$ and $\alpha = .85$. We then consider three alternative values for the precision of information: $\kappa_x/\kappa_\theta = 1$, which is our baseline; $\kappa_x/\kappa_\theta = \infty$, which corresponds to the extreme of perfect information (as in the standard Calvo model); and $\kappa_x/\kappa_\theta = 0$, which corresponds to the alternative extreme, that of no information about the current shock other than the prior (i.e., the past shock).

Figure 4

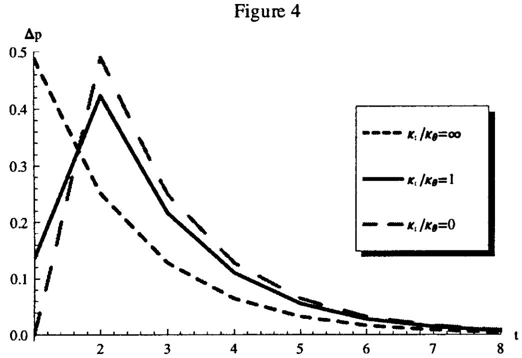

Figure 5

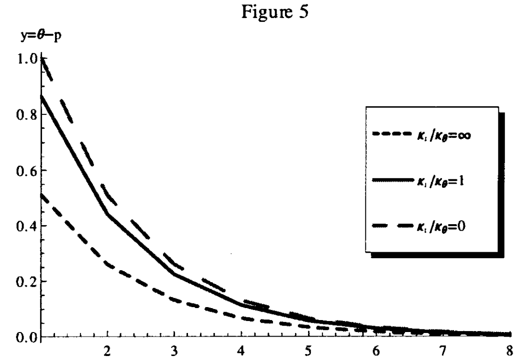

Figure 4 illustrates how the incompleteness of information affects the dynamics of inflation, relative to the complete-information Calvo model. First, the instantaneous impact effect of a monetary shock on inflation is increasing in κ_x/κ_θ. As the noise in private information increases, prices react less initially to a nominal disturbance. Second, as the precision of private information decreases, the second-period inflation becomes higher and higher. As the past nominal demand now becomes common knowledge, prices with low sensitivity to the monetary shock last period greatly increase in the second period to reflect this new information. Except for sufficiently high values of κ_x/κ_θ, this is where inflation reaches its peak. Last, although the decay rate of inflation is constant after this date, because of the high inflation experienced in the second period, lower κ_x/κ_θ leads to a higher level of inflation for all subsequent periods.

From Figure 5 one then sees that the impact effect of a monetary shock on output is decreasing in the precision of private information. Of course, this is simply the mirror image of what happens to prices. Furthermore, like the impulse responses for inflation, real output is higher for lower levels of κ_x/κ_θ for all subsequent periods.

It is interesting here to note that incomplete information has lasting effects on the levels of inflation and real GDP even though the shocks become common knowledge after just one period. This is precisely because of the interaction of incomplete information with price staggering and with strategic complementarity: by the time the shock becomes common knowledge, some firms have already set their price on the basis of incomplete information about the shock; strategic complementarity then guarantees that the firms that now have access to full information will still find it optimal to respond to the shock as if themselves had incomplete information.

Note that, except for high values of κ_x/κ_θ, the peak of output occurs before the peak in inflation. This is in contrast to the standard Calvo model which predicts strong price increases during the period in which the shock is realized and therefore typically has inflation peaking before output. A similar observation has been made in Woodford (2003), but with two important differences: Woodford (2003) abstracts from price staggering; it also assumes that past shocks and past outcomes never become known, thus appearing to require an implausibly slow degree of learning about the underlying shocks. Here, we show how the empirically appealing property that inflation peaks after real output can be obtained even with quite fast learning, provided one interacts incomplete information with price staggering.

In the present model, the peak of inflation happens at most one period after the innovation in θ. This particular property is an artifact of the assumption that the shock becomes common knowledge exactly one period after the innovation takes place. If we extend the model so that the shock becomes common knowledge after, say, 4 periods, then the peak in inflation can occur as late as 4 periods after the shock. The more general insight is that inflation can start low if firms initially have little information about the innovation and can rise in the early phases of learning, but once firms have accumulated enough information about the shock then inflation will begin to fall. In other words, the dynamics of learning are essential for the dynamics of inflation only as long as the firms remain sufficiently uncertain about the shock; but once the firms have learned enough about the shock, the subsequent dynamics of inflation are determined primarily by the Calvo mechanics.

3.4 Uncertainty about precisions

The analysis so far has focused on a Gaussian specification for the information structure that is quite standard in the pertinent literature. Under this specification, the response of prices to nominal shocks was determined by three parameters: (i) the degree of price rigidity; (ii) the degree of strategic complementarity; and (iii) the precision of information about the underlying nominal shock. In this section we show that, under a plausible variation of the information structure, knowledge of these parameters need not suffice for calibrating the degree of price inertia. The key insight is that the precision of information about the underlying shock pins down the response of the firms' forecasts of this shock, but not necessarily the response of their forecasts of the forecasts of other firms (i.e., their higher-order beliefs); and what matters for the response of equilibrium prices to the shock is not only the former but also the latter.

Apart from serving as an example for this more general insight, the variant considered here has its own appeal in that it introduces a plausible source of uncertainty: it allows firms to face uncertainty regarding the precision of information that other firms may have regarding nominal demand. In particular, we introduce a second aggregate state variable, which permits us to capture uncertainty about the average precision of available information in the cross-section of the economy.

This new state variable is modeled as a binary random variable, $s_t \in \{h, l\}$, which is i.i.d. over time and independent of the nominal shock θ_t, and which takes each of the two possible values h and l with probability 1/2. Let $\gamma, \kappa_h, \kappa_l$ be scalars, commonly known to all firms, with $1/2 < \gamma < 1$ and $0 < \kappa_l < \kappa_h$. The "type" of a firm is now given by the pair $(x_{i,t}, \kappa_{i,t})$, where $x_{i,t} = \theta_t + \varepsilon_{i,t}$ is the particular signal the firm receives about the current nominal shock, $\varepsilon_{i,t} \sim \mathcal{N}(0, 1/\kappa_{i,t})$ is the noise in this signal, and $\kappa_{i,t}$ is its precision. The latter is specific to the firm and is contingent on the new state variable s_t as follows: when this state is $s_t = h$, the precision of firm i is $\kappa_{i,t} = \kappa_h$ with probability γ and $\kappa_{i,t} = \kappa_l$ with probability $1 - \gamma$; and, symmetrically, when the state is $s_t = l$, the precision of firm i is $\kappa_{i,t} = \kappa_l$ with probability γ and $\kappa_{i,t} = \kappa_h$ with probability $1 - \gamma$. Finally, because these realizations are independent across the firms, γ is also the fraction of the population whose signals have precision κ_s when the state is s, for $s \in \{h, l\}$.

Note that a firm knows his own $\kappa_{i,t}$, but not the underlying state s_t. A firm's $\kappa_{i,t}$ thus serves a double role: it is both the precision of the firm's own information about the nominal shock θ_t and a noisy signal of the average precision in the cross-section of the economy. Therefore, the key difference from the baseline model is the property that firms face an additional source of informational heterogeneity: they face uncertainty regarding how informed other firms might be about the nominal demand shock. Note then that the coefficient γ parameterizes the level of this heterogeneity: when $\gamma = 1$, all firms have the same precisions, and this fact is common knowledge; when instead $\gamma \in (1/2, 1)$, different firms have different precisions, and each firm is uncertain about the distribution of precisions in the rest of the population.

Furthermore, note that knowing the state variable s_t would not help any firm improve his forecast of the nominal shock θ_t. This is simply because belief of a firm about the nominal shock depends only on its own precision (which the firm knows), not on the precisions of other firms (which the firm does not know). Nevertheless, the firm would love to know s_t because this could help him improve his forecast of the forecasts and actions of other firms in equilibrium. Indeed, since an firm's own expectation of θ_t depends on both his x_{it} and his κ_{it}, it is a safe guess that the equilibrium choice of the firm also depends on both x_{it} and κ_{it} and therefore that the aggregate price level depends both on θ_t and on κ_t. It then follows that firms face uncertainty about the aggregate price level, not only because they do not know the underlying innovation in θ_t, but also because they don't know how precisely other firms are informed about this shock. Finally, note that, while the uncertainty about θ_t matters for individual pricing behavior, and hence for aggregate prices, even when firms' pricing decisions are strategically independent ($\alpha = 0$), the uncertainty about s_t matters only when their pricing decisions are interdependent ($\alpha \neq 0$). This highlights the distinctive nature of the additional source of uncertainty that we have introduced in this section.

To better appreciate this point, it is useful to study the stochastic properties of the hierarchy of beliefs about θ. Let E_t^1 denote the cross-sectional average of $\mathbb{E}_{it}[\theta_t]$ conditional on the current state for the precisions being s_t. Next, for any $k \geq 2$, let E_t^k denote the cross-sectional average of $\mathbb{E}_{it}[E_t^{k-1}]$; that's the k^{th}-order average beliefs. Clearly, all these average beliefs are functions of the current and past nominal shocks and the current precision state s_t. Finally, let $\bar{E}_t^k$ denote the expectation of the k^{th}-order average belief conditional on the nominal shocks alone (that is, averaging across the two possible s_t states). It is easy to check that

$$\bar{E}_t^k = \eta_k \theta_t + (1 - \eta_k)\theta_{t-1},$$

for some constant η_k. The constant $\eta_k \equiv \partial \bar{E}_t^k / \partial \theta_t$ thus identifies the sensitivity of the k^{th}-order average belief to the underlying nominal shock. As anticipated in the Introduction, the response of the price level to the underlying nominal shock is determined by the sensitivities $\{\eta_k\}$.[9]

The behavior of the hierarchy of beliefs is illustrated in Figures 6 and 7. In Figure 6, we focus on the impact of the precision of information when this precision is common knowledge. We thus restrict $\kappa_h = \kappa_l = \kappa$ (in which case γ becomes irrelevant) and consider how the sensitivities of the beliefs to the shock vary with κ (which now identifies the common precision of information). The following qualitative properties are then worth emphasizing. First, the hierarchy of beliefs about θ_t converges to the common prior expectation, θ_{t-1}, as the signals become uninformative: for all k, $\eta_k \to 0$ as $\kappa \to 0$. Second, the beliefs converge to the true underlying state, θ_t, as the signals become perfect: for all k, $\eta_k \to 1$ as $\kappa \to \infty$. Finally, whenever the signals are informative but not perfect, higher-order beliefs are more anchored towards the prior than lower-order beliefs: for any $\kappa \in (0, \infty)$, $1 > \eta_1 > \eta_2 > ... > 0$.

[9] The discussion in the Introduction had abstracted from price rigidities (i.e., it had imposed $\lambda = 0$), but the insight is clearly more general.

Figure6

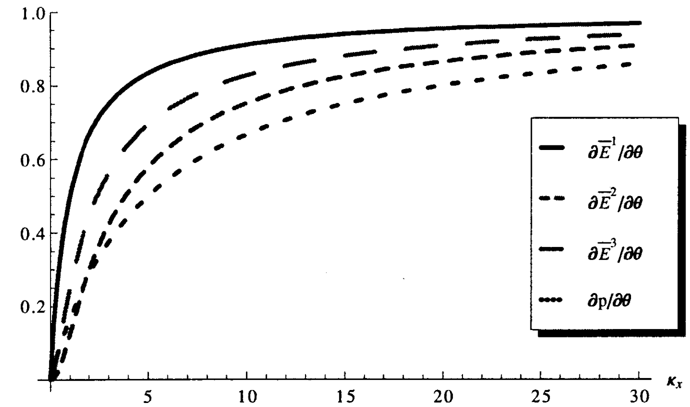

Figure 7

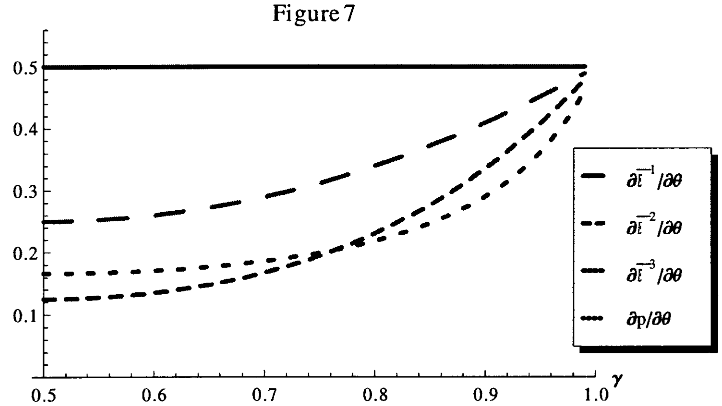

In Figure 7, we turn our focus to the impact of the uncertainty regarding the precision of others' information. In particular, we let $\kappa_h > \kappa_l$ and consider how the beliefs vary with the coefficient γ (which parameterizes the heterogeneity of information regarding the underlying precision state). One can then observe that η_1 is invariant to γ, while η_2 and η_3 increase with γ. That is, the sensitivity of first-order beliefs to the nominal shock is independent of γ, while the sensitivities of higher-order beliefs increase with γ.

Along with the fact that the price level depends not only on first-order but also on higher-order beliefs, this suggests that γ should affect the response of the price level to the nominal shock even though it does not affect the response of first-order beliefs. Indeed, following similar steps as in the baseline model, the equilibrium can be solved as follows.[10]

[10] As mentioned earlier, in both our baseline model and in all three variants of it, we identify the equilibrium

Proposition 38 *(i) There exists an equilibrium in which the pricing strategy of a firm is given by*

$$p_{i,t} = \begin{cases} b_1 p_{t-1} + b_{2,h} x_{i,t} + b_{3,h}\theta_{t-1} & \text{if } \kappa_{i,t} = \kappa_h \\ b_1 p_{t-1} + b_{2,h} x_{i,t} + b_{3,h}\theta_{t-1} & \text{if } \kappa_{i,t} = \kappa_l \end{cases}$$

while the aggregate price level is given by

$$p_t = \begin{cases} c_1 p_{t-1} + c_{2,h}\theta_t + c_{3,h}\theta_{t-1} & \text{if } s_t = h \\ c_1 p_{t-1} + c_{2,h}\theta_t + c_{3,h}\theta_{t-1} & \text{if } s_t = l \end{cases}$$

for some coefficients $(b_1, b_{2,h}, b_{2,l}, b_{3,h}, b_{3,l})$ *and* $(c_1, c_{2,h}, c_{2,l}, c_{3,h}, c_{3,l})$.

(ii) Let $c_2 \equiv \frac{1}{2}(c_{2,h} + c_{2,l})$ *and* $c_3 \equiv \frac{1}{2}(c_{3,h} + c_{3,l})$ *be the mean sensitivity of the price level to the current and past nominal shock, averaging across the precision states. The equilibrium value of* c_1 *does not depend on* γ *and is identical to that in the baseline model, while the equilibrium values of* c_2 *and* c_3 *depend on* γ *if and only if* $\alpha \neq 0$.

This result is also illustrated in Figure 7, which plots the coefficient $c_2 \equiv \mathbb{E}[\partial p_t / \partial \theta_t]$ as a function of γ. As γ decreases, the sensitivity of the first-order beliefs to the current nominal shock stays constant, while the sensitivity of the price level decreases. As anticipated, this is because a lower γ decreases the sensitivity of second- and higher-order beliefs.

To recap, the example of this section has highlighted how, even if one were to fix the sensitivity of the firms forecasts to the underlying nominal shock, one could still have significant freedom in how higher-order beliefs, and thereby equilibrium prices, respond to the shock. This is important for understanding the quantitative implications of incomplete information: to estimate the degree of price inertia caused by incomplete information, one may need direct or indirect information, not only about the firm's expectations about the underlying nominal shocks, but also about their higher-order expectations.

Finally, it is interesting to note how a variant of the model introduced here could generate

as the fixed point of the best-response condition (3.3). For the current variant, this is with some abuse, since the non-Gaussian nature of the information structure implies that this condition is not exact: it is only a log-linear approximation. However, the properties that first-order beliefs do not depend on γ, while higher-order beliefs and hence equilibrium outcomes do depend on γ, do not hinge on this approximation. Finally, keep in mind that condition (3.3) is exact in either the baseline model or the two other variants that we consider subsequently.

the possibility that all firms are perfectly informed about the nominal shock, are free to adjust their prices fully, and yet find it optimal to adjust only partly. To see this, suppose that when the precision state is $s = h$ all firms get $\kappa_i = \infty$ (which means that their signals are perfectly informative); but when $s = l$, some firms get $\kappa_i = \infty$ and others get $\kappa_i = 0$ (which means that the signal is completely uninformative). Under this scenario, when the precision state is $s = h$, all firms are perfectly informed. However, this fact is not common knowledge. This is because each firm cannot tell whether she is perfectly informed because the state is $s = h$ or because the state is $s = l$ but she was among the lucky ones to receive the perfectly precise signals. As a result, each firm must assign positive probability to the event that some other firms might not be informed and hence might not adjust their prices. But then because of strategic complementarity every firm will find it optimal not to adjust fully to the shock. It follows that there exist events where all firms are perfectly informed about the shock and nevertheless do not fully adjust their prices.

