Essays in Cooperation and Repeated Games

by

Daniel Barron

B.A. Economics, Indiana University-Bloomington (2008)
B.S. Mathematics, Indiana University-Bloomington (2008)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

Signature of Author ....

Department of Economics
May 15, 2013

Certified by ....

Robert Gibbons
Sloan Distinguished Professor of Management and Economics
Thesis Supervisor

Certified by ....

Juuso Toikka
Assistant Professor of Economics
Thesis Supervisor

Accepted by ....

Michael Greenstone
3M Professor of Environmental Economics
Chairman, Departmental Committee on Graduate Studies

Essays in Cooperation and Repeated Games

by

Daniel Vincent Barron

Submitted to the Department of Economics
on May 15, 2013 in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy in
Economics

ABSTRACT

This dissertation explores cooperation when formal contracts and legal institutions are imperfect. The first chapter (co-authored with Isaiah Andrews) considers how a principal allocates business among a group of agents to motivate them in the context of a repeated game with imperfect private monitoring. If players are impatient, the optimal relational contract dynamically allocates future business among agents depending on past performance. An optimal allocation rule favors an agent who performs well, even if he later performs poorly. An agent loses favor only if he is unable to produce and his replacement performs well. The principal may allows some relationships to deteriorate into persistent shirking in order to better motivate other agents. We find conditions under which the principal either does or does not benefit by concealing information from the agents.

The second chapter proves that approximately Pareto efficient outcomes can be sustained in a broad class of games with imperfect public monitoring and Markov adverse selection when players are patient. Consider a game in which one player's utility evolves according to an irreducible Markov process and actions are imperfectly observed. Then any payoff in the interior of the convex hull of all Pareto efficient and min-max payoffs can be approximated by an equilibrium payoff for sufficiently patient players. The proof of this result is partially constructive and uses an intuitive "quota mechanism" to ensure approximate truth-felling. Under stronger assumptions, the result partially extends to games where one player's private type determines every player's utility.

The final chapter explores how firms might invest to facilitate their relationships with one another. Consider a downstream firm who uses relational contracts to motivate multiple suppliers. In an applied model with imperfect private monitoring, this chapter shows that the suppliers might "put the relationship first:" they invest to flexibly produce many of the products required by the downstream firm, rather than cutting costs by specializing. A downstream firm that relies on relational contracts tends to source from fewer suppliers, each of whom can inefficiently manufacture many different products required by that firm.

Thesis Supervisor: Robert Gibbons
Title: Sloan Distinguished Professor of Management and Economics

Contents

Acknowledgments

I am tremendously indebted to my advisor Robert Gibbons, who has been absolutely instrumental in my personal and professional development during my time at MIT. I can only aspire to his breadth and depth of knowledge, his keen understanding of the literature, and his consistent insight and constant encouragement through the process. His mentorship shaped my understanding of applied theory and its role within economics, for which I can only offer my sincere gratitude.

I would like to deeply thank Juuso Toikka for providing support, insight, and detailed advising and feedback. From broad and wide-ranging conversations over beers after the theory seminar to focused discussions of papers and proofs, Juuso was always very generous with his time and very insightful in his advice. Glenn Ellison provided instant, honest, deep, and structured feedback to my papers. He makes sharp thinking and concise writing look easy, and I can only hope to emulate his clarity and wisdom. I thank Bengt Holmstrom, Parag Pathak, Alessandro Bonatti, and many others who offered valuable insights and teaching.

I benefited greatly from co-authoring with Robert Gibbons, Ricard Gil, Kevin Murphy, and especially Isaiah Andrews, who was sounding board, discipline device, and font of insight as we developed and wrote what ended up being my job market paper.

I am very lucky to be surrounded by exceptionally brilliant, humble, and fun colleagues. Thanks to my roommates Isaiah, Conrad Miller, Adam Sacarny, and Yuhei Miyauchi. Greg Leiserson for incisive commentary and fantastic hiking experiences. Mike Powell for great discussions encompassing relational contracts, the job market, and computer games. Xiao Yu Wang for random talks and general awesomeness. Brad Shapiro for always being ready to hang out over dinner, games, or fine whiskey. Marco DiMaggio for being a wonderful roommate generals year, and sharing the best food ever. Sebastian DiTella for many study sessions, introducing me to excellent television, and being (almost) as nervous as I was. Sally Hudson for hiking, grilling, and insight into the Huskers mentality. Brad Larsen and Chris Palmer for many great conversations, and instructing me in the ways of empirics. Heartfelt thanks to Chris Walters, Gabriel Carroll, David Chan, Eliza Forsythe, Amir Kermani, Juan Pablo Xandri, Nils Wernerfelt, Miikka Rokkanen, Sarah Venables, Ben Golub, Iuliana Pascu, Alexander Wolitzky, Anton Kolotilin, Luis Zermeno, Jennifer Peck, Jason Abaluck, Robby Akerlof, Hongyi Li, and too many others to list. You made this journey amazing.

The members of the Lutheran-Episcopal Ministry at MIT and Christ Church Cambridge deserve special thanks for providing a core of comfort and support at MIT. Thanks to MITOC for nurturing my addiction to hiking and climbing, and to the folks at SSG-Boston for training in a variety of archaic weapons.

I could not have accomplished anything without the support and love of my parents, John and Cathy. Thanks Mom and Dad for encouraging me to go into Economics, for much sage advice along the journey, and for always being there with a supportive and loving ear. Thanks also to my siblings, Jeff, Julie, and Mike, and their families. You were there for me growing up and at every step along the way.

Finally, I am tremendously grateful to have Emily in my life. Thank you so much for your humor, your comfort during stressful times, your enthusiasm, intelligence, and love. I am so lucky to have met you, and your presence makes the future immeasurably brighter.

Chapter 1

The Allocation of Future Business

With Isaiah Andrews

1.1 Introduction

When formal contracts are unavailable, costly, or imperfect, individuals rely on informal relationships with one another to sustain cooperation. The resulting relational contracts are pervasive both within and between firms. For example, informal agreements typically complement incomplete formal contracts in the "just-in-time" supply chains used by Chrysler, Toyota, and others. Similarly, managers rely on long-term relationships to motivate employees and divisions.[1] Regardless of the setting, a principal who uses relational contracts must have an incentive to follow through on her promises to the agents. If the parties interact repeatedly, then these promises are made credible by the understanding that if they are ever broken, the relationship sours and surplus is lost.

When a principal interacts with several agents, how she chooses to allocate business among them critically determines the strength of her relational contracts with each of them. In this paper, we consider a repeated game in which the principal requires a single product in each period. Only a subset of the agents is able to produce at any given time, and the principal chooses one agent from this subset to exert private effort that stochastically

[1] See Sako (2004) and Liker and Choi (2004) on Toyota, Hohn (2010) on Chrysler, and Baker, Gibbons, and Murphy (2002) for many other examples and references.

determines the profitability of the final product. Importantly, each agent observes only his own output and pay, so each has only limited information about the principal's other relationships. By promising additional future business to one agent, the principal increases the total surplus created by that bilateral relationship, which in turn allows the principal to give that agent more effective incentives today. Ideally, the principal would be able to promise enough future business to each agent to motivate him. However, the same future business cannot be promised to different agents; therefore, if the parties are not very patient, the principal must prudently allocate business over time to ensure that she can credibly reward an agent's strong performance.

We show that future business is dynamically allocated in response to past performance, and that this allocation decision plays fundamentally different roles following high or low output. An agent is motivated by his expected future payoff, which is determined by wages and bonuses. The principal can only credibly promise a large payoff if it is accompanied by a substantial amount of future business. Therefore, an optimal allocation rule *rewards success* by promising to favor a high-performing agent with future business, which serves as a form of relational capital to make a large reward credible. In contrast, the principal can immediately punish a low-performing agent using transfer payments, regardless of how much future business is allocated to that agent. As a result, an optimal allocation rule *tolerates failure*, as the principal prefers to use short-term transfers to punish poor performance rather than weakening a long-term favored relationship by withdrawing future business. More formally, we characterize a dynamic relational contract that rewards success, tolerates failure, and induces first-best effort whenever any equilibrium does. In this equilibrium, an agent rises to favored status when he performs well, then is gradually displaced by more recent high performers. When players are too impatient to attain first-best, the principal sometimes asks one agent to shirk in order to promise enough future business to the *other* agent to credibly reward high performance. Considering the game with two agents and making a realistic but substantive equilibrium restriction,[2] we show that one (non-unique) optimal relational contract continues to reward success and tolerate failure. In this equilibrium, the principal eventually permanently favors one agent, while her relationship with the other agent devolves

[2]Analogous to belief-free equilibrium; see definition 5 for more details.

into persistent perfunctory performance. In other words, this relational contract optimally prescribes a permanent inefficiency in order to solve a short-term moral hazard problem.

Methodologically, we consider a repeated game with imperfect private monitoring—agents do not observe the output produced by or the payments made to their counterparts—so we cannot rely on standard recursive methods and instead develop alternative tools. This monitoring structure implies both that agents cannot coordinate to jointly punish deviations and that the principal may be able to exploit agents' limited information about past play to better motivate effort. In our relational contract that attains first-best whenever any equilibrium does, the agents would be willing to work hard even if they learned the true history of past play. In contrast, if first-best is unattainable, we show that the principal typically finds it optimal to exploit the private monitoring structure by concealing information from the agents.

Broadly, this paper considers how individuals can solve short-term moral hazard problems by making long-term policy commitments—in this case, by promising to allocate future business in a certain way. For instance, consider a firm with two divisions. Each division might learn of a profitable project in each period, but the firm can support only one project at a time. If projects are very profitable, costs are low, or resources to launch projects are abundant, we will show that the firm can allocate resources to always motivate both divisions, so that one division might be temporarily favored but neither becomes permanently dominant. When resources or profitable projects are scarce, then in order to properly motivate one division, the firm allows that division to become dominant for a long period of time or potentially permanently following a strong performance. Thus, this paper speaks to the observation by Cyert and March (1963) that distinct coalitions within a company motivate and reward one another by making long-term "policy commitments." In our setting, persistent policies that promise to favor one agent over another are important motivational tools that make monetary payments credible.

The allocation of future business plays a pivotal role in many different real-world environments. For example, companies frequently allocate business among long-term suppliers in response to past performance. Toyota's *keiretsu* is a classic example of relational supply chain management. Asanuma (1989) notes that Toyota ranked its suppliers in four different tiers. Better-ranked suppliers are given first priority on valuable contracts, and suppliers move between tiers in response to past performance.[3] Similarly, Farlow et al (1996) report that Sun Microsystems regularly divides a project among suppliers based in part on their past performance. In fact, Krause et al (2000) present a survey of many different firms which finds that suppliers are frequently motivated by "increased volumes of present business and priority consideration for future business." Similarly, firms allocate tasks and promotions to reward successful employees.

The seminal papers on relational contracts by Bull (1987), MacLeod and Malcomson (1989), and Levin (2003) have spurred a large and growing literature; Malcomson (2012) has an extensive survey. Levin (2002) analyzes multilateral relational contracts, arguing that bilateral relationships may be more robust to misunderstandings between players. In our setting and unlike Levin (2002), bilateral relationships are the natural result of private monitoring and the allocation of business plays a central role. Segal (1999) considers a static setting in which the principal can privately offer a formal contract to each of a group of agents, which is a monitoring assumption similar to our own.

In a model related to ours, Board (2011) exploits the notion that future business is a scarce resource to limit the number of agents with whom a principal optimally trades. He shows that the principal separates potential trading partners into "insiders" and "outsiders," trades efficiently with the insiders, and is biased against outsiders. While agents in his model do not observe one another's choices, the principal perfectly observes actions and so the dynamics are driven by neither moral hazard nor private monitoring, which are the central forces in our paper. In the context of repeated procurement auctions, Calzolari and Spagnolo (2010) similarly argue that an auctioneer might want to limit participation in the auction to bolster relational incentives with the eventual winner, and they discuss the interaction between relational incentives and bidder collusion. Li, Zhang, and Fine (2011) show that a principal who is restricted to cost-plus contracts may use future business to motivate agents, but the dynamics of their equilibrium are not driven by relational concerns and are quite different from ours. In research related to the present paper, Barron (2012) considers how

[3]Toyota is known for encouraging its suppliers to communicate and collectively punish a breach of the relational contract. Our baseline model does not allow such communication, but we show in Section 1.7 that qualitatively similar dynamics hold with communication (albeit under stronger assumptions).

suppliers optimally invest to maximize the value of their relational contracts and shows that such suppliers might opt to be "generalists" rather than specializing in a narrow range of products.

Our model is also related to a burgeoning applied theory literature using games with private monitoring. Kandori (2002) provides a nice overview of the theory of such games. Fuchs (2007) considers a bilateral relational contracting problem and shows that efficiency wages are optimal if the principal privately observes the agent's output. Wolitzky (2012a,b) considers enforcement in games where each agent observes an action either perfectly or not at all. In our model, output—which is an imperfect signal of an agent's private effort—is similarly observed by some but not all of the players. Ali and Miller (2012) analyze what networks best sustain cooperation in a repeated game, but do not consider the allocation of business among players. Harrington and Skrzypacz (2011) discuss collusion when firms privately observe both prices and quantities, deriving an equilibrium that resembles the actions of real-world cartels.

The rest of the paper proceeds as follows. The next section lays out the repeated game, devoting special attention to the private monitoring structure that drives many of our results. We prove a set of necessary and sufficient conditions for equilibrium in Section 1.3. A producing agent is willing to work hard only if he believes that *following high output*, he would produce a sufficiently large amount of surplus in subsequent periods. Of course, the principal's allocation decision determines how much surplus each agent expects to produce at any given time and so determines when these conditions hold. Section 1.4 fully characterizes a non-stationary relational contract that attains first-best whenever any Perfect Bayesian Equilibrium does. Turning to the game with two agents and parameters such that first-best is unattainable, Section 1.5 shows that non-stationary allocation rules are typically optimal. Restricting attention to equilibria that provide *ex post* incentives to the agents, we characterize one optimal allocation rule. In this relational contract, the principal's relationship with one of the agents eventually becomes "perfunctory," in the sense that the agent is allocated production infrequently and perpetually shirks. Which relationship breaks down depends on past performance: when one agent performs well, the *other* relationship is likely to suffer. In Section 1.6, we consider equilibria that do not provide *ex post* incentives. So long as first-best is attainable, it can be attained using *ex post* incentives. In contrast, if parameters are such that first-best is unattainable, we show that the principal can typically benefit by concealing information from the agents. The baseline model makes the stark assumption that agents cannot communicate; in Section 1.7, we allow agents to send messages to one another and show that under an intuitive additional assumption, results that are qualitatively similar to our baseline model continue to hold. Section 1.8 concludes.

Omitted proofs are in Appendix 1.9. Additional results referenced in the text are in Supplementary Appendices ?? and ??, available at http://economics.mit.edu/grad/dbarron.

1.2 Model and Assumptions

1.2.1 Model

Consider a repeated game with $N + 1$ players, denoted $\{0, 1, ..., N\}$. We call player 0 the principal ("she"), while every other player $i \in \{1, ..., N\}$ is an agent ("he"). In each round, the principal requires a single good that could be made by any one of a subset of the agents. This good can be thought of as a valuable input to the principal's production process that only some of the agents have the capacity to make in each period. After observing which agents can produce the good, the principal allocates production to one of them, who chooses whether to accept or reject production and how much effort to exert. Critically, utility is transferable between the principal and each agent but *not* between two agents. At the very beginning of the game, the principal and each agent can "settle up" by transferring money to one another; these payments are observed only by the two parties involved.

Formally, we consider the infinite repetition $t = 0, 1, ...$ of the following stage game with common discount factor δ:

1. A subset of *available agents* $\mathcal{P}_t \in 2^{\{1,..,N\}}$ is publicly drawn with probability $F(\mathcal{P}_t)$.

2. The principal publicly chooses a single agent $x_t \in \mathcal{P}_t \cup \{\emptyset\}$ as the *producer*.

3. $\forall i \in \{1, ..., N\}$, the principal transfers $w_{i,t}^A \geq 0$ to agent i, and agent i simultaneously transfers $w_{i,t}^P \geq 0$ to the principal. $w_{i,t}^A$ and $w_{i,t}^P$ are observed *only* by the principal and

agent i. Define $w_{i,t} = w_{i,t}^A - w_{i,t}^P$ as the *net transfer to agent* i.[4]

4. Agent x_t accepts or rejects production, $d_t \in \{0,1\}$. d_t is observed only by the principal and x_t. If $d_t = 0$, then $y_t = 0$.[5]

5. If $d_t = 1$, agent x_t privately chooses an effort level $e_t \in \{0,1\}$ at cost ce_t.

6. Output $y_t \in \{0, y_H\}$ is realized and observed by agent x_t and the principal, with $Prob\{y_t = y_H | e_t\} = p_{e_t}$ and $1 \geq p_1 > p_0 \geq 0$.

7. $\forall i \in \{1, ..., N\}$, transfers $\tau_{i,t}^A, \tau_{i,t}^P$ are simultaneously made to and from agent i, respectively. $\tau_{i,t}^A, \tau_{i,t}^P$ are observed only by the principal and agent i, with the *net transfer to agent* i denoted $\tau_{i,t} = \tau_{i,t}^A - \tau_{i,t}^P$.

Let $1_{i,t}$ be the indicator function for the event $\{x_t = i\}$. Then discounted payoffs are

$$U_0 = (1-\delta)\sum_{t=0}^{\infty}\delta^t\left(d_t y_t - \sum_{i=1}^{N}(w_{i,t} + \tau_{i,t})\right) \equiv \sum_{t=0}^{\infty}\delta^t(1-\delta)u_0^t$$

and

$$U_i = (1-\delta)\sum_{t=0}^{\infty}\delta^t\left(w_{i,t} + \tau_{i,t} - d_t c e_t 1_{i,t}\right) \equiv \sum_{t=0}^{\infty}\delta^t(1-\delta)u_i^t$$

for the principal and agent i, respectively.

This monitoring structure implies that the principal observes every variable except effort, whereas agents do not see any of one another's choices.[6] Several features of this model allow us to cleanly discuss the role of future business in a relational contract. First, agents cannot communicate with one another. While stark, this assumption implies both that punishments are bilateral and that agents' information plays an important role in equilibrium.[7] In Section 1.7, we give conditions under which some of our results hold when agents can communicate.

[4] By convention, $w_{i,t}^P = 0$ if $w_{i,t} > 0$ and $w_{i,t}^A = 0$ if $w_{i,t} \leq 0$.

[5] Results analogous to those in Sections 1.3 - 1.6 hold for outside option $\bar{u} > 0$ so long as the principal can reject production by agent x_t in step 4 of the stage game without this rejection being observed by the other agents.

[6] In particular, agents cannot pay one another. In Appendix ??, we show that a stationary contract would typically be optimal if they could. Thus, our model may be best-suited for settings in which the agents primarily interact with one another *via* the principal.

[7] In the absence of private monitoring we could restrict attention to equilibria with (some definition of) bilateral punishments. It is non-trivial to formally define such a behavioral restriction in this model due to the role of the allocation rule.

Second, the wage is paid *before* the agent accepts or rejects production. One way to interpret $d_t = 0$ is as a form of shirking that guarantees low output, rather than explicit rejection.[8] Third, we assume that some agents are unable to produce in each round. Between firms, suppliers might lack the time or appropriate capital to meet their downstream counterpart's current needs;[9] within firms, a division might be unavailable because it has no new project that requires resources from headquarters. Finally, the principal cannot "multisource" by allocating production to several agents in each round. While this restriction substantially simplifies the analysis, the allocation of business would remain important with multisourcing *so long as* the principal profits from only one agent's output in each period.[10]

1.2.2 Histories, Strategies, and Continuation Payoffs

Define the *set of histories* h^T *at the start of round* T as

$$\mathcal{H}_B^T = \left\{\mathcal{P}_t, x_t, \{w_{i,t}^A\}, \{w_{i,t}^P\}, d_t, e_t, y_t, \{\tau_{i,t}^A\}, \{\tau_{i,t}^P\}\right\}_{t=0}^T$$

and denote by $\mathcal{N}$ the different nodes of the stage game. A *partial history* $(h^{t-1}, n_t) \in \mathcal{H}_B^T \times \mathcal{N}$ denotes that h^{t-1} occurred in the first $t-1$ rounds and n_t is the node in period t.[11] The set of all histories is

$$\mathcal{H} = \{\emptyset\} \cup \left\{\bigcup_{T=1}^{\infty} \left\{(h^{T-1}, n_T) \mid h^{T-1} \in \mathcal{H}_B^T, n_T \in \mathcal{N}\right\}\right\}.$$

For each player $i \in \{0, 1, ..., N\}$, $\mathcal{I}_i : \mathcal{N} \rightrightarrows \mathcal{N}$ describes i's information sets in the stage game. Private histories at the start of each round for player i are $h_i^T \in \mathcal{H}_{i,B}^T$,[12] and player

[8] Indeed, the results in Sections 1.3 - 1.6 hold if the agents accept or reject production before the wage is paid, but can costlessly choose effort "$e_t = -1$" that guarantees $y_t = 0$.

[9] For example, Jerez et al (2009) report that one of Infosys' partners sources a product from the market only if Infosys "does not have the capability" to supply it.

[10] If multisourcing is possible, the choice of *which agent's output to use* when *multiple* agents produce y_H plays a similar role to x_t in the baseline model.

[11] A partial history may also be written in terms of actions: $(h^{t-1}, \mathcal{P}_t, x_t)$ denotes history h^{t-1} followed by $\{\mathcal{P}_t, x_t\}$.

[12] $\forall i \in \{1, ..., N\}$,

$$\mathcal{H}_{i,B}^T = \left\{\mathcal{P}_t, x_t, w_{i,t}^A, w_{i,t}^P, 1\{x_t = i\}d_t, 1\{x_t = i\}e_t, 1\{x_t = i\}y_t, \tau_{i,t}^A, \tau_{i,t}^P\right\}_{t=0}^T.$$

i's private histories are

$$\mathcal{H}_i = \{\emptyset\} \cup \left\{ \bigcup_{T=1}^{\infty} \left\{ (h_i^{T-1}, \mathcal{I}_i(n_T)) \mid h_i^{T-1} \in \mathcal{H}_{i,B}^T, n_T \in \mathcal{N} \right\} \right\}.$$

Two histories h^t, $\hat{h}^t$ are *indistinguishable by agent* i if $h_i^t = \hat{h}_i^t$.

Strategies for player i are denoted $\sigma_i \in \Sigma_i$ with strategy profile $\sigma = (\sigma_0, ..., \sigma_N) \in \Sigma = \Sigma_0 \times ... \times \Sigma_N$. Let u_i^t be player i's stage-game payoff in round t.