Of course, in this last example the possibility that firms are perfectly informed and yet do not respond perfectly to the shock can occur only with probability strictly less than one: the will also be events where some firms are relative uninformed and nevertheless find it optimal to respond quite a bit to the shock because they expect that other firms will be more informed. That is, this example cannot generate situations where in all events firms are perfectly informed and nevertheless expect other firms to be less informed. Indeed, based on the results of Kajii and Morris (1997) regarding the robustness of complete-information equilibria to the introduction of incomplete information, one can safely guess that if the common prior assigns probability near 1 to the set of events where the firms are nearly perfectly informed about the underlying nominal shock, then with probability near 1 the response of prices to the underlying nominal shock will be nearly the same as in the common-knowledge benchmark. Nevertheless, the results of this section do highlight how quantifying the response of higher-order beliefs is essential for quantifying the response of prices to nominal shocks.

3.5 Heterogeneous priors

In this section we study how heterogeneous priors regarding the signals firms receive can affect the behavior of higher-order beliefs and thereby the response of prices to the underlying nominal shocks. In particular, we consider a system of heterogeneous priors that induces firms to behave in equilibrium as if they lived in a world where other firms were less informed about the underlying nominal shocks. This sustains a partially self-fulfilling equilibrium in which firms react little to the underlying shock, even if they have nearly perfect information about it.

Apart from the introduction of heterogeneous priors, the setup is identical to our baseline Calvo model of Section 3. As there, θ_t follows an exogenous random-walk process. In any given period, a firm may change its price with probability $1-\lambda$, in which case the price it chooses is a weighted average of all future target prices:

$$p_{i,t} = (1-\beta\lambda)\sum_{j=0}^{\infty}(\beta\lambda)^j\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_{t+j} + \alpha\mathbb{E}_{i,t}p_{t+j}\right] \tag{3.9}$$

As in the baseline model, each period firms learn perfectly the nominal demand of the previous period, θ_{t-1}, and receive a private signal of the current period's nominal demand:

$$x_{i,t} = \theta_t + \varepsilon_{i,t}.$$

However, firms disagree on the stochastic properties of the noise in their signals.

In particular, each firm believes that its own signal is an unbiased signal of θ_t. Specifically, firm i believes the error in its own private signal is drawn from the following distribution:

$$\varepsilon_{i,t} \sim \mathcal{N}(0, 1/\kappa_x).$$

At the same time, each firm believes that the private signals of all other firms are biased. Specifically, firm i believes that the errors in the private signals of all other firms are drawn independently from the following distribution:

$$\varepsilon_{j,t} \sim N\left(\delta_{i,t}, 1/\kappa_x\right)\ \forall j \neq i$$

where $\delta_{i,t}$ is the bias that firm i believes to be present in the private signals of other firms. Finally, we assume that the perceived biases are negatively correlated with the innovation in the fundamental (the nominal shock). Specifically, for all i and all t,

$$\delta_{i,t} = \delta_t \equiv -\chi \nu_t = -\chi(\theta_t - \theta_{t-1}),$$

where $\chi \in [0,1]$ is a parameter that controls the correlation of the perceived bias with the innovation in the nominal shock. Finally, these perceptions are commonly understood and mutually accepted: the firms have agreed to disagree.[11]

To understand the difference between the baseline model (which had assumed a common prior) and the current model (with allows for heterogeneous priors), it is useful to consider the beliefs of each firm about the *average* expectation of θ_t. In either model, each firm's own (first-order) expectation of the fundamental is

$$\mathbb{E}_{i,t}\theta_t = \frac{\kappa_x}{\kappa_x + \kappa_\theta} x_t + \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \theta_{t-1}$$

In the baseline model, this implied that each firm believed that the average first-order expectation in the rest of the population satisfied

$$\bar{E}_t^1 = \frac{\kappa_x}{\kappa_x + \kappa_\theta} \theta_t + \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \theta_{t-1}.$$

That is, the firm's second-order expectation was given by

$$\mathbb{E}_{i,t}\bar{E}_t^1 = \frac{\kappa_x}{\kappa_x + \kappa_\theta} \mathbb{E}_{i,t}\theta_t + \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \theta_{t-1}.$$

In contrast, now that firms have heterogeneous priors, each firm believes the average first-order expectation satisfies

$$\bar{E}_t^1 = \frac{(1-\chi)\,\kappa_x}{\kappa_x + \kappa_\theta} \theta_t + \frac{\kappa_\theta + \chi \kappa_x}{\kappa_x + \kappa_\theta} \theta_{t-1},$$

[11] We have used a related heterogeneous-priors specification in Angeletos and La'O (2009b), albeit within a different context and for different purposes. We refer the reader to that paper for a more thorough discussion on the modeling role of heterogeneous priors: they are convenient, but are not strictly necessary for the type of effects we document. For example, in that paper we document how a certain type of sunspot-like fluctuations can obtain with either heterogeneous priors or a common prior, but the former maintain a higher level of tractability.

That is, the firm's second-order expectation is now given by

$$\mathbb{E}_{i,t}\bar{E}_t^1 = \frac{(1-\chi)\,\kappa_x}{\kappa_x+\kappa_\theta}\mathbb{E}_{i,t}\theta_t + \frac{\kappa_\theta + \chi\kappa_x}{\kappa_x+\kappa_\theta}\theta_{t-1},$$

Therefore, the heterogeneous priors introduced in this section do not affect first-order beliefs, but they do affect second- and higher-order beliefs: the higher χ is, the more each firm believes that the beliefs of others will be less sensitive to innovations θ, even though its own belief is not affected.

We now examine how this affects equilibrium behavior. We conjecture once again an equilibrium in which the price set by a firm in period t is a linear function of $(p_{t-1}, \theta_{t-1}, x_t,)$:

$$p_{i,t} = P(p_{t-1}, \theta_{t-1}, x_{i,t}) = b_1 p_{t-1} + b_2 x_{i,t} + b_3\theta_{t-1} \tag{3.10}$$

for some coefficients b_1, b_2, b_3. Accordingly, firm i expects that the price level will satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\int P(p_{t-1}, \theta_{t-1}, x)dF_t(x) \tag{3.11}$$

where F_t is the cross-sectional distribution of signals as perceived by the typical firm. Our assumption regarding the heterogeneous priors implies that each firm thinks that the cross-sectional mean of the signals in the rest of the population is $(1-\chi)\theta_t + \chi\theta_{t-1}$. It follows that each firm expects the price level to satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\,P(p_{t-1}, \theta_{t-1}, (1-\chi)\theta_t + \chi\theta_{t-1})$$

or equivalently

$$p_t = c_1 p_{t-1} + c_2\theta_t + c_3\theta_{t-1} \tag{3.12}$$

where

$$c_1 = \lambda + (1-\lambda)b_1, \quad c_2 = (1-\lambda)b_2(1-\chi) \quad c_3 = (1-\lambda)(b_3 + b_2\chi). \tag{3.13}$$

But now recall from the baseline model that, no matter what are the coefficients (c_1, c_2, c_3), the best response of a firm to (3.12) is to set a price as in (3.10), with the coefficients (b_1, b_2, b_3) defined by the solution to (3.8). It follows that the equilibrium values of the coefficients (b_1, b_2, b_3)

and (c_1, c_2, c_3) are now given by the joint solution of (3.8) and (3.13).

Note that χ enters only the conditions for c_2 and c_3 in (3.13), not the condition for c_1. It follows that the equilibrium value of c_1 (and hence also that of b_1) remains the same as in our baseline model (or, equivalently, as in the standard Calvo model). Moreover, the price process continues to be homogeneous of degree one, so that $c_1 + c_2 + c_3 = 1$. Finally, the equilibrium value of c_2 now satisfies

$$c_2 = \frac{\lambda(1 - c_1)}{\lambda + c_1 \frac{\kappa_\theta + \chi\kappa_x}{\kappa_x(1-\chi)}}.$$

Comparing this last condition with the corresponding condition for the baseline model, one observes that the equilibrium values of (c_1, c_2, c_3) for the present model coincide with those of the baseline model if the precision of information in that model is adjusted to the value $\tilde{\kappa}_x$ defined by

$$\frac{\tilde{\kappa}_x}{\kappa_\theta} \equiv \frac{(1-\chi)\,\kappa_x}{\kappa_x + \chi\kappa_\theta}, \tag{3.14}$$

which is clearly decreasing in χ. This observation in turn establishes a certain isomorphism between the present model and the baseline one: in the heterogeneous-prior economy, firms expect the price level to respond to the underlying nominal shock in the same way as in a common-prior economy that is identical to the one in our baseline model except for the fact that the precision of available information is decreased from κ_x to $\tilde{\kappa}_x$.

Along with the fact that, because of strategic complementarity ($\alpha > 0$), the incentive of a firm to respond to its own information is lower the lower the expected response of the price level, we conclude that our heterogeneous-priors specification reduces the response of each firm to its own information about the underlying nominal shock.

Proposition 39 *In the equilibrium of the heterogeneous-priors economy, firms respond to their information in the same way as in an a common-prior economy in which the precision of information that other firms have is a decreasing function of χ. By implication, the sensitivity b_2 of a firm's price to its own signal about the underlying nominal shock is also decreasing in χ for given precision.*

To further appreciate this result, for the moment allow us to abstract from price rigidities ($\lambda = 0$) and consider the limit as $\kappa_x \to \infty$, meaning that each firm is (nearly) perfectly informed about the shock. In the common-prior world, this would have guaranteed that prices move one-to-one with the nominal shock and hence that the nominal shock has no real effect. But now let $\chi \to 1$ along with $\kappa_x \to \infty$ in such a way that the quantity $\tilde{\kappa}_x$ defined in (3.14) stays bounded away from ∞. In this limit, each firm is perfectly informed about the shock but expects other firms to respond as if they were imperfectly informed; firms therefore find it optimal to adjust their prices less than one-to-one in response to the nominal shock, thereby causing the shock to have a real effect, even though they are perfectly informed about the shock and there is no price rigidity.

The preceding analysis has focused on how heterogeneous priors affect the instantaneous response of prices to the underlying shock: in the model considered above, the dynamics of prices after the initial shock is driven solely by the Calvo mechanics, much alike the baseline model. However, heterogeneous priors can also affect these dynamics, to the extent that the perceived bias is persistent over time. To see this, suppose that bias δ_t follows an autoregressive process of the following form:

$$\delta_t = -\chi v_t + \rho \delta_t,$$

for some $\rho \in [0, 1)$. One can the easily extend the preceding analysis to show the following.

Proposition 40 *There exist coefficients* (b_1, b_2, b_3, b_4)*, which depend on* (χ, ρ)*, such that the equilibrium strategy of firm* i *is given by*

$$p_{i,t} = P(p_{t-1}, x_{it}, \theta_{t-1}, \delta_{t-1}) \equiv b_1 p_{t-1} + b_2 x_{it} + b_3 \theta_{t-1} + b_4 \delta_{t-1}$$

At this point, it is important to recognize that so far we have used the model to make predictions only about what the *firms* expect the price level to do and how they respond to their information about this shock—we have not used the model to make predictions about what an outside observer, or an "econometrician", should expect the price level to do. This is where heterogeneous priors make things delicate. To analyze the equilibrium strategy of the firm as a function of their signals (their "types"), we do not need to take a stand on whether the firms signals are "truly" biased or not; we only need to postulate the system of their beliefs and then we can use the model to make predictions about how these beliefs will map into behavior. In contrast, to analyze the impulse response of the aggregate price level to the underlying shock, it no more suffices to characterize the mapping from beliefs to behavior; we also need to specify a mapping from the underlying nominal shock to the cross-sectional distribution of beliefs. In particular, we must now take a stand on how the cross-sectional average of signals relates to the underlying nominal shock in the eyes of the outside observer.

One possibility, then, is to assume that the "econometrician" believes that the signals are biased, so that the cross-sectional average of x_{it} is $\theta_t + \delta_t$. But another equally plausible possibility is to assume that the "econometrician" believes that the signals are unbiased, so that the cross-sectional average of x_{it} is θ_t. Under the former scenario, the econometrician's law of motion of the price level coincides with the one in the minds of the firms. In particular, it is given by

$$p_t = (1-\lambda)p_{t-1} + \lambda P(p_{t-1}, \theta_t + \delta_t, \theta_{t-1}, \delta_{t-1}),$$

with the function P given by Proposition 40. Under latter scenario, instead, the econometrician's law of motion for the price level differs from the one in the mind of the firms. In particular, it is now given by

$$p_t = (1-\lambda)p_{t-1} + \lambda P(p_{t-1}, \theta_t, \theta_{t-1}, \delta_{t-1}),$$

with the function P given once again by Proposition 40. Evidently, the only difference in these two law of motions in that the cross-sectional average of x_{it} is assumed to be $\theta_t + \delta_t$ in the first case and θ_t in the latter case.

For the remainder of the analysis, we assume the latter scenario, keeping though in mind three properties. First, under this scenario the econometrician predicts that firms react to δ_t not because their signals are biased, but only because they believe that other firms will do so. Second, the two scenarios deliver similar qualitative properties as long as $\alpha > 0$. And third, the quantitative difference between the two scenarios vanishes as $\alpha \to 1$. Both of these properties are direct implications of strategic complementarity.

Assuming the second scenario and combining the stochastic process of θ_t and δ_t with the law of motion for the price level, we conclude that the dynamics of the economy, as seen from the perspective of the econometrician, are given by the following:

$$\begin{bmatrix} \theta_t \\ \delta_t \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \rho & 0 \\ \lambda b_2 + \lambda b_3 & \lambda b_4 & 1 - \lambda + \lambda b_1 \end{bmatrix} \begin{bmatrix} \theta_{t-1} \\ \delta_{t-1} \\ p_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ -\chi \\ c_2 \end{bmatrix} v_t$$

with the coefficients (b_1, b_2, b_3, b_4) being determined as in Proposition 40. One can then use this system, along with the fact that real output is $y_t = \theta_t - p_t$, to simulate the impulse responses of inflation and output to a positive innovation in v_t (of a size equal to the standard deviation of v_t).

Figure 8: inflation

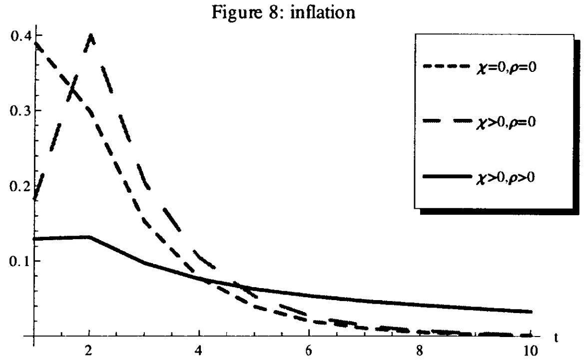

Figure 9: output

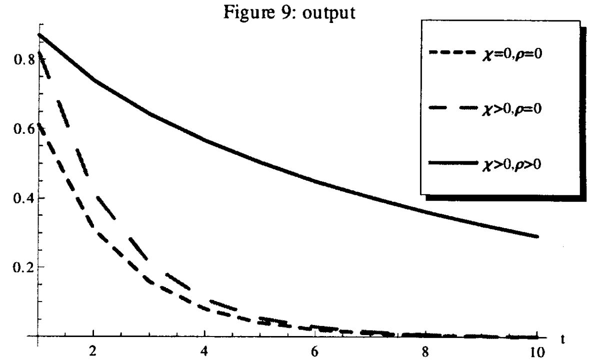

Figures 8 and 9 illustrate these impulse responses for different values of χ and ρ. The baseline (common-prior) model corresponds to $\chi = \rho = 0$. As anticipated, letting $\chi > 0$ but keeping $\rho = 0$ affects the impact effect of the innovation but not its persistence. In particular, in the period that the the innovation in nominal demand materializes, the response of inflation is dampened by letting $\chi > 0$, and by implication the positive effect on real output is amplified. But as long as $\rho = 0$ the dynamics following this initial period are determined solely by the Calvo propagation mechanism and hence the persistence is the same as in the baseline model. This is because any discrepancy between either first-order or higher-order beliefs and the true value of the shock vanishes after the initial period. In contrast, letting $\rho > 0$ permits the discrepancy to persist in higher order beliefs even after it has vanished in first-order beliefs, thereby contributing to additional persistence in the real effects of the nominal shock.

To sum up, heterogeneous priors can help rationalize significant inertia the response of prices to changes in nominal demand simply by inducing inertia in the response of higher-order expectations. This is true no matter how high is the firms' precision of information, that is, the sensitivity of first-order beliefs to the nominal shock. Moreover, this insight is not specific to the contemporaneous response of beliefs to the shock, but also to the entire dynamic adjustment of the beliefs. Indeed, in work that not reported here because of space limitations, we have obtained similar results in a heterogeneous-priors variant of Woodford (2003), in which the shock never becomes common knowledge, thus allowing both first- and high-order beliefs converge to the truth only slowly over time. But whereas in Woodford (2003) the rate of convergence of higher-order beliefs is tightly connected to that of first-order beliefs, heterogeneous priors permits us to break this connection, so that higher-order beliefs and prices may converge very slowly to their complete-information values even if first-order beliefs converge very fast.

Finally, note that the present model shares a bit of the flavor of the model with uncertain precisions considered in the previous section: there we focused on the possibility that firms may face uncertainty about the precision of other firms' information about the shock; here we showed how heterogeneous priors can induce firms to behave as if they expected other firms to have less precise information than themselves. In either case, the key is the behavior of higher-order beliefs as opposed to first-order beliefs.[12]

3.6 Heterogenous priors and cost-push shocks

In the preceding two sections we highlighted how higher-order beliefs can induce inertia in the response of prices to nominal shocks in the economy. In so doing, we focused on the role of higher-order beliefs for the propagation of certain structural shocks, namely nominal shocks. In this section, building on Angeletos and La'O (2009b), we highlight how higher-order beliefs can *themselves* be the source of a certain type of fluctuations in the price level and real output—

[12]The similarity of the results of our two models is reminiscent of a point made in Lipman (2003). That paper shows that, if one fixes a specific hierarchy of beliefs, one cannot tell apart a common prior from heterogeneous priors from properties of any finite order of beliefs. This in turn suggests that in certain cases it may be possible to replicate, or approximate, the equilibrium behavior that obtains for any particular hierarchy of beliefs with either a common prior or heterogeneous priors. At the same time, because the reduced-form game that characterizes the general equilibrium of our economy admits a unique rationalizable outcome, and hence also a unique correlated equilibrium, the results of Kajii and Morris (1997) suggest that the complete-information equilibrium is robust to the introduction of incomplete information. It then follows that, with a common prior, one needs sufficient "noise" to make the predictions of the model under incomplete information sufficiently different.

fluctuations that resemble the ones generated by cost-push, or mark-up, shocks.

For this purpose, the model of the previous section is modified as follows. Firms continue to have heterogeneous priors about their signals, but the perceived bias is no more correlated with the underlying nominal shock. Rather, the bias follows an independent stochastic process, given by

$$\delta_t = \rho\delta_{t-1} + \omega_t \tag{3.15}$$

where ω_t is a Normally distributed shock that is i.i.d. across time and independent of θ_τ for all τ. Following similar steps as in the previous section, one can then show the following.

Proposition 41 *There exist coefficients* $(b_1, b_2, b_3, b_4, b_5)$ *such that the equilibrium strategy of firm* i *is given by*

$$p_{i,t} = b_1 p_{t-1} + b_2 x_{i,t} + b_3\theta_{t-1} + b_4\delta_{t-1} + b_5\omega_t$$

Figure 10: inflation

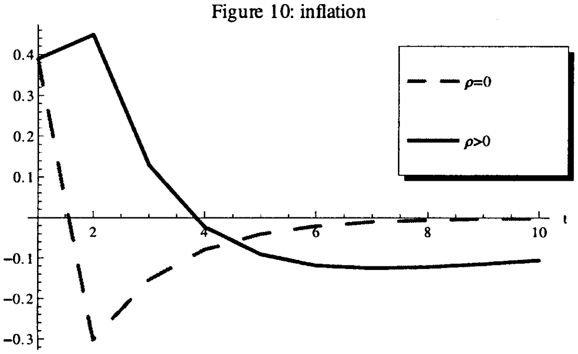

Figure 11: output

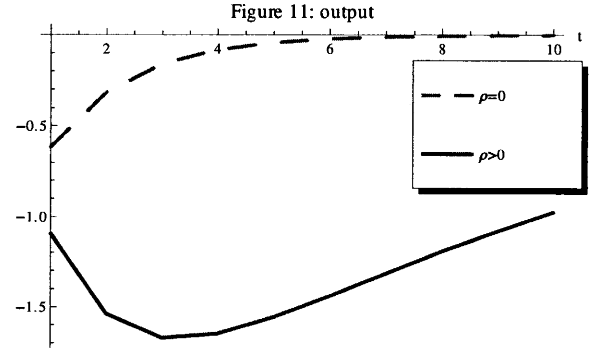

Taking once again the perspective of an econometrician who believes that the signals are unbiased, the impulse responses of the price level can be obtained from the following dynamic system:

$$\begin{bmatrix} \theta_t \\ \delta_t \\ p_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \rho & 0 \\ \lambda b_2 + \lambda b_3 & \lambda b_4 & 1 - \lambda + \lambda b_1 \end{bmatrix} \begin{bmatrix} \theta_{t-1} \\ \delta_{t-1} \\ p_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ \lambda b_2 \end{bmatrix} v_t + \begin{bmatrix} 0 \\ 1 \\ \lambda b_5 \end{bmatrix} \omega_t$$

Figures 10 and 11 illustrate the impulse responses of inflation and real output to a positive innovation in ω_t. Such a shock causes firms to raise their prices even though aggregate nominal demand hasn't change. As a result, inflation increases and output contracts. The resulting fluctuations thus resembled to those often identified as the impact of "cost-push" shocks. When $\rho = 0$, this cost-push-like shock is transitory; the moderately persistent effects on real output are then due merely to the Calvo mechanics. When instead $\rho > 0$, the bias is itself persistent, which contributes to additional persistence in the real effects of the shock.

3.7 Concluding Remarks

this chapter studied how the combination of incomplete information and infrequent price adjustment may dampen the response of prices to nominal shocks. A series of variant models was considered that progressively shifted focus from the stickiness of prices and the precision of available information about the shocks to the dynamics of higher-order expectations. We thus sought to highlight that quantifying the degree of price stickiness and the speed of learning (i.e., the rate at which first-order beliefs adjust to the shock) does not suffice for quantifying the rate at which higher-order beliefs adjust, and therefore also does not suffice for quantifying the rate at which prices adjust. We also illustrated how the distinct role of higher-order beliefs could be readily accommodated within the Calvo framework without any sacrifice in analytical tractability—which in turn may pave the way to bringing the ideas of this chapter closer to the data.

We thus hope that more effort will be devoted to quantifying the behavior of expectations and their implications for the degree of price inertia at the macro level. How can this be done? One possibility, at least in principle, would be to start conducting surveys of the higher-order beliefs of economic agents. We find this both impractical and unnecessary. As further argued in Angeletos and La'O (2009a, 2009b), in our eyes dispersed information—and higher-order beliefs—are primarily modeling devices for capturing the uncertainty that economic agents may face about aggregate economic activity beyond the one that they face about the underlying fundamentals. Indeed, in macro models and games alike, higher-order beliefs matter only to the extent that they impact forecasts of the equilibrium actions of other agents. Furthermore, in most macro models this is typically summarized in forecasts of few macroeconomic variables, such as the price level and the level of aggregate output or employment. Therefore, we do not think that it is essential to collect data on the details of higher-order beliefs. Rather, we believe that data on forecasts of economic activity can provide more direct guidance for quantifying the type of effects we document here.

Turning to potential implications for monetary policy, we wish to make the following points. If one interprets the exogenous shock in our analysis as an innovation in monetary policy, one may conclude that our results justify strong real effects for exogenous changes to monetary policy. We would not necessarily favor such an interpretation. The theory we have presented here does not imply the same price inertia with respect to all shocks in the economy, not even all the monetary ones. Rather, it is absolutely crucial whether there is non-trivial lack of common knowledge about the shock under consideration. But then note that information about changes in monetary policy is readily available and closely followed, not only from participants in financial markets, but also from the general public when it seems to matter most. Moreover, it is commonly understood that this is the case; there is a lot of communication in the market regarding monetary policy; and financial prices adjust within seconds to changes in monetary policy. In our eyes, these are indications that assuming common knowledge about the innovations to monetary policy might not be a terrible benchmark after all.

At the same time, we suspect a significant lack of common knowledge for a variety of other "structural" shocks hitting the economy, such aggregate productivity shocks, financial shocks, labor-market shocks, and so on. We then expect incomplete information to have more bite on monetary policy in a different dimension: the interest-rate rule followed by the central bank can affect how much the incompleteness of information about these shocks impacts equilibrium outcome. In particular, the central bank can control the degree of strategic complementarity in pricing decision by designing it response to realized inflation and output; in so doing, it can also mitigate the inertia effects we have documented here. Further exploring this possibility is left for future work.[13]

[13]The insight that the response of policy to macroeconomic outcomes can impact the decentralized use of information, and thereby the response of the economy to both the fundamentals and noise draws from Angeletos and Pavan (2009). See also Angeletos and La'O (2008) and Lorenzoni (2009) for some recent work on the design of optimal monetary policy when information is dispersed.

Appendix

Proof of Proposition 37. The condition that determines c_1 can be restated as

$$f(c_1, \lambda, \beta) = \alpha, \tag{3.16}$$

where

$$f(c_1, \lambda, \beta) \equiv \frac{(1 - \beta\lambda c_1)(c_1 - \lambda)}{c_1(1 - \lambda)(1 - \beta\lambda)}.$$

As mentioned in the main text, there are two solutions to this equation—one with $c_1 > 1$ and $c_1 \in (\lambda, 1)$—and we focus on the non-explosive one.

Clearly, the aforementioned equation is independent of κ_x and κ_θ, implying that c_1 is independent of the information structure. Moreover,

$$\frac{\partial c_1}{\partial \lambda} = -\frac{\partial f/\partial \lambda}{\partial f/\partial c_1} = \frac{c_1(1-c_1)(1-\beta c_1)(1-\beta\lambda^2)}{\lambda(1-\lambda)(1-\beta\lambda)(1-\beta c_1^2)} > 0,$$

$$\frac{\partial c_1}{\partial \alpha} = \frac{1}{\partial f/\partial c_1} = \frac{\lambda(1-\beta c_1^2)}{c_1^2(1-\lambda)(1-\beta\lambda)} > 0.$$

It follows that c_1 is increasing in both λ and α. Next, recall from the main text that c_2 satisfies

$$c_2 = \frac{\lambda(1-c_1)}{\lambda + c_1\frac{\kappa_\theta}{\kappa_x}}.$$

Since c_1 is independent of $(\kappa_x, \kappa_\theta)$, it is immediate that c_2 is increasing in κ_x/κ_θ; and since $c_3 = 1 - c_1 - c_2$, it is immediate that c_3 is decreasing in κ_x/κ_θ. Moreover, since the above expression for c_2 is independent of α for given c_1 and is decreasing in c_1, and since c_1 is itself increasing in α, it follows that c_2 is decreasing in α. Finally, the fact that c_2 is non-monotonic in λ and that c_3 is non-monotone in α can be establish by numerical example.

Proof of Proposition 38. The equilibrium can be characterized in a similar fashion as in the baseline model. First, by aggregating the strategy of the firms, we infer that the coefficients

$(c_1, c_{2,h}, c_{2,l}, c_{3,h}, c_{3,l})$ must solve the following system:

$$\begin{aligned}
c_1 &= \lambda + (1-\lambda)b_1 \\
c_{2,s} &= (1-\lambda)\left[\gamma b_{2,s} + (1-\gamma)b_{2,-s}\right] \\
c_{3,s} &= (1-\lambda)\left[\gamma b_{3,s} + (1-\gamma)b_{3,-s}\right]
\end{aligned}$$

where we use the convention that $-s = l$ when $s = h$ and $-s = h$ when $s = l$. Next, by taking the firms' best response, we infer that the coefficients $(b_1, b_{2,h}, b_{2,l}, b_{3,h}, b_{3,l})$ solve the following system:

$$\begin{aligned}
b_1 &= ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\, c_1 \\
b_{2,s} &= \Big[\, (1-\beta\lambda)(1-\alpha) + ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\,(\gamma c_{2,s} + (1-\gamma)c_{2,-s}) \\
&\quad + (\beta\lambda)\left(\tfrac{1}{2}(b_{2,l} + b_{2,h}) + (b_{3,l} + b_{3,h})\right) \Big] \frac{\kappa_x}{\kappa_x + \kappa_\theta} \\
b_{3,s} &= \Big[\, (1-\beta\lambda)(1-\alpha) + ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\,(\gamma c_{2,s} + (1-\gamma)c_{2,-s}) \\
&\quad + (\beta\lambda)\left(\tfrac{1}{2}(b_{2,l} + b_{2,h}) + (b_{3,l} + b_{3,h})\right) \Big] \frac{\kappa_\theta}{\kappa_x + \kappa_\theta} \\
&\quad + ((1-\beta\lambda)\alpha + (\beta\lambda)b_1)\,(\gamma c_{3,s} + (1-\gamma)c_{3,-s}).
\end{aligned}$$

Clearly, c_1 and b_1 continue to be determined by the same equations as in the baseline model. Once again, we focus on the solution with $c_1 \in (0,1)$. Given this solution, the remainder of the conditions consist a linear system, which admits a unique solution for the coefficients $(b_{2,s}, b_{3,s}, c_{2,s}, c_{3,s})_{s \in \{h,l\}}$.

Proof of Proposition 39. In the main text of the article we showed that the equilibrium values of (c_1, c_2, c_3) and (b_1, b_2, b_3) are determined by the solution to two systems of equations, referenced as conditions (8) and (13) in the article; that c_1 and b_1 continue to be determined as in the baseline model and are thus independent of χ; and that c_2 satisfies

$$c_2 = \frac{\lambda(1-c_1)}{\lambda + c_1 \frac{\kappa_\theta + \chi\kappa_x}{\kappa_x(1-\chi)}},$$

which is decreasing in χ. Along with the fact that $c_2 = (1-\lambda)b_2(1-\chi)$, we get that

$$b_2 = \frac{\lambda(1-c_1)}{(1-\lambda)\left[\lambda + c_1\frac{\kappa_\theta}{\kappa_x} + (c_1-\lambda)\chi\right]},$$

which is also decreasing in χ, since $c_1 \in (\lambda, 1)$.

Proof of Propositions 40 and 41. To nest both the model of Section 5 and that of Section 6, we let the bias be given by

$$\delta_t = -\chi v_t + \omega_t + \rho\delta_{t-1}.$$

We then conjecture an equilibrium in which the price set by a firm in period t is a linear function of $(p_{t-1}, x_t, \theta_{t-1}\delta_{t-1}, \omega_t)$:

$$p_{i,t} = P(p_{t-1}, x_{i,t}, \theta_{t-1}, \delta_{t-1}, \omega_t) = b_1 p_{t-1} + b_2 x_{i,t} + b_3\theta_{t-1} + b_4\delta_{t-1} + b_5\omega_t \tag{3.17}$$

for some coefficients b_1, b_2, b_3, b_4, b_5.

Our specification of the heterogeneous priors implies that each firm thinks that the cross-sectional mean of the signals in the rest of the population is $\bar{x}_t = \theta_t + \delta_t = (1-\chi)\theta_t + \chi\theta_{t-1} + \rho\delta_{t-1} + \omega_t$. It follows that each firm expects the price level to satisfy

$$p_t = \lambda p_{t-1} + (1-\lambda)\, P(p_{t-1}, \bar{x}_t, \theta_{t-1}, \delta_{t-1}, \omega_t)$$

or equivalently

$$p_t = c_1 p_{t-1} + c_2\theta_t + c_3\theta_{t-1} + c_4\delta_{t-1} + c_5\omega_t \tag{3.18}$$

where

$$\begin{aligned} c_1 &= \lambda + (1-\lambda)\, b_1, \quad c_2 = (1-\lambda)\, b_2\,(1-\chi), \quad c_3 = (1-\lambda)\,(b_2\chi + b_3) \\ c_4 &= (1-\lambda)\,(b_2\rho + b_4), \quad c_5 = (1-\lambda)\,(b_2 + b_5) \end{aligned} \tag{3.19}$$

Next, note that we may write the firm i's best response, which is given by condition (9) in the article, as follows:

$$p_{i,t} = (1-\beta\lambda)\left[(1-\alpha)\,\mathbb{E}_{i,t}\theta_t + \alpha\mathbb{E}_{i,t}p_t\right] + (\beta\lambda)\,\mathbb{E}_{i,t}p_{i,t+1}.$$

Given that the firm expects the price level to evolve according to (3.18), for the firm's best response to be consistent with our conjecture (3.17), it must be that the coefficients $(b_1, b_2, b_3, b_4, b_5)$ solve the following system:

$$\begin{aligned}
b_1 &= \left[(1-\beta\lambda)\,\alpha + (\beta\lambda)\,b_1\right] c_1 \\
b_2 &= \left[(1-\beta\lambda)\,(1-\alpha+\alpha c_2) + (\beta\lambda) + (\beta\lambda)\,(b_1c_2 + b_2 + b_3 - \chi b_4)\right]\frac{\kappa_x}{\kappa_x+\kappa_\theta} \\
b_3 &= \left[(1-\beta\lambda)\,(1-\alpha+\alpha c_2) + (\beta\lambda) + (\beta\lambda)\,(b_1c_2 + b_2 + b_3 - \chi b_4)\right]\frac{\kappa_\theta}{\kappa_x+\kappa_\theta} \\
&\quad + (1-\beta\lambda)\,\alpha c_3 + (\beta\lambda)\,b_1c_3 + (\beta\lambda)\,b_4\chi \\
b_4 &= (1-\beta\lambda)\,\alpha c_4 + (\beta\lambda)\,b_1c_4 + (\beta\lambda)\,b_4\rho \\
b_5 &= (1-\beta\lambda)\,\alpha c_5 + (\beta\lambda)\,b_1c_5 + (\beta\lambda)\,b_4
\end{aligned} \tag{3.20}$$

Combining conditions (3.19) and (3.20) gives us a system of equations which characterizes the equilibrium values for $(b_1, b_2, b_3, b_4, b_5, c_1, c_2, c_3, c_4, c_5)$.

It is immediate that the conditions that determine c_1 and b_1 are identical to those in the baseline model. As in the baseline model, we ignore the solution that has $c_1 > 1$ and focus on the solution that has $c_1 \in (0, 1)$. Given this solution, the remaining conditions define a linear system, which has a unique solution for the remaining coefficients.

Chapter 4

Predatory Trading and Credit Freeze

4.1 Introduction

The inter-bank lending market and the discount window of the Fed are two facilities which allow banks to borrow short-term in order to meet temporary liquidity needs. However, these opportunities are not always availed by traders. Financial institutions often appear reluctant to borrow, even at times when liquidity is most needed. In this chapter I study how strategic interactions among banks may deter financial institutions from raising money in times of temporary financial distress.