Definition 1 *$\forall i \in \{0, ..., N\}$, player i's continuation surplus $U_i : \Sigma \times \mathcal{H} \to \mathbb{R}$ given strategy σ and history (h^{t-1}, n_t) is*

$$U_i(\sigma, (h^{t-1}, n_t)) = E_\sigma \left[\sum_{t'=1}^{\infty} (1-\delta)\delta^{t'-1} u_i^{t+t'} \mid (h^{t-1}, n_t) \right] \tag{1.1}$$

and the payoff to the principal from agent i's production, $U_0^i : \Sigma \times \mathcal{H} \to \mathbb{R}$, is

$$U_0^i(\sigma, (h^{t-1}, n_t)) = E_\sigma \left[\sum_{t'=1}^{\infty} (1-\delta)\delta^{t'-1} (1\{x_{t+t'} = i\} d_{t+t'} y_{t+t'} - w_{i,t+t'} - \tau_{i,t+t'}) | (h^{t-1}, n_t) \right]. \tag{1.2}$$

Intuitively, U_0^i is the "stake" the principal has in her relationship with agent i: it is the expected output produced by agent i but not earned by him. Agents do not know the true history, so their beliefs about continuation surplus are *expectations* conditional on their private history: $E_\sigma \left[U_i(\sigma, (h^{t-1}, n_t)) | h_i^{t-1}, \mathcal{I}_i(n_t) \right]$.

Definition 2 *The i-*dyad surplus *$S_i : \Sigma \times \mathcal{H} \to \mathbb{R}$ is the total expected surplus generated by agent i:*

$$S_i(\sigma, (h^{t-1}, n_t)) = U_0^i(\sigma, (h^{t-1}, n_t)) + U_i(\sigma, (h^{t-1}, n_t)). \tag{1.3}$$

Dyad surplus plays a critical role in the analysis. Intuitively, S_i is *agent i's contribution to total surplus*—it is the surplus from those rounds when agent i is allocated production.

The principal's history is

$$\mathcal{H}_{0,B}^T = \left\{ \mathcal{P}_t, x_t, \{w_{i,t}^A\}, \{w_{i,t}^P\}, d_t, y_t, \tau_{i,t}^A, \tau_{i,t}^P \right\}_{t=0}^{T}.$$

We will show that S_i is a measure of the collateral that is available to support incentive pay for agent i. We typically suppress the dependence of S_i, U_i, and U_0^i on σ.

A *relational contract* is a Perfect Bayesian Equilibrium (PBE) of this game, where the set of PBE payoffs is $PBE(\delta)$ for discount factor $\delta \in [0,1)$.[13] We focus on the *optimal* relational contract, which maximizes *ex ante total* surplus: $\max_{v \in PBE(\delta)} \sum_{i=0}^{N} v_i$. This is equivalent to maximizing the principal's surplus because utility is transferable between the principal and each agent,[14] but it is not equivalent to maximizing agent i's surplus.

Three assumptions are maintained throughout the paper, unless explicitly stated otherwise.

Assumption 1 *Agents* $\{1,...,N\}$ *are symmetric:* c, p_e, *and* y_H *are identical for each agent, and* $\forall \mathcal{P}, \mathcal{P}' \in 2^{\{1,...,N\}}$, $F(\mathcal{P}) = F(\mathcal{P}')$ *if* $|\mathcal{P}| = |\mathcal{P}'|$.

Assumption 2 *Full support: for any non-empty* $\mathcal{P} \subseteq \{1,...,N\}$, $F(\mathcal{P}) > 0$.

Assumption 3 *High effort is costly but efficient:* $y_H p_1 - c > y_H p_0$, $c > 0$.

1.3 The Role of Future Business in Equilibrium

In this section, we prove two essential lemmas that form the foundation of our analysis. The first clarifies the punishments that can be used to deter deviations, and the second identifies the role of S_i in the relational contract.

As a benchmark, Proposition 1 shows that when y_t is contractible, the optimal formal contract generates first-best total surplus $V^{FB} = (1 - F(\emptyset))(y_H p_1 - c)$ regardless of the allocation rule.

[13] A Perfect Bayesian Equilibrium is an assessment consisting of both a strategy σ^* and belief system for each player $\mu^* = \{\mu_i^*\}_{i=0}^{N}$. Beliefs for player i $\mu_i^* : \mathcal{H}_i \to \Delta(\mathcal{H})$ assign a distribution over true histories (h^{t-1}, n_t) at each of i's information sets $(h_i^{t-1}, \mathcal{I}_i(n_t))$. Given these beliefs, σ_i^* is a best response at each information set. Bayes Rule is used to update beliefs μ_i^* whenever possible. When Bayes Rule does not apply (i.e., the denominator is 0), μ_i^* assigns positive weight only to histories that are consistent with $(h_i^{t-1}, \mathcal{I}_i(n_t))$, but is otherwise unconstrained.

[14] Proof sketch: Consider an optimal PBE σ^*. At the beginning of the game, agent i pays a transfer equal to $U_i(\sigma^*, \emptyset)$ to the principal. The equilibrium proceeds as in σ^* if agent i pays this transfer. If agent i does not pay, then in every future period $w_{i,t} = \tau_{i,t} = 0$, agent i chooses $d_t = 0$, and the principal chooses an allocation rule that maximizes her continuation payoff given these actions. These strategies form an equilibrium if the principal believes i choose $e_t = 0$ and agent i believes the principal offers $w_{i,t}^A = \tau_{i,t}^A = 0$ in every future period following a deviation.

Proposition 1 *If output y is contractible and the principal offers a take-it-or-leave-it contract after $\mathcal{P}_t$ is realized, then $\exists$ a PBE with surplus V^{FB}.*

Proof:

The following is an equilibrium contract of the stage game: $\tau_i(0) = 0$, $\tau_i(1) = \frac{c}{p_1 - p_0}$, and $w_i = c - p_1 \frac{c}{p_1 - p_0}$. Under this contract, agent x_t accepts and chooses $e = 1$. Any allocation rule that satisfies $x_t \neq \emptyset$ if $\mathcal{P}_t \neq \emptyset$ is incentive-compatible and generates V^{FB}. ■

Since all players are risk-neutral, the principal can use a formal contract to costlessly induce first-best effort. This result stands in stark contrast to the rest of the analysis, in which future business is the critical determinant of a relationship's strength. In this setting, the incompleteness of formal contracts is a necessary prerequisite for the allocation of business to play an important role.

Agents' beliefs about the true history can evolve in complicated ways, so our next step is to derive intuitive incentive constraints that depend on dyad-specific surplus S_i. In the online Appendix, we show that any PBE payoff can be attained by an equilibrium in which agents do not condition on their past effort choices. Apart from effort, every variable is observed by the principal and at least one agent; deviations in these variables can be punished as harshly as possible.[15] Lemma 1 demonstrates that if the principal or agent i reneges on a promised payment, the harshest punishment is the bilateral breakdown of their relationship.

Lemma 1 *Fix equilibrium σ^* and an on-path history $(h^t, n_{t+1}) \in \mathcal{H}$. Consider two histories $h^{t+1}, \tilde{h}^{t+1} \in \mathcal{H}_B^{t+1}$ that are successors to (h^t, n_{t+1}) such that $h^{t+1} \in supp\{\sigma | h^t, n_t\}$ but $\tilde{h}^{t+1} \notin supp\{\sigma | h^t, n_t\}$, and suppose that $\tilde{h}_j^{t+1} = h_j^{t+1}$, $\forall j \notin \{0, i\}$. In the continuation game, the payoffs of the principal and agent i satisfy*

$$E_{\sigma^*}\left[U_i(\tilde{h}^{t+1}) | \tilde{h}_i^{t+1}\right] \geq 0, \tag{1.4}$$

$$U_0(\tilde{h}^{t+1}) \geq \max_{\hat{h}^{t+1} | \forall j \neq i, \hat{h}_j^{t+1} = \tilde{h}_j^{t+1}} \left[\sum_{j \neq i} U_0^j(\hat{h}^t)\right] \geq \sum_{j \neq i} U_0^j(h^t). \tag{1.5}$$

[15] Analogous to Abreu (1988).

Proof:

If (1.4) is not satisfied, then agent i could choose $w_{i,t}^P = \tau_{i,t}^P = 0$, $e_t = 0$, and $d_t = 0$ in each period, earning strictly higher surplus. Contradiction.

Define $\mathcal{H}_i(\tilde{h}^{t+1}) = \left\{\hat{h}^{t+1} | \forall j \neq i, \hat{h}_j^{t+1} = \tilde{h}_j^{t+1}\right\}$. If (1.5) is not satisfied, let

$$\hat{h}^{t+1} \in \arg \max_{h^{t+1} \in \mathcal{H}_i(\tilde{h}^{t+1})} U_0(h^{t+1}) - U_0^i(h^{t+1}).$$

Recursively define a feasible continuation strategy for the principal from $\tilde{h}^{t+1}$. $\forall t' \geq 1$, the principal plays $\sigma_0^*(\hat{h}^{t+t'})$, with the sole exception that $w_{i,t+t'}^A = \tau_{i,t+t'}^A = 0$. Let $h^{t+t'}$ be the observed history at the end of round $t+t'$. The principal chooses $\hat{h}^{t+t'+1}$ according to the distribution of length $t+t'+1$ histories induced by $\sigma_0^*(\hat{h}^{t+t'})$, conditional on the event $\hat{h}_j^{t+t'+1} = h_j^{t+t'+1}$ $\forall j \neq i$. This conditional distribution is well-defined because the specified strategy differs from $\sigma_0^*(\hat{h}^{t+t'})$ only in variables that are observed only by the principal and agent i, and $\hat{h}_j^{t+t'} = h_j^{t+t'}$ $\forall j \neq i$ by construction.

Under this strategy, agents $j \neq i$ cannot distinguish $\tilde{h}^{t+t'}$ and $\hat{h}^{t+t'}$ for any $t' \geq 1$, so the principal earns at least $E_{\sigma^*}\left[\sum_{j \neq i}^N -(w_{j,t+t'} + \tau_{j,t+t'}) | \hat{h}^{t+t'}\right]$ if $x_{t+t'} = i$ and $E_{\sigma^*}\left[u_0^{t+t'} + w_{i,t+t'} + \tau_{i,t+t'} | \hat{h}^{t+t'}\right]$ if $x_{t+t'} \neq i$. Agent i is awarded production with the same probability as under $\sigma_0^*(\hat{h}^{t+1})$, so the principal's payment is bounded from below by $U_0(\hat{h}^{t+1}) - U_0^i(\hat{h}^{t+1})$. ■

Intuitively, following a deviation that is observed by only agent i and the principal, the harshest punishment for both parties is the termination of trade in the i^{th} dyad. Terminating trade holds agent i at his outside option, which is his min-max continuation payoff. However, the principal remains free to trade with the other agents because the deviation was not observed by these agents. In particular, the principal can act as if the true history is any $\hat{h}^t$, so long as that history is consistent with the beliefs of agents $j \neq i$. By choosing $\hat{h}^t$ to maximize her surplus, the principal exploits the other agents' limited information to mitigate the harm of i's refusal to trade.

Building on Lemma 1, Lemma 2 proves that agent i's beliefs about his dyad-specific surplus S_i determine whether he works hard in equilibrium. Following high output, an agent can be credibly promised no more than S_i total surplus. Following low output, transfers can hold an agent at his outside option 0. As a result, equilibrium conditions are summarized in a set of intuitive incentive constraints that must be satisfied whenever an agent both *works hard* and *produces high output.*

Lemma 2 *1. Let σ^* be a PBE. Then*

$$
\begin{gathered}
(1-\delta)\frac{c}{p_1-p_0} \leq \delta E_{\sigma^*}\left[S_i(h^{t+1}) \mid \left(h_i^{t-1}, \mathcal{I}_i(n_t)\right)\right], \\
\forall i \in \{1,...,N\}, \forall h^{t-1} \in \mathcal{H}, n_t \text{ on the equilibrium path immediately after } y_t \\
\text{such that } x_t = i,\ e_t = 1,\ y_t = y_H.
\end{gathered}
\tag{1.6}
$$

2. Let σ be a strategy that generates total surplus V and satisfies (1.6). Then $\exists$ a PBE σ^* that generates the same total surplus V and joint distribution of $\{\mathcal{P}_t, x_t, d_t e_t, y_t\}_{t=1}^{T}$ as σ for all T.

Proof:

We prove Statement 1 here and defer the proof of Statement 2 to Appendix 1.9.

For any t and $h^{t+1} \in \mathcal{H}_B^{t+1}$, define

$$
D_i(h^{t+1}) = U_0(h^{t+1}) - \max_{\hat{h}^{t+1} \mid \forall j \neq i, \hat{h}_j^{t+1} = h_j^{t+1}} \left[U_0(\hat{h}^{t+1}) - U_0^i(\hat{h}^{t+1})\right]. \tag{1.7}
$$

By Lemma 1, the principal is punished by no more than $D_i(h^{t+1})$ if she deviates from a history h^{t+1} to $\hat{h}^{t+1}$ in round t and it is only observed by agent i. Note that $D_i(h^{t+1}) \leq U_0^i(h^{t+1})$ by definition.

Fix history $(h^{t-1}, \mathcal{P}_t, x_t = i, \{w_{i,t}\}, d_t = 1, e_t = 1, y_t) = (h^{t-1}, n_t)$, and let $\tau_i(y_t) = E_{\sigma^*}\left[\tau_{i,t} \mid h_i^{t-1}, \mathcal{I}_i(n_t)\right]$, $U_i(y_t) = E_{\sigma^*}\left[U_i \mid h_i^{t-1}, \mathcal{I}_i(n_t)\right]$ be agent i's expected transfer and continuation payoff following output y_t. On the equilibrium path, agent i chooses $e_t = 1$ only if

$$
\begin{aligned}
&p_1[(1-\delta)\tau_i(y_H) + \delta U_i(y_H)] + (1-p_1)[(1-\delta)\tau_i(0) + \delta U_i(0)] - (1-\delta)c \geq \\
&\quad p_0[(1-\delta)\tau_i(y_H) + \delta U_i(y_H)] + (1-p_0)[(1-\delta)\tau_i(0) + \delta U_i(0)].
\end{aligned}
\tag{1.8}
$$

Let $\tilde{h}^t$ be the history reached by choosing $\tau_{i,t} = 0$ rather than the scheduled transfer in round

t. For bonus $\tau_{i,t}$ to be paid in equilibrium, we must have

$$\begin{aligned}(1-\delta)E_\sigma\left[\tau_{i,t}|h^{t-1},n_t,y_t\right] &\leq \delta E_\sigma\left[D_i(\tilde{h}^t)|h^{t-1},n_t,y_t\right]\\ -(1-\delta)E_\sigma\left[\tau_{i,t}|h^{t-1},n_t,y_t\right] &\leq \delta E_\sigma\left[U_i(h^t)\mid h_i^{t-1},\mathcal{I}_i(n_t,y_t)\right]\end{aligned} \tag{1.9}$$

or else the principal or agent i would choose not to pay $\tau_{i,t}$. Plugging these constraints into (1.8) and applying the definition of S_i results in (1.6). ■

Define $\tilde{S} = \frac{1-\delta}{\delta}\frac{c}{p_1-p_0}$ as the minimum dyad-specific surplus that agent x_t must believe he receives following high effort and output. Statement 1 is analogous to MacLeod and Malcomson's (1989, Proposition 1) and Levin's (2003, Theorem 3) arguments that reneging constraints on transfers (1.9) can be aggregated across principal and agent.

To prove Statement 2, we construct an equilibrium σ^* that uses the same allocation rule, accept/reject decisions, and effort choices as the given strategy σ. In order to make these actions part of an equilibrium, we have to solve three problems. First, the principal must be willing to follow the proposed allocation rule x_t^*. Second, each agent i must be willing to work hard whenever condition (1.6) holds. Third, because (1.6) depends on agent i's information, we must ensure that each agent knows the appropriate amount of information about the true history.

We use wages and bonus payments to solve these three problems. In each period of σ^*, the producing agent x_t is paid an efficiency wage $w_{x_t,t} = d_t^* y_H p_{e_t^*} + d_t^* e_t^* \frac{c}{p_1-p_0}$. This payment ensures that the principal earns 0 in every period and so is willing to implement the chosen allocation rule, while agent x_t earns the entire expected surplus produced in that period. Agent x_t can also infer the desired accept/reject decision d_t^* and effort level e_t^*. Following the agent's accept or reject decision and effort, no bonus is paid unless the agent accepted production, was supposed to work hard, and produced $y_t = 0$. In that case, agent x_t pays his continuation surplus back to the principal, earning 0.[16] So long as the original strategies satisfy (1.6), these payments communicate exactly enough information to ensure that x_t is

[16]These transfers may seem unusual, since the principal pays a large up-front wage to the producing agent and is repaid if output is low. In the proof, these transfers ensure both that the principal is willing to implement the desired allocation rule, and that any deviations are immediately detected and punished as harshly as possible. Other transfer schemes may also work, but all relational contracts must contend with the same constraint (1.6).

willing to both work hard when asked and repay the principal if he produces low output. A deviation in any variable other than effort is observed by the principal and punished with the termination of trade.

Agent i's expected dyad-specific surplus S_i determines whether he can be induced to work hard, and S_i is determined in turn by the equilibrium allocation rule. From (1.6), it is clear that agent i is only willing to work hard if S_i is sufficiently large following $y_t = y_H$, and hence the principal's allocation rule matters only when the agent both *work hards* and *produces high output*. This result highlights the interaction between *wages and bonus payments* and the *allocation of business*. Following low output, an agent can be held at his outside option through short-term transfers, regardless of how much business he is promised in the future. In contrast, following high output an agent can be credibly promised only S_i continuation surplus. Because the agent's *beliefs*—which are typically coarser than the true history—determine whether he works hard, the principal must consider both the allocation rule and agent expectations. Therefore, private monitoring potentially shapes both off-equilibrium punishments and actions on the equilibrium path.

Because the condition (1.6) depends only on the allocation rule, accept/reject decisions, and effort choices, we define an optimal relational contract in terms of these variables in the rest of the analysis. Moreover, a corollary of this lemma is that the set of equilibrium payoffs would not be changed if the principal could send a private, costless message to each agent. Any such communication can be replicated in the baseline game using appropriately-chosen wage and bonus payments.[17] Together, Lemmas 1 and 2 underscore that favoring an agent with future business *following high effort and output* makes it easier to motivate him, at the cost of reducing the future business that can be promised to other agents. This trade-off shapes the optimal dynamic relational contract.

[17]Of course, the set of equilibrium outcomes *would* change substantially if agents were allowed to directly communicate. This case is studied in Section 1.7.

1.4 An Optimal Allocation Rule if First-Best is Attainable

This section considers relational contracts that induce first-best effort. Stationary allocation rules are effective only when players are patient, but we introduce a simple non-stationary equilibrium that attains first-best for a strictly larger set of parameters than any stationary allocation rule. In this equilibrium, agents are favored with additional future business when they perform well and remain favored if they later perform poorly, but fall from favor when they are unable to produce and their replacement performs well.

Lemma 2 implies that S_i determines the bonuses that can be credibly promised to agent i. If δ is close to 1, players care tremendously about the future and this collateral is abundant regardless of the allocation rule. Define a stationary allocation rule as one in which actions do not depend on previous rounds along the equilibrium path; then a relational contract with a stationary allocation attains first-best when players are patient.[18]

Proposition 2 *There exists a stationary equilibrium with surplus* V^{FB} *if and only if*

$$(1-\delta)\frac{c}{p_1-p_0} \leq \delta\frac{1}{N}V^{FB}. \tag{1.10}$$

Proof:

Consider the following stationary allocation rule, accept/reject decision, and effort choice: in each round, $\text{Prob}_\sigma\{x_t = i|\mathcal{P}_t\} = \left\{\frac{1}{|\mathcal{P}_t|} \text{ if } i \in \mathcal{P}_t, 0 \text{ otherwise}\right\}$, while agents choose $d_t = e_t = 1$. Agents are symmetric, so $x_t = i$ with *ex ante* probability $\frac{1}{N}(1-F(\emptyset))$. By Lemma 2, these actions are part of an equilibrium that attains first-best if and only if

$$(1-\delta)\frac{c}{p_1-p_0} \leq \delta\frac{1}{N}(1-F(\emptyset))(y_Hp_1-c),$$

which proves that (1.10) is a sufficient condition to induce first-best with stationary contracts.

To prove that (1.10) is also necessary, note that for any stationary allocation rule there

[18]Given a stationary allocation rule, the optimal contract within each dyad is very similar to that in Levin (2003), Theorem 6.

exists some i such that $\text{Prob}\{x_{Stat} = i|\mathcal{P} \neq \emptyset\} \leq \frac{1}{N}$ because $\sum_i \text{Prob}\{x_{Stat} = i|\mathcal{P} \neq \emptyset\} \leq 1$. If (1.10) does not hold, then i always chooses $e_t = 0$. But $F(\{i\}) > 0$, so first-best cannot be attained. ■

In a stationary relational contract, the allocation rule does not evolve in response to past performance. While stationary relational contracts attain V^{FB} only if players are patient, a non-stationary allocation rule can induce first-best effort even when (1.10) does not hold. We introduce the *Favored Producer Allocation*—illustrated in Figure 1—and prove that it induces first-best effort whenever any equilibrium does.

In the Favored Producer Allocation, the principal ranks the agents from $1, ..., N$ and awards production to the "most favored" available agent—the $i \in \mathcal{P}_t$ with the lowest rank. If that agent produces low output, the rankings remain unchanged, while he immediately becomes the most favored agent following high output. This allocation rule rewards success because an agent who produces y_H is immediately promised a larger share of future business, and it is tolerant of failure in the sense that a favored agent remains so even if he performs poorly. Once an agent is favored he loses that favor only when another agent performs well and replaces him.[19] In each round, every agent is the sole available producer with positive probability and so every agent has the opportunity to become favored. The resulting dynamics resemble a tournament in which the most recent agent to perform well "wins" favored status.

Formally, the Favored Producer Automaton can be described as follows.

Definition 3 *Let ϕ be an arbitrary permutation of $\{1, ..., N\}$ and ϕ^I be the identity mapping. The* Favored Producer Allocation *is defined by:*

1. *Set $\phi_1 = \phi^I$.*

2. *In each round t, $x_t = \arg\min_{i \in \{1,..,N\}} \{\phi_t(i) | i \in \mathcal{P}_t\}$.*

3. *If $y_t = 0$: $\phi_{t+1} = \phi_t$. If $y_t = y_H$: $\phi_{t+1}(x_t) = 1$, $\phi_{t+1}(i) = \phi_t(i) + 1$ if $\phi_t(i) < \phi_t(x_t)$, and $\phi_{t+1}(i) = \phi_t(i)$ otherwise.*

[19] Interestingly, there is some anecdotal evidence that downstream firms tend not to withdraw business from a poorly-performing supplier. For instance, Kulp and Narayanan (2004) report that one supplier "thought it unlikely that Metalcraft would pull business if a given supplier's score dropped below acceptable levels."