Financial markets are often modeled as interactions between small traders in perfectly competitive markets taking prices as given. However, in reality these markets are not devoid of large players with market impact. For this reason, a recent literature has begun to emphasize strategic behavior among large financial institutions. "Predatory trading" is one such strategic interaction. Brunnermeier and Pedersen (2005) define predatory trading as "trading that induces and/or exploits the need of other investors to reduce their positions." That is, predatory trading is a strategy in which a trader can profit by trading against another trader's position, driving an otherwise solvent but distressed trader into insolvency. The forced liquidation of the distressed trader leads to price swings from which the predator can then profit. Brunnermeier and Pedersen (2005) provide a framework to study this type of interaction, and show how predatory trading can in fact induce a distressed trader's need to liquidate.

In this chapter I explore how predatory trading may affect the incentives of banks to seek loans in times of financial distress. In general, a distressed bank or trader may wish to raise money in order to temporarily bridge financial short-falls. However, in an environment in which banks have private information about their own finances, searching for extra capital from outside lenders may become a signal of financial weakness. Traders can then exploit this information, and predatorily trade against funds that they infer to be sufficiently week. Therefore, the mere act of searching for loans may expose a distressed firm to predatory trading and possible insolvency.

The key assumption behind this result is that there exists asymmetric information among traders–that is, ex ante, traders have private information about their own balance sheet that is not available to other traders. Within an asymmetric information environment, actions undertaken by banks to relieve financial distress may convey information about its underlying financial state. Hence, in deciding whether to search for a loan, a distressed bank faces a trade-off between the financial cushion provided by a loan and the information this act reveals. In equilibrium, I find that some distressed funds who would otherwise seek to recapitalize may be reluctant to search for extra capital in the presence of potential predators. Predatory trading may therefore deter banks and financial institutions from raising funds in times when they need it the most.

Anecdotal Evidence. The findings of this chapter support certain anecdotal evidence about strategic trading and the reluctance of financial institutions to find loans in times of distress. One of the most often-cited examples of predatory trading is the alleged front-running against the infamous hedge fund Long-Term Capital Management (LTCM) in the fall of 1998. After realizing losses in a number of markets, it is reported that LTCM began searching for capital from a number of Wall Street banks, most notably Goldman Sachs & Co. LTCM alleges that with this information Goldman then traded heavily against LTCM's positions in credit-default swaps, front-running LTCM's eventual unwinding. *Business Week* writes,[1]

> ..if lenders know that a hedge fund needs to sell something quickly, they will sell

[1] "The Wrong Way to Regulate Hedge Funds," *Business Week*, February 26, 2001, p. 90.

> the same asset–driving the price down even faster. Goldman Sachs & co. and other counterparties to LTCM did exactly that in 1998. Goldman admits it was a seller but says it acted honorably and had no confidential information.

Similarly, in *When Genius Failed: The rise and fall of Long-Term Capital Management*, Lowenstein writes,[2]

> As it scavenged for capital, Long-Term had been forced to reveal bits and pieces and even the general outline of its portfolio... Meriwether bitterly complained to the Fed's Peter Fisher that Goldman, among others, was "front-running", meaning trading against it on the basis of inside knowledge. Goldman, indeed, was an extremely active trader in mid-September, and rumors that Goldman was selling Long-Term's positions in swaps and junk bonds were all over Wall Street.

Furthermore, an interesting study by Cai (2007), uses a unique data set of audit transactions to examine the trading behavior of market makers in the Treasury bond futures market during LTCM's collapse. Cai finds strong evidence of predatory front running behavior by market makers, based on their informational advantages.[3]

The findings of this chapter also provide a possible explanation as to why financial firms may not obtain loans in times of financial shortfall. During the 2008-2009 financial crisis, sources report that Lehman Brothers was reluctant to publicly raise liquidity, a month or so before its collapse. *The Wall Street Journal* writes,[4]

> As the credit crunch deepened, the Fed had set up a new lending facility for investment banks. Although the central bank doesn't reveal who borrows from it, the market generally figures it out, and there's a stigma associated with it. Lehman

[2] According to the author, *When Genius Failed: The Rise and Fall of Long-Term Capital Management* was based on interviews with former employees and partners of the firm, as well as interviews conducted at the major Wall street invesment banks.

[3] Although identities are concealed in the transactions dataset, Cai finds one large clearing firm (coded "PI7") with large customer orders during the crisis period which closely match various features of LTCM's trades executed through Bear Stearns, including trade size, pattern and timing. More importantly, Cai finds that market makers traded on their own accounts in the same direction just one or two minutes before before PI7 customer orders were executed.

[4] "The Two Faces of Lehman's Fall." *The Wall Street Journal,* October 6, 2008.

> didn't do so over the summer, because it didn't want to be seen as needing Fed money, says one person familiar with the matter.

The *WSJ* further reports that Lehman eventually tried to secretly raise funds from the European Central Bank:

> In the weeks before it collapsed, Lehman Brothers Holdings Inc. went to great lengths to conceal how fast it was careening toward the financial precipice. The ailing securities firm quietly tapped the European Central Bank as a financial lifeline.

Eventually, any funds Lehman could acquire were apparently not enough, and the investment bank declared bankruptcy on September 15, 2008. The *Associated Press* writes,[5]

> If the mortgage meltdown is like a financial hurricane, then think of Lehman Brothers as a casualty that waited too long to cry for help.

Related Literature. Brunnermeier and Pedersen (2005) provides the basic framework for predatory trading used in this chapter. Brunnermeier and Pedersen (2005) show that if a distressed firm is forced to liquidate a large position, other traders have the incentive to trade in the same direction, in order to profit from large price swings. Furthermore, they show that predatory trading can even induce the distressed trader's need to liquidate. In their analysis, the predator is perfectly informed of the distressed trader's balance sheet, whereas in this chapter I relax this assumption and allow traders to have private information about their own finances. This is motivated by the observation that banks often know more about their own balance sheet (and portfolio) than other institutions.

this chapter more generally emphasizes the importance of considering non-competitive markets in which large strategic traders do not take prices as given. Strategic trading based on private information about security fundamentals is studied by Glosten and Milgrom (1985) and Kyle (1985), while speculative trading by investors with no knowledge of fundamental values, but who do possess superior knowledge of the trading environment is studied by Madrigal (1996) and Vayanos (2001). Allen and Gale (1992), on the other hand, study stock price manipulation in which an investor buys and sells shares, incurring profits by convincing others that he

[5] "Financial hurricane victim Lehman waited too long." *The Associated Press*, September 14, 2008.

is informed. Finally, Carlin, Lobo, and Viswanathan (2007) offer a complementary theory of predatory trading: they show how predation is a manifestation of a breakdown in cooperation between market participants.

this chapter furthermore examines how lending problems may arise from the strategic interactions among banks. In this way, this chapter is related to Acharya, Gromb, and Yorulmazer (2009), who study market power in the interbank lending market. They show that during crises episodes, the profits a surplus bank may gain from buying fire-sale assets and increasing market share may lead to a lower willingness to supply interbank loans. Similarly, this chapter is related to the literature on the role of the central bank during episodes of aggregate liquidity shortages or interbank-lending market breakdown, see for example Allen and Gale (1998), Holmstrom and Tirole (1998), Diamond and Rajan (2005), and Gorton and Huang (2006). Finally, Bolton and Scharfstein (1990) show that an optimal lending contract may leave a firm unable to fully counter predation risk. They consider product market predation, not financial market predation. Finally, while all of these papers emphasize the provision of liquidity by banks and central banks, i.e. the suppliers of funds, they do not consider the signal value of searching for liquidity by distressed financial firms and how that endogenously affects the demand for funds.

this chapter is organized as follows. Section 4.2 describes the model. In Section 4.3, I define the equilibrium of the economy and analyze the optimal decision for each type of trader. Section 4.4 examines the benchmark case in which there is no predator. Section 4.5 characterizes the equilibria in the full model with predatory trading. Section 4.6 concludes.

4.2 The Model

There are 3 periods: $t \in \{1, 2, 3\}$ and two tradeable assets: a riskless bond and a risky asset. The risk-free rate is normalized to 0. The risky asset has an aggregate supply Q and a final payoff z at time $t = 3$, where z is a random variable with an expected value of $\mathbb{E}_i z = \bar{z} > 0$. The price of the risky asset at any time t is denoted s_t.

There are two strategic traders, the distressed trader and the (potential) predator, which are denoted by $i \in \{d, p\}$. Both traders are risk neutral and seek to maximize their expected profit at time $t = 3$, which I denote as $\tilde{w}_i$. Each strategic trader is large, and hence, his trading impacts the equilibrium price. Traders can be thought of as hedge funds or the proprietary trading desks of large investment banks. Let $x_{i,t}$ denote trader i's holding of the stock at time t. Each strategic trader has a given initial endowment, $x_{i,1}$, of the risky asset and is restricted to hold $x_{i,t} \in [-\bar{x}, \bar{x}]$.[6] For simplicity I assume that each trader's initial endowment is equal to its maximum long position, that is, $x_{p,1} = x_{d,1} = \bar{x}$.

In addition to the two large strategic traders, the market is populated by long-term investors. The long-term traders are price-takers and have at each point in time an aggregate demand curve given by

$$Y(s_t) = \frac{1}{\lambda}(\bar{z} - s_t). \tag{4.1}$$

This demand schedule has two important attributes. First, it is downward sloping: in order for long-term traders to hold more of the risky asset, they must be compensated in terms of lower prices.[7] Second, the long-term traders' demand depends only on the current price s_t, that is, they do not attempt to profit from future price swings.[8]

The market clearing price solves $Q = Y(s_t) + x_{p,t} + x_{d,t}$. Market clearing implies that the equilibrium stock price is given by $s_t = \bar{z} - \lambda[Q - (x_{p,t} + x_{d,t})]$. Due to the constraint on asset holdings, strategic traders cannot take unlimited positions. Assuming the case of limited capital, i.e. $2\bar{x} < Q$, the equilibrium price is always lower than the fundamental value: $s_t < \bar{z}$, $\forall t$. Therefore, strategic traders can expect positive profits from holding the asset until time $t = 3$.

In addition to the risky asset, each strategic trader is endowed with a non-tradeable investment. At time $t = 3$, this investment, if not liquidated, yields a payoff of u, where u is a random variable with an expected value of $\mathbb{E}_i u = \bar{u} > 0$. This investment is non-tradeable in the following sense: it cannot be sold by the trader at any point in time before the investment has materialized in the last stage. I let $v_{i,t}$ represent the paper value at time t of this investment. For example, if the trader is an investment bank, $v_{i,t}$ may be thought of as the value of

[6] This position limit can be interpreted more broadly as a risk limit or a capital constraint.

[7] This could be due to risk aversion or due to institutional frictions that make the risky asset less attractive for long-term traders. For instance, long-term traders may be reluctant to buy complicated derivatives such as asset-backed securities.

[8] Long-term investors may be interpreted as pension funds and individual investors. Under this interpretation, long-term investors may not have sufficient information, skills, or time to predict future price changes.

investments made by the lending side of the bank which, perhaps due to agency reasons, cannot be securitized.

The paper value of the distressed's investment is subject to liquidity shocks, such that $v_{d,t}$ is not necessarily equal to $\bar{u}$ at every point in time. In particular, $v_{d,t}$ at any point in time takes one of three values: $v_{d,t} \in V_d \equiv \{v_l, v_m, v_h\}$, where without loss of generality $v_l < v_m < v_h$. The realizations of $v_{d,t}$ are however independent of u, so that the trader's expected final payoff from his non-tradeable investment is always given by $\bar{u}$. On the other hand, the predator's valuation of non-tradeable assets is constant over time, and equal to its expected payoff: $v_{p,t} = \bar{u}$, $\forall t$.

At any time t, a trader's "mark-to-market" wealth is given by $w_{i,t} = x_{i,t}s_t + v_{i,t}$. If the trader survives to period 3, its expected payoff from holding its portfolio is $\mathbb{E}_i[\tilde{w}_i] = \mathbb{E}_i[w_{i,3}] = x_{i,3}\bar{z} + \bar{u}$. Let $\bar{w}$ denote the maximum expected wealth of a trader's portfolio, that is, $\bar{w} \equiv \bar{x}\bar{z} + \bar{u}$. However, if at any time before the last period a trader's wealth drops below some threshold level $\underline{w}$, then the trader must liquidate all assets at fire sale prices. This assumption of forced liquidation could be due to margin constraints, risk management, or other considerations in connection with low wealth. Let $L < \bar{w}$ be the fire-sale value of the entire portfolio if the trader is forced to liquidate before the last stage, and let $\Delta \equiv \bar{w} - L$ denote the difference between the expected payoff from the portfolio and its fire sale value. One may think of Δ as the penalty the trader incurs for liquidating prematurely.

Timing and Information. Before stage 1, Nature draws initial value $v_{d,1} \in V_d = \{v_l, v_m, v_h\}$ of the distressed's non-tradeable holdings according to the following distribution

$$v_{d,1} = \begin{cases} v_l & \text{with probability} \quad q_l \\ v_m & \text{with probability} \quad q_m \\ v_h & \text{with probability} \quad q_h \end{cases}$$

where $q_l + q_m + q_h = 1$. For simplicity, I let $q_l = q_m = q_h = 1/3$. One may think of this as the initial "type" of the distressed trader. That is, the distressed is initially a low type if $v_{d,1} = v_l$, a medium type if $v_{d,1} = v_m$, and a high type if $v_{d,1} = v_h$.

Stage 1. In stage 1, the distressed trader learns his initial type $v_{d,1}$ (or valuation of his non-tradeable investment), but the distressed's type is not observed by the predator.[9] This can be interpreted as investors conducting a valuation of the financial firm, but this value is not released publicly. Once it observes $v_{d,1}$, the distressed trader then has the option to search for additional resources from an outside lender. I let $a_d \in A_d \equiv \{S, NS\}$ denote the action taken by the distressed, where S denotes the decision to "search" for a loan, and NS denotes the decision to "not search". For this reason, I refer to stage 1 as the "loan-seeking stage".

If the distressed decides to search, he receives a loan which may increase his wealth. Before deciding to search, however, the distressed does not know the value of this loan. In particular, I assume that the loan is stochastic, and with some probability may bring the distressed trader into a higher valuation state for stage 2. For example, if the distressed initially is a medium type ($v_{d,1} = v_m$) and decides to search for a loan, with probability π_{mh} he receives a loan which makes him a high-type firm for stage 2 ($v_{d,2} = v_h$). On the other hand, with probability π_{mm} the distressed does not receive a loan and stays a medium type for stage 2. More formally, if the distressed searches, his stage 2 type $v_{d,2}$ is determined by the following transition matrix

$$\pi\left(v_{d,2}|v_{d,1}\right) = \begin{bmatrix} \pi_{ll} & \pi_{lm} & \pi_{lh} \\ 0 & \pi_{mm} & \pi_{mh} \\ 0 & 0 & 1 \end{bmatrix} \tag{4.2}$$

where $\pi_{ll} + \pi_{lm} + \pi_{lh} = 1$ and $\pi_{mm} + \pi_{mh} = 1$. Like the bond, the loan has zero interest but does have a fixed cost of $c > 0$ which is incurred in the last period. Finally, if the distressed decides to not search for a loan, his type remainds constant; that is, $v_{d,2} = v_{d,1}$.

After the distressed decides whether or not to search for a loan, the value $v_{d,2}$ is realized. This value is again observed by the distressed but not by the predator.

Stage 2. Although the predator does not observe the distressed's type $v_{d,2}$ directly, the predator does however observe whether the distressed decided to search or not. Specifically, the predator observes a_d. After observing a_d, the predator then decides whether or not to predatorily trade against the distressed. I let $a_p \in A_p \equiv \{P, NP\}$ denote the action taken by the potential predator, where P denotes the decision to "predate", and NP denotes the action

[9]The predator's type v_p is constant and common knowledge throught out the game.

to "not predate". For this reason, I refer to this stage as the "predatory phase".

If the predator decides to predatorily trade, then the strategic traders engage in a "predation war". The results of this predation war are derived from Brunnermeier and Pedersen (2005). The mechanics of this predation war are not the main focus of this chapter. For this reason, in this section I only present the important (reduced-form) results that are pertinent to understanding the model. A more detail description of the predation war is given in the Appendix.

If the predator decides to predate, then the strategic traders engage in a "predation war" in which both traders sell the risky asset as fast as possible. This predation war continues until one of the traders is forced to leave the market. The trader who is forced to leave the market is the trader whose wealth falls below the minimum wealth threshold $\underline{w}$ first–that is, the trader with the lower amount of wealth will be forced to leave the market. The predator therefore wins the predation war and the distressed loses if and only if $v_p > v_{d,2}$. In this case, the predator buys back up to its optimal position $\bar{x}$, receives strictly positive profits $m > 0$ and moves on to stage 3, while the distressed trader is forced into liquidation and receives final payoff $\tilde{w}_d = L$. On the other hand, if the predator loses and the distressed wins the predation war, the predator must liquidate its assets at fire sale prices and receives final payoff $\tilde{w}_p = L$, while the distressed buys back up to its optimal position $\bar{x}$ and continues on to stage 3.[10]

If the predator decides not to predate, then there is no predation war. Both traders move on to the next period with no change to their current or expected wealth.

Stage 3. In stage 3, conditional on making it to this stage (either not engaging or winning the predation war in stage 2), the predator receives the final realized wealth from his portfolio, $\tilde{w}_p = x_{p,3}z + u$.