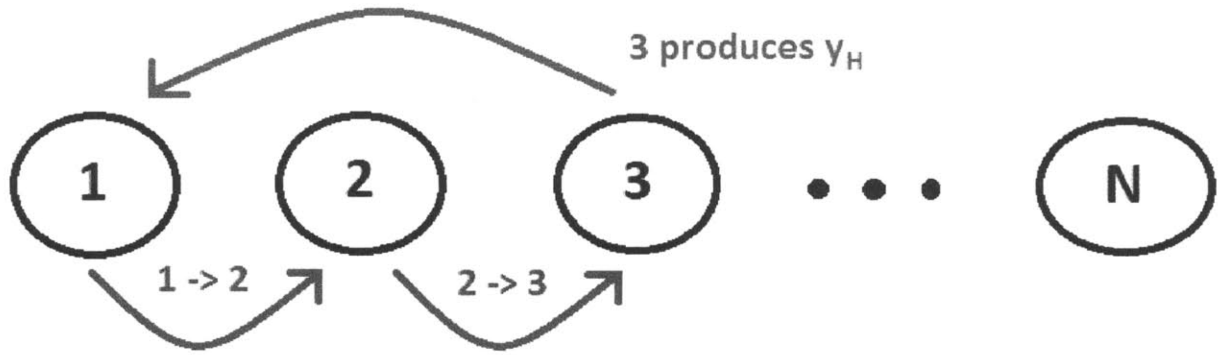

Figure 1.1: The Favored Producer Allocation. In this picture, the agents ranked 1 and 2 are not available, so production is given to the third-ranked agent. This agent produces y_H and so displaces the top-ranked agent.

4. On the equilibrium path, $d_t = e_t = 1$ iff (1.6) holds, and otherwise $d_t = 1$, $e_t = 0$.

Proposition 3 proves that a relational contract using the Favored Producer Allocation attains first-best whenever any equilibrium does.

Proposition 3 *Suppose V^{FB} is attainable in a PBE. Then the Favored Producer Allocation is part of a PBE that generates V^{FB}. Moreover, $\exists$ a non-empty, open $\Delta \subseteq [0,1]$ such that if $\delta \in \Delta$, the Favored Producer Allocation attains first-best but no stationary equilibrium does.*

Proof:

See Appendix 1.9.

Before discussing the proof, let's consider why the Favored Producer Allocation rule is (a) tolerant of failure and (b) rewards success. For (a), consider a harsher allocation rule that withdraws business from agent i whenever he produces $y_t = 0$. Typically, i is favored because he produced high output at some point in the past. Thus, this harsher allocation rule would tighten (1.6) when agent i produced y_H. At the same time, the harsher allocation rule would *not* relax this constraint for the other agents to the same extent because (1.6) depends only on histories following *high output.* An optimal allocation rule is tolerant of failure precisely because the promise of future business does not simply serve as a direct incentive for effort, but determines what incentive payments are credible. Transfers ensure that regardless of his favored status, an agent earns 0 if he produces poorly. An agent can be credibly promised S_i after high output, which is larger if that agent is also favored in the future.

To show (b), compare the Favored Producer Allocation to a stationary relational contract. Supposing that first-best is attained, the allocation rule determines *what fraction* of the first-best surplus is produced by each agent at every history. In a (symmetric) stationary equilibrium, an agent who performs well has $S_i = \frac{1}{N}V^{FB}$. In contrast, an agent who produces y_H in the Favored Producer Allocation has $S_i > \frac{1}{N}V^{FB}$, since this agent is more likely to produce in each future period and so is likely to retain favored status.

The Favored Producer Allocation is simple enough that we can explicitly calculate when first-best can be attained. Fixing $F(\emptyset)$, the region for which first-best is attained is increasing in the probability of both agents are available. Intuitively, the allocation of business plays a more important role in the relational contract and first best is easier to attain when the probability that *both* agents are available $F(\{1,2\})$ is large.[20] This comparative static has important implications for real-world relationships: for example, frequent (independent) disruptions in a supply chain decrease both total surplus and the probability that both agents are available and so *a fortiori* make it difficult to motivate effort.

We now turn to the proof of Proposition 3. Because PBE strategies depend on private histories—which grow increasingly complex over time—this game is not amenable to a recursive analysis. Instead, our analysis focuses on the relatively easy-to-compute *beliefs at the beginning of the game.* Using this technique, we derive necessary conditions for first-best to be attained in a PBE and show that the Favored Producer Allocation attains first-best whenever these conditions hold. We will focus on intuition for the proof here; details may be found in Appendix 1.9.

The basic idea of the proof is to sum up the obligation owed to each agent—that is, the amount of future business that has been promised to an agent in expectation in any given round. By Lemma 2, agent x_t must believe his dyad-specific surplus is at least $\tilde{S}$ whenever he chooses $e_t = 1$ and produces $y_t = y_H$. We relax this constraint so that it must hold only the *first time* each agent produces y_H. This obligation is paid off when the principal

[20] At $F(\{1,2\}) = 1$ (which is ruled out by Assumption 2), the Favored Producer Allocation collapses to a stationary equilibrium and so stationary equilibria are optimal.

allocates business to that agent in the future. Therefore, we define the *expected obligation* (denoted Ω_t^i) owed to agent i at time t as the net *ex ante* expected future business that must been promised to agent i to motivate first-best effort in rounds $1, ..., t$.

Definition 4 *Given strategies σ, define*

$$\beta_{i,t}^L = Prob_\sigma \left\{ \{\nexists t' < t \text{ s.t. } x_{t'} = i, y_{t'} = y_H\} \cap \{x_t = i\} \right\},$$
$$\beta_{i,t}^H = Prob_\sigma \left\{ \{\exists t' < t \text{ s.t. } x_{t'} = i, y_{t'} = y_H\} \cap \{x_t = i\} \right\}.$$

Then the expected obligation owed to agent i at time t, Ω_t^i, *is*

$$\Omega_t^i \equiv \frac{\Omega_{t-1}^i}{\delta} + \beta_{i,t}^L p_1 \delta \tilde{S} - \beta_{i,t}^H (1-\delta)(y_H p_1 - c) \tag{1.11}$$

with initial condition $\Omega_0^i \equiv 0$.

Given a proposed equilibrium with first-best effort σ, Ω_t^i is a state variable that tracks how much expected surplus is "owed" to agent i in round t. At $t = 1$, agent i is allocated production with probability $\beta_{i,1}^L$, produces y_H with probability p_1, and must be promised $\delta\tilde{S}$ future surplus *in expectation* following y_H to satisfy (1.6). Therefore, $\Omega_1^i \equiv \beta_{i,t}^L p_1 \delta \tilde{S}$ equals the expected continuation surplus that must be promised to agent i following the first period. At $t = 2$, agent i is still owed this initial amount—now worth $\frac{\Omega_1^i}{\delta}$ due to discounting—and accumulates an additional obligation $\beta_{i,2}^L p_1 \delta \tilde{S}$ from those histories in which he produces y_H for the first time in this round. If i y_H in $t = 1$ and is allocated production at $t = 2$—which occurs with probability $\beta_{i,2}^H$—then his obligation can be "paid off" at a rate equal to the expected surplus from production, $(1-\delta)(y_H p_1 - c)$. Putting these pieces together, the expected obligation promised to agent i in round 2 is

$$\frac{\Omega_1^i}{\delta} + \beta_{i,2}^L p_1 \delta \tilde{S} - \beta_{i,2}^H (1-\delta)(y_H p_1 - c),$$

which equals Ω_2^i. A similar intuition applies for each t.

If $\Omega_t^i \to \infty$ as $t \to \infty$, then agent i incurs obligation faster *in expectation* than it could possibly be paid off by promises of future business. In Appendix 1.9, we show that $\lim_t \Omega_t^i < \infty$ is a *necessary* condition for the proposed strategy σ to be an equlibrium.

Because Ω_t^i is constructed by taking an expectation across histories, this necessary condition is independent of agent beliefs. While concealing information from agent i might alter the incentive constraint (1.6), it cannot *systematically* trick agent i into expecting more future business than is possible to provide *in expectation.*

Given this result, it suffices to find conditions under which $\lim_{t\to\infty}\Omega_t^i = \infty$ for *every* set of $\{\beta_{i,t}^L, \beta_{i,t}^H\}_{i,t}$ that are consistent with a feasible allocation rule. Because the agents are symmetric, it turns out that obligation is minimized when we treat all agents who have produced y_H in the past symmetrically, and likewise with all those who have not yet produced y_H. Moreover, those agents who have already produced y_H should be allocated business whenever possible, since doing so pays off their obligation and also minimizes the obligation incurred by the other agents. Therefore, we can exactly pin down the $\{\beta_{i,t}^L, \beta_{i,t}^H\}_{i,t}$ that minimize Ω_t^i for every agent.

To complete the proof, we show that if the Favored Producer Allocation does not attain first-best, then $\lim_{t\to} \Omega_t^i = \infty$ and hence first-best is unattainable by *any* PBE. Intuitively, the Favored Producer Allocation gives an agent who has produced y_H the maximal fraction of future surplus, subject to the constraints that $e_t = 1$ in every period and (1.6) is satisfied. Given that agents cannot be tricked into believing that total surplus is larger than V^{FB} this allocation minimizes the probability that non-favored agents are awarded production, which maximizes the likelihood that a favored agent remains favored. Even when a favored agent is replaced, his punishment is mild because he retains priority whenever more-favored agents are unable to produce.

In one important sense, agents' private information does not play a role in the Favored Producer Allocation: (1.6) is satisfied at the true history h^t on the equilibrium path, so each agent would be willing to follow his strategy even if he learned the true history. We refer to any equilibrium with this property as a *full-disclosure equilibrium.* Formally, full-disclosure equilibria provide *ex post* incentives to each player and so are belief-free.[21] In the online Appendix, we show that a strategy profile generates the same total surplus as a belief-free equilibrium if and only if (1.6) holds at the true history, so we define full-disclosure

[21] As introduced by Ely and Valimaki (2002) and Ely, Horner, and Olszewski (2005). We use a definition of belief-free equilibrium given in the online Appendix.

equilibrium in terms of this condition.

Definition 5 *A PBE σ^* is a* full-disclosure equilibrium *(FDE) if* $\forall\ i \in \{1, ..., N\}$,

$$\begin{gathered}(1-\delta)\frac{c}{p_1 - p_0} \leq \delta E_{\sigma^*}\left[S_i(h^t)|h^{t-1}, n_t\right] \\ \forall i \in \{1, ..., N\}, \forall h^{t-1} \in \mathcal{H}, n_t \text{ on the equilibrium path immediately after } y_t \\ \text{such that } x_t = i,\ e_t = 1,\ y_t = y_H.\end{gathered} \tag{1.12}$$

If (1.12) does not hold, then σ^* conceals information from the agents. *The set of full disclosure equilibrium payoffs is $FD(\delta) \subseteq PBE(\delta) \subseteq \mathbb{R}^{N+1}$.*

The sole difference between (1.12) and (1.6) is that (1.12) conditions on the true history rather than agent i's (coarser) information set. Because full-disclosure relational contracts provide *ex post* incentives for effort, each agent is willing to follow the equilibrium even if he learns additional information about the true history. That is, these relational contracts are robust to the agent learning more information about past play in some (unmodeled) way, for instance by reading trade journals or the newspaper. Proposition 3 implies that if first-best is attainable, then there exists an equilibrium that induces first-best using *ex post* incentives.

Corollary 1 *Suppose $\exists$ a PBE σ^* that generates surplus V^{FB}. Then $\exists$ a full-disclosure equilibrium σ^{FD} that generates V^{FB}.*

Proof:

Condition (1.12) holds by direct computation in the Favored Producer Allocation.■

In this section, we have explicitly characterized a relational contract that attains first-best whenever any Perfect Bayesian Equilibrium does. This allocation rule rewards success and resembles a tournament: agents compete for a temporary favored status that lasts until another agent produces high output. It is also tolerant of failure: rankings are unaffected when an agent produces low output. In this way, the principal ensures that she can credibly promise a large relational bonus to an agent that performs well.

1.5 Unsustainable Networks and Relationship Breakdown

We now turn to the case with two agents and consider relational contracts if first-best is unattainable. In this case, non-stationary relational contracts typically dominate stationary contracts, but the principal cannot always induce every agent to work hard. In particular, the principal might forego providing effective incentives to one agent to ensure that another agent is promised enough future business to motivate him. In the class of full-disclosure relational contracts with two agents, a variant of the Favored Producer Allocation turns out to be (non-uniquely) optimal: if an agent produces high output, he sometimes enters an exclusive relationship in which the principal permanently favors him with future business. Once this occurs, the other agent shirks.

So long as first-best is unattainable but some cooperation is possible, Proposition 4 proves that every optimal relational contract tailors the allocation of business to past performance.

Proposition 4 *Suppose V^{FB} cannot be attained in a PBE and $\tilde{S} \leq \sum_{\mathcal{P}|i\in\mathcal{P}} F(\mathcal{P})(y_H p_1 - c)$. Then for any stationary equilibrium σ^{Stat}, there exists an equilibrium σ^* that generates strictly higher surplus.*

Proof:

Let x_{Stat} be an optimal stationary allocation rule. Because $\tilde{S} \leq \sum_{\mathcal{P}|i\in\mathcal{P}} F(\mathcal{P})(y_H p_1 - c)$, it must be that $\tilde{S} \leq (y_H p_1 - c)\text{Prob}\{x_{Stat} = i\}$ holds for $i \in \mathcal{M}_{Stat}$, where $0 < |\mathcal{M}_{Stat}| < N$ agents. Only $i \in \mathcal{M}_{Stat}$ choose $e = 1$ in equilibrium. Consider the non-stationary equilibrium that chooses a set of agents $\mathcal{M}(\mathcal{P}_1)$ with $|\mathcal{M}(\mathcal{P}_1)| = |\mathcal{M}_{Stat}|$ and $\mathcal{M}(\mathcal{P}_1) \cap \mathcal{P}_1 \neq \emptyset$, then allocates production to the agents in $\mathcal{M}(\mathcal{P}_1)$ as in $\mathcal{M}_{Stat}$. For $t > 1$, this non-stationary equilibrium generates the same surplus as the stationary equilibrium; for $t = 1$, it generates strictly higher surplus, since $\text{Prob}\{\mathcal{P}_1 \cap \mathcal{M}_{Stat} = \emptyset\} > 0$ by assumption 2. ■

Whenever one agent works hard and produces high output in the relational contract, the principal must promise that agent a large amount of future business. One way to satisfy this promise is to forego providing incentives to the *other agent* at some histories. For this reason, shirking occurs on the equilibrium path if first-best cannot be attained. Our next goal is to characterize *when* shirking will occur, and show that a relational contract may optimally entail one relationship devolving into perfunctory performance.

For the rest of the section, we restrict attention to full-disclosure relational contracts. This is a substantial restriction that allows us to use recursive techniques to characterize equilibrium. These relational contracts also have the valuable property that they do not rely on subtle features of an agent's beliefs. In a real-world environment, it might be difficult to completely stop an agent from learning unmodeled additional information about the true history; full-disclosure relational contracts are robust to this additional information.[22]

Among full-disclosure relational contracts, it turns out that a simple variant of the Favored Producer Allocation is non-uniquely optimal. In this relational contract, the principal's relationship with an agent might eventually become perfunctory: while both agents work hard at the beginning of the game, as $t \to \infty$ it is almost surely the case that one of the agents chooses $e_t = 0$ whenever he produces. Hence, the principal sacrifices one relationship in order to provide adequate incentives in the other. The principal continues to rely on this perfunctory relationship when no better alternatives exist because $y_H p_0 \geq 0$, but she offers no incentive pay and has low expectations about output.

Definition 6 *Let $N = 2$. The (q_1, q_2)-Exclusive Dealing allocation rule is:*

1. *Begin the game in state G_1. In state G_i, $Prob_\sigma\{x_t = i | i \in \mathcal{P}_t\} = 1$, $Prob_\sigma\{x_t = -i | \mathcal{P}_t = \{-i\}\} = 1$, and both agents choose $e_t = 1$. If $y_t = y_H$, transition to ED_{x_t} with probability $q_{x_t} \geq 0$, otherwise transition to state G_{x_t}. If $y_t = 0$, stay in G_i.*

2. *In state ED_i, $Prob_\sigma\{x_t = i | i \in \mathcal{P}_t\} = 1$ and $Prob_\sigma\{x_t = -i | \mathcal{P}_t = \{-i\}\} = 1$. If $x_t = i$, $e_t = 1$; otherwise, $e_t = 0$. Once in ED_i, remain in ED_i.*

We refer to continuation play in state ED_i as exclusive dealing.

In (q_1, q_2)-Exclusive Dealing, each agent faces the possibility that his relationship with the principal breaks down at some time in the future. Before breakdown occurs, the alloca-

[22] In the context of a collusion model with adverse selection rather than moral hazard, Miller (2012) argues that *ex post* incentives are natural for this reason.

tion rule is the same as in the Favored Producer Allocation and both agents are expected to choose $e = 1$. Once agent i enters exclusive dealing—which happens with probability q_i whenever i produces y_H—agent $-i$ stops exerting effort and his relationship with the principal becomes perfunctory. Like the Favored Producer Allocation, (q_1, q_2)-Exclusive Dealing both rewards success and tolerates failure: high output is rewarded by both favored status and the possibility of a permanent relationship, while low output does not change the allocation rule but leaves the door open for the other agent to win exclusive dealing.

Proposition 5 shows that (q^*, q^*)-Exclusive Dealing is optimal among full-disclosure equilibria for appropriate $q^* \in [0, 1]$. A critical caveat is that it is not uniquely optimal: there exist other allocation rules that work, including some that do not involve any permanent break-down of a relationship.[23]

Proposition 5 *Let $N = 2$. $\exists q^* \in [0, 1]$ such that the (q^*, q^*)-Exclusive Dealing equilibrium is an optimal full-disclosure equilibrium.*

Proof:

See Appendix 1.9.

While the proof is lengthy, the intuition for this result is straightforward. Consider the continuation game immediately after agent 1 works hard and produces $y_t = y_H$. Total surplus is higher if agent 2 works hard. However, agent 2 is only willing to work hard if he is promised a substantial fraction of future business, which makes it more difficult to give agent 1 $\tilde{S}$ dyad-specific surplus. It turns out that this benefit and cost of agent 2 working hard both scale at the rate $\delta^{t'-t}$ over time. Therefore, an optimal relational contract simply maximizes the sum of discounted probabilities that agent 2 works hard, subject to satisfying agent 1's incentive constraint. In particular, there exists a $q^* \in [0, 1]$ such that (q^*, q^*)-Exclusive Dealing solves this problem.

Exclusive dealing is certainly not uniquely optimal, but it is interesting that an optimal relational contract might require the permanent break-down of one relationship in order to

[23] For example, an allocation rule that grants *temporary* exclusive dealing to a high performer for K periods immediately following y_H is also optimal (subject to integer constraints on K).

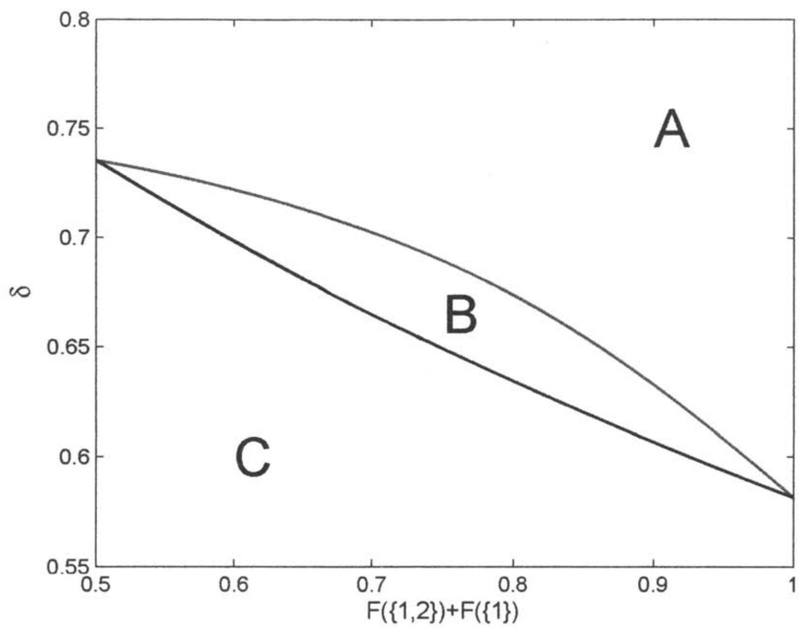

Figure 1.2: Optimal equilibria for $y_H = 10$, $p_1 = 0.9$, $p_0 = 0$, $c = 5$, $F(\emptyset) = 0$. (A) the Favored Producer Allocation can attain first-best; (B) (q^*, q^*)-Exclusive Dealing is an optimal FDE, and non-stationary equilibria strictly dominate stationary equilibria; (C) no effort can be supported in equilibrium.

credibly motivate the *other* agent.[24] This result is consistent with the observation by Cyert and March (1963) that individuals may use policy commitments to compensate one another for past actions: the principal commits to a long-term inefficient policy to solve a short-run moral hazard problem. Figure 2 illustrates the implications of Propositions 3, 4, and 5 in a game with two agents.

Agents are *ex ante* identical in the baseline model, so the *identity* of the agent whose relationship sours does not affect total continuation surplus. If agent 1 is instead more productive than agent 2—so that high output for agent 1 is $y_H + \Delta_y > y_H$—then which relationship breaks down influences long-run profitability. In the online Appendix, we show that the proof of Proposition 5 extends to this asymmetric case for some parameters, implying that (q_1, q_2)-Exclusive dealing is optimal (albeit with $q_1 \neq q_2$). If Δ_y is not too large,

[24]As in Proposition 1 of Board (2012), this result can be interpreted as saying it is optimal for the principal to separate agents into "insiders" and "outsiders" and be biased against the "outsiders." Unlike Board, we focus on a moral hazard problem, which implies that the allocation rule is tolerant of failure and looks like a tournament between the agents.

the principal might optimally enter an exclusive relationship with *either* agent, so that *ex ante* identical principal-agent networks exhibit *persistent differences in long-run productivity.* In a market of such principal-agents groups, an outside observe would observe persistent differences in productivity, even though each group was playing an *ex ante* identical optimal relational contract.[25]

1.6 When is it Optimal to Conceal Information?

Corollary 1 shows that concealing information is unnecessary when first-best can be attained, and Proposition 5 shows that Exclusive Dealing is an optimal full-disclosure equilibria otherwise. In this section, we prove that providing full-disclosure incentives is typically costly if first-best is unattainable: the principal might do even better by not providing *ex post* incentives and keeping agents in the dark.[26]

Full-disclosure relational contracts have the great advantage of being simple and providing robust incentives for effort. In contrast, the relational contract we construct to show that concealing information can be optimal is more complicated and relies on relatively subtle features of the agents' beliefs. As a result, a full-disclosure relational contract may be easier to implement in practice, even if the principal could theoretically earn a higher surplus by concealing information. Nevertheless, some firms appear to be secretive about their relational scorecards—Farlow et al (1996) report that Sun Microsystems used to withhold the details of its relational scorecards from suppliers.[27]

Proposition 6 proves that concealing information is optimal whenever first-best is unattainable but agents can be strictly motivated to exert effort.