In addition, the distressed trader, conditional on making to this stage (either not engaging or winning the predation war in stage 2), is subject to an exogenous income shock. This income shock has two outcomes, either it results in stage 3 wealth below the threshold $\underline{w}$, forcing the distressed trader to liquidate, or it results in stage 3 wealth above the threshold.

The distressed in period 3, has valuation $v_{d,3}$ equal to its valuation in the previous period,

[10] Note that the distressed does not make profits from winning the pedation war. This may be interpreted as the distressed isn't trying, or does not have the skills, to profit from the exit of the predator. Thus, in the event that the distressed wins a predation war, it receives zero gains: $m_d = 0$.

$v_{d,2}$. The probability of hitting the lower bound on wealth after the income shock depends on the distressed's current type. In particular, the probability that the trader's wealth after the income shock is above the threshold is increasing in $v_{d,3}$. If the distressed trader is low type, then his wealth after the income shock is above the threshold $\bar{w}$ with probability p_l. If the trader is medium type, then his wealth after the income shock is above the threshold with probability p_m. Finally, if the trader is high type, then his wealth after the income shock is above the threshold with probability p_h. I assume $p_l < p_m < p_h$, so that the high type has the lowest probability of hitting the lower bound on wealth, and the low type has the highest probability of hitting the lower bound.

If the distressed hits the lower bound on wealth after the income shock, he is forced to liquidate all assets at fire-sale prices and receives final payoff $\tilde{w}_d = L$. If instead the distressed's wealth is above the threshold after the income shock, then he receives the final payoffs from holding the portfolio, $\tilde{w}_d = x_{d,3}z + u$. If the distressed took out a loan in stage 1, he now pays the fixed cost c to the outside lender.

Remarks and Assumptions. The extensive form of this game is presented in Figure 4.1 for stages 1 and 2. I omit the last stage for simplicity and because no player moves in the last stage. Therefore, what is omitted in this figure is the stochastic income shock in stage 3 and the realization of payoffs for each player. From this figure, one can see that Nature (denoted by "N") moves first, choosing the distressed's non-tradable investment value v_l, v_m, v_h with probabilities q_l, q_m, q_h respectively. After observing its own type, the distressed decides whether to search (S) or not search (NS) for a loan. If the distressed searches, Nature then draws another type based on the transition matrix given in (4.2). Finally, the predator, although he does not observe the distressed's type, does get to observe the distressed's action. Given this information set, the predator then decides whether to predate (P) or not (NP).

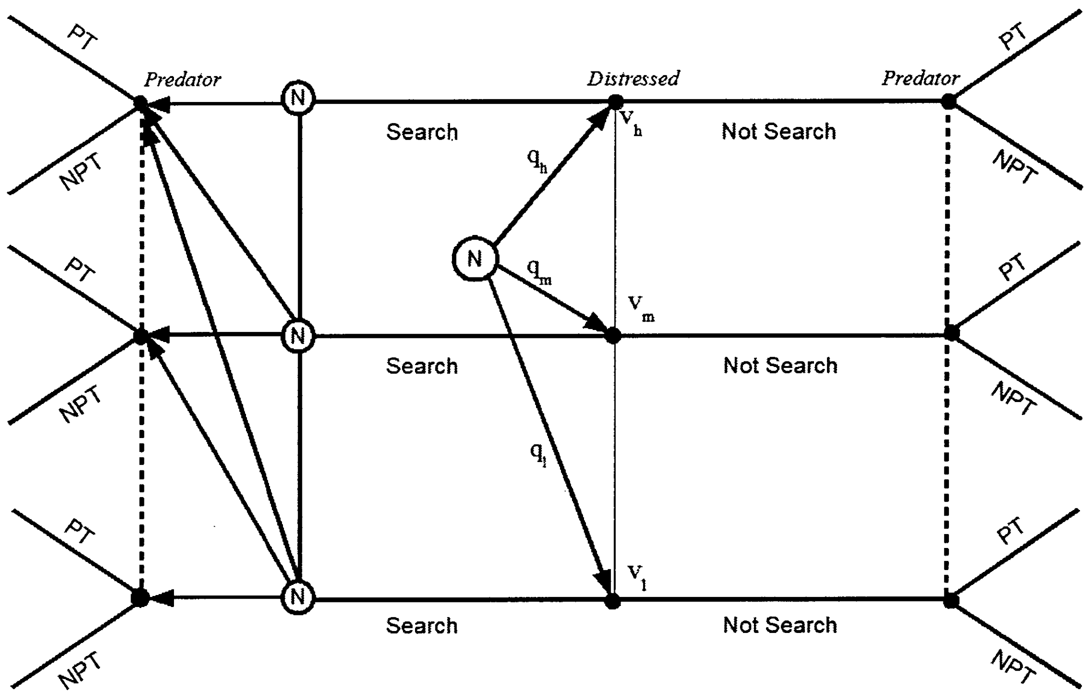

Figure 4.1

This figure therefore illustrates the two main decisions nodes in the game: first, the decision of the distressed to search for a loan, and second, the decision of the predator to predate.

The key decision for the distressed trader occurs in stage 1, the loan seeking stage. In this stage, after observing his initial value (or type), the distressed trader decides whether or not to search for a loan. In making this decision, there are two future risks that the distressed trader faces: predatory trading risk in stage 2 and exogenous income risk in stage 3. Searching for a loan is the distressed trader's only way of potentially protecting itself against these risks. In stage 3, the lower the trader's valuation, $v_{d,3}$, the greater the probability of hitting the lower bound on wealth. For this reason a loan would be desirable. However, the main caveat of searching for a loan is its possible signal value–that is, the potential predator sees whether the distressed searched for a loan, and hence infers some information from this action. Therefore, in deciding whether to search for a loan, a distressed bank faces a trade-off between the financial cushion provided by a loan and the information it conveys.

The key decision for the predator occurs in stage 2. While the predator does not observe the distressed's type, he does observe whether or not the distressed searched for a loan. This is motivated by the following. Financial firms must contact outside lenders, counterparties, or central banks when seeking loans. Although any loan amount received may not be observed by the market, the act of seeking liquidity is likely to become public. Thus, potential predators may be able to infer information from this action. In the next section, I show how the predator forms beliefs about the distressed's type optimally via Bayes rule.

Finally, note that is optimal for each trader to always hold $\bar{x}$ of the risky asset, unless engaged in a predation war. This corner solution is due to the long-term investor's demand curve and to the fact that traders have limited capital, so that the equilibrium price is always lower than the fundamental value.

Assumptions. I make the following assumptions on parameter values. First,

$$v_m < v_p < v_h.$$

That is, if the predator predates in stage 2, he succeeds if $v_{d,2} = v_l$ or v_m, but fails if $v_{d,2} = v_h$. Note that even after receiving a loan, the distressed firm may not be a high type. Thus, this assumption implies that a loan will not always bring the firm into the range where it is not subject to predation risk. Bolton and Scharfstein (1990) show that an optimal financial contract may leave an agent cash constrained even if the agent is subject to predation risk.[11]

Second, I assume that

$$p_h = 1 \quad \text{and} \quad p_l = 0.$$

This simply states that in stage 3, the high type never hits the lower bound on wealth while the low type always hits the lower bound on wealth. This will imply that the distressed's type space has dominance regions. That is, for the two extreme types–the low type and the high type–searching and not searching, respectively, are strictly dominant. The dominant strategies

[11] They consider product market predation, not financial market predation. Furthermore, they do not consider the signal value of searching for liquidity when information is asymmetric.

will be proven and shown in the following section.

Third, I assume that

$$0 < c < \Delta \quad \text{and} \quad 0 < m < \Delta.$$

This assumption states that the gain incurred from successfully predating is less than the loss due to a forced liquidation, and that the fixed cost from obtaining a loan is less than the loss due to a forced liquidation.

Finally, I assume the following condition for the transition probabilities.

$$0 < \pi_{lh} < \pi_{mh}$$

This assumption states that, conditional on searching for a loan, the probability of the medium type becoming a high type is greater than the probability of the low type becoming a high type.

4.3 Equilibrium Definition

Both strategic traders are risk-neutral and expected payoff maximizers. There are two stages in this game where the traders make choices, and the choices each trader makes may affect their final payoff $\tilde{w}_i$. In this section I define the equilibrium in this game and characterize the decision rules for each agent.

Note that in this game, each agent–the distressed and the predator–may face the risk of liquidating prematurely and receiving final payoff $\tilde{w}_i = L$. For the predator, this could be the outcome of the predation war. For the distressed, this could either be the outcome of the predation war, or the outcome of the exogenous income shock in stage 3. In terms of final outcomes of the game, I say that a particular trader "survives" if he is never forced to liquidate. That is, "survival" refers to the event that the trader makes it through the entire game without liquidating and receives the final value of holding its portfolio, $\tilde{w}_i = x_{i,3}z + u$.

In stage 1, the distressed trader, after observing its initial type, $v_{d,1}$, decides whether or not to search for a loan. To make this decision the distressed trader forms beliefs about his survival probability that depend not only on his chosen action and its initial type, but also on the strategy of the predator. Given an initial type $v_{d,1}$, let $\alpha(a_d, v_{d,1}|r_p)$ denote the distressed's belief it survives if it chooses action a_d, conditional on the predator following strategy r_p. Thus, given initial type $v_{d,1}$, the distressed's expected payoff conditional on not searching is given by

$$\mathbb{E}_{d,1}[\tilde{w}_d|NS, v_{d,1}, r_p] = \bar{w}\alpha(NS, v_{d,1}|r_p) + (\bar{w} - \Delta)(1 - \alpha(NS, v_{d,1}|r_p)), \tag{4.3}$$

since he gets expected payoff $\bar{w}$ if he survives, and liquidation value $L = \bar{w} - \Delta$ otherwise. On the other hand, given initial type v_d, the distressed's expected payoff conditional on searching is given by

$$\mathbb{E}_{d,1}[\tilde{w}_d|S, v_{d,1}, r_p] = (\bar{w} - c)\alpha(S, v_{d,1}|r_p) + (\bar{w} - \Delta)(1 - \alpha(S, v_{d,1}|r_p)) \tag{4.4}$$

since he gets expected payoff $\bar{w}$ if he survives, minus the fixed cost of the loan, and liquidation value $L = \bar{w} - \Delta$ otherwise.

Likewise, in stage 2, the predator, after observing the action of the distressed, a_d, decides whether or not to predate. To make this decision the predator forms beliefs about his survival probability that depend not only on his chosen action, but also on the observed action of the distressed and the distressed's strategy. Given an observed action a_d, let $\beta(a_p, a_d|r_d)$ denote the predator's belief he survives if he chooses action a_p, conditional on the distressed following strategy r_d. If the distressed chooses to predate, then the survival probability is merely the posterior probability that $v_{d,2} < v_p$, i.e. $\beta(P, a_d|r_d) = \Pr[v_{d,2} < v_p|a_d, r_p]$. On the other hand, if the predator does not predate, then $\beta(NP, a_d|r_d) = 1$, for any a_d, r_d. Therefore, given observed action a_d, the predator's expected payoff conditional on predatorily trading is given by

$$\mathbb{E}_{p,2}[\tilde{w}_p|P, a_d, r_d] = (\bar{w} + m)\beta(P, a_d|r_d) + (\bar{w} - \Delta)(1 - \beta(P, a_d|r_d)) \tag{4.5}$$

since he gets expected payoff $\bar{w}$ if it survives, plus profits m, and liquidation value $L = \bar{w} - \Delta$ otherwise. On the other hand, given observed action a_d, the predator's expected payoff conditional on not predatorily trading is given by

$$\mathbb{E}_{p,2}[\tilde{w}_p|NP, a_d, r_d] = \bar{w} \tag{4.6}$$

A strategy of the distressed trader is merely a mapping from the distressed's type space to an action, that is $r_d : V_d \to A_d$. A strategy for the predator is merely a mapping from its information set to an action, that is, $r_p : A_d \to A_p$. The equilibrium of this game is then defined as follows.

Definition 42 *An equilibrium is a strategy for the distressed $r_d : V_d \to A_d$, a strategy for the predator $r_p : A_d \to A_p$, and a belief system $\alpha : A_d \times V_d \to [0,1]$ and $\beta : A_p \times A_d \to [0,1]$*

(i)For each $v_{d,1} \in V_d$, the distressed of initial type v_d searches for a loan if and only if his expected payoff from doing so is greater than his expected payoff from not searching

$$r_d\left(v_{d,1}\right) = S \quad \textit{if and only if} \;\; \mathbb{E}_{d,1}\left[\tilde{w}_d | S, v_{d,1}, r_p\right] > \mathbb{E}_{d,1}\left[\tilde{w}_d | NS, v_{d,1}, r_p\right], \tag{4.7}$$

conditional on the predator following strategy r_p.

(ii)For each $a_d \in A_d$, the predator who observes a_d predates if and only if his expected payoff from doing so is greater than his expected payoff from not predating

$$r_p\left(a_d\right) = P \quad \textit{if and only if} \;\; \mathbb{E}_{p,2}\left[\tilde{w}_p | P, a_d, r_d\right] > \mathbb{E}_{p,2}\left[\tilde{w}_p | NP, a_d, r_d\right] \tag{4.8}$$

conditional on the distressed following strategy r_d

(iii) The survival belief of the distressed, α, is based on the predator following strategy r_p, and the survival belief of the predator β is formed using Bayes rule and based on the distressed following strategy r_d.

Conditions (i) and (ii) of Definition 42 require that the strategies of the distressed and the predator are sequentially rational given their beliefs. Condition (iii) states that the belief system must be consistent given the strategy profile of the players. Thus, the equilibrium definition is that of a standard perfect-Bayesian equilibrium, in which the distressed is the sender and the predator is the receiver. Finally, I prove shortly that in this game there are no out-of-equilibrium beliefs.

Decision rule for the distressed trader. I first consider the decision for the distressed trader in stage 1. The expected payoffs for the distressed from searching and from not searching are given in (4.4) and (4.3), respectively. Combining these with the distressed's decision rule stated in (4.7), it follows that optimal action for the distressed trader may be expressed as follows.

Lemma 43 *Given initial valuation $v_{d,1}$ and conditional on the predator following strategy r_p, the distressed trader searches for a loan if and only if*

$$\frac{\alpha\left(NS, v_{d,1}|r_p\right)}{\alpha\left(S, v_{d,1}|r_p\right)} < \frac{\Delta - c}{\Delta} \tag{4.9}$$

where $\frac{\Delta - c}{\Delta} \in (0, 1)$.

The above Lemma gives a simple cut-off rule, in terms of the distressed's beliefs, for when it is optimal for the distressed to search for a loan. To see this, note that the left-hand side of equation (4.9) is merely the ratio of the probability of survival from not searching to the probability of survival from searching. Lemma 43 states that the distressed trader will search if and only if this ratio is sufficiently low. The cut-off for this ratio, i.e. the right-hand side of equation (4.9), is a constant which is decreasing in the fixed cost of obtaining a loan, c, but increasing in the liquidation penalty Δ. Therefore, the distressed finds it optimal to search if the probability of surviving from not searching is low relative to that of searching. That is, searching is optimal if it greatly increases the distressed's chances of survival. However, a higher fixed cost of the loan or a lower liquidation penalty makes it less likely that the distressed will find it optimal to search.

Lemma 43 gives a simple decision rule for the distressed trader, given his initial type. Using this decision rule, it is now clear that for the two extreme types–the low type v_d and the high type v_h–searching and not searching, respectively, are strictly dominant. This is stated in the following lemma.

Lemma 44 *For any strategy of the predator, the low type always finds it optimal to search. Likewise, for any strategy of the predator, the high type always finds it optimal to not search.*

The proof of Lemma 44 is simple. Consider first the low-type's decision. For any strategy of the predator, if the low type decides not to search, his probability of survival is zero, since no matter what happens in stage 2, the exogenous income shock in stage 3 will force the firm to liquidate ($p_l = 0$). On the other hand, for any strategy of the predator if the low type decides to search for a loan, his probability of survival is strictly positive. Condition (4.9) is hence satisfied for all r_p. Therefore, no matter the strategy of the predator, the low type always finds it optimal to search.

Second, consider the decision of the high type in stage 1. For any strategy of the predator, the high type's probability of survival, whether it searches or not, is always equal to 1. This is due to the fact that neither the predator in stage 2 nor the income shock in stage 3 can force the high type to liquidate. Therefore, according to condition (4.9), the high type always finds it optimal to not search.

Lemma 44 clarifies the type of equilibria that may exist in this game. The property that the low type always searches and the high type never searches, i.e. that there are dominance regions in the type space, implies that any possible equilibrium in this game must be a separating (or semi-separating) equilibrium. Any action observed by the predator is consistent with an equilibrium path, and hence no off-the-equilibrium path beliefs need be specified.

Finally, note that Lemma 44 also contributes to understanding the signalling nature in this game. Because the predator does not observe the type of the distressed, it can only infer information from the distressed's action. From Lemma 44, we see that regardless of the predator's strategy, it is strictly dominant for the low type to search, and it is strictly dominant for the high type to not search. Therefore, when the medium type makes its decision whether or not to search, part of its trade-off is whether to be pooled with the low types or to be pooled with the high types, and in this way convey information to the predator.

Decision rule of the predator. I now consider the decision for the predator in stage 2. The expected payoffs for the predator from trading and from not trading are given in (4.5) and (4.6), respectively. Combining these with the predator's decision rule stated in (4.8), it follows that optimal action for the predator may be expressed as follows.

Lemma 45 *Given observed action a_d and conditional on the distressed following strategy r_d,* *the predator predatorily trades if and only if*

$$\frac{1-\beta\left(P, a_d | r_d\right)}{\beta\left(P, a_d | r_d\right)} < \frac{m}{\Delta} \tag{4.10}$$

where $\frac{m}{\Delta} \in (0,1)$.

Much like Lemma 43, Lemma 45 gives a simple cut-off rule, in terms of the predator's beliefs, for when it is optimal for the predator to predatorily trade. To see this, note that the left-hand side of equation (4.9) is merely the ratio of the probability of failing to the probability of succeeding, if the trader decides to predatorily trade. Lemma 45 states that the predator will predate if and only if this ratio is sufficiently low. The cut-off for this ratio, i.e. the right-hand side of equation (4.9), is a constant which is increasing in the gain from winning a predation war, m, but decreasing in the liquidation penalty Δ. Therefore, the predator finds it optimal to predate if the probability of surviving conditional on predatorily trading is high enough. However, a lower gain from predatorily trading or a higher liquidation penalty makes it less likely that the predator will find it optimal to predate.

The decision rules stated in Lemmas 43 and 45 greatly simplify the equilibrium analysis. In what follows, in Section 4.4 I first look at the equilibrium when there is no predator, and then in Section 4.5 I turn to the equilibrium of the full game with predation.

4.4 Equilibrium without Predator

In this section I analyze the equilibrium in a benchmark case in which there is no predator. That is, I consider a setting identical to that described in Section 4.2, but without stage 2, the predatory stage. Within this predator-less setting, I need only to consider the optimal strategy of the distressed.

In this environment, Lemma 44 continues to hold; that is, it is optimal for the low types to search and for the high types to not search. Thus, in terms of the distressed's strategy, one needs only to find what is optimal for the medium type. Although there is no predation risk, the medium type still faces income shock risk. Hence if the medium type decides not to search, his probability of survival is given by $\alpha\left(NS, v_m\right) = p_m$. On the other hand, if the medium decides to search for a loan, there is some probability he becomes a high type; for this reason the probability of surviving is given by $\alpha(S, v_m) = \pi_{mm}p_m + \pi_{mh}$. Combining these probabilities with the decision rule in (4.9), I find that the medium-type searches if and only if

$$\frac{p_m}{\pi_{mm}p_m + \pi_{mh}} < \frac{\Delta - c}{\Delta}$$

Therefore, depending on parameter values, the medium type could find either choice optimal in the benchmark with no predator.

For the remainder of this chapter, I focus on the case where the medium type searches for a loan in the absence of predators. In other words, I impose the following condition.

Condition 46 $\frac{p_m}{\pi_{mm}p_m+\pi_{mh}} < \frac{\Delta-c}{\Delta}$

The above condition is imposed for the following reason. The focus of this chapter is to study the effect of predatory trading on the incentives of financial firms to seek liquidity in times of distress. For this reason, it seems reasonable to consider an environment in which there is a clear incentive for a bank to seek out a loan in the absence of predators. Condition 46 imposes that in the absence of predation risk, searching for a loan greatly increases the distressed's chances of avoiding the lower wealth bound, making it preferable for him search. Another way to read this condition is that the liquidation penalty Δ is sufficiently high relative to the cost of a loan that the medium type prefers to search. In any case there is a clear incentive for the medium type to seek out a loan when there are no predators.[12]

Using this condition, the following proposition characterizes the optimal strategy for the distressed in the benchmark with no predatory trading.

Proposition 47 *When there are no predators, under condition 46, the distressed follows a strategy in which the low and medium types search for a loan and the high type does not search.*

[12] Another way to justify this condition is to imagine there were a continuum of types. Then there would exist a type, strictly greater than the low type that would search.

4.5 Equilibrium with Predator

I now study the equilibrium (or equilibria) of the full game with predatory trading, as laid out in Section 4.2. As this is a signalling game, there can in principle be multiple equilibria. In order to characterize the set of all possible equilibria, I consider the entire set of possible strategies of one of the traders. Here I choose to focus on the set of strategies of the predator. For each of the predator's strategies, I determine under what conditions the strategy may be compatible with an equilibrium. By systematically considering each of the predator's strategies, this procedure allows me to characterize the set of all possible equilibria.

A strategy of the predator is merely a mapping from its information set to an action, $r_p : A_d \rightarrow A_p$. Since A_d and A_p each contain only two elements, there are only four possible strategies for the predator. I label these strategies as $\{r_p^1, r_p^2, r_p^3, r_p^4\}$ and describe each as follows:

- Let r_p^1 be strategy in which the predator never predates

$$r_p^1(a_d) \equiv \begin{cases} NP & \text{if} \quad a_d = S \\ NP & \text{if} \quad a_d = NS \end{cases}$$

- Let r_p^2 be strategy in which the predator always predates

$$r_p^2(a_d) \equiv \begin{cases} P & \text{if} \quad a_d = S \\ P & \text{if} \quad a_d = NS \end{cases}$$

- Let r_p^3 be strategy in which the predator predates if and only if he observes that the distressed did not search.

$$r_p^3(a_d) \equiv \begin{cases} NP & \text{if} \quad a_d = S \\ P & \text{if} \quad a_d = NS \end{cases}$$

- Finally, let r_p^4 be strategy in which the predator predates if and only if he observes that the distressed did search.

$$r_p^4\left(a_d\right) \equiv \begin{cases} P & \text{if} \quad a_d = S \\ NP & \text{if} \quad a_d = NS \end{cases}$$

In this section, I consider each strategy $r_p \in R_p \equiv \left\{r_p^1, r_p^2, r_p^3, r_p^4\right\}$ separately. For a given proposed strategy r_p of the predator, I find the best response of the distressed trader. That is, I find the survival probabilities of the distressed based on the belief that the predator is following strategy r_p, and given these survival beliefs I find the optimal strategy of the distressed. Allow me to denote this strategy as $r_d' = BR_d\left(r_p\right)$, where BR_d signifies that is is the best response of the distressed. Next, given the distressed's best response strategy r_d', I then find the best response of the predator. That is, I find the survival probabilities of the predator based on the belief that the distressed is following strategy r_d', and then find the optimal strategy of the predator based on these survival beliefs. Allow me to denote this strategy as $r_p' = BR_p\left(r_d'\right)$, where BR_p signifies that is is the best response of the predator.

I then characterize under what conditions the predator's best response strategy coincides with the proposed strategy. If $r_p' = r_p$, then there exists a fixed point in the traders' best responses: $r_d' = BR\left(r_p'\right)$ and $r_p' = BR\left(r_d'\right)$. In this case, the strategy profile $\left\{r_p', r_d'\right\}$ and corresponding survival beliefs therefore constitute an equilibrium. On the other hand, if under no conditions the fixed point exists, I then conclude that an equilibrium in which the predator follows the proposed strategy does not exist.

I now consider each strategy.

The predator follows strategy r_p^1. Suppose the predator follows a strategy in which he never predates.

Distressed trader's best response. If the predator follows a strategy in which he never predates, in terms of the distressed trader's decision making process it is as if the predator did not exist. In other words, the distressed never faces any predation risk. The distressed trader will thus follow the same strategy outline above in Section 4.4 for the benchmark case in which there is no predator. Proposition 47 implies that the distressed's best response to r_p^1 is a strategy in which the low and medium types search, while the high type does not search.

Predatory trader's best response. The predator forms its survival beliefs based on the presumption that the distressed is following a strategy in which the low and medium types search, while the high type does not search. If the predator observes that the distressed did not search, the predator infers that he must be facing a high type. Hence, the predator knows that if he predates, his probability of surviving is zero. On the other hand, if he does not predatorily trade he receives a final payoff of $\bar{w}$. Therefore, if the predator observes that the distressed did not search, it is optimal for the predator to not predate.

On the other hand, if the predator observes that the distressed searched for a loan, then the probability of the predator surviving a predation war is given by

$$\beta\left(P, S|r_d'\right) = \frac{\left(\pi_{ll} + \pi_{lm}\right) q_l + \pi_{mm} q_m}{q_l + q_m}$$

This is simply the probability that the distressed's wealth after seeking a loan is less than the wealth of the predator. Combining this with the predator's decision rule in (4.10), the predator finds it optimal to predate if and only if

$$\kappa_1 < \frac{m}{\Delta}$$

where

$$\kappa_1 \equiv \frac{\left(1 - \pi_{ll} - \pi_{lm}\right) + \left(1 - \pi_{mm}\right)}{\pi_{ll} + \pi_{lm} + \pi_{mm}} \tag{4.11}$$

where I have used the fact that $q_l = q_m$. Therefore the proposed equilibrium in which the predator never predates exists if and only if the above condition is not satisfied, that is, when $\frac{m}{\Delta} < \kappa_1$.

The predator follows strategy r_p^2. Suppose the predator follows a strategy in which he always predates.

Distressed trader's best response. From Lemma 44, we know that the low type chooses to search and the high type chooses to not search. Now consider the optimal choice of the medium type. Given that the predator is following a strategy in which he always predates, if the medium type chooses to not search, then he will be engaged in a predation war which he will surely lose, since $v_{d,2} = v_m < v_p$. Hence, the medium type's probability of survival from not searching, $\alpha\left(NS, v_m | r_p^2\right)$, is equal to zero. On the other hand, if the medium type chooses to search, then he still faces a predation war. In this case, however, there is some positive probability that the distressed receives a loan that makes it a high type and hence wins the predation war. Therefore, the medium type's probability of survival from searching for a loan is given by $\alpha\left(S, v_m | r_p^2\right) = \pi_{mh}$, which is strictly greater than zero. According to the distressed's decision rule (4.9), it is optimal for the medium type to search. The distressed's best response to r_p^2 is therefore a strategy in which the low and medium types search, while the high type does not search.

Predatory trader's best response. The predator forms its survival beliefs based on the presumption that the distressed is following a strategy in which the low and medium types search, while the high type does not search. If the predator observes that the distressed did not search, the predator infers that he must be facing a high type. Hence, the predator knows that if he predates, his probability of surviving is zero. On the other hand, if he does not predatorily trade he receives a final payoff of $\bar{w}$. Therefore, if the predator observes that the distressed did not search, it is optimal for the predator to not predate. Thus, given the strategy of the distressed, under no conditions does the best response of the predator coincide with r_p^2. Therefore, no equilibrium exists in which the predator follows a strategy in which he always predates.

The predator follows strategy r_p^3. Suppose the predator follows a strategy in which he predates if and only if the distressed does not search.

Distressed trader's best response. The low type searches and the high type does not. The medium type forms his beliefs based on the strategy of the predator. Given that the predator is following a strategy in which it predates if and only if the distressed does not search, if the medium type chooses to not search, then he will be engaged in a predation war which he will surely lose, since $v_{d,2} = v_m < v_p$. Hence, the medium type's probability of survival from not searching, $\alpha\left(NS, v_m | r_p^3\right)$, is equal to zero. On the other hand, given the predator's strategy, if the medium type chooses to search, then he will not face any predation risk and the only risk he faces is the exogenous income shock in stage 3. In this case, his probability of survival is given by $\alpha\left(S, v_m | r_p^3\right) = \pi_{mm} p_m + \pi_{mh}$, which is strictly greater than zero. According to the distressed's decision rule (4.9), it is optimal for the medium type to search. The distressed's best response to r_p^3 is therefore a strategy in which the low and medium types search, while the high type does not search.

Predatory trader's best response. The predator forms its survival beliefs based on the presumption that the distressed is following a strategy in which the low and medium types search, while the high type does not search. If the predator observes that the distressed did not search, the predator infers that he must be facing a high type. Hence, the predator knows that if he predates, his probability of surviving is zero. On the other hand, if he does not predatorily trade he receives a final payoff of $\bar{w}$. Therefore, if the predator observes that the distressed did not search, it is optimal for the predator to not predate. Thus, given the strategy of the distressed, under no conditions does the best response of the predator coincide with r_p^3. Therefore, no equilibrium exists in which the predator follows a strategy in which he predates if and only if the distressed does not search.

The predator follows strategy r_p^4. Finally suppose the predator follows a strategy in which he predates if and only if the distressed searches.

Distressed trader's best response. The low type searches and the high type does not. The medium type forms his beliefs based on the strategy of the predator. Given that the predator is following a strategy in which he predates if and only if the distressed searches, if the medium type chooses to not search, then he will not face any predation risk and the only risk he faces is the exogenous income shock in stage 3. Hence, the medium type's probability of survival from not searching is given by $\alpha\left(NS, v_m | r_{p,4}\right) = p_m$. On the other hand, given the predator's strategy, if the medium type chooses to search, then he will be engaged in a predation war with the predator in stage 2, which the distressed will lose if he is still a medium type, but will win if he receives a loan that makes him a high type. Therefore, the medium type's probability of survival from searching for a loan is given by $\alpha\left(S, v_m | r_{p,4}\right) = \pi_{mh}$. According to the distressed's decision rule (4.9), the medium type searches if and only if

$$\frac{p_m}{\pi_{mh}} < \frac{\Delta - c}{\Delta}$$

Therefore, depending on parameter values, the medium type could find either choice optimal. If $\frac{p_m}{\pi_{mh}} < \frac{\Delta - c}{\Delta}$, then conditional on the predator's strategy it is optimal for the medium type to search. In this case, the distressed's best response to r_p^4 is a strategy in which the low and medium types search, while the high type does not search. On the other hand, if $\frac{p_m}{\pi_{mh}} > \frac{\Delta - c}{\Delta}$, then conditional on the predator's strategy, it is optimal for the medium type to not search. In this case, the distressed's best response to r_p^4 is a strategy in which the low type searches for a loan, and the medium and high types do not.

Predatory trader's best response. To characterize the best response of the predator, I consider separately the two cases: first, the case in which $\frac{p_m}{\pi_{mh}} < \frac{\Delta - c}{\Delta}$ and second the case in which $\frac{p_m}{\pi_{mh}} > \frac{\Delta - c}{\Delta}$.

First, suppose that $\frac{p_m}{\pi_{mh}} < \frac{\Delta - c}{\Delta}$. In this case, the predator forms his survival beliefs based on the presumption that the distressed is following a strategy in which the low and medium types search, while the high type does not search. Note that this strategy of the distressed is identical to the distressed's best response to r_p^1. Using the findings of that discussion, one may infer that if the predator observes that the distressed does not search, it is optimal for the predator to not predate. On the other hand, if the predator observes that the distressed does search for a loan, then the predator finds it optimal to predate if and only if $\kappa_1 < \frac{m}{\Delta}$, where κ_1 is given in (4.11). In this case, the predator's best response coincides with $r_{p,}^4$. Therefore the proposed equilibrium in which the predator predatorily trades if and only if the distressed searches exists whenever $\kappa_1 < \frac{m}{\Delta}$ and $\frac{p_m}{\pi_{mh}} < \frac{\Delta - c}{\Delta}$.

Second, suppose that $\frac{p_m}{\pi_{mh}} > \frac{\Delta - c}{\Delta}$. In this case, the predator forms his survival beliefs based on the presumption that the distressed is following a strategy in which the low type searches for a loan, but the medium and high types do not. If the predator observes that the distressed does not search, he infers that it must be facing either a medium or high type. In this case, the probability of the predator surviving a predation war is given by $\beta\left(P, NS|r_d'\right) = \frac{q_m}{q_m + q_h}$. Combining this with the predator's decision rule in (4.10), I find that it is optimal for the predator to not predate. On the other hand, if the predator observes that the distressed does search for a loan, then the probability of the predator surviving a predation war is given by $\beta\left(P, S|r_d'\right) = \pi_{ll} + \pi_{lm}$. Combining this with the predator's decision rule in (4.10), the predator finds it optimal to predate if and only if

$$\kappa_2 < \frac{m}{\Delta}$$

where[13]

$$\kappa_2 \equiv \frac{\pi_{lh}}{\pi_{ll} + \pi_{lm}} \tag{4.12}$$

Therefore the proposed equilibrium in which the predator predatorily trades if and only if the distressed searches exists whenever $\kappa_2 < \frac{m}{\Delta}$ and $\frac{p_m}{\pi_{mh}} > \frac{\Delta - c}{\Delta}$.

4.5.1 Equilibrium Characterization

Given the above analysis, I first state the following non-existence result.

Lemma 48 *No equilibrium exists in which the predator follows a strategy in which he always predates. Moreover, no equilibrium exists in which the predator follows a strategy in which he predates if and only if he observes that the distressed did not search.*

This lemma is useful in that it implies that in any equilibrium, the predator is either playing a strategy in which he never predates, or one in which he predates if and only if he observes that the distressed searched for a loan. In the former case, the equilibrium of the game will be similar to that in the benchmark with no predator. In the latter case, the fact that the predator predates if and only if he observes that the distressed searched for a loan, implies that the presence of predators creates an incentive for the distressed to refrain from searching.

I now characterize the set of all possible equilibria in this game in Propositions 49 and 50. In Proposition 49, I first consider the case in which the ratio p_m/π_{mh} is low, that is, when the medium type's probability of surviving the exogenous income shock is low relative to the transition probability of becoming a high type after searching for a loan.

Proposition 49 *Suppose* $\frac{p_m}{\pi_{mh}} < \frac{\Delta - c}{\Delta}$.