Proposition 6 *Let $N = 2$ and suppose that first-best cannot be attained in a PBE but that*

$$\frac{c}{p_1 - p_0} < \frac{\delta}{1-\delta}\left(F(\{1\}) + F(\{1,2\})\right)\left(y_H p_1 - c\right). \tag{1.13}$$

[25]Persistent performance differences among seemingly similar companies are discussed in Gibbons and Henderson (2012).

[26]Using the definition of belief-free equilibria given in Appendix B, the equilibrium used to prove this result is weakly belief-free in a sense analogous to Kandori (2011), but not belief-free.

[27]Sun did eventually reveal the details of these scorecards, but only so that their suppliers could adopt the same scorecard to manage *their* own (second-tier) relationships.

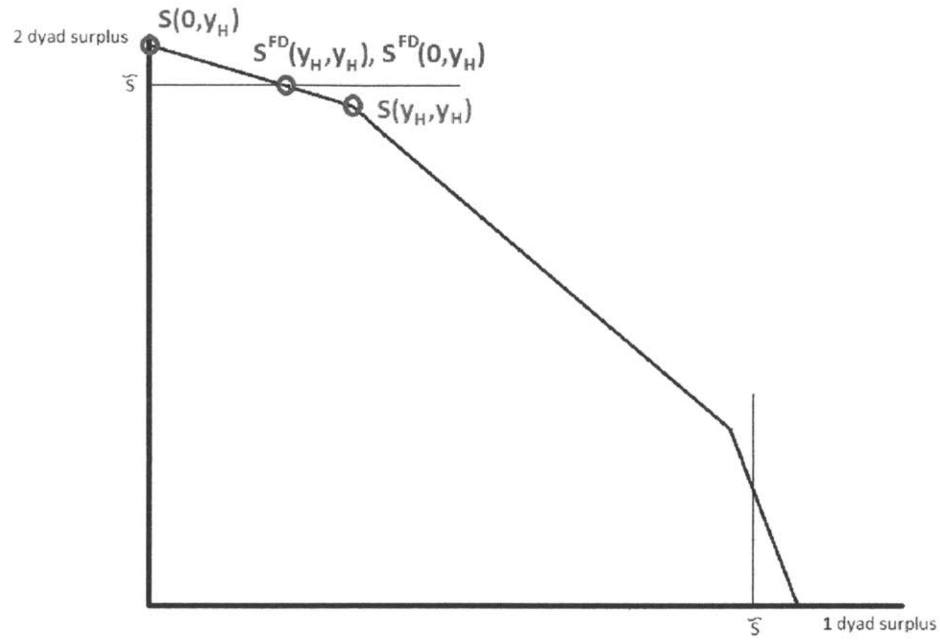

Figure 1.3: Set of full-disclosure dyad-specific surpluses when $V^{FB} \notin PBE(\delta)$

Let σ^* be a full-disclosure equilibrium; then is it not an optimal PBE.

Proof:

See Appendix 1.9.

The proof of Proposition 6 constructs an equilibrium that dominates (q^*, q^*)-Exclusive Dealing by concealing information. In such an equilibrium, the full-disclosure constraint (1.12) need not hold in every history. If first-best is unattainable, the optimal relational contract specifies inefficient continuation play, and these inefficiencies can be mitigated by allowing the principal to conceal information from the agents.

More precisely, Figure 3 illustrates the set of dyad-specific surpluses that can be sustained in a full-disclosure equilibrium. Consider a (q^*, q^*)-Exclusive Dealing equilibrium, which we know to be an optimal full-disclosure equilibrium from Proposition 5. Suppose agent 1 is allocated production and produces y_1 in $t = 1$. In $t = 2$, agent 2 is allocated production, works hard, and produces y_H. Let $(S_1^{FD}(y_1, y_H), S_2^{FD}(y_1, y_H))$ be the vector of dyad-specific surpluses for the agents following this history. Because this is a full-disclosure relational contract, it must be that $S_2^{FD}(y_1, y_H) \geq \tilde{S}$ for each $y_1 \in \{0, y_H\}$.

Now, notice that the larger is $S_2^{FD}(y_H, y_H)$, the harder it is to satisfy agent 1's constraint (1.6) in $t = 1$, since $S_1^{FD}(y_H, y_H)$ enters that constraint. In contrast, $S_1^{FD}(0, y_H)$ is irrelevant for agent 1's incentives because (1.6) only matters following high output. Therefore, consider a non-full-disclosure equilibrium in which agent 2 is informed of y_1 only *after* he chooses $e_2 = 1$, and let $(S_1(y_1, y_H), S_2(y_1, y_H))$ be the vector of dyad-specific surpluses in this alternative strategy profile. For agent 2 to work hard, it need only be the case that $E\left[S_2(y_1, y_H)|h_2^t\right] \geq \tilde{S}$. In particular, we can set $S_2^{FD}(y_H, y_H) < \tilde{S}$, which in turn relaxes (1.6) for agent 1 in $t = 1$. This slack can then be used to implement a more efficient continuation payoff when $y_1 = y_H$.

Proposition 6 illustrates that it is sometimes optimal for the principal to refrain from telling one agent about her obligations to other agents. This result is related to the long literature on correlated equilibria such as Aumann (1974) and Myerson (1986), as well as to recent results by Rayo and Segal (2010), Kamenica and Gentzkow (2011), and Fong and Li (2010). Like those papers, the principal in this model can determine an agent's information and thereby create slack in their incentive constraints, which can in turn be used to induce high effort in other histories. Unlike those papers, the information concealed by the principal concerns the past performance of other agents and is only valuable in the context of the larger equilibrium.

1.7 Communication

The baseline model makes the stark assumption that agents cannot send messages to one another and so are unable to multilaterally punish deviations. While this assumption is realistic in dispersed supply chains, agents within an organization or a close-knit group of suppliers may be able to coordinate their actions. In this section, we show that the allocation of future business remains an important tool even if agents can communicate, so long as they also earn rents.[28] To consider how joint punishment affects our results, we first define an augmented game with communication between agents.

[28]The baseline model and this extension provide two examples of settings in which the allocation of business affects the strength of each relationship. More generally, so long as the surplus at stake in each dyad depends on how much business is allocated to that dyad, the allocation of business will matter for incentives.

Definition 7 *The* augmented game with communication *is identical to the baseline repeated game, except that each player simultaneously chooses a publicly-observed message $m_i \in M$ at the beginning of each round, where $|M|$ is large but finite.*

Agents can use the message space M to share information with one another and coordinate to punish the principal. In the augmented game, the allocation of business would be irrelevant so long as $\tilde{S} \leq V^{FB}$ and the principal earned all the rents from production. On the other hand, the allocation rule remains an important motivational tool if agents keep some of the profits they produce, a notion formalized in Assumption 4.

Assumption 4 *Fix $\gamma \in [0,1]$. An equilibrium satisfies* γ-rent-seeking *if at any history $(h^{t-1}, \mathcal{P}_t, x_t)$ on the equilibrium path, agent x_t earns*

$$u_{x_t}^t = \gamma \sum_{i=0}^{N} u_i^t,$$

the principal earns $u_0^t = (1-\gamma)\sum_{i=0}^{N} u_i^t$, and all $i \notin \{0, x_t\}$ earn $u_i^t = 0$.

Assumption 4 is similar to an assumption made in Halac (2012) and can be viewed as a reduced-form model of bargaining power. In round t, the producing agent x_t earns a fraction γ of the total surplus produced in that round. As a result, the principal earns only a fraction $1-\gamma$ of the surplus produced in each relationship. Because agents do not pay bonuses to one another, the maximum surplus at stake in agent i's relationship is the sum of his and the principal's surplus. In particular, the rents earned by one agent cannot be used to make incentive pay to *another* agent credible.

More precisely, Lemma 3 shows that agent i is only willing to work hard if *the sum of his and the principal's surpluses exceeds $\tilde{S}$.*

Lemma 3 *The following condition holds in any PBE σ^*:*

$$\tilde{S} \leq \delta E_{\sigma^*}\left[U_0(h^{t+1}) + U_i(h^{t+1}) \mid (h_i^{t-1}, \mathcal{I}_i(n_t))\right]$$
$$\forall i \in \{1, ..., N\}, \forall h^{t-1} \in \mathcal{H}, n_t \text{ on the equilibrium path immediately after } y_t \quad (1.14)$$
$$\text{such that } x_t = i,\ e_t = 1,\ y_t = y_H.$$

Proof:

This proof is similar to that of Statement 1 of Lemma 2, except that transfers must instead satisfy

$$\begin{aligned} \tau_i(y_t) &\leq \tfrac{\delta}{1-\delta} E_\sigma \left[U_0(h^{t+1}) | h^{t-1}, n_t, y_t \right] \\ -\tau_i(y_t) &\leq \tfrac{\delta}{1-\delta} E_\sigma [U_i(h^{t+1}) \mid h_i^t, \mathcal{I}_i(n_t, y_t)] \end{aligned}$$

where n_t is a node following $e_t = 1$ and output y_t. If these conditions are not satisfied, then either the principal or agent i would strictly prefer to deviate to $\tau_{i,t} = 0$ and be min-maxed. Plugging these expressions into (1.8) yields (1.14). ■

Condition (1.14) is similar to (1.6), except that the right-hand side includes the principal's *total* expected continuation surplus rather than just her surplus from dyad i. The agents can coordinate to hold the principal at her outside option following a deviation, so her entire continuation surplus can be used to support incentive pay in each period. However, $U_0(h^{t+1}) + U_i(h^{t+1})$ does *not* include the continuation surplus for the *other* agents.

As in the original game, promising future business to agent i increases $U_i(h^{t+1})$ and so relaxes (1.14) for i while tightening this constraint for the other agents. Unlike Lemma 2, the principal may have an incentive to deviate from an optimal allocation rule, so (1.14) is only a necessary condition for equilibrium. Nevertheless, Lemma 4 shows that a version of the Favored Producer Allocation continues to be an equilibrium.

Lemma 4 *Let Assumption 4 hold. If (1.14) is satisfied under the Favored Producer Allocation for $d_t = e_t = 1$, $\forall t$, then there exists an equilibrium that uses the Favored Producer Allocation and generates V^{FB} total surplus.*

Proof:

See Appendix 1.9.

In the proof of Lemma 4, every player simultaneously reports every variable except e_t in each period. All of these variables are observed by at least two players, so any lie is immediately detected and punished by complete market breakdown. These messages effectively make the monitoring structure public; the Favored Producer Allocation remains a relational contract because it is a full-disclosure equilibrium.

Allocating business to i relaxes his constraint (1.14) but tightens this constraint for the other agents. Using the same logic as in Proposition 3, Proposition 7 proves that a relational contract using the Favored Producer Allocation attains first-best whenever any PBE does.

Proposition 7 *Consider the game in Definition 7, and suppose Assumption 4 holds. If any PBE has total surplus V^{FB}, then the equilibrium from Lemma 4 also has total surplus V^{FB}. If $\gamma > 0$, $\exists$ a non-empty interval $\Delta \in [0,1]$ such that if $\delta \in \Delta$ this equilibrium attains V^{FB} but stationary equilibria do not.*

Proof:

See Appendix 1.9.

Under Assumption 4, $U_0(h^t) = (1-\gamma)V^{FB}$ at any on-path history h^t in a relational contract that attains first-best. By (1.14), agent i must expect to earn at least $\tilde{S} - (1-\gamma)V^{FB}$ in order to be willing to work hard. Using this intuition, we can define the *residual obligation* $\hat{\Omega}_t^i$ as the amount of expected dyad-specific surplus that must be given to agent i for him to work hard:

$$\hat{\Omega}_t^i \equiv \frac{\hat{\Omega}_{t-1}^i}{\delta} + \beta_t^L p_1 \delta(\tilde{S} - (1-\gamma)V^{FB}) - \beta_t^H \left[(1-\delta)\gamma(y_H p_1 - c)\right]. \qquad (1.15)$$

The proof of this result is then similar to that of Proposition 3.

When the fraction of surplus γ earned by an agent is small, the principal earns more surplus and thus (1.14) is easier to satisfy. Intuitively, the rent-seeking activities of one agent have a *negative externality* on the principal's relationship with *other* agents. Rent-seeking by i makes the principal more willing to renege on the other agents, since she loses less surplus in the punishment following a deviation. Agent i does not internalize this negative externality because his relationship with the principal is determined by $U_0(h^t) + U_i(h^t)$ and so is only affected by how rent is shared in *other* dyads.[29]

[29] The observation that rent-seeking by one agent can impose a negative externality on other agents' relationships has been made by Levin (2002).

1.8 Conclusion

In the absence of formal contracts, a principal must carefully allocate a limited stream of business among her agents in order to motivate them. We have shown that agents only work hard if they are promised sufficient future business following *high* output, and so the optimal allocation rule (1) rewards success and (2) tolerates failure. When first-best is attainable, an optimal relational contract resembles a tournament, where the prize is a (temporary) larger share of future business and an agent "wins" if he is allocated business and performs well. When first-best cannot be attained, the optimal full-disclosure equilibrium may involve exclusive dealing with a high-performing agent while other relationships deteriorate. The principal can mitigate these inefficiencies by concealing information about the history of play from the agent. Thus, a downstream firm (or boss) who interacts with multiple suppliers (or workers, or divisions) must carefully consider both the rule she uses to allocate business and the amount of information she reveals.

Like much of the relational contracting literature, one shortcoming of our model is that competition does not pin down the division of rents between players. In some realistic cases, suppliers might be expected to compete away their rents, so that a downstream firm would opt to cultivate multiple sources in order to secure better prices. There are potentially interesting interactions between rent-seeking and relationship cultivation, since an agent's incentives depend critically on his beliefs about his future surplus. Nevertheless, those companies with very close supplier relations tend to source from a small number of suppliers, who earn substantial rents. Toyota even goes so far as to enshrine "[carrying] out business....without switching to others" in their 1939 Purchasing Rules (as noted by Sako (2004)).

More broadly, this paper presents one setting where the principal chooses seemingly inefficient actions (eg, exclusive dealing with one agent) in order to cultivate a close relationship. In repeated games with impatient players, short-run moral hazard problems can sometimes be solved only by committing to persistently inefficient continuation play. Thus, organizations exhibit tremendous inertia: inefficient policies persist in a relational contract,[30] while seemingly welfare-improving changes undermine existing agreements and inhibit cooperation.

[30] See Gibbons and Henderson (2012) for much more on persistent performance differences among seemingly similar enterprises.

1.9 Appendix: Proofs[31]

Lemma 2

Statement 2: Sufficiency

Given a strategy σ satisfying (1.6), we construct an equilibrium σ^*.

Let $\mathcal{H}(\sigma)$ be the set of on-path histories for generic strategies σ, with element $h^t(\sigma) \in \mathcal{H}(\sigma)$. Define an *augmented history* as an element $(h^t; \tilde{h}^t) \in \mathcal{H} \times \mathcal{H}$. The set $\mathcal{H}^{Aug}$ will be defined to relate on-path histories in σ^* to histories in σ. To simplify notation, let $\mathcal{N}^e$ be the set of nodes in the extensive-form stage game in which e was just chosen, before Nature chooses y.

Constructing Equilibrium Strategies: We recursively construct the set $\mathcal{H}^{Aug}$ and σ^*.

First, we construct $\mathcal{H}^{Aug}$:

1. If $t = 0$, then the initial augmented history is $(\emptyset, \emptyset) \in \mathcal{H}^{Aug}$. If $t > 0$, let $(h^{t-1}(\sigma); h^{t-1}(\sigma^*)) \in \mathcal{H}^{Aug}$.

2. For every history $(h^{t-1}(\sigma), n_t)$, $n_t \in \mathcal{N}^e$, that is on-path for σ, define the augmented history

$$\left((h^{t-1}(\sigma), n_t); (h^{t-1}(\sigma^*), n_t^*)\right) \in \mathcal{H}^{Aug} \tag{1.16}$$

   where $n_t^* \in \mathcal{N}^e$ is defined as follows. Note that actions with * are part of n_t^*, while those without are part of n_t.

   (a) $\mathcal{P}_t^* = \mathcal{P}_t$.

   (b) If $p_0 > 0$ or $d_t > 0$, then $x_t^* = x_t$, $d_t^* = d_t$, $e_t^* = e_t$. If $p_0 = 0$ and $d_t = 0$, then $x_t^* = x_t$, $d_t^* = 1$, and $e_t^* = 0$.[32]

[31] We frequently refer to "all histories on the equilibrium path" such that some condition holds. Formally, interpret "all histories on the equilibrium path" as "almost surely on the equilibrium path."

[32] Note: if $p_0 = 0$, then $\{d_t = 0\}$ and $\{d_t = 1, e_t = 0\}$ generate the same surplus for every player and $w^*_{x_t,t} = 0$ for both of these events. As a result, $d_t^* = 1$ and $e_t^* = 0$ in σ^* whenever $d_t = 0$ in σ. In this case, an augmented history $(h^t(\sigma), h^t(\sigma^*))$ may have rounds in which $d_{t'}^* = 1$ and $e_{t'}^* = 0$ but $d_{t'} = 0$; however, σ^* still forms an equilibrium that generates the same total surplus and distribution over $\{\mathcal{P}_t, x_t, d_t e_t, y_t\}_{t=1}^T$

(c) $w^*_{x_t,t}$ satisfies

$$w^*_{x_t,t} = d_t \left[(1-e_t) y_H p_0 + e_t \left(y_H p_1 + (1-p_1) \frac{c}{p_1 - p_0} \right) \right]$$

and $w^*_{i,t} = 0$, $\forall i \neq x_t$.

3. For any successor history $h^t(\sigma)$ to $(h^{t-1}(\sigma), n_t)$ in the support of σ, define the augmented history

$$\left(h^t(\sigma); (h^t(\sigma^*)\right) \in \mathcal{H}^{Aug}$$

by setting $y_t = y^*_t$, $\tau^*_{i,t} = 0 \ \forall i \neq x_t$, and

$$\tau^*_{x_t,t} = -1\{w^*_{x_t,t} \geq y_H p_1, y^*_t = 0\} \frac{c}{p_1 - p_0}$$

Next, we recursively define the candidate equilibrium σ^* in which the principal tracks an augmented history.

1. At $t = 0$, the principal chooses augmented history $(\emptyset, \emptyset) \in \mathcal{H}^{Aug}$. For any $t > 0$, let $(h^{t-1}(\sigma), h^{t-1}(\sigma^*)) \in \mathcal{H}^{Aug}$ be the unique augmented history on the equilibrium path at the beginning of period t.

2. When $\mathcal{P}_t$ is realized, the principal chooses $(h^{t-1}(\sigma), n_t)$ for $n_t \in \mathcal{N}^e$ according to $\Psi^e_\sigma(h^{t-1}(\sigma), \mathcal{P}_t)$, which is the conditional distribution given by σ over $n_t \in \mathcal{N}^e$, given $(h^{t-1}(\sigma), \mathcal{P}_t)$. For this $(h^{t-1}(\sigma), n_t)$, there is a unique augmented history $((h^{t-1}(\sigma), n_t), (h^{t-1}(\sigma^*), n^*_t)) \in \mathcal{H}^{Aug}$.

3. Given this unique augmented history $((h^{t-1}(\sigma), n_t), (h^{t-1}(\sigma^*), n^*_t))$, the principal chooses allocation x^*_t and wages $\{w^*_{i,t}\}$.

4. Agent x_t accepts production iff $w^*_{x_t,t} \geq y^i_H p_0$, chooses $e_t = 1$ iff $w^*_{x_t,t} = y_H p_1 + (1 - p_1)\frac{c}{p_1 - p_0}$, and pays transfer $\tau^*_{i,t}$ following output y^*_t.

$\forall T < \infty$ as σ.

5. Given successor history $(h^{t-1}(\sigma), n_t, y_t^*)$, define $\Psi_\sigma(h^{t-1}(\sigma), n_t, y_t^*)$ as the conditional distribution of histories $h^t(\sigma)$ given $(h^{t-1}(\sigma), n_t, y_t^*)$, which is well defined on the equilibrium path. The principal randomly chooses history $h^t(\sigma)$ according to $\Psi_\sigma(h^{t-1}(\sigma), n^e, y_t^*)$, resulting in the augmented history $(h^t(\sigma), h^t(\sigma^*))$, where $h^t(\sigma^*)$ is the realized history under σ^*. Along the equilibrium path, $(h^t(\sigma), h^t(\sigma^*)) \in \mathcal{H}^{Aug}$ by construction.

6. Following any deviation in a variable other than e_t, the principal thereafter chooses $x_t = \min\{i | i \in \mathcal{P}_t\}$ and $w_{i,t}^A = \tau_{i,t}^A = 0$. Agents who observe the deviation choose $d_t = 0$.

Uniqueness of $(h^t(\sigma), h^t(\sigma^*))$ for each $h^t(\sigma) \in \mathcal{H}(\sigma)$: Given an augmented history $(h^{t-1}(\sigma), h^{t-1}(\sigma^*)) \in \mathcal{H}^{Aug}$ for every on-path $h^{t-1}(\sigma)$, then this recursive structure identifies every $h^t(\sigma) \in \text{supp}\{\sigma | h^{t-1}(\sigma)\}$ with a unique augmented history $(h^t(\sigma), h^t(\sigma^*))$. The unique initial history is $(\emptyset, \emptyset)$, and so every on-path $h^t(\sigma)$ is linked to a unique $(h^t(\sigma), h^t(\sigma^*))$ by induction.

Payoff-Equivalence of Old and New Strategies: We claim that σ^* generates the same distribution over $\{\mathcal{P}_s, x_s, d_s e_s, y_s\}_{s=1}^t, \forall t$, and same total surplus as σ. Define $\chi_\sigma(h^t(\sigma))$ as the *ex ante* distribution over $h^t(\sigma)$ under strategies σ and $\chi_{\sigma,\sigma^*}(h^t(\sigma), h^t(\sigma^*))$ as the distribution over *augmented* histories $(h^t(\sigma), h^t(\sigma^*))$. Then:

1. $\forall\, t \geq 0$ and $h^t(\sigma) \in \mathcal{H}(\sigma)$, if $(h^t(\sigma), h^t(\sigma^*)) \in \mathcal{H}^{Aug}$, then

$$\chi_{\sigma,\sigma^*}(h^t(\sigma), h^t(\sigma^*)) = \chi_\sigma(h^t(\sigma))$$

   and so σ and σ^* generate the same distribution over $\{\mathcal{P}_s, x_s, d_s e_s, y_s\}_{s=1}^t, \forall t$.