(i) if $\kappa_1 < \frac{m}{\Delta}$, *then there exists a unique equilibrium such that the distressed trader follows a strategy in which the low and medium types search for a loan while the high type does not search, and the predator follows a strategy in which he predates if and only if the distressed searches.*

(ii) if $\frac{m}{\Delta} < \kappa_1$, *then there exists a unique equilibrium such that the distressed trader follows a strategy in which the low and medium types search for a loan while the high type does not search, and the predator follows a strategy in which he never predatorily trades.*

[13]Note that, given the assumptions on the parameter values, $\kappa_2 < \kappa_1$.

Thus, when the ratio p_m/π_{mh} is sufficiently low, the distressed trader always behaves as he does in the benchmark with no predator. That is, in any equilibrium the distressed follows a strategy in which the low and medium types search, and the high type does not search. Note that when $\frac{m}{\Delta} < \kappa_1$, the equilibrium is almost identical to the benchmark without predation, except that in this case it is optimal for the predator to never predatorily trade.

On the other hand, when $\kappa_1 < \frac{m}{\Delta}$, the predator predatorily trades if and only if he observes that the distressed searched. When the predator observes that the distressed searched, he knows that the distressed must be initially a low or medium type. Using Bayes rule, the predator can then form posterior beliefs over the probability of winning in a predation war. Under these beliefs, it is optimal for the predator to predate. The key insight here is that by searching for a loan, the distressed is signalling that he is either a low or medium type. Hence by taking measures to increase its financial viability, the distressed is in effect signalling its financial weakness.

Furthermore, note that the predator's strategy provides an incentive for the distressed to refrain from searching. However, the increase in survival probability the medium type gains from searching is high enough to compensate for the increased predation risk. Thus, the medium type still finds it optimal to search.

I now consider the case in which the ratio p_m/π_{mh} is high, that is, when the medium type's probability of surviving the exogenous income shock is high relative to the transition probability of becoming a high type after searching for a loan.

Proposition 50 *Suppose* $\frac{p_m}{\pi_{mh}} > \frac{\Delta - c}{\Delta}$.

(i) if $\kappa_2 < \kappa_1 < \frac{m}{\Delta}$, *then there exists a unique equilibrium such that the distressed trader follows a strategy in which the low type searches for a loan while the medium and high types do not search, and the predator follows a strategy in which he predates if and only if the distressed searches.*

(ii) if $\kappa_2 < \frac{m}{\Delta} < \kappa_1$, *there exist two (pure-strategy) equilibria:*[14]

[14] *Of course, there also exists a third equilibria in mixed strategies. In this equilibrium, the distressed strategy will be one in which the low type searches, the high type does not search, and the medium type randomizes between searching and not searching, while the predator's strategy is one in which it randomizes between predatorily trading and not predatorily trading.*

(a) the distressed trader follows a strategy in which the low type searches for a loan while the medium and high types do not search, and the predator follows a strategy in which he predates if and only if the distressed searches.

(b) the distressed trader follows a strategy in which the low and medium types search for a loan while the high type does not search, and the predator follows a strategy in which he never predatorily trades.

(iii) if $\frac{m}{\Delta} < \kappa_2 < \kappa_1$, *then there exists a unique equilibrium such that the distressed trader follows a strategy in which the low and medium types search for a loan while the high type does not search, and the predator follows a strategy in which he never predatorily trades.*

Therefore, when the ratio p_m/π_{mh} is sufficiently high, under certain parameter values there exists an equilibrium in which the distressed trader follows a strategy such that the low type searches while the medium and high types do not search, and the predator follows a strategy in which it predates if and only if the distressed searches. This equilibrium is interesting because the medium type finds it optimal to not search, and hence plays a different action than he would in the benchmark without predation.

In this equilibrium, the predator finds it optimal to predate if he observes that the distressed searched. Of course, this is because the predator knows that only the low types search, and hence the predator has a high chance of winning in a predation war. Thus, even though the predator cannot directly observe types, searching for a loan is a strong signal that the distressed has a very weak financial status. For this reason the predator finds it optimal to predate.

The predator's equilibrium strategy provides a strong incentive for the distressed to refrain from searching. Consider the decision of the medium type distressed trader. If the medium type chooses to not search, then he will not face any predation risk; the only risk he faces is in the exogenous income shock in stage 3. On the other hand, if the medium type chooses to search for a loan, he then engages in a predation war in stage 2, which he can win only if he receives a large enough loan to become a high type. Therefore, when the probability of surviving the income shock as a medium type is high relative to the transition probability of becoming a high type after searching for a loan, that is, when the ratio p_m/π_{mh} is sufficiently high, then the medium type finds it optimal to not search. In other words, the medium type prefers to pool himself with high types by not searching and consequently facing greater income risk, over pooling himself with low types but consequently facing predation risk. One can think of this as a financially weak firm that tries to ride out a temporary financial shortfall on its own, without signalling any weakness to predators by seeking outside liquidity.

This equilibrium clearly illustrates how predator trading may adversely affect the incentives of banks to seek loans in times of financial distress. In the benchmark without predation, the medium-type distressed trader searches for a loan in order to protect itself against exogenous income risk. However, when there are predators who cannot directly observe the wealth of traders, actions undertaken by these traders to relieve financial distress may convey information about their underlying financial state. For this reason, predators have the incentive to predate when they see a large trader searching for a loan. Thus, in deciding whether or not to search for a loan, the distressed firms face a trade-off between the financial cushion provided by a loan and the information this act reveals. There are equilibria in which medium-type distressed funds who would otherwise seek to recapitalize may be reluctant to search for loans in the presence of predators.

Finally, this analysis may have further implications for regulation and policy. Policies that may break the separating equilibria found in this analysis, i.e. lead to pooling equilibria, may weaken the adverse signal value of searching for funds. For example, an interesting policy to consider would be one in which the government forces all firms to obtain a loan. If all traders–low types, medium types, and high types–are forced to take a loan, then the predator cannot use this information to infer underlying financial states. However, under this policy one would need to then consider the cost of saving funds or banks that may in fact be insolvent.

Another policy to consider is one which increases the probability of successful search–that is, a policy that increases π_{mh}. This policy would then shift the equilibrium from the one in Proposition 50 in which the medium type does not search for a loan, to an equilibrium in Proposition 49 in which the medium type does find it optimal to search. In this case, the distressed trader would act exactly as it would in the benchmark without predation. This would then be an interesting new perspective on central bank policy during crises episodes–as opposed to the lender of last resort rationale, the central bank may be important as an institution that improves and facilitates inter-bank lending.

4.6 Concluding remarks

this chapter analyzes how predatory trading may affect the incentives of banks to seek loans in times of financial distress. I find that when a distressed trader is more informed than other traders about its own balances, searching for extra capital from lenders can become a signal of financial need, thereby opening the door for predatory trading and possible insolvency. I find equilibria in which some distressed traders who would like borrow short-term in order to meet temporary liquidity needs, may be reluctant to do so in the presence of potential predators. Predatory trading may therefore deter banks and financial institutions from raising funds in times when they need it the most.

4.7 Appendix: The Predation War

In this appendix I provide a more detailed analysis of the predation war discussed in Section 4.2. In stage 2, if the predator decides to predatorily trade, then the strategic traders engage in a "predation war". The results of this predation war are derived from Brunnermeier and Pedersen (2005).

Suppose that time is continuous within this stage and denoted by $\tau \in [0, \bar{\tau}]$. That is, traders may now trade continuously in the asset. Let $x_i(\tau)$ denote the position of trader i in the asset at time τ, and let $s(\tau)$ denote the price of that asset. At the beginning of stage 2, each trader has an initial position , $x_i(0) = \bar{x}$, of the risky asset. Within stage 2 the trader can now continuously trade the asset by choosing his trading intensity, $a_i(\tau)$. Hence, at time τ the trader's position in the risky asset is given by

$$x_i(\tau) = x_i(0) + \int_0^{\tau} a_i(u)du$$

As mentioned previously, each strategic trader is restricted to hold $x_i(\tau) \in [-\bar{x}, \bar{x}]$. Finally, I consider the case of limited capital, such that $2\bar{x} < Q$. In addition to the two large strategic traders, the market is populated by long-term investors, whose aggregate demand curve is given in (4.1).

Furthermore, it is assumed that traders cannot sell infinitely fast. Strategic traders can as a whole can trade at most $A \in R$ shares per time unit at the current price. That is, at any moment τ, the aggregate amount of trading must satisfy $a_d(\tau) + a_p(\tau) \leq A$. This implies that if both strategic traders are trading, the greatest intensity at which each trader may trade is $A/2$.[15]

Trader i's within-stage objective is to maximize his expected wealth at the end of the stage. His earnings from investing in the risky asset is given by the final value of his stock holdings, $x_i(\bar{\tau})\, \dot{z}$, minus the cost of buying shares. That is, a strategic trader's objective is to choose a

[15] Brunnermeier and Pedersen (2005) assume that strategic traders can as a whole trade at most $A \in R$ shares per time unit at the current price. Rather than simply assuming that orders beyond A cannot be executed, they assume that traders suffer temporary impact costs if orders exceed this bound. They then show that it is optimal for traders to trade as fast as possible without incurring this cost.

trading process so as to maximize

$$\max E\left[x_i(\bar{\tau})\,z+\int_0^{\bar{\tau}} a_i(\tau)s(\tau)\,d\tau\right]$$

Due to limited capital of the strategic traders, $s(\tau)<\bar{z}$ at any time, and hence, any optimal trading strategy satisfies $x_i(\bar{\tau})=\bar{x}$ for the surviving trader. That is, any surviving trader ends up with the maximum capital in the arbitrage position. The qualitative results presented in this appendix will then depend on the following: (i) strategic traders have limited capital, that is $2\bar{x}<Q$, otherwise $s(\tau)=\bar{z}$ and (ii) markets are illiquid in the sense that large trades move prices $(\lambda>0)$ and traders cannot trade arbitrarily fast $(A<\infty)$.

Let τ_d and τ_p denote the amount of time it takes for the distressed trader and the predator to hit their lower bounds on wealth, respectively, if both were trading simultaneously at their highest intensity. That is,

$$\tau_d\equiv\frac{w_{d,2}(0)-\underline{w}}{A/2}\qquad\text{and}\qquad\tau_p\equiv\frac{w_{p,2}(0)-\underline{w}}{A/2}$$

where $w_{d,2}=\bar{x}s(0)+v_{d,2}$ and $w_{p,2}=\bar{x}s(0)+v_p$. In equilibrium, both traders sell as fast as possible until one of the traders is forced to leave the market. Specifically, both traders trade at constant speed $-A/2$ from from time 0 to time τ^*, where $\tau^*\equiv\min\{\tau_d,\tau_p\}$. Therefore, the pivotal time τ^* is determined by the wealth of the trader who is closest to the threshold; in other words, the trader who begins the period with lower wealth (i.e. the lower v) is the trader who is forced to leave the market. I assume that $\underline{w}$ is high enough such that at least one trader hits the lower bound.

More precisely, the trader which is forced to leave the market trades according to the following process

$$a_i(\tau)=\begin{cases}-A/2 & \text{for}\quad \tau\in[0,\tau^*]\\ 0 & \text{for}\quad \tau\geq\tau^*\end{cases}$$

While the surviving trader trades according to the following process

$$a_i(\tau) = \begin{cases} -A/2 & \text{for} \quad \tau \in [0, \tau^*] \\ A & \text{for} \quad \tau \in \left[\tau^*, \tau^* + \frac{\bar{x}-x(\tau^*)}{A}\right] \\ 0 & \text{for} \quad \tau \geq \tau^* + \frac{\bar{x}-x(\tau^*)}{A} \end{cases}$$

Thus, both traders trade as fast as they can at constant speed $-A/2$ for τ^* periods, at which point one trader is forced to leave the market. This liquidation strategy is known by both strategic traders. At time τ^*, the surviving trader then buys at a constant rate back up to the original arbitrage position $\bar{x}$; this takes $\frac{\bar{x}-x(\tau^*)}{A}$ periods. From then on, the surviving trader remains in this position.

Finally, the equilibrium price follows the following trajectory

$$s(\tau) = \begin{cases} s(0) - \lambda A \tau & \text{for} \quad \tau \in [0, \tau^*] \\ s(0) - \lambda 2\bar{x} + \lambda A (\tau - \tau^*) & \text{for} \quad \tau \in \left[\tau^*, \tau^* + \frac{\bar{x}-x(\tau^*)}{A}\right] \\ \bar{z} - \lambda(\bar{x} - Q) & \text{for} \quad \tau \geq \tau^* + \frac{\bar{x}-x(\tau^*)}{A} \end{cases}$$

The simultaneous selling by both strategic traders leads to price "overshooting." This implies that the surviving trader may yield a gain from winning the predation war. This gain is given by

$$m = \int_0^{\overline{T}} a_i(\tau) s(\tau)\, d\tau.$$

This is because the surviving trader sells his assets for an average price that is higher than the price at which he buys them back after the other trader has left the market. Therefore, the predator has an incentive to predate in order to profit from the price swings that occur in the wake of the liquidation. Furthermore, the overshooting price due to simultaneous selling makes liquidation excessively costly for the trader who is ultimately forced to leave the market.

Bibliography

[1] Acharya, Viral, Denis Gromb, and Tanju Yorulmazer, (2009). “Imperfect Competition in the Interbank Market for Liquidity as a Rationale for Central Banking,” NYU mimeo.

[2] Adam, Klaus (2007), “Optimal monetary policy with imperfect common knowledge?,” Journal of Monetary Economics 54, 276-301.

[3] Allen, Franklin, and Douglas Gale (1992). “Stock-price manipulation,” Review of Financial Studies 5, 503–529.

[4] Allen, Franklin, and Douglas Gale (2000). “Financial Contagion,” Journal of Political Economy, 108, 1-31.

[5] Alvarez, Fernando and Robert Shimer (2008), "Search and Rest Unemployment," NBER Working Papers 13772.

[6] Amador, Manuel, and Pierre-Olivier Weill (2007), “Learning from Private and Public Observations of Others' Actions,” Stanford University/UCLA mimeo.

[7] Amador, Manuel, and Pierre-Olivier Weill (2008), “Learning from Prices: Public Communication and Welfare,” Stanford University/UCLA mimeo.

[8] Amato, Jeffery, and Hyun Song Shin (2006), “Imperfect Common Knowledge and the Information Value of Prices,” Economic Theory 27, 213-241.

[9] Angeletos, George-Marios, and Jennifer La'O (2009a), “Noisy Business Cycles,” In D. Acemoglu, K. Rogoff, and M. Woodford (eds.), NBER Macroeconomics Annual 2009, Volume 24, Forthcoming.

[10] Angeletos, George-Marios, and Jennifer La'O (2009b), "Incomplete Information, Higher-Order Beliefs, and Price Inertia," Journal of Monetary Economics, Volume 56:S, pp. 19-37

[11] Angeletos, George-Marios, and Jennifer La'O (2009c), "Sentiments," MIT mimeo.

[12] Angeletos, George-Marios, and Jennifer La'O (2008), "Dispersed Information over the Business Cycle: Optimal Fiscal and Monetary Policy with Dispersed Information," MIT mimeo.

[13] Angeletos, George-Marios, Guido Lorenzoni, and Alessandro Pavan (2009), "A Rational Theory of Irrational Exuberance," MIT/Northwestern mimeo.

[14] Angeletos, George-Marios, and Alessandro Pavan (2009), "Policy with Dispersed Information," Journal of the European Economic Association 7, 11-60.

[15] Angeletos, George-Marios, and Alessandro Pavan (2007), "Efficient Use of Information and Social Value of Information," Econometrica 75:4, 1103-1142.

[16] Angeletos, George-Marios, and Alessandro Pavan (2007b), "Socially Optimal Coordination: Characterization and Policy Implications," Journal of the European Economic Association 5, 585-593.

[17] Angeletos, George-Marios, and Alessandro Pavan (2004), "Transparency of Information and Coordination in Economies with Investment Complementarities," American Economic Review 94, 91-98.

[18] Bacchetta, Philippe, and Eric van Wincoop (2005), "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?," American Economic Review 95.

[19] Barro, Robert (1976), "Rational Expectations and the Role of Monetary Policy," Journal of Monetary Economics 2, 1-32.

[20] Barro, Robert and Robert King (1984). "Time Seperable Preferences and Intertemporal Substitution Models of Business Cycles." Quarterly Journal of Economics 99: 817-839.

[21] Barsky, Robert, and Eric Sims (2009), "News Shocks," University of Michigan mimeo. 3.

[22] Basu, Susanto, John Fernald, and Miles Kimball (2006), "Are Technology Improvements Contractionary?," American Economic Review 96, 1418-1448.

[23] Beaudry, Paul, and Franck Portier (2004), "Exploring Pigou's Theory of Cycles," Journal of Monetary Economics 51, 1183-1216.

[24] Beaudry, Paul, and Franck Portier (2006), "Stock Prices, News, and Economic Fluctuations," American Economic Review 96, 1293-1307.

[25] Benhabib, Jess, Richard Rogerson, and Randall Wright (1991), "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," Journal of Political Economy 99, 1166-1197.

[26] Bernanke, Ben, Mark Gertler, and Simon Gilchrist (1999), "The Financial Accelerator in a Quantitative Business Cycle Framework," in The Handbook of Macroeconomics, ed. by J. B. Taylor, and M. Woodford, pp. 1341–1393. Elsevier Science B.V., Amsterdam, North Holland.

[27] Bernanke, Ben, and Ilian Mihov (1995), "Measuring Monetary Policy," The Quarterly Journal of Economics 113, 869-902.