2. σ generates the same total surplus in period t surplus at history $h^t(\sigma)$ as σ^* does in period t at augmented history $(h^t(\sigma), h^t(\sigma^*))$.

We tackle each claim separately.

Claim 1: By induction. For $t = 1$, $\mathcal{H} = \emptyset$ and the result follows. Consider the conditional distribution over $h^t(\sigma)$ given $h^{t-1}(\sigma)$ and compare it to the conditional distribution over $(h^t(\sigma), h^t(\sigma^*))$ given $(h^{t-1}(\sigma), h^{t-1}(\sigma^*))$. The distribution $F(\mathcal{P})$ is exogenous. Conditional on the realized $\mathcal{P}_t$, for $p_0 > 0$ actions $(x_t^*, d_t^*, e_t^*) \sim \Psi_\sigma(h^{t-1}(\sigma), \mathcal{P}_t)$, while for $p_0 = 0$ the analogous statement holds for $(x_t^*, d_t^* e_t^*)$. Since each history $(h^{t-1}(\sigma), \mathcal{P}_t, n^e)$ corresponds to a unique augmented history $((h^{t-1}(\sigma), \mathcal{P}_t, n^e), (h^{t-1}(\sigma), \mathcal{P}_t, n^{e*})$, the augmented history $((h^{t-1}(\sigma), \mathcal{P}_t, n^e), (h^{t-1}(\sigma), \mathcal{P}_t, n^{e*}))$ has the same distribution as $(h^{t-1}(\sigma), \mathcal{P}_t, n^e)$ conditional on the $(t-1)$-period history. The distribution of y_t depends only on d_t and e_t, so the distribution of y_t^* given this augmented history is the same as the distribution of y_t given $(h^{t-1}(\sigma), \mathcal{P}_t, n^e)$. Given this, the distribution over t-period augmented histories $(h^t(\sigma), h^t(\sigma^*))$ is identical to the distribution over $h^t(\sigma)$, which proves Claim 1.

Claim 2: The actions that affect total surplus are (x_t^*, d_t^*, e_t^*) (or $(x_t^*, d_t^* e_t^*)$ for $p_0 = 0$), which have an identical distribution under σ and σ^* by claim 1. This proves Claim 2.

Optimality of Principal's Actions: Next, we show that the principal has no profitable deviation. Under σ^*, the principal earns no more than 0 at every h^t, so it suffices to show that the principal cannot earn strictly positive surplus by deviating. Fix history h^t. The principal could deviate from σ^* in one of three variables: $x_t^*, \{w_{i,t}^*\}_{i=1}^N$ or $\{\tau_{i,t}^*\}_{i=1}^N$. Following the deviation, the principal earns 0 whenever she allocates production to an agent that observed the deviation. Thus, the principal has no profitable deviation in $\{\tau_{i,t}^*\}_{i=1}^N$, since $\tau_{i,t}^* \leq 0$ and so a deviation would be myopically costly. A deviation in $w_{i,t}^*$ for $i \neq x_t$ would be unprofitable for the same reasons, since $w_{i,t}^* = 0$. The principal would earn $d_t y_t - w_{x_t,t}$ following a deviation in $w_{x_T,t}^*$. If $w_{x_t,t} < y_H p_0$ then $d_t^* = 0$ and the deviation is unprofitable. If $w_{x_t,t} \in [y_H^i p_0, y_H^i p_1)$, then $d_t^* = 1$ and $e_t^* = 0$, so the principal's deviation earns $y_H p_0 - w_{x_t,t} < 0$. If $w_{x_t,t} > y_H p_1$, then the principal earns no more than

$$y_H p_1 - w_{x_t,t} + (1 - p_1) \frac{w_{x_t,t} - y_H p_1}{1 - p_1} = 0$$

so this deviation is not profitable. Because the principal can never earn more than 0 regardless of her allocation decision, she has no profitable deviation from x_t^*, proving the claim.

Optimality of Agent's Actions: Finally, we argue that each agent has no profitable deviation. This argument entails tracking the agent's private information at each history.

Let $h^{t-1}(\sigma), \hat{h}^{t-1}(\sigma) \in \mathcal{H}(\sigma)$, and let $(h^{t-1}(\sigma), h^{t-1}(\sigma^*)), (\hat{h}^{t-1}(\sigma), \hat{h}^{t-1}(\sigma^*)) \in \mathcal{H}^{Aug}$.

Claim 1: Suppose that agent i cannot distinguish $(h^{t-1}(\sigma), n_t)$ and $(\hat{h}^{t-1}(\sigma), \hat{n}_t)$. Then agent i cannot distinguish $(h^{t-1}(\sigma^*), n_t^*))$ and $(\hat{h}^{t-1}(\sigma^*), \hat{n}_t^*)$. Proof by induction. For $t = 0$, the result holds trivially. Suppose the result holds for all $t < T$, and consider a history in round T. Suppose towards contradiction that i can distinguish $(h^{T-1}(\sigma^*), n_T^*)$ and $(\hat{h}^{T-1}(\sigma^*), \hat{n}_T^*)$. From the inductive step, there is a variable in round T that differs between n_T^* and $\hat{n}_T^*$. Because n_T and $\hat{n}_T$ are indistinguishable by assumption, $\mathcal{P}_T^*$, x_T^* cannot differ, and so d_t^* and e_t^* cannot differ either.

Thus, it must be that $w_{i,T}^* \neq \hat{w}_{i,T}^*$ or $\tau_{i,T}^* \neq \hat{\tau}_{i,T}^*$. But $\hat{\tau}_{i,T}^*$ is determined entirely by $\hat{w}_{i,T}^*$, so it must be that $w_{i,T}^* \neq \hat{w}_{i,T}^*$, which implies that $x_t^* = i$. If $d_T^* = 0$, then $w_{i,T}^* = \hat{w}_{i,T}^* = 0$; if $d_T^* = 1$ and $e_T^* = 0$, then $w_{i,T}^* = \hat{w}_{i,T}^* = y_H p_0$; and if $d_T^* = e_T^* = 1$, then $w_{i,T}^* = \hat{w}_{i,T}^* = y_H p_1 + (1 - p_1)\frac{c}{p_1 - p_0}$. Thus, contradiction obtains.

Claim 2: Agent i has no profitable deviation from $\tau_{i,t}^*$. Agent i plays a myopic best-response and earns 0 continuation surplus immediately following any deviation. Hence, we need only check a deviation when $\tau_{i,t}^* < 0$, which occurs when $x_t^* = i$, $e_t^* = 1$, and $y_t^* = 0$ on the equilibrium path. Let $(h^{t-1}(\sigma^*), n_t^*)$ be such a history, and consider agent i's beliefs.

Agent i knows the true history is consistent with $(h_i^{t-1}(\sigma^*), \mathcal{I}_i(n_t^*))$ and infers $e_t^* = 1$ because $w_{i,t}^* \geq y_H p_1$. Because σ and σ^* are payoff equivalent,

$$S_i(\sigma, h^t(\sigma)) = S_i(\sigma^*, h^t(\sigma^*)). \tag{1.17}$$

Furthermore,

$$U_i(\sigma^*, h^{t+1}(\sigma^*)) = S_i(\sigma^*, h^{t+1}(\sigma^*)) \tag{1.18}$$

by construction of σ^*. By the previous claim, the strategies σ^* induce a *coarser partition* than σ over the set of augmented histories. That is,

$$\left\{(\hat{h}^{t-1}(\sigma^*),\hat{n}_t^*) \mid ((\hat{h}^{t-1}(\sigma),\hat{n}_t),(\hat{h}^{t-1}(\sigma^*),\hat{n}_t^*)) \in \mathcal{H}^{Aug},\ (\hat{h}_i^{t-1}(\sigma),\mathcal{I}_i(\hat{n}_t)) = (h_i^{t-1}(\sigma),\mathcal{I}_i(n_t))\right\} \subseteq \left\{(h^{t-1}(\sigma^*),n_t^* | h_i^{t-1}(\sigma^*),\mathcal{I}_i(n_t^*)\right\}. \tag{1.19}$$

Because (1.6) holds by assumption, $E_\sigma\left[S_i(\sigma,h^{t+1}(\sigma))|h_i^{t-1}(\sigma),\mathcal{I}_i(n_t)\right] \geq \frac{1-\delta}{\delta}\frac{c}{p_1-p_0}$. Combined with (1.17) and (1.19), we

$$E_{\sigma^*}\left[S_i(\sigma^*,h^{t+1}(\sigma^*))|h_i^{t-1}(\sigma^*),\mathcal{I}_i(n_t^*)\right] = E_{\sigma^*}\left[E_\sigma\left[S_i(\sigma,h^{t+1}(\sigma))|h_i^{t-1}(\sigma),\mathcal{I}_i(n_t)\right]|((h^{t-1}(\sigma),n_t),h^{t-1}(\sigma^*),n_t^*) \in \mathcal{H}^{Aug}\right] \geq \frac{1-\delta}{\delta}\frac{c}{p_1-p_0} \tag{1.20}$$

Agent i has a profitable deviation if and only if

$$0 > (1-\delta)\tau_{i,t}^* + \delta E_{\sigma^*}\left[S_i(\sigma^*,h^{t+1}(\sigma^*))|h_i^{t-1}(\sigma^*),\mathcal{I}_i(n_t^*)\right]$$

but $\tau_{i,t}^* = \frac{c}{p_1-p_0}$, so (1.20) implies that this inequality never holds. Thus, agent i has no profitable deviation from $\tau_{i,t}^*$.

Claim 3: Agent i has no profitable deviation from d_t^*, e_t^*, or $w_{i,t}^*$. A deviation in d_t^* or $w_{i,t}^*$ is myopically costly and leads to a continuation payoff of 0, so the agent never has a profitable deviation in these variables. Similarly, agent i has no profitable deviation from $e_t^* = 0$. Given $\tau_{i,t}^*$, agent i has a profitable deviation from $e_t^* = 1$ iff

$$-(1-p_1)\frac{c}{p_1-p_0} - c < -(1-p_0)\frac{c}{p_1-p_0}$$

which cannot hold.

Thus, σ^* is a Perfect Bayesian Equilibrium that satisfies the conditions of Lemma 2.■

Proposition 3

We present two lemmas to prove Proposition 3.

Statement of Lemma A1

Consider strategies σ that attain the first-best total surplus V^{FB}, and suppose $\exists\, i \in \{1, .., N\}$ such that

$$\limsup_{t\to\infty} \Omega_t^i = \infty.$$

Then σ is not a equilibrium.

Proof of Lemma A1

Given a σ that attains first-best with $\limsup_{t\to\infty} \Omega_i^t = \infty$ for some $i \in \{1, ..., N\}$. Because σ attains first-best, $d_t = e_t = 1$ whenever $\mathcal{P}_t \neq \emptyset$. Towards contradiction, suppose that σ is an equilibrium.

Let

$$b_i(h_i^t) = 1\{x_t = i, y_t = y_H\} * 1\{\forall t' < t, x_{t'} = i \Rightarrow y_{t'} = 0\} \tag{1.21}$$

be the indicator function for the event that $x_t = i$, $y_t = y_H$ in period t, and $y_{t'} = 0\ \forall t' < t$ with $x_t = i$. Lemma 2 implies that (1.6) must be satisfied the *first time* that $x_t = i$ and $d_t = e_t = 1$, which implies

$$b_i(h_i^t)\delta\tilde{S} \leq (1-\delta) b_i(h_i^t) E_\sigma\left[\sum_{t'=1}^{\infty} \delta^{t'} 1_{i,t+t'}(y_H p_1 - c) \mid h_i^t\right]. \tag{1.22}$$

A fortiori, (1.22) must hold in expectation:

$$E_\sigma\left[b_i(h_i^t)\right] * \delta\tilde{S} \leq E_\sigma\left[(1-\delta) b_i(h_i^t) E_\sigma\left[\sum_{t'=1}^{\infty} \delta^{t'} 1_{i,t+t'}(y_H p_1 - c) \mid h_i^t\right]\right].$$

Because σ attains first-best, $E\left[b_i(h_i^t)\right] = \text{Prob}\{x_t = i, y_t = y_H, x_{t'} = i \rightarrow y_{t'} = 0\} = p_1\beta_t^L$. Dividing by δ^{K-t} and summing across $t = 1, ..., K$ yields

$$\sum_{k=1}^{K} \frac{p_1\beta_k^L}{\delta^{K-k}}\delta\tilde{S} \leq (1-\delta)V^{FB}\sum_{k=1}^{K} E_\sigma\left[\sum_{t'=1}^{\infty}\frac{1}{\delta^{K-k-t'}}1_{i,k+t'} \mid b_i(h_i^{k+t'}) = 1\right]\text{Prob}_\sigma\{b_i(h_i^k) = 1\}. \tag{1.23}$$

This infinite sum is dominated by $\sum_{t'=1}^{\infty}\delta^{t'}(y_H p_1 - c)$ and so converges absolutely. Switching the order of summation on the right-hand side of (1.23) and, then re-index using $\psi = k + t'$. Because the events of (1.21) are disjoint, $\sum_{k=1}^{K}\text{Prob}\{b_i(h_i^k) = 1\} = \text{Prob}\{\exists k \leq K \text{ s.t. } x_k =$

$i, y_k = y_H\}$. Therefore, (1.23) can be written

$$\sum_{\psi=1}^{K} \frac{p_1 \beta_\psi^L}{\delta^{K-\psi}} \delta \tilde{S} \leq (1-\delta) V^{FB} \sum_{\psi=1}^{\infty} \frac{1}{\delta^{K-\psi}} \text{Prob}_\sigma \{x_\psi = i, \exists t' \leq \min\{K, \psi - 1\} \text{ s.t. } x_{t'} = i, y_{t'} = y_H\}. \quad (1.24)$$

Now, the tail for $\psi > K$ on the right-hand side of (1.24) is dominated by $\sum_{k=K}^{\infty} K \frac{1}{\delta^{K-k}} = \frac{1}{1-\delta}$ and so absolutely converges. Hence, (1.24) implies that $\exists\, C \in \mathbb{R}$ such that $\forall K \in \mathbb{N}$,

$$\sum_{\psi=1}^{K} \frac{p_1 \beta_\psi^L}{\delta^{K-\psi}} \delta \tilde{S} - (1-\delta) V^{FB} \sum_{\psi=1}^{K} \frac{1}{\delta^{K-\psi}} \text{Prob}_\sigma \{x_\psi = i, \exists t' \leq \min\{K, \psi - 1\} \text{ s.t. } x_{t'} = i, y_{t'} = y_H\} \leq C.$$

By definition, $\text{Prob}_\sigma \{x_\psi = i, \exists t' \leq \min\{K, \psi - 1\} \text{ s.t. } x_{t'} = i, y_{t'} = y_H\} = \beta_\psi^H$, so in any equilibrium that attains first-best, $\exists\, C \in \mathbb{R}$ such that $\forall K \in \mathbb{N}$,

$$\sum_{\psi=1}^{K} \frac{p_1 \beta_\psi^L}{\delta^{K-\psi}} \delta \tilde{S} - (1-\delta) V^{FB} \sum_{\psi=1}^{K} \frac{1}{\delta^{K-\psi}} \beta_\psi^H \leq C \quad (1.25)$$

Contradiction. So σ cannot be an equilibrium. ■

Statement of Lemma A2

Suppose $\exists$ strategies σ that attain the first-best total surplus V^{FB} such that $\forall i \in \{1, ..., N\}$, $\limsup_{t\to\infty} \Omega_t^i < \infty$. $\forall h^t$, let

$$E(h^t) = \left\{ i \mid h^t, \exists t' < t \text{ s.t. } x_{t'} = i, y_{t'} = y_H \right\}$$

and $\forall t$,

$$E_i^t = \left\{ h^t | i \in E(h^t) \right\}.$$

Then $\exists$ strategies $\hat{\sigma}$ that also attain V^{FB} such that:

1. Obligation is finite: $\forall i \in \{1, ..., N\}$, $\limsup_{t\to\infty} \hat{\Omega}_t^i < \infty$.
2. Ex ante, the allocation rule treats agents who have produced y_H symmetrically: $\forall t, i \in \{1, ..., N\}$

$$\text{Prob}\{x_t = i | E_i^t\} = \tfrac{1}{N} \textstyle\sum_{j=1}^{N} \text{Prob}\left\{x_t = j | E_j^t\right\}$$
$$\text{Prob}\left\{x_t = i | \left(E_i^t\right)^C\right\} = \tfrac{1}{N} \textstyle\sum_{j=1}^{N} \text{Prob}\left\{x_t = j | \left(E_j^t\right)^C\right\}.$$

3. Agents who have produced y_H are favored: $\forall t, h^t$, $\text{Prob}\{x_t \in E(h^t) | E(h^t) \cap \mathcal{P}_t \neq \emptyset\} = 1$.

Proof of Lemma A2

Given σ, define the strategy profile $\tilde{\sigma}$ in the following way. At $t = 0$, draw a permutation $\rho : \{1, ..., N\} \to \{1, ..., N\}$ uniformly at random. Then play in $\tilde{\sigma}$ is identical to σ, except that agent i is treated as agent $\rho(i)$.

Because agents are *ex ante* symmetric, $\tilde{\sigma}$ generates first-best surplus. If $\limsup_{t\to\infty} \max_i \Omega_i^t < \infty$ under σ, then

$$\limsup_{t\to\infty} \tilde{\Omega}_i^t = \limsup_{t\to\infty} \frac{1}{N} \sum_{i=1}^{N} \Omega_i^t < \infty$$

$\forall i \in \{1, .., N\}$ under $\tilde{\sigma}$. Define $E(h^t) = \{i \in \{1, ..., N\} | \exists t' < t \; s.t. \; x_{t'} = i, y_{t'} = y_H\}$ as the agents who have produced y_H at least once in history h^t. Then $\forall i \in \{1, ..., N\}$,

$$\text{Prob}_{\tilde{\sigma}}\{x_t = i | h^t \; s.t. \; i \in E(h^t)\} = \frac{1}{N} \sum_{j=1}^{N} \text{Prob}_{\sigma}\{x_t = j | h^t \; s.t. \; j \in E(h^t)\}.$$

Therefore, every $i \in E(h^t)$ is treated symmetrically, as is every $j \notin E(h^t)$.

Next, define strategies $\hat{\sigma}$ in the following way. Like $\tilde{\sigma}$, draw a permutation ρ uniformly at random at $t = 0$. In each round, the allocation rule assigns production to an agent in $E(h^t)$ whenever possible:

$$\hat{x}_t \in \left\{ i \;\middle|\; \begin{array}{c} \mathcal{P}_t \cap E(h^t) = \emptyset \to i = \arg\min_{i' \in \mathcal{P}_t} \rho(i) \\ \mathcal{P}_t \cap E(h^t) \neq \emptyset \to i = \arg\min_{i' \in \mathcal{P}_t \cap E(h^t)} \rho(i) \end{array} \right\}. \tag{1.26}$$

Agents always choose $\hat{d}_t = \hat{e}_t = 1$. Transfers are $\hat{w}_{i,t} = \hat{\tau}_{i,t} = 0$.

We claim that if $\limsup_{t\to\infty} \max_i \tilde{\Omega}_i^t < \infty$ under $\tilde{\sigma}$, then $\limsup_{t\to\infty} \max_i \hat{\Omega}_i^t < \infty$ under $\hat{\sigma}$. Under both strategies,

$$\text{Prob}_{\hat{\sigma}}\{x_t = i | h^t \; s.t. \; i \in E(h^t)\} = \text{Prob}_{\hat{\sigma}}\{x_t = j | h^t \; s.t. \; j \in E(h^t)\} \tag{1.27}$$

for all i, j, so $\exists \; \tilde{\Omega}^t, \hat{\Omega}^t$ such that $\forall i \in \{1, ..., N\}$, $\tilde{\Omega}_i^t = \tilde{\Omega}^t$ and $\hat{\Omega}_i^t = \hat{\Omega}^t$.

Denote by $E_t^m \subseteq \mathcal{H}$ the event that m agents have produced y_H by time t; then

$$\phi_t^m \equiv \frac{1}{m}\text{Prob}_{\tilde{\sigma}}\{x_t \in E(h^t)|E_t^m\} \leq \frac{1}{m}\text{Prob}_{\hat{\sigma}}\{x_t \in E(h^t)|E_t^m\}. \tag{1.28}$$

Because $\hat{\sigma}$ and $\tilde{\sigma}$ are symmetric, $\text{Prob}\{i \in E_t(h^t)|E_t^m\} = \frac{m}{N}$ for both strategies. Then we can explicitly write β_t^H and β_t^L for $\tilde{\sigma}$:

$$\beta_t^H = \sum_{m=1}^{N} \frac{m}{N}\text{Prob}_{\tilde{\sigma}}\{E_t^m\}\phi_t^m$$

$$\beta_t^L = \sum_{m=0}^{N-1} \text{Prob}_{\tilde{\sigma}}\{E_t^m\}\left(\frac{1-F(\emptyset)}{N} - \frac{m}{N}\phi_t^m\right).$$

Now,

$$\text{Prob}_{\tilde{\sigma}}\{E_t^0\} = \left[F(\emptyset) + (1-F(\emptyset))(1-p_1)\right]\text{Prob}_{\tilde{\sigma}}\{E_{t-1}^0\}$$

and

$$\text{Prob}_{\tilde{\sigma}}\{E_t^m\} =$$
$$\phi_{t-1}^m\text{Prob}_{\tilde{\sigma}}\{E_{t-1}^m\} + (1-F(\emptyset)-\phi_{t-1}^m)(1-p_1)\text{Prob}_{\tilde{\sigma}}\{E_{t-1}^m\} + (1-F(\emptyset)-\phi_{t-1}^{m-1})p_1\text{Prob}_{\tilde{\sigma}}\{E_{t-1}^{m-1}\}$$

because $\tilde{\sigma}$ attains first-best.

To show that $\limsup_{t\to\infty}\tilde{\Omega}^t > \limsup_{t\to\infty}\hat{\Omega}^t$, by (1.28) we need only show that obligation decreases in ϕ_t^m. In (1.25), obligation is a smooth function of β_t^L and β_t^H, and hence $\{\phi_t^m\}_{m,t}$, so we must show $\frac{\partial\Omega_t}{\partial\phi_{t'}^m} \leq 0\ \forall t,t' \in \mathbb{N}, m \in \{1,...,N\}$.

Using (1.11) and repeatedly substituting for Ω^{t-1} yields

$$\frac{\partial\Omega^t}{\partial\phi_{t'}^m} = \frac{1}{\delta^{t-t'-1}}\frac{\partial}{\partial\phi_{t'}^m}\Omega^{t'+1} + \sum_{s=1}^{t-t'-1}\frac{1}{\delta^s}\left(\frac{\partial}{\partial\phi_{t'}^m}\beta_{t-s}^L p\delta\tilde{S} - \frac{\partial}{\partial\phi_{t'}^m}\beta_{t-s}^H(1-\delta)V^{FB}\right). \tag{1.29}$$

$$\frac{\partial\beta_t^L}{\partial\phi_{t'}^m} = \begin{cases} \sum_{k=1}^{N}\left(\frac{1-F(\emptyset)}{N} - \frac{k}{N}\phi_t^k\right)\frac{\partial\text{Prob}_{\tilde{\sigma}}\{E_t^k\}}{\partial\phi_{t'}^m} & \text{if } t' < t \\ -\frac{m}{N}\text{Prob}_{\tilde{\sigma}}\{E_{t'}^m\} & \text{if } t' = t \\ 0 & t' > t \end{cases}$$

$$\frac{\partial \beta_t^H}{\partial \phi_{t'}^m} = \begin{cases} \sum_{k=1}^N \frac{k}{N}\phi_t \frac{\partial \text{Prob}_{\tilde{\sigma}}\{E_t^k\}}{\partial \phi_{t'}^m} & \text{if } t' < t \\ \frac{m}{N}\text{Prob}_{\tilde{\sigma}}\{E_{t'}^m\} & \text{if } t' = t \\ 0 & t' > t \end{cases} .$$

Because $\tilde{\sigma}$ is symmetric and $\lim_{t\to\infty} \text{Prob}_{\tilde{\sigma}}\{E_t^N\} = 1$, $\beta_t^H + \beta_t^L \equiv \frac{1-F(\emptyset)}{N}$ and $\sum_{t=1}^\infty \beta_t^L \equiv \frac{1}{p_1}$. Differentiating the first identity gives us that

$$\frac{\partial}{\partial \phi_{t'}^m}\beta_t^L \equiv -\frac{\partial}{\partial \phi_{t'}^m}\beta_t^H \tag{1.30}$$

while differentiating the second gives us $\frac{\partial}{\partial \phi_t^m}\beta_t^L = -\sum_{s=t+1}^\infty \frac{\partial}{\partial \phi_t^m}\beta_s^L$. Hence

$$\frac{\partial}{\partial \phi_{t-1}^m}\Omega_t = \left(p_1\delta\tilde{V} + (1-\delta)V^{FB}\right)\frac{\partial}{\partial \phi_{t-1}^m}\beta_{t-1}^L. \tag{1.31}$$

The fact that $\frac{\partial}{\partial \phi_t^m}\beta_t^L = -\sum_{s=t+1}^\infty \frac{\partial}{\partial \phi_t^m}\beta_s^L$ (and $\frac{\partial}{\partial \phi_t^m}\beta_t^L < 0$) implies that

$$\frac{\partial}{\partial \phi_{t'}^m}\beta_{t'}^L < -\sum_{s=t'+1}^{\infty} \delta^{s-t'}\frac{\partial}{\partial \phi_{t'}^m}\beta_s^L < -\sum_{s=t'+1}^{t} \delta^{s-t'}\frac{\partial}{\partial \phi_{t'}^m}\beta_s^L \tag{1.32}$$

where we have used the fact that $\frac{\partial}{\partial \phi_t^m}\beta_s^L > 0$ for $s > t$. Together, (1.30) and (1.32) imply that

$$\begin{aligned} &\textstyle\sum_{s=1}^{t-t'-1} \frac{1}{\delta^s}\left(\frac{\partial}{\partial \phi_{t'}^m}\beta_{t-s}^L p_1\delta\tilde{S} - \frac{\partial}{\partial \phi_{t'}^m}\beta_{t-s}^H(1-\delta)V^{FB}\right) = \\ &\left(p_1\delta\tilde{S} + (1-\delta)V^{FB}\right)\textstyle\sum_{s=1}^{t-t'-1}\frac{1}{\delta^s}\frac{\partial}{\partial \phi_{t'}^m}\beta_{t-s}^L < -\left(p_1\delta\tilde{S} + (1-\delta)V^{FB}\right)\frac{\partial}{\partial \phi_{t'}^m}\beta_{t'}^L. \end{aligned} \tag{1.33}$$

If we plug this expression into (1.29) and apply (1.31), we finally conclude

$$\frac{\partial \Omega_t}{\partial \phi_{t'}^m} = \left(p_1\delta\tilde{S} + (1-\delta)V^{FB}\right)\frac{\partial}{\partial \phi_{t'}^m}\beta_{t'}^L + \left(p_1\delta\tilde{S} + (1-\delta)V^{FB}\right)\sum_{s=1}^{t-t'-1}\frac{1}{\delta^s}\frac{\partial}{\partial \phi_{t'}^m}\beta_{t-s}^L < 0$$

precisely as we wanted to show. ■

Completing Proof of Proposition 3

By Lemma A1, we need only show that $\lim_{t\to\infty}\Omega_t = \infty$ if the Favored Producer Automaton does not attain first-best. Define $F^m = \sum_{\mathcal{P}|\mathcal{P}\cap\{1,\ldots,m\}\neq\emptyset} F(\mathcal{P})$ as the probability that at least one of a set of m agents can produce.

Under the allocation rule (1.26), we can explicitly calculate $\phi^m = \frac{F^m}{m}$ as the probability that an agent is given production if he has produced y_H in the past and $m-1$ other agents have also produced y_H in the past. Plugging β_t^L and β_t^H into (1.25) yields

$$\Omega^t = \sum_{k=0}^{t-1}\frac{1}{\delta^k}\left(\sum_{m=0}^{N}\text{Prob}\{E_{t-1-k}^m\}\left(p_1\delta\tilde{S}\left(\frac{1-F(\emptyset)-F^m}{N}\right)-(1-\delta)V^{FB}\frac{F^m}{N}\right)\right).$$

This expression is positive and strictly increasing in $\tilde{S}$ because $1-F(\emptyset)-F^m \geq 0$ with equality iff $m = N$. Therefore, if $\exists\tilde{S}^*$ such that $\lim_{t\to\infty}\Omega_t = C$ for some $C\in\mathbb{R}$, then $\lim_{t\to\infty}\Omega_t = \infty$ for any $\tilde{S} > \tilde{S}^*$.

Any two strategies with identical $\{\beta_t^L\}_t$ and $\{\beta_t^H\}_t$ have identical $\{\Omega^t\}_t$. Whenever possible, the Favored Producer Automaton allocates production to agents who have already produced, so has the same obligation as $\hat{\sigma}$ from Lemma A2. Let S_G be agent i's surplus immediately after he produced y_H in the Favored Producer Automaton. If obligation converges when $\tilde{S} = S_G$, and production is allocated as in the Favored Producer Automaton, then it diverges for $\tilde{S} > S_G$ and the proposition is proven.

In the Favored Producer Automaton, at an h^t where $x_t = i$, $y_t = y_H$, agent i's future surplus is

$$\delta S_G = E\left[(1-\delta)V^{FB}\sum_{t'=1}^{\infty}\delta^{t'}b_{t+t'}|h^t\right] \tag{1.34}$$

where $b_{t+t'} = 1\{x_{t+t'} = i, \exists t'' < t+t' \text{ s.t. } x_{t''} = i, y_{t''} = y_H\}$. Let $\zeta_t = 1\{x_t = i; \forall t' < t, x_{t'} = i \to y_{t'} = 0\}$. Because (1.34) holds at *every* history h^t where $x_t = i$ and $y_t = y_H$, we can write

$$\delta S_G = E\left[(1-\delta)V^{FB}\sum_{t'=1}^{\infty}\delta^{t'}b_{t+t'}|h^t \text{ s.t. } y_t = y_H, \zeta_t = 1\right] \tag{1.35}$$

for any h^t where $x_t = i$, $y_t = y_H$. Because $e = 1$ at every on-path history, $\text{Prob}\{y_t = y_H|\zeta_t = 1\} = p_1$. The event $\{\zeta_t = 1, y_t = y_H\}$ can occur only once, so we can manipulate (1.35) by taking expectations and applying the Law of Iterated Expectations and dominated convergence to rewrite

$$\delta S_G \sum_{t=1}^{\infty} \delta^t \beta_t^L = (1-\delta) V^{FB} E\left[\sum_{t=1}^{\infty} \delta^t b_t\right]. \tag{1.36}$$

Using dominated convergence to interchange summation and integration and noting that $E[b_t] = \beta_t^H$, we have

$$p_1 \delta S_G \sum_{k=1}^{\infty} \delta^k \beta_k^L - (1-\delta) V^{FB} \sum_{k=1}^{\infty} \delta^k \beta_k^H = 0.$$

Hence, by definition of Ω_t, if production is allocated according to the Favored Producer Automaton, then $\lim_{t\to\infty} \delta^t \Omega_t = 0$ when $\tilde{S} = S_G$.

Finally, suppose $\tilde{S} = S_G$. Then by (1.25),

$$\Omega_t = \frac{1}{\delta^t}(1-\delta) V^{FB} \sum_{k=t+1}^{\infty} \delta^k \beta_k^H - \frac{1}{\delta^t} p_1 \delta S_G \sum_{k=t+1}^{\infty} \delta^k \beta_k^L.$$

In any first-best equilibrium, $\beta_k^L \to 0$ as $t \to \infty$, so the second term in this expression vanishes as $t \to \infty$. $\lim_{t\to\infty} \beta_k^H = \frac{1}{N}$ due to symmetry, so

$$\lim_{t\to\infty} \Omega_t = V^{FB} \delta \frac{1}{N}$$

so obligation converges at $\tilde{S} = S_G$. Thus, for $\tilde{S} > S_G$, obligation diverges to ∞, which implies by Lemma A1 that first-best cannot be attained. But for $\tilde{S} \leq S_G$, the Favored Producer Automaton attains first-best.

Finally, we argue that $\exists$ open $\Delta \subseteq [0,1]$ such that for $\delta \in \Delta$, the Favored Producer Allocation attains first best but a stationary equilibrium does not. Let δ_{Stat} solve $(1-\delta_{Stat}) \frac{c}{p_1 - p_0} = \delta_{Stat} \frac{1}{N} V^{FB}$. Proposition 2 implies that a stationary equilibrium attains first-best iff $\delta > \delta_{Stat}$, and Assumption 3 implies that $\delta_{Stat} > 0$. Since both sides of (1.6) are continuous in δ, by Lemma 2 it suffices to show that for $\delta = \delta_{Stat}$, the Favored Producer Automaton satisfies $E_{\sigma^*}\left[S_i\left(h^{t+1}\right) \middle| \left(h_i^{t-1}, \mathcal{I}_i\left(n_t\right)\right)\right] > \frac{1}{N} V^{FB}$ at any (h^{t-1}, n_t) immediately following $y_t = y_H$ such that $x_t = i$ and $e_t = 1$. At any such history, i knows with certainty that he will be assigned rank 1 next period. Further, if we take $r_{i,t}$ to denote the rank of player i at the start of round t in the Favored Producer Automaton, then for any on-path history h^t, $E_{\sigma^*}\left[S_i\left(h^t\right)\middle| h^t\right] = E_{\sigma^*}\left[S_i\left(h^t\right)\middle| r_{i,t}\right]$. By Assumptions 1 and 2, $E_{\sigma^*}\left[S_i\left(h^t\right)\middle| r_{i,t}\right]$ is strictly decreasing in $r_{i,t}$ since $r_{i,t} < r_{j,t}$ implies both that $\text{Prob}\left\{x_t = i\right\} > \text{Prob}\left\{x_t = j\right\}$ and that $\text{Prob}\left\{r_{i,t+t'} < r_{j,t+t'}\right\} > \frac{1}{2}$ for all $t' \geq 0$. Since $\sum_{i=1}^{N} E_{\sigma^*}\left[S_i\left(h^{t+1}\right)\middle| r_{i,t+1}\right] = V^{FB}$, $E_{\sigma^*}\left[S_i\left(h^{t+1}\right)\middle| r_{i,t+1} = 1\right] > \frac{1}{N}V^{FB}$, and so (1.6) is slack in the Favored Producer Allocation for $\delta = \delta_{Stat}$. This proves the result.■

Proposition 5

In Appendix B, we prove that the payoff set $FD(\delta)$ is recursive. Let $FD^{\mathcal{P}}(\delta)$ be the Pareto frontier of $FD(\delta)$. By definition of a full-disclosure equilibrium, agent x_t chooses $e_t = 1$ and produces $y_t = y_H$ immediately preceding history (h^{t-1}, n_t) if and only if

$$(1-\delta)\frac{c}{p_1 - p_0} \leq \delta E\left[S_{x_t}(h^t)|h^{t-1}, n_t\right]. \tag{1.37}$$

Let σ^* be an optimal full-disclosure equilibrium, and let $V^{Eff} = S_1(\sigma^*, \emptyset) + S_2(\sigma^*, \emptyset)$ be the total surplus produced in σ^*. Define $\alpha \equiv F(\{i\}) + F(\{1,2\})$, and note that $F(\{i\}) = F(\emptyset) - \alpha$.

Claim 1: Either $V^{Eff} = y_H p_0$, or $\forall h^t$ on the equilibrium path such that in each $t' \leq t$, either $e_{t'} = 0$ or $y_{t'} = 0$, $\text{Prob}_{\sigma^*}\left\{e_{t+1} = 1|h^t\right\} = 1$. Let h^t be a history such that such that $y_{t'} = 0\ \forall t' < t$.

We first argue that $\sum_{i=1}^{2} S_i(h^t) = \max_{v \in FD(\delta)} v_1 + v_2$. If not, then $\exists$ a FDE $\tilde{\sigma}$ with strictly higher surplus than $\sigma^*(h^t)$. Consider the alternative strategy $\hat{\sigma}$ that is identical to σ^* except at history h^t, at which $\hat{\sigma}(h^t) = \tilde{\sigma}$. Then $\hat{\sigma}$ generates more surplus than σ^* and satisfies (1.37). Contradiction.

If $(1-\delta)\frac{c}{p_1 - p_0} > \alpha(y_H p_1 - c)$, then (1.6) cannot hold and $V^{Eff} = y_H p_0$. Otherwise, Proposition 4 gives a full-disclosure equilibrium with surplus strictly larger than $y_H p_0$. In this case, $\exists\ h^t$ such that either $e_{t'} = 0$ or $y_{t'} = 0\ \forall t' < t$ and $\exists i \in \{1,2\}$ such that $\text{Prob}_{\sigma^*}\left\{e_{t+1} = 1|h^t, x_{t+1} = i\right\} > 0$. Let (h^t, n_{t+1}) be the history immediately following $x_{t+1} =$

i, $e_{t+1} = 1$, and $y_{t+1} = y_H$. Then

$$(1-\delta)(y_H p_1 - c) + \delta p_1 \sum_{j=1}^{2} E\left[S_j(h^t)|h^{t-1}, n_t\right] + (1-p_1)\delta V^{Eff} > (1-\delta)y_H p_0 + \delta V^{Eff} \quad (1.38)$$

since otherwise $V^{Eff} = (1-\delta)y_H p_0 + \delta V^{Eff}$, which contradicts $V^{Eff} > y_H p_0$. Let $S(h^{t+1})$ be the continuation payoffs following (h^t, n_{t+1}).

Consider a strategy σ that is identical to σ^* except after history h^t. Let $\text{Prob}_\sigma\{e_{t+1} = 1|h^t, \mathcal{P}_{t+1} \neq \emptyset\} = 1$. Following h^t, if $y_{t+1} = 0$, σ continues with a FDE generating V^{Eff}. if $y_t = y_H$, continuation dyad surplus is $S(h^{t+1})$ if $x_t = i$ and $(S_2(h^{t+1}), S_1(h^{t+1}))$ if $x_t \neq i$. These continuation payoffs are feasible because players are symmetric. Then σ satisfies (1.6) and generates strictly higher surplus than σ^* by (1.38). Contradiction.

Set-Up

Consider a full-disclosure equilibrium σ, and let (h^{t-1}, n_t) be a history satisfying $x_t = i$, $e_t = 1$, and $y_t = y_H$. Let $\hat{S} \in \mathbb{R}^2_+$ be a vector of dyad-specific surpluses in the continuation game. Let h^t be an equilibrium successor history to (h^{t-1}, n_t).

We consider a FDE that maximizes total continuation surplus such that i-dyad surplus is no less than $\hat{S}_i$. Such an equilibrium must solve

$$\max_{\{u_{1,t+t'}, u_{2,t+t'}\}} \sum_{t'=1}^{\infty} \delta^{t'}(1-\delta)E\left[u_{1,t+t'} + u_{2,t+t'} \,|\, h^{t-1}, n_t\right] \quad (1.39)$$

subject to the constraint that $E\left[u_{1,t+t'} + u_{2,t+t'}\right]_{t'=1}^{\infty}$ be feasible and $\forall i \in \{1,2\}$,

$$\sum_{t'=1}^{\infty} \delta^{t'}(1-\delta)E\left[u_{i,t+t'} \,|\, h^{t-1}, n_t\right] \geq \hat{S}_i \quad (1.40)$$

$$\begin{gathered}\sum_{s=\hat{t}}^{\infty} \delta^{s-\hat{t}}(1-\delta)E\left[u_{x_{t+s},t+s}\middle| h^{t+\hat{t}-1}, n_{t+\hat{t}}\right] \geq \tilde{S} \\ \forall \hat{t},\ (h^{t+\hat{t}-1}, n_{t+\hat{t}}) \text{ s.t. } e_{t+\hat{t}} = 1, y_{t+\hat{t}} = y_H\end{gathered} \quad (1.41)$$

Where (1.41) follows from (1.12).

Define

$$L_i^{t+t'} = \left\{\forall 1 \leq \hat{t} \leq t',\, x_{t+\hat{t}} = -i \Rightarrow (e_{t+\hat{t}} = 0 \text{ OR } y_{t+\hat{t}} = 0), h^t\right\} \quad (1.42)$$

as the event that $\left\{x_{t+\hat{t}}=-i, e_{t+\hat{t}}=1, y_{t+\hat{t}}=y_H\right\}$ has *not yet occurred* in any round after h^t, and let $L_{-i}^{t+t'}=\left\{L_i^{t+t'}\right\}^C$ be the complement of this event. Define $\lambda_{t+t'}=\text{Prob}_\sigma\left\{e_{-i,t+t'}=1|x_{t+t'}=-i, L_i^{t+t'}\right\}$, $\beta_{t+t'}=\text{Prob}_\sigma\left\{x_{t+t'}=i|L_i^t\right\}$, and $b_{t+t'}=\left\{x_{t+t'}=-i, e_{-i,t+t'}=1, L_i^{t+t'}\right\}$, and note that

$$\text{Prob}_\sigma\left\{L_{-i}^{t+t'}\right\}=\sum_{k=1}^{t'-1}\text{Prob}_\sigma\left\{b_{t+t'}\right\}p_1. \tag{1.43}$$

Define $(S_{1,t+t'}^*, S_{2,t+t'}^*)=S_{t+t'}^*$ by

$$S_{i,t+t'}^*=\sum_{s=t'}^{\infty}\delta^s(1-\delta)E\left[u_{x_{t+s},t+s}\middle| h^t, b_{t+t'}, y_{t+t'}=y_H\right]$$

and note that

$$E\left[u_{j,t+t'}|L_{-i}^{t+t'}\right]\text{Prob}_\sigma\left\{L_{-i}^{t+t'}\right\}=\sum_{s=1}^{t'}E\left[u_{j,t+t'}|b_{t+s}, y_{t+s}=y_H\right]\text{Prob}_\sigma\{b_{t+s}, y_{t+s}=y_H\}$$

because the events $\{b_{t+s}, y_{t+s}=y_H\}$ for $s\in\{1,...,t'\}$ partition the event $L_{-i}^{t+t'}$. Consider a relaxation of the problem (1.39)-(1.41) that ignores constraint (1.41) for player i, and replaces this constraint for $-i$ with the weaker set of constraints

$$\Sigma_{s=\hat{t}}^{\infty}\delta^{s-\hat{t}}(1-\delta)E\left[u_{x_{t+s},t+s}\middle| h^{t-1}, b_{t+\hat{t}}, y_{t+\hat{t}}=y_H\right]=S_{-i,t+\hat{t}}^*\geq\tilde{S}$$
$$S_{t+\hat{t}}^*\in FD(\delta)$$

The first constraint is a relaxation of the true constraint (1.41), while the second is implied by (1.41). Manipulating the problem (1.39)-(1.41) yields the following relaxed problem:

$$\max_{\{x_{t+t'},e_{t+t'}S_{t+t'}^*\}}\sum_{t'=1}^{\infty}\delta^{t'}(1-\delta)\sum_{j=1}^{2}\left[E\left[u_{j,t+t'}|L_i^{t+t'}\right]\text{Prob}_\sigma\left\{L_i^{t+t'}\right\}+\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}, y_{t+t'}=y_H\}S_{j,t+t'}^*\right] \tag{1.44}$$

subject to:

1. Promise-keeping constraints for agent $j\in\{1,2\}$

$$\sum_{t'=1}^{\infty}\delta^{t'}(1-\delta)\left[E\left[u_{j,t+t'}|L_i^{t+t'}\right]\text{Prob}_\sigma\left\{L_j^{t+t'}\right\}+\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}, y_{t+t'}=y_H\}S_{j,t+t'}^*\right]\geq\hat{S}_j. \tag{1.45}$$

2. Constraint (1.37) holds if event $L_{-i}^{t+t'}$ occurs:

$$S^*_{-i,t+t'} \geq \tilde{S},\ \forall t' \geq 1. \tag{1.46}$$

3. Continuation play after $L_{-i}^{t+t'}$ must be a full-disclosure equilibrium:

$$S^*_{t+t'} \in FD(\delta),\ \forall t' \geq 1. \tag{1.47}$$

Claim 2: Without loss of generality, $S^*_{1,t+t'} = S^*_1$ and $S^*_{2,t+t'} = S^*_2$, $\forall t' \geq 1$. Let σ solve (1.44)-(1.47) but not satisfy Claim 2. Consider strategies identical to σ, except that whenever event $b_{t+t'}$ occurs, continuation play generates surplus $S^* = (S^*_1, S^*_2)$, where $S^* \in \arg\max_{V \in FD(\delta)} \sum_{i=0}^{2} V_i$ subject to

$$V_i = \frac{1}{\sum_{t'=0}^{\infty} \delta^t \text{Prob}_\sigma\{b_{t+t'}, y_{t+t'} = y_H\}} \sum_{t'=0}^{\infty} \delta^t \text{Prob}_\sigma\{b_{t+t'}, y_{t+t'} = y_H\} S^*_{i,t+t'}, i \in \{1, 2\}.$$

S^*_i is the limit of a sequence of convex combinations of FDE payoffs. $FD(\delta)$ is bounded, closed, and convex as a corollary of Lemma 2, proven in Appendix B. Therefore, $S^* \in FD(\delta)$. Finally, if the original $\{S^*_{j,t+t'}\}$ solved the relaxed problem (1.44)-(1.47), then so does S^*, which proves the claim.