[28] Bils, Mark, and Pete Klenow (2004), "Some Evidence on the Importance of Sticky Prices," Journal of Political Economy 112, 947-985.

[29] Blanchard, Olivier, and Danny Quah (1989), "The Dynamic Effects of Aggregate Demand and Supply Disturbances," American Economic Review 79, 655-673.

[30] Bolton, Patrick, and David S. Scharfstein, (1990). "A theory of predation based on agency problems in financial contracting," American Economic Review 80, 93–106.

[31] Brunnermeier, Markus K. and Lasse H. Pedersen (2005). "Predatory Trading." The Journal of Finance 60(4): 1825-1863

[32] Caballero, R. and A. Krishnamurthy (2001) "International and Domestic Collateral Constraints in a Model of Emerging Market Crises," in Journal of Monetary Economics, 48, 3, December.

[33] Cai, Fang (2003). "Was there front running in the LTCM crisis?" International Finance Discussion Papers, Board of Governors of the Federal Reserve System, 758

[34] Calvo, Guillermo (1983), "Staggered Prices in a Utility Maximizing Framework," Journal of Monetary Economics 12, 383-398.

[35] Carlin, Bruce, Miguel Lobo, and S. Viswanathan (2007). "Episodic Liquidity Crises: Cooperative and Predatory Trading," The Journal of Finance 62(5): 2235-2274.

[36] Carlstrom, Charles and Timothy Fuerst (1997) "Agency costs, net worth, and business fluctuations: a computable general equilibrium analysis," American Economic Review 87:893-910.

[37] Chari, V.V., Patrick Kehoe, and Ellen McGrattan (2007), "Business Cycle Accounting," Econometrica 75, 781-836.

[38] Chari, V.V., Christiano, Lawrence and Patrick Kehoe (2009), "The Behavior of Small and Large Firms over the Business Cycle," mimeo.

[39] Chen, Kaiji and Zheng Song (2009), "Financial Frictions on Capital Allocation: A Transmission Mechanism of TFP Fluctuations," mimeo.

[40] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (1999), "Monetary Policy Shocks: What Have We Learned and to What End?," in J. B. Taylor & M. Woodford , eds., Handbook of Macroeconomics, Elsevier.

[41] Christiano, Lawrence, Martin Eichenbaum, and Robert Vigfusson (2003), "What Happens After A Technology Shock?" Board of Governors of the Federal Reserve System, International Finance Discussion Papers No. 768.

[42] Christiano, Lawrence, Cosmin Ilut, Roberto Motto, and Massimo Rostagno (2008), "Monetary policy and stock market boom-bust cycles," ECB Working Paper No. 955.

[43] Christiano, Lawrence, Roberto Motto, and Massimo Rostagno (2008) "Financial Factors in Economic Fluctuations," Manuscript, Northwestern University.

[44] Cochrane, J. (1994) "Shocks," in Carnegie-Rochester Conference Series on Public Policy, 41, December.

[45] Coibion, Olivier, and Yuriy Gorodnichenko (2008), "What can survey forecasts tell us about informational rigidities?" UC Berkeley mimeo.

[46] Collard, Fabrice, and Harris Dellas (2005a), "Technology shocks and employment," CNRS-GREMAQ/University of Bern mimeo.

[47] Collard, Fabrice, and Harris Dellas (2005), "Misperceived vs Unanticipated Money: A Synthesis," CNRS-GREMAQ/University of Bern mimeo.

[48] Cooley, T., R. Marimon, and V. Quadrini (2001) "Aggregate Consequences of Limited Contract Enforceability," Mimeo.

[49] Cooley, Thomas, Ramon Marimon, and Vincenzo Quadrini (2004) "Aggregate Consequences of Limited Contract Enforceability," Journal of Political Economy, 112(4), 817–847.

[50] Cordoba, J.C., and M. Ripoll (2004): "Credit Cycles Redux," International Economic Review, 45(4), 1011–1046.

[51] Diamond, Douglas, and Raghuram Rajan (2005), "Liquidity Shortages and Banking Crises," The Journal of Finance 60(2), 615-647.

[52] Fazzarri, Steven, R. Glenn Hubbard and Bruce Petersen (1988), "Financing constraints and corporate investment," Brookings Papers on Economic Activity 1988(1):141-195.

[53] Francis, Neville, and Valerie Ramey (2003a), "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," University of California, San Diego, mimeo.

[54] Gale, Douglas and Martin Hellwig (1985), "Incentive-compatible debt contracts: the one-period problem", Review of Economic Studies 52:647-664.

[55] Galí, Jordi (1999), "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," American Economic Review 89, 249-271.

[56] Galí, Jordi (2003), "New Perspectives on Monetary Policy, Inflation, and the Business Cycle," in M. Dewatripont, L. Hansen, and S. Turnovsky (eds.), Advances in Economic Theory, vol. III, 151-197, Cambridge University Press.

[57] Galí, Jordi, J. David López-Salido, and Javier Vallés (2003), "Technology shocks and monetary policy: Assessing the FedÕs performance, " Journal of Monetary Economics 50, 723-743.

[58] Galí, Jordi, and Pau Rabanal (2004), "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?," NBER Macroeconomics 2004.

[59] Gertler, Mark and Simon Gilchrist (1994) "Monetary policy, business cycles, and the behavior of small manufacturing firms," Quarterly Journal of Economics 59:309-340.

[60] Gertler, Mark, Simon Gilchrist, and Fabio Natalucci (2007) "External Constraints on Monetary Policy," Journal of Money, Credit and Banking, 39, 295–330.

[61] Gertler, Mark, R. Glenn Hubbard, and Anil Kashyap (1991) "Interest Rate Spreads, Credit Constraints, and Investment Fluctuations: An Empirical Investigation," in Financial Markets and Financial Crises, edited by R. Glenn Hubbard, University of Chicago Press, 1991, pp. 11-31.

[62] Gilchrist, Simon, and C.E Himmelberg (1995), "Evidence on the role of cash flow for investment," Journal of Monetary Economics 36:541- 572.

[63] Gilchrist, Simon, and John Leahy (2002), "Monetary Policy and Asset Prices," Journal of Monetary Economics 49, 75-97.

[64] Glosten, L. and P. Milgrom (1985). "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," Journal of Financial Economics 13, 71-100.

[65] Lowenstein, Roger (2001). When Genius Failed: The Rise and Fall of Long-Term Capital Management. Random House.

[66] Goodfriend, Marvin, and Robert King (2001), "The Case for Price Stability," NBER Working Paper 8423.

[67] Goodfriend, Marvin, and Robert King (1997), "The New Neoclassical Synthesis and the Role of Monetary Policy," NBER Macroeconomics Annual 1997, Cambridge, MIT Press.

[68] Gorton, Gary, and Lixin Huang (2006). "Banking Panics and the Endogeneity of Central Banking," Journal of Monetary Economics, 53, 1613-1629.

[69] Graham, Liam, and Stephen Wright (2008), "Information, Heterogeneity and Market Incompleteness," UCL/Birkbeck College mimeo.

[70] Hall, Robert E. (1997), "Macroeconomic Fluctuations and the Allocation of Time," Journal of Labor Economics, 15, S223-S250.

[71] Hansen, Gary (1985), "Indivisible Labor and the Business Cycle," Journal of Monetary Economics 16, 309-337.

[72] Hellwig, Christian (2002), "Public Announcements, Adjustment Delays, and the Business Cycle," UCLA mimeo.

[73] Hellwig, Christian (2005), "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," UCLA mimeo.

[74] Hellwig, Christian, and Laura Veldkamp (2008), "Knowing What Others Know: Coordination Motives in Information Acquisition," Review of Economic Studies, forthcoming.

[75] Hellwig, Christian, and Venky Venkateswaran (2008), "Setting The Right Prices For The Wrong Reasons," UCLA mimeo.

[76] Holmstrom, Bengt, and Jean Tirole (1998), "Private and Public Supply of Liquidity," Journal of Political Economy, 106, 1–40.

[77] Hubbard, R. Glenn, Anil Kashyap and Toni Whited (1995), "Internal finance and firm invesmlent," Journal of Money, Credit and Banking 27:683 701.

[78] Iacoviello, Matteo (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," American Economic Review, 95(3), 739–764.

[79] Jaimovich, Nir, and Sergio Rebelo (2009), "Can News About the Future Drive the Business Cycle?," American Economic Review, 99(4), 1097-1118.

[80] Kajii, Atsushi, and Stephen Morris (1997), "The Robustness of Equilibria to Incomplete Information," Econometrica 65, 1283-1309.

[81] Kashyap, Anil, Owen Lamont, and Jeremy Stein (1994) "Credit conditions and the cyclical behavior of inventories," Quarterly Journal of Economics 109:565 592.

[82] King, Robert (1982), "Monetary Policy and the Information Content of Prices," Journal of Political Economy 90, 247-279.

[83] King, Robert and Sergio Rebelo (2000), "Resusitating Real Business Cycles," in J. B. Taylor & M. Woodford, eds., Handbook of Macroeconomics, Elsevier.

[84] Klenow, J. Peter and Jonathan L. Willis (2007), "Sticky information and sticky prices," Journal of Monetary Economics 54(Supplement), 79-99.

[85] Kiyotaki, Nobuhiro. (1998). "Credit and business cycles," Japanese Economic Review 49 (March): 18–35.

[86] Kiyotaki, Nobuhiro, and Moore, John. (1997). "Credit cycles," Journal of Political Economy 105 (April): 211–48.

[87] Kocherlakota, Narayana. (2000) "Creating Business Cycles Through Credit Constraints," in Federal Reserve Bank of Minneapolis Quarterly Review, 24, No. 3, Summer.

[88] Krishnamurthy, A. (2003). "Collateral Constraints and the Amplification Mechanism," Journal of Economic Theory. 111(2): 277-292.

[89] Krusell, Per and Anthony Smith (1998) "Income and Wealth Heterogeneity in the Macroeconomy," Journal of Political Economy, vol. 106(5), 867-896.

[90] Kydland, Finn, and Edward Prescott (1982), "Time to Build and Aggregate Fluctuations," Econometrica 50, 1345-70.

[91] Kyle, Albert S (1985). "Continuous auctions and insider trading," Econometrica 53, 1315–1335.

[92] Liu, Zheng, Pengfei Wang, and Tao Zha (2009), "Asset Prices Credit Constraints, and Macroeconomic Fluctuations," mimeo.

[93] Lipman, Barton (2003), "Finite Order Implications of Common Priors," Econometrica 71, 1255-1267.

[94] Lorenzoni, Guido (2008), "A Theory of Demand Shocks," American Economic Review, forthcoming.

[95] Lorenzoni, Guido (2009), "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information," Review of Economic Studies, forthcoming.

[96] Lucas, Robert E., Jr. (1972), "Expectations and the Neutrality of Money," Journal of Economic Theory 4, 103-124.

[97] Lucas, Robert E., Jr. (1975), "An Equilibrium Model of the Business Cycle," Journal of Political Economy 83, 1113-1144.

[98] Luo, Yulei (2008), "Consumption Dynamics under Information Processing Constraints," Review of Economic Dynamics 11, 366-385.

[99] Mackowiak, Bartosz, and Mirko Wiederholt (2008), "Optimal Sticky Prices under Rational Inattention," American Economic Review, forthcoming.

[100] Mackowiak, Bartosz, and Mirko Wiederholt (2009), "Business Cycle Dynamics under Rational Inattention," European Central Bank/Nothwestern University mimeo.

[101] Madrigal, Vicente, (1996), "Non-fundamental speculation," Journal of Finance 51, 553–578.

[102] Mankiw, N. Gregory and Ricardo Reis (2002), "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," Quarterly Journal of Economics 117, 1295-1328.

[103] Mankiw, N. Gregory and Ricardo Reis (2006), "Pervasive Stickiness," American Economic Review 96, 164-169.

[104] McGrattan, Ellen (2004), "Comment on Gali and Rabanal," Federal Reserve Bank of Minneapolis, Research Department Staff Report 338.

[105] Mertens, Thomas (2009), "Excessively Volatile Stock Markets: Equilibrium Computation and Policy Analysis," NYU Stern mimeo.

[106] Moscarini, Giuseppe (2004), "Limited Information Capacity as a Source of Inertia," Journal of Economic Dynamics and Control 28, 2003-2035.

[107] Morris, Stephen, and Hyun Song Shin (2002), "The Social Value of Public Information," American Economic Review 92, 1521-1534.

[108] Morris, Stephen, and Hyun Song Shin (2006), "The Inertia of Forward-Looking Expectations," American Economic Review 96, 1521-1534.

[109] Nimark, Kristoffer (2008), "Dynamic Pricing and Imperfect Common Knowledge," Journal of Monetary Economics 55, 365-382.

[110] Pastor, Lubos, and Robert F. Stambaugh, (2003), "Liquidity risk and expected stock returns," Journal of Political Economy 111, 642–685.

[111] Phelps, Edmund (1970), "Introduction: The New Microeconomics in Employment and Inflation Theory," in A.A. Alchian, C.C. Holt et al., Microeconomic Foundations of Employment and Inflation Theory, New York: Norton.

[112] Porapakkarm, Ponpoje, and Eric Young (2008), "Information Heterogeneity in the Macroeconomy,"University of Virginia mimeo.

[113] Prescott, Edward (1986), "Theory Ahead of Business Cycle Measurement," Federal Reserve Bank Minneapolis Quarterly Review 10, 9-22.

[114] Prescott, Edward and José-Víctor Ríos-Rull (1992), "Classical Competitive Analysis of Economies with Islands," Journal of Economic Theory 57, 73-98.

[115] Rajan, Raghuram and Luigi Zingales (1998), "Financial Dependence and Growth," American Economic Review, 88(3), pp. 559-86.

[116] Reis, Ricardo (2006), "Inattentive Producers," Review of Economic Studies 73, 793-821.

[117] Reis, Ricardo (2008), "Optimal Monetary Policy Rules in an Estimated Sticky-Information Model," American Economic Journal: Macroeconomics, forthcoming.

[118] Rogerson, Richard (1988), "Indivisible Labor, Lotteries and Equilibrium," Journal of Monetary Economics 21, 3-16.

[119] Rondina, Giacomo (2008), "Incomplete Information and Informative Pricing," UCSD mimeo.

[120] Rotemberg, Julio (2003), "Stochastic technical progress, smooth trends, and nearly distinct business cycles," American Economic Review 93, 1543-1559.

[121] Rotemberg, Julio, and Michael Woodford (1991), "Markups and the Business Cycle," in O.J. Blanchard and S. Fischer (eds.), NBER Macroeconomics Annual 1991, Cambridge, MIT Press.

[122] Rotemberg, Julio, and Michael Woodford (1997), "An Optimization-Based Econometric Model for the Evaluation of Monetary Policy," in O.J. Blanchard and S. Fischer (eds.), NBER Macroeconomics Annual 1997, Cambridge, MIT Press.

[123] Rotemberg, Julio, and Michael Woodford (1999), "The Cyclical Behavior of Prices and Costs.," in J. B. Taylor & M. Woodford (eds.), Handbook of Macroeconomics, Elsevier.

[124] Shapiro, Matthew and Mark Watson (1988) "Sources of Business Cycle Fluctuations." NBER Macroeconomics Annual 3:111-148.

[125] Sharpe, Steven (1994). "Financial market imperfections, firm leverage, and the cyclicality of employment," American Economic Review 84:1060-1074.

[126] Shimer, Robert (2009), "Convergence in Macroeconomics: The Labor Wedge," American Economic Journal: Macroeconomics 1, 280-297.

[127] Sims, Christopher (2003), "Implications of Rational Inattention," Journal of Monetary Economics 50, 665-690.

[128] Sims, Christopher (2006), "Rational Inattention: Beyond the Linear-Quadratic Case," American Economic Review 96, 158-163.

[129] Sims, Eric (2009)."Expectations Driven Business Cycles: An Empirical Evaluation" May 2009

[130] Stock, James H. & Watson, Mark W., (1999). "Business cycle fluctuations in us macroeconomic time series," Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 1, pages 3-64 Elsevier.

[131] Townsend, Robert M. (1979), "Optimal contracts and competitive markets with costly state verification," Journal of Economic Theory 21:265-293.

[132] Townsend, Robert M. (1983), "Forecasting the Forecasts of Others," Journal of Political Economy 91, 546-588.

[133] Van Nieuwerburgh, Stijn, and Laura Veldkamp (2008), "Information Acquisition and Under-Diversification," NYU Stern mimeo.

[134] Van Nieuwerburgh, Stijn, and Laura Veldkamp (2006), "Learning Asymmetries in Real Business Cycles," Journal of Monetary Economics 53, 753-772.

[135] Vayanos, Dimitri (2001). "Strategic trading in a dynamic noisy market," The Journal of Finance 56, 131–171.

[136] Veldkamp, Laura (2006), "Media Frenzies in Markets for Financial Information," American Economic Review 96, 577-601.

[137] Veldkamp, Laura, and Justin Woolfers (2006), "Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement," Journal of Monetary Economics 54, 37-55.

[138] Wen, Yi (2001), "Demand-Driven Business Cycles: Explaining Domestic and International Comovements,"Cornell University mimeo.

[139] Woodford, Michael (2003a), "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps, Princeton University Press.

[140] Woodford, Michael (2003b), Interest and Prices: Foundations of a Theory of Monetary Policy, Princeton: Princeton University Press.

[141] Woodford, Michael (2008), "Information-Constrained State Dependent Pricing," Columbia mimeo.