Return to Derivations

Recall that $x_t = i$, $e_t = 1$, and $y_t = y_H$ in round t of h^{t+1}. Therefore, $\hat{S}_i \geq \tilde{S}$. Relax the problem (1.44)-(1.47) by ignoring (1.45) for agent $-i$. Furthermore, while $\{\beta_{t+t'}\}_{t'}$ and $\{\text{Prob}_\sigma\{b_{t+t'}\}\}$ are closely related to one another, we relax the problem further by allowing these two sets of variables to be chosen independently. Total surplus increases and all constraints are relaxed when $\beta_{t+t'}$ increases, so it is optimal to choose $\beta_{t+t'} = \alpha\ \forall t'$.

Under this further simplification, and following much manipulation (omitted here for brevity), the relaxed program may be written

$$\max_{\{\text{Prob}_\sigma\{b_{t+t'}\}\},S_1^*,S_2^*} \begin{array}{c} \alpha\delta\left(y_H(p_1-p_0)-c\right)+\delta y_H p_0 + \\ \left((1-\delta)\left(y_H\left(p_1-p_0\right)-c\right)+\delta p_1\left(S_1^*+S_2^*\right)\right)\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}\} - \\ \left(\delta\alpha\left(y_H(p_1-p_0)-c\right)-\delta y_H p_0 p_1\right)\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}\} \end{array} \tag{1.48}$$

subject to the promise-keeping constraint

$$\delta\hat{S}_i \leq \alpha\delta\left(y_H p_1 - c\right) + \delta p_1\left(S_i^* - \alpha\left(y_H p_1 - c\right)\right)\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}\} \tag{1.49}$$

and the equilibrium constraints

$$S_{-i}^* \geq \tilde{S}, \quad S^* \text{a full-disclosure eq'm} \tag{1.50}$$

Effectively, this problem has three choice variables, $\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}\}$, S_1^*, and S_2^*.

$S_i^* \leq \alpha(y_H p_1 - c)$ because $\alpha(y_H p_1 - c)$ is the maximum continuation surplus that can be given to agent i, and S^* is optimally on the Pareto frontier of $FD(\delta)$. Therefore,

$$\delta\alpha(y_H(p_1-p_0)-c)+\delta y_H p_0 p_1 \leq \delta\alpha(y_H p_1 - c) + \delta(1-\alpha)y_H p_0 \leq S_1^* + S_2^*$$

because $\delta\alpha(y_H p_1 - c) + \delta(1-\alpha)y_H p_0$ is the value of exclusive dealing with a single agent and thus the smallest Pareto-efficient payoff. Thus,

$$\left((1-\delta)\left(y_H\left(p_1-p_0\right)-c\right)+\delta p_1\left(S_1^*+S_2^*\right)-\delta\alpha\left(y_H(p_1-p_0)-c\right)-\delta y_H p_0 p_1\right) \geq 0.$$

The maximal feasible value of $\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}\}$ that satisfies (1.49) solves this problem.

Next, we show that (1.49) binds at the optimum. Some value of $\sum\delta^{t'}\text{Prob}_\sigma\{b_{t+t'}\}$ makes (1.49) bind, so it suffices to show that some strategy attains this value. Let S_G and S_B be dyad-specific payoffs in the good and bad states of the Favored Producer Automaton, respectively. Because first-best cannot be attained, $\tilde{S} > S_G$. $S_{-i}^* \geq \tilde{S}$, so $S_i^* < V^{FB} - \tilde{S} < S_B$. Because $\hat{S}_i \geq \tilde{S}$, (1.49) is violated when $\beta_{t+t'} = \alpha$, $\Psi_{t+t'} = 1$ always; otherwise, the Favored Producer Automaton would induce high effort and thus attain the first-best. Additionally, (1.49) must be satisfied when $\sum \delta^{t'} \text{Prob}_\sigma\{b_{t+t'}\} = 0$, since otherwise $\hat{S}_i > \alpha(y_H p_1 - c)$, which is not feasible.

Consider the following class of strategies. Agent i chooses $e_{t+t'} = 1$. With probability q, agent $-i$ chooses $e_{t+t'} = 0$, $\forall t'$; we will henceforth refer to this continuation strategy as "exclusive dealing with agent i." Otherwise, agent $-i$ chooses $e_{t+t'} = 1$ $\forall t'$ until $b_{t+t'} = 1$. In either case, $\beta_{t+t'} = \alpha$ until $b_{t+t'} = 1$. Once $b_{t+t'} = 1$, continuation play generates fixed surplus S^*. For this class of strategies, agent i's dyad-specific surplus V_i^q is given by

$$S_i^q = \alpha\left[(1-\delta)(y_H p_1 - c) + \delta(q\alpha(y_H p_1 - c) + (1-q)S_i^q\right] + (1-\alpha)\left[\delta p_1 S_i^* + \delta(1-p_1)S_i^q\right].$$

Rearranging this expression, we immediately show that S_i^q is continuous in q, with $\lim_{q\to 1} S_i^q = \alpha(y_H p_1 - c)$ and $\lim_{q\to 0} S_i^q \leq V_G$. By the Intermediate Value Theorem, there exists a q^* such that $S_i^{q^*} = \hat{S}_i$. $\beta_{t+t'} = \alpha$ in this class of strategies, so the value of $\sum \delta^{t'} \text{Prob}_\sigma\{b_{t+t'}\}$ corresponding to q^* makes (1.49) bind.

Finally, the Pareto frontier is downward-sloping, so if (1.49) binds then the promise $\hat{S}_{-i}$ must also be satisfied or else $\hat{S}$ is infeasible. Fixing continuation play S^*, the class of equilibria characterized by (q_1, q_2)-Exclusive Dealing solves the relaxed problem. It remains to show that $q_1 = q_2 = q^*$ at each relevant history h^{t+1}.

Claim 3: Fix the vector of dyad-specific surpluses (S_1, S_2) at history h^{t+1}. $\exists$ a solution $(\bar{S}_1, \bar{S}_2)$ to (1.48) subject to (1.49) and (1.50) such that $(S_1, S_2) \leq (\bar{S}_1, \bar{S}_2)$ and $\bar{S}_{-i}^* = \tilde{S}$. Define $S^{Eff,i} = (S_1^{Eff,i}, S_2^{Eff,i})$ by

$$S^{Eff,i} = \arg\max_{V \in FD(\delta)} V_i$$

subject to

$$S_1^{Eff,i} + S_2^{Eff,i} = V^{Eff}.$$

There are two cases to consider: (1) $S_i^{Eff,i} \geq \tilde{S}$, or (2) $S_i^{Eff,i} < \tilde{S}$.

Consider case (2). In this case, $S_1^* + S_2^*$ is decreasing in S_{-i}^* because the payoff frontier is convex. Let

$$\bar{S}^* = \arg\max_{v \in FD(\delta)} v_2$$

subject to

$$v_1 = \tilde{S}.$$

Define $Q = \sum_{t'=1}^{\infty} \delta^{t'} \text{Prob}_\sigma\{b_{t+t'}\}$. Let (Q, S^*) satisfy (1.49), and consider an alternative $(\bar{Q}, \bar{S}^*)$, where $\bar{Q}$ chosen so that (1.49) binds. Now $\tilde{S} \leq S_{-i}^*$, so $\bar{S}_1^* + \bar{S}_2^* \geq S_1^* + S_2^*$ because $FD(\delta)$ is convex. Moreover, $\bar{S}_i^* \geq S_i^*$ because the Pareto frontier of $FD(\delta)$ is downward-sloping and $\bar{S}_{-i}^* \leq S_{-i}^*$, so (1.49) continues to be satisfied. If $\bar{S}_i^* = S_i^*$, then $\bar{S}_{-i}^* \geq S_{-i}^*$, $Q = \bar{Q}$, and $(S_1, S_2) \leq (\bar{S}_1, \bar{S}_2)$ follows immediately. If $\bar{S}_i^* > S_i^*$, then $\bar{Q} > Q$, which further increases total surplus (1.48). Moreover, because (1.49) binds in both solutions, $S_i^* = \bar{S}_i^* = \hat{S}_i$, and so $\bar{S}_{-i}^* \geq S_{-i}^*$ because $\bar{S}_1^* + \bar{S}_2^* \geq S_1^* + S_2^*$. This proves the claim in case (2).

Consider case (1). I claim that case (1) contradicts the assumption that first-best cannot be attained. By claim 1, $e = 1$ in the first round, and so we must have

$$V^{Eff} = (1-\delta)(y_H p_1 - c) + \delta V^{Eff}$$

which implies that $V^{Eff} = y_H p_1 - c = V^{FB}$. Contradiction.

Final Steps

By claim 3, the optimal equilibrium awards player i with exactly $\tilde{S}$ surplus following ($x_t = i$, $e_t = 1$, $y_t = y_H$) regardless of the rest of the history h^t. Together with previous results, we conclude that (q^*, q^*)-Exclusive Dealing with q^* chosen so (1.6) binds is a solution to the relaxed problem. This is an equilibrium, proving the claim.■

Proposition 6

It suffices to show that $\exists$ a relational contract that conceals information and is strictly better than the optimal full-disclosure equilibrium. Define $\alpha \equiv F(\{1\}) + F(\{1,2\})$.

By Proposition 5, $\exists q^* \in [0,1]$ such that q^*-Exclusive Dealing is an optimal full-disclosure equilibrium. By (1.13), $q = 1$ would gives players strict incentives to work hard; because first-best cannot be attained, $q^* \in (0,1)$. For any $\epsilon > 0$ such that $q^* - \epsilon \in (0,1)$, define

$$\phi(\epsilon) \equiv \frac{(1 - q^* + \epsilon)p_1}{(1 - q^* + \epsilon)p_1 + (1 - p_1)} < p_1.$$

Note that $\exists\ \bar{q}(\epsilon), \underline{q}(\epsilon)$ such that $1 \geq \bar{q}(\epsilon) > \underline{q}(\epsilon) \geq 0$ and $\phi(\epsilon)\underline{q}(\epsilon) + (1 - \phi(\epsilon))\bar{q}(\epsilon) = q^*$.

We construct a relational contract that conceals information; by Lemma 2 and Corollary ??, it suffices to specify an information partition, allocation rule, accept/reject decision, and effort choice. In $t = 1$, $x_1 \in \mathcal{P}_1$ is chosen randomly and $d_1 = e_1 = 1$. Without loss of generality, let $x_1 = 1$. In $t = 2$, agent 2 remains uninformed of y_1 with probability 1 if $y_1 = 0$ and with probability $1 - q^* + \epsilon$ if $y_1 = y_H$. If $1 \in \mathcal{P}_2$ then $x_2 = 1$, otherwise agent 2 is given production whenever possible. If $x_2 = 1$, then $d_2 = e_2 = 1$. If $x_2 = 2$, then $d_2 = 1$, and $e_2 = 1$ if and only if agent x_2 was not informed of y_1. From $t = 3$ onwards, the continuation equilibrium is chosen from the full-disclosure Pareto frontier. Let $V^i \in FD(\delta)$ be the full-disclosure payoff that maximizes agent i's payoff among all **optimal** full-disclosure payoffs (that is, those which maximize total surplus), and $V^{Ex,i} \in FD(\delta)$ be the continuation payoff when i is given exclusive dealing: $V_i^{Ex,i} = \alpha(y_H p_1 - c)$ and $V_{-i}^{Ex,i} = (F(\emptyset) - \alpha)y_H p_0$. If $x_1 = 1$, then the continuation payoff $V^{Eq}(y_1, y_2) \in FD(\delta)$ is chosen identically to the optimal full-disclosure equilibrium. If $x_1 = 2$ and $y_1 = y_H$ was revealed to 2, then the continuation payoff is $V^{NEq,L}(y_1, y_2) = V^{Ex,x_1}$. If $x_1 = 2$ and y_1 was not revealed, then the continuation payoff is $V^{NEq,H}(y_1, y_2)$, where $V^{NEq,H}(0,0) = V^{NEq,H}(y_H, 0) = V^{x_1}$, $V^{NEq,H}(y_H, y_H) = \underline{q}(\epsilon)V^{Ex,x_2} + (1 - \underline{q}(\epsilon))V^{x_2}$, and $V^{NEq,H}(0, y_H) = \bar{q}(\epsilon)V^{Ex,x_2} + (1 - \bar{q}(\epsilon))V^{Ex,x_2}$.

By Lemma 2, we need only show that (1.6) holds for the agents at each history. For $t \geq 3$, the continuation equilibrium is full-disclosure and so this inequality holds by definition. In round $t = 2$ and $x_2 = 1$, this constraint is

$$\tilde{S} \leq (q^* V_{x_1}^{Ex,x_1} + (1-q^*)V_{x_1}^{x_1}) \tag{1.51}$$

which holds because q^*-Exclusive Dealing is an equilibrium. If $x_2 = 2$ and $y_1 = y_H$ is revealed, then $e_2 = 0$ and so (1.6) does not apply. If y_1 is concealed, then agent 2 believes $\text{Prob}_\sigma \{y_1 = y_H | h_2^2\} = \phi(\epsilon)$. Therefore, agent 2's expected continuation surplus following $y_2 = y_H$ is

$$\phi(\epsilon)V^{NEq,H}(y_H, y_H) + (1-\phi(\epsilon))V^{NEq,L}(0, y_H) = q^* V_{x_2}^{Ex,x_2} + (1-q^*)V_{x_2}^{x_2} \geq \tilde{S}$$

where equality is follows from the definition of $V^{NEq,H}$, $\bar{q}$, and $\underline{q}$ and the inequality is a result of agent symmetry and (1.51).

Next, we check 1's incentives in $t = 1$. Following $y_1 = y_H$, 1 believes he receives continuation surplus

$$S_1(y_H, \epsilon) = \begin{array}{c} (1-\delta)\alpha(y_H p_1 - c) + (q^* - \epsilon)\delta V_1^{Ex,1} + \\ (1 - q^* + \epsilon)\alpha\delta \left(p_1 \left[(1-q^*)V_1^1 + q^* V_1^{Ex,1}\right] + (1-p_1)V_1^1\right) + \\ (1-\alpha)(1-q^*+\epsilon)\delta \left(p_1(\underline{q}(\epsilon)V_1^{Ex,2} + (1-\underline{q}(\epsilon))V_1^2) + (1-p_1)V_1^1\right). \end{array} \tag{1.52}$$

Because $\underline{q}(\epsilon) < q^*$ and $V_1^{Ex,2} < V_1^2$,

$$p_1(\underline{q}(\epsilon)V_1^{Ex,2} + (1-\underline{q}(\epsilon))V_1^2) + (1-p_1)V_1^1 > p_1(q^* V_1^{Ex,2} + (1-q^*)V_1^2) + (1-p_1)V_1^1$$

so

$$S_1(y_H, \epsilon) > \begin{array}{c} (1-\delta)\alpha(y_H p_1 - c) + (q^* - \epsilon)\delta V_1^{Ex,1} + \\ (1 - q^* + \epsilon)\alpha\delta \left(p_1 \left[(1-q^*)V_1^1 + q^* V_1^{Ex,1}\right] + (1-p_1)V_1^1\right) + \\ (1-\alpha)(1-q^*+\epsilon)\delta \left(p_1(q^* V_1^{Ex,2} + (1-q^*)V_1^2) + (1-p_1)V_1^1\right). \end{array}$$

The right-hand side of this expression is continuous in ϵ and equals $q^* V_1^{Ex,1} + (1-q^*)V_1^1$ when $\epsilon = 0$. Hence, $S_1(y_H, 0) > \tilde{S}$ by (1.51), and so $\exists \epsilon > 0$ such that $S_1(y_H, \epsilon) \geq \tilde{S}$.

It remains to show that an equilibrium with such an ϵ generates strictly higher surplus than q^*-Exclusive Dealing. In $t = 1$, total surplus in both equilibria is $(1-\delta)(y_H p_1 - c)$. In $t = 2$, the equilibrium that conceals information generates strictly higher surplus, since

$$\begin{gathered}(\alpha + (1-\alpha)(1 - p_1 + (1 - q^* + \epsilon)p_1))\,(y_H p_1 - c) > \\ (\alpha + (1-\alpha)(1 - p_1 + (1 - q^*)p_1))\,(y_H p_1 - c)\end{gathered}$$

because $\epsilon > 0$. In the continuation game starting at $t = 3$, the equilibrium that conceals information has total surplus

$$\begin{gathered}\alpha E\left[V_1^{Eq}(y_1, y_2) + V_2^{Eq}(y_1, y_2)|e_1 = e_2 = 1\right] + (1-\alpha)p_1(q^* - \epsilon)(V_1^{Ex,1} + V_2^{Ex,1}) + \\ (1-\alpha)(1 - p_1(q^* - \epsilon))\left(q^*(V_1^{Ex,2} + V_2^{Ex,2}) + (1-q^*)(V_1^2 + V_2^2)\right).\end{gathered}$$

On the other hand, q^*-Exclusive dealing has continuation payoffs

$$\begin{gathered}\alpha E\left[V_1^{Eq}(y_1, y_2) + V_2^{Eq}(y_1, y_2)|e_1 = e_2 = 1\right] + (1-\alpha)p_1 q^*(V_1^{Ex,1} + V_2^{Ex,1}) + \\ (1-\alpha)(1 - p_1 q^*)\left(q^*(V_1^{Ex,2} + V_2^{Ex,2}) + (1-q^*)(V_1^2 + V_2^2)\right).\end{gathered}$$

Comparing these payoffs, we find that the concealing-information equilibrium dominates the optimal full-disclosure equilibrium so long as

$$-\left(V_1^{Ex,1} + V_2^{Ex,1}\right) + \left(q^*(V_1^{Ex,2} + V_2^{Ex,2}) + (1-q^*)(V_1^2 + V_2^2)\right) \geq 0$$

which holds because $V_1^{Ex,1} + V_2^{Ex,1} = V_1^{Ex,2} + V_2^{Ex,2} < V_1^2 + V_2^2$. Thus, we have found an equilibrium that is strictly better than the optimal full-disclosure equilibrium. ■

Lemma 4

It suffices to show that in each round, players' strategies are best-responses in the game where all variables except e are publicly observed, since messages reveal the true history along the equilibrium path. By construction, the agent is willing to choose $e = 1$ if the IC constraint holds. Given that agent surplus is $U_i = \gamma V_G$ following high output and the principal's payoff is $(1-\gamma)V^{FB}$ in any first-best equilibrium, we have

$$p_1[\delta(1-\gamma)V^{FB} + \delta\gamma V_G] - (1-\delta)c \geq p_0[\delta(1-\gamma)V^{FB} + \delta\gamma V_G]$$

which can be rearranged to yield the condition

$$(1-\delta)\frac{c}{p_1 - p_0} \leq \delta\left[(1-\gamma)V^{FB} + \gamma V_G\right]$$

In each round, the agent awarded production earns $\gamma(1-\delta)(y_H p_1 - c)$ and the principal earns $(1-\gamma)(1-\delta)(y_H p_1 - c)$, so Assumption 4 is satisfied. If players follow the message strategy, the principal is willing to follow the allocation rule and pay $\tau(y_H) = (1-\gamma)\frac{\delta}{1-\delta}V^{FB}$, since any deviation is immediately revealed by messages and mutually min-maxed, leaving the principal a payoff of 0. Similarly, agent i is willing to pay $\frac{\delta}{1-\delta}\gamma V_G$ if favored and $\frac{\delta}{1-\delta}\gamma V_B$ otherwise.

Finally, we must check that messages are incentive compatible. Any unilateral deviation is immediately observed because the producing agent and principal send identical messages in each period. Thus, any deviation is punished by mutual min-maxing, yielding a continuation payoff of 0. Hence, no player has an incentive to deviate from the prescribed message profile and these strategies form an equilibrium. ■

Proposition 7

Claim 1: Obligation is the Correct Notion

We closely follow the proof of Proposition 3. Given strategy σ that yields first-best surplus, define the *residual expected obligation for player i* as

$$\hat{\Omega}_t^i = \frac{\hat{\Omega}_{t-1}^i}{\delta} + \beta_t^L p_1 \delta(\tilde{S} - (1-\gamma)V^{FB}) - \beta_t^H\left[(1-\delta)\gamma(y_H p_1 - c)\right]. \quad (1.53)$$

We first claim that if $\limsup_{t\to\infty} \hat{\Omega}_t^i = \infty$, then the strategy σ cannot form an equilibrium that attains first-best.

Under Assumption 4, continuation surpluses for agent i and the principal at on-path history h^t are

$$U_i(h^t) = \gamma V^{FB} E_\sigma \left[\sum_{t'=0}^{\infty} \delta^{t'}(1-\delta)1_{i,t+t'} \mid h^t \right],$$

$$U_0(h^t) = (1-\gamma)V^{FB}$$

respectively. Therefore, the necessary condition (1.14) can be rewritten

$$(1-\delta)\frac{c}{p_1 - p_0} \leq \delta E_\sigma \left[(1-\gamma)V^{FB} + \gamma V^{FB} \sum_{t'=0}^{\infty} \delta^{t'}(1-\delta)1_{i,t+t'} \mid (h_{x_t}^{t-1}, n_t, y_H) \right] \quad (1.54)$$

where $n_t = (\mathcal{P}_t, x_t, w_{i,t}, d_t = 1, e_t = 1) \in \mathcal{N}_e$. As before, consider the relaxed game where (1.54) must only hold the first time each agent i produces y_H. Define $b_i(h_i^t)$ as in (1.21); then we can write (1.54) as

$$b_i(h_i^t)p_1\delta\tilde{S} \leq b_i(h_i^t)p_1\delta E_\sigma \left[(1-\gamma)V^{FB} + \gamma V^{FB} \sum_{t'=0}^{\infty} \delta^{t'}(1-\delta)1_{i,t+t'} \mid (h_{x_t}^t, n_t, y_H) \right]$$

which can be rearranged to yield

$$b_i(h_i^t)p_1\delta(\tilde{S} - (1-\gamma)V^{FB}) \leq b_i(h_i^t)p_1\delta E_\sigma \left[\gamma V^{FB} \sum_{t'=0}^{\infty} \delta^{t'}(1-\delta)1_{i,t+t'} \mid (h_{x_t}^t, n_t, y_H) \right].$$

This expression is nearly identical to the corresponding inequality (1.22), with the sole exception that the obligation incurred upon production is $\tilde{S} - (1-\gamma)V^{FB}$ rather than $\tilde{S}$. A stationary equilibrium attains first-best if $\tilde{S} - (1-\gamma)V^{FB} \leq 0$. If $\tilde{V} - (1-\gamma)V^{FB} > 0$, then an argument identical to Lemma A1 proves the desired result.

Claim 2: Business Allocated to an Agent who has Already Produced

Next, we argue that Lemma A2 holds for $\hat{\Omega}_t^i$: that is, $\limsup_{t\to\infty} \hat{\Omega}_t^i < \infty$ for first-best strategy σ only if $\limsup_{t\to\infty} \hat{\Omega}_t^i < \infty$ for first-best strategy $\hat{\sigma}$ that (1) randomizes agent labels at the beginning of the game, and (2) awards production to an agent that has already produced y_H whenever possible. Result (1) follows by the same argument as in Lemma A2.

For Result (2), $\tilde{S} - (1-\gamma)V^{FB} > 0$ is independent of the allocation rule, so the only difference between this argument and the proof of Lemma A2 is that the final term in (1.53) is multiplied by γ. Performing the same set of manipulations, we can show

$$\frac{\partial \hat{\Omega}_t}{\partial \phi_{t'}^m} = \begin{array}{c} \left(p\delta(\tilde{S} - (1-\gamma)V^{FB}) + (1-\delta)\gamma V^{FB}\right) \frac{\partial}{\partial \phi_{t'}^m} \beta_{t'}^L + \\ \left(p\delta(\tilde{S} - (1-\gamma)V^{FB}) + (1-\delta)\gamma V^{FB}\right) \sum_{s=1}^{t-t'-1} \frac{1}{\delta^s} \frac{\partial}{\partial \phi_{t'}^m} \beta_{t-s}^L \end{array}$$

which is negative by (1.33), proving the desired claim.

Claim 3: $\lim_{t\to\infty} \hat{\Omega}_t = \infty$ if the Favored Producer Automaton does not attain first-best

Finally, we prove the proposition by arguing that $\hat{\Omega}_t \to \infty$ whenever the Favored Producer Automaton does not generate first-best surplus. We follow Proposition 3: in this context, the Favored Producer Automaton gives agent i continuation surplus

$$\delta V_1 = \gamma E\left[(1-\delta)V^{FB} \sum_{t'=1}^{\infty} \delta^{t'} b_{t+t'} | h^t\right]$$

as in (1.34). Manipulating this expression to resemble (1.36), we have

$$p_1 \delta S_G \sum_{k=1}^{\infty} \delta^k \beta_k^L - (1-\delta)\gamma V^{FB} \sum_{k=1}^{\infty} \delta^k \beta_k^H = 0.$$

In order for the Favored Producer Automaton to attain first-best, it must be that $S_G \leq \tilde{S} - (1-\gamma)V^{FB}$. Plugging $S_G = (\tilde{S} - (1-\gamma)V^{FB})$ into a non-recursive expression for residual obligation, we yield

$$\delta^t \hat{\Omega}_t + p_1 \delta(\tilde{S} - (1-\gamma)V^{FB}) \sum_{k=t+1}^{\infty} \delta^k \beta_k^L - (1-\delta)\gamma V^{FB} \sum_{k=t+1}^{\infty} \delta^k \beta_k^H = 0$$

and hence

$$\hat{\Omega}_t = \frac{1}{\delta^t}(1-\delta)\gamma V^{FB} \sum_{k=t+1}^{\infty} \delta^k \beta_k^H - \frac{1}{\delta^t} p_1 \delta(\tilde{S} - (1-\gamma)V^{FB}) \sum_{k=t+1}^{\infty} \delta^k \beta_k^L.$$

In any first best equilibrium $\beta_k^L \to 0$ as $t \to \infty$, so the second term in this expression vanishes as $t \to \infty$. $\lim_{t\to\infty} \beta_k^H = \frac{1}{N}$ due to symmetry, so

$$\lim_{t\to\infty} \hat{\Omega}_t = V^{FB}\delta\gamma\frac{1}{N}$$

and this *lower bound* on obligation converges at $\tilde{S} - (1-\gamma)V^{FB} = S_G$, which proves that $\hat{\Omega}_t \to \infty$ for any larger $\tilde{S}$, as desired.■

Chapter 2

Attaining Efficiency with Public Monitoring and Markov Adverse Selection

2.1 Introduction

Many economically interesting settings entail both moral hazard and adverse selection problems. Collusive firms have private information about their costs and imperfectly observe their opponents' pricing decisions; a boss motivates her workers without knowing the opportunity cost of their effort or the profitability of the activity; a regulator enforces pricing rules to maximize welfare, but does not observe all actions taken by a monopolist or underlying market demand. Moreover, private information about players' utility is frequently persistent, so players have different beliefs about both present and future payoffs at any given time. The purpose of this paper is to take one step toward the analysis of these situations by proving that nearly efficient payoffs can be obtained with patient players in games when one player has Markov private information about her payoff and monitoring is imperfect but public. More precisely, any payoff in the interior of the convex hull of Pareto efficient and min-max payoffs can be approximated by an equilibrium for patient players. The proof partially constructs a class of equilibria and sheds intuitive light on the underlying structure of these games.

In recent years, a growing body of research has focused on dynamic Bayesian games with persistent private types. When utility is transferable, players can commit to actions, and private information is Markov, Athey and Segal (2012) construct a dynamic team mechanism that implements an efficient allocation rule and can be replicated in an equilibrium of the game without commitment. In work that is closely related to this paper, Escobar and Toikka (2012)—henceforth referred to as ET—prove a similar efficiency result for a general class of games with observable actions and Markov private types. Building upon an idea proposed by Jackson and Sonnenschein (2007) and applying it dynamic mechanism design problems with Markov private types, ET show that this mechanism can be replicated by an equilibrium using Fudenberg and Maskin's (1986) "carrot-and-stick" punishments. I also use Jackson and Sonnenschein's idea, albeit in a simpler mechanism design problem with one-sided private information. Because monitoring is imperfect and hence harsh punishments would lead to inefficient payoffs on the equilibrium path, I cannot use ET's techniques to support this mechanism as an equilibrium. Instead, I modify Fudenberg, Levine, and Maskin's (1994)—henceforth FLM—and Fudenberg and Levine's (1994) analysis of repeated games with imperfect public monitoring to show that continuation play can be used to *approximate transfers* when players are patient.[1] Renault, Solan, and Vieille (2012, henceforth RSV) show that efficient payoffs can be attained in dynamic sender-receiver games with perfect monitoring where the sender's type affects both players' utility. Adapting their argument, I partially extend my result to some games in which player 1's type affects everyone's payoff.

Two fundamental problems complicate the application of FLM to a game with Markov adverse selection. First, the player with private information must report in such a way that the resulting payoffs are approximately efficient. I initially focus on this adverse selection problem, demonstrating that nearly efficient payoffs can be obtained via a dynamic mechanism without transfers when players are patient and can *commit* to actions as a function of the reported type. Given this result, I turn to the problem of enforcing the correct actions in equilibrium. Here, the difficulty is that a player's private information informs her beliefs

[1]In independently-developed work, Horner, Takahashi, and Vieille (not yet available) extend Fudenberg and Levine (1994) to games with Markov private types.

about *continuation* payoffs. If naively designed, a punishment harsh enough to deter *all* types from deviating might lead to quite inefficient continuation play, even when players are patient. I use the fact that the private type in each round is not very informative about payoffs *in the distant future* to sidestep this problem. Because types evolve according to a Markov process, all players have approximately the same beliefs about the distribution of types in the far future, so rewards and punishments in these distant periods can be precisely targeted to provide incentives without being excessively harsh. In order to identify a potential deviator, the distribution of public signals must satisfy a "pairwise full rank" condition that is qualitatively similar to, but somewhat stronger than, the condition required by FLM.

The proposed equilibrium breaks the infinite-horizon game into sets of K blocks of T periods each. Within each T period block, actions correspond to a fixed allocation rule and the player with private information is given a "quota" for the number of times she can report each type. This technique of breaking an equilibrium into finite-length blocks is used by ET in games with Markov private types; earlier work by Radner (1981, 1985); and Matsushima (2004), Sugaya (2012), and others in the context of games with imperfect private monitoring. For T sufficiently large, the *distribution of private information* within the block—for example, a player's private types—closely matches the true distribution; in my setting, this implies that an agent who reports truthfully is unlikely to exceed an appropriately-chosen quota until the final rounds in the block.[2]

At the end of each block, a count is made of the number of times each report-signal pair occurs, and this count determines play in blocks separated from the present round by jK other blocks, for $j \in \mathbb{N}$. Thus, the messages and signals in block 1 influence the targeted allocation rules in blocks $1+K$, $1+2K$, and so on. In this way, a deviator cannot exploit his private information to deviate only when his expected punishment would be mild. Essentially, each sequence of blocks $1, 1+K, 1+2K, ...$ is treated as a separate game, where the private types in any two blocks are approximately independent. Continuation payoffs are chosen to give linear incentives, so as in Holmstrom and Milgrom (1987) players have constant incentives at every history within each block to play the desired action. As players

[2]For games with imperfect private monitoring, the private information is instead the *private signal*. Loosely speaking, aggregating these signals across many periods increases precisions and ensures that players rarely need to engage in inefficient punishment.

grow patient, the size of each block T increases so that the empirical distribution of types in that block more closely matches the invariant distribution. K also increases to ensure that private information in one block has a minimal impact on players' expected rewards or punishments for actions taken in that block.

In addition to providing a bound on equilibrium payoffs when players are very patient, this proof sheds intuitive light on the interaction between moral hazard and Markov adverse selection. Critically, *delay can be beneficial* if private information decays over time; by postponing a punishment, a player can more precisely target a desired level of utility and mitigate the need for unduly harsh penalties. Hence, an institution that implements punishments after a delay might be more efficient than one that immediately penalizes deviations. Second, adverse selection problems are mitigated in a dynamic setting. In the equilibrium presented in this paper, the informed player loses the option of sending an advantageous message in the future whenever she chooses to send the same message today; her incentive to tell the truth arises from this tension between her current and future selves. Intertemporal trade-offs are imposed by a host of real-world institutions, from Ostrom's (1990) analysis of the graduated sanctions that regulate public goods in many communities to employment contracts that allow a maximum number of sick days each year.

The rest of the paper is organized as follows. In Section 2.2, I introduce the model and state the main result as Theorem 1. The proof begins in Section 2.3, where I prove that approximate truth-telling can be guaranteed when players are sufficiently patient and can *commit* to actions. I turn to the game without commitment in Section 2.4, introduce a notion of enforceability, and use this definition to prove the main result. Section 2.5 extends the efficiency result to some games in which the private information affects everyone's payoff, and Section 2.6 concludes with discussion. Full proofs may be found in online appendices, available at http://economics.mit.edu/grad/dbarron.

2.2 Model and Definitions

Consider an infinite-horizon dynamic game Γ with N players. Player 1 has a private type $\theta_t \in \Theta = \{\theta_1, ..., \theta_{|\Theta|}\}$ with $|\Theta| < \infty$ that evolves according to a Markov process with initial distribution $\nu \in \Delta(\Theta)$ and transition probability $P(\theta_{t+1}|\theta_t)$. To facilitate players adapting to the state of the world, I assume that player 1 can send a public message $m_t \in M_t$ after learning θ_t. After observing this message, players simultaneously choose unobserved actions $a_i \in A_i$ with profile $a = (a_1, ..., a_N) \in A = A_1 \times ... \times A_N$, which influences a public signal $y \in Y$ with $|Y| < \infty$ via the distribution $F(y|a_1, ..., a_N)$. To prevent players from drawing inferences about the action profile from their payoff, I assume that the utility of player i depends only on his private action a_i, the public signal y, and - in the case of player 1 - the type θ: $u_1(a_1, y, \theta)$ and $u_i(a_i, y)$ for $i \neq 1$. In Section 2.5, I consider games in which θ affects every player's utility. Payoffs are weighted by the common discount factor δ, and player i seeks to maximize his average discounted payoff

$$U = \sum_{t=0}^{\infty} \delta^t (1 - \delta) u_{i,t}$$

where $u_{i,t}$ is the expected utility of player i in round t.

In summary, the stage game in each period is:

1. A public randomization device $\xi_t \sim U[0, 1]$ is realized.

2. Type $\theta_t \in \Theta$ is drawn according to transition probability $P(\theta_t|\theta_{t-1})$, with $|\Theta| < \infty$.

3. Player 1 observes θ_t and sends public message $m_t \in M$.

4. After observing m_t, each player i simultaneously chooses an action $a_i \in A_i$.

5. Public signal y is realized according to distribution $F(y|a)$.

6. Payoffs are realized: $u_1(a_1, y, \theta)$ for player 1, $u_i(a_i, y)$ for $i \in \{2, ..., N\}$.

Define the *expected stage-game payoff for player* 1 *and* $i \neq 1$, respectively, as

$$g_1(a, \theta) = E_y[u_1(a_1, y, \theta)|a]$$

and

$$g_i(a) = E_y[u_i(a_i, y)|a].$$

To simplify notation, I sometimes include θ as a argument of player i's utility, even when $i \neq 1$: $g_i(a, \theta)$. Thus, player i's average payoff in the dynamic game is

$$\sum_{t=0}^{\infty} \delta^t (1-\delta) g_i(a_t, \theta_t).$$

Because $A \times \Theta$ is finite, $g_i : A \times \Theta \to \mathbb{R}$ is uniformly bounded; without loss of generality, let $|g_i| \leq 1$ for all i. Player 1 is referred to as "she," while all other players are "he."

Informally, my main result states that as players grow patient (δ increases), some Pareto efficient payoffs can be approximated arbitrarily closely by an equilibrium payoff. More precisely, any point on the interior of the convex hull of Pareto-efficient and min-max payoffs can be approximately attained for patient players. To formalize this result, I first make an assumption on the transition probability $P(\cdot|\cdot)$.

Assumption 5 *$P(\theta_t|\theta_{t-1})$ has a unique stationary distribution $\pi \in \Delta(\Theta)$.*

Assumption 1 holds if $P(\cdot|\cdot)$ is irreducible. Under this condition, the long-run distribution of types is π, regardless of the current information. Using this fact, I consider the set of "long-run" expected payoffs that are attainable under the distribution π.

Definition 8 *The set of* stationary payoffs *is*

$$V = co\left\{v \in \mathbb{R}^N | \exists \alpha : \Theta \to \Delta(A) \ s.t. \ \forall i, \ E_\pi \left[g_i(\alpha(\theta), \theta)\right] = v_i\right\}.$$

Then the set of Pareto efficient stationary payoffs *is*

$$PF = \{v \in V | \forall v' \in V, \ if \ v' \geq v, \ then \ v' = v\}.$$

V is the set of payoffs that can be attained using an allocation rule $\alpha : \Theta \to \Delta(A)$ when the true type θ is drawn according to the distribution π. Dutta (1995) proves that V is the limit (in the Hausdorff metric) of the set of feasible payoffs in the dynamic game as $\delta \to 1$, where this convergence is uniform in the prior distribution ν. I focus on the set of equilibrium payoffs for very patient players, so V is the natural comparison set.

There are several different ways to define min-max payoffs for each player. Following ET, the *stationary min-max profile for player i* specifies a constant, *pure-strategy* action for every player except for i, who is allowed to best-respond to the actions of the other players. This is not the weakest possible notion of min-max, since the action profile can neither depend on θ nor be mixed. It will also be helpful to define the *max-max payoff for player i*, which is the maximum possible payoff i can earn when types are drawn according to π. I will typically refer to these payoffs as the *min-max* and *max-max* payoffs for i.

Definition 9 *The* stationary min-max payoff for player i *is*

$$g_i^m = \min_{a_{-i} \in A_{-i}} \max_{\alpha_i : \Theta \to A_i} E_\pi[g_i(\alpha_i(\theta), a_{-i}, \theta)]$$

with strategy profile $\alpha^{i,m}$. *The* max-max payoff for player i *is*

$$g_i^M = \max_{\alpha : \Theta \to A} E_\pi[g_i(\alpha(\theta), \theta)].$$

with strategy profile $\alpha^{i,M}$.[3] *Let*

$$V^m = \left\{ v \in V | v = E_\pi \left[g(\alpha(\theta), \theta) \right] \text{ where } \alpha \in \{\alpha^{i,m}, \alpha^{i,M}\} \text{for some } i \in \{1, ..., N\} \right\}$$

be the set of payoffs corresponding to min-max and max-max action profiles, and define the individually rational convex hull of Pareto efficient and min-max payoffs *as*

$$V^* = \{ v \in V | v \in co(PF \cup V^m) \; s.t. \; v_i \geq g_i^m, \forall i \}.$$

I prove that every payoff in the interior of V^* can be approximated in an equilibrium as players grow patient. This set need not correspond to the entire set of feasible and individually rational utilities, since it only includes those payoff vectors that can be written as a convex combination of Pareto-efficient, min-max, and max-max payoffs.[4] Note that only

[3]Note that $\alpha^{i,m}(\theta)$ and $\alpha^{i,M}(\theta)$ are independent of θ if $i \neq 1$.

[4]Notice that this payoff set is slightly different from that of Escobar and Toikka (2012), whose result is for the convex hull of Pareto-efficient and *constant* actions. I conjecture that my result can be extended to include the set of all constant actions without incident.

player 1's min-max and max-max action profiles depend on θ, so $\alpha^{i,m}$ and $\alpha^{i,M}$ are constant for $i \neq 1$.

Histories in this game include both the private actions and public signals observed in each round. For player 1, a history also includes the sequence of past types $\{\theta_1, ..., \theta_t\}$.

Definition 10 *Define player i's* private history in round t *as*

$$h_i^t = \{\xi_{t'}, \theta_{t'} 1\{i = 1\}, m_{t'}, a_{i,t'}, y_{t'}\}_{t' \leq t}.$$

The set of all player i's private histories of length T is denoted $\mathcal{H}_i^T$, and $\mathcal{H}_i = \bigcup_{T \in \mathbb{N}} \mathcal{H}_i^T$.

The public history in round t *consists of variables that everyone observes:*

$$h^t = \{\xi_{t'}, m_{t'}, y_{t'}\}_{t' \leq t}.$$

The set of public histories of length T is $\mathcal{H}^T$, with $\mathcal{H} = \bigcup_{T \in \mathbb{N}} \mathcal{H}^T$.

My solution concept is the Perfect Bayesian Equilibrium (PBE). Let $\mu_i = (\mu_i^t)_{t \geq 1}$ be player *i's* sequence of beliefs $\mu_i^t : \mathcal{H}_i^t \to \Delta(\Theta_1)$, and consider an *assessment* (σ, μ) where $\sigma = (\sigma_i)_{i \in \{1,...,N\}}$ is a strategy profile $\sigma_i : \mathcal{H}_i \to A_i$. (σ, μ) is a PBE if it is sequentially rational and μ is computed from σ using Bayes rule whenever possible.

For approximate efficiency to be attainable among patient players, continuation payoffs must be tailored to punish some players while rewarding others. In order to identify which players should be punished following a realization of y, the distribution of signals must satisfy a *pairwise full rank* condition. While not required for the result, I also assume for convenience that $F(y|a)$ has full support.

Assumption 6 *Let $\alpha : \Theta \to A$ be a Pareto efficient, min-max, or max-max strategy profile. Then $\forall \theta \in \Theta$, $F(y|\alpha(\theta))$ has full support on Y and satisfies* pairwise full rank, *defined as* *the following: for any two players $i, j \in \{1, ..., N\}$, define the matrix of signal distributions*

$$\Pi_{ij}^{y}(a) = \begin{bmatrix} F(y_1|a) & \cdots & F(y_{|Y|}|a) \\ F(y_1|a_i = a_i^1, a_{-i}) & \cdots & F(y_{|Y|}|a_i = a_i^1, a_{-i}) \\ \vdots & \cdots & \vdots \\ F(y_1|a_i = a_i^{|A_i|}, a_{-i}) & \cdots & F(y_{|Y|}|a_i = a_i^{|A_i|}, a_{-i}) \\ F(y_1|a_j = a_j^1, a_{-j}) & \cdots & F(y_{|Y|}|a_j = a_j^1, a_{-j}) \\ \vdots & \cdots & \vdots \\ F(y_1|a_j = a_j^{|A_j|}, a_{-j}) & \cdots\cdots & F(y_{|Y|}|a_j = a_j^{|A_j|}, a_{-j}) \end{bmatrix}$$

where $\{a_i^1, ..., a_i^{|A_i|}\} = A_i$. Then F satisfies pairwise full rank *at action profile a if $\forall i, j \in \{1, ..., N\}$, rank$(\Pi_{ij}^{y}(a)) = |A_i| + |A_j| - 1$, which is the maximal rank for this matrix.*

Pairwise full rank encompasses two intuitive conditions. First, given a_{-i}, each of player i's actions lead to a different conditional distribution $F(\cdot|\tilde{a}_i, a_{-i})$. Hence, any unilateral deviation by player i can be detected and punished. Furthermore, deviations by player i are statistically distinguishable from those by player j, so that a transfer scheme between i and j can be designed to deter deviations by either of these players. This second feature is critical, since it implies that agents can be punished using transfers without a budget breaker or "burning money."

Assumption 6 is somewhat stronger than the condition assumed by FLM, who requires only that a single action profile satisfies pairwise full rank. So long as there exists an action profile with pairwise full rank, FLM show that a dense subset of (potentially mixed-strategy) action profiles also satisfy pairwise full rank. Unlike FLM, the construction presented in this paper relies on pure-strategy actions, so pairwise full rank must be assumed to hold at every action profile.

It is sometimes convenient to consider a T-period truncation of dynamic game. This finite-horizon game is used in Section 2.3 to demonstrate that a mechanism without transfers can implement approximate truth-telling.

Definition 11 *Define the* T-round dynamic game *as the truncated game that ends after the* *stage game is played for T rounds, with payoffs*

$$\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t g_i(a_t,\theta_t)$$

I will refer to the game that has not been truncated as the infinite-horizon dynamic game.

Utility is discounted in the T-round game so that the resulting average payoff is comparable for any T. To preview, T-round blocks are the fundamental components in my constructed equilibrium; agents will be induced to conform to an equilibrium in each block using changes in continuation play at the end of each block.

The next sections are dedicated to proving the following efficiency result.

Theorem 1 *Let Assumptions 5 and 6 hold. Suppose $W \subseteq int(V^*)$ is a smooth*[5] *set. Then $\forall \epsilon > 0$, $\exists\ \delta^* < 1$ such that $\forall \delta \geq \delta^*$, $\forall w \in W$, there exists a sequential equilibrium of the infinite-horizon dynamic game that generates payoff v^* satisfying $||w - v^*|| < \epsilon$.*

2.3 The Mechanism

As the first step of the analysis, I consider a mechanism design problem in which players commit to actions as a function of the history of messages and public signals in the T-round dynamic game.

Definition 12 *The* game with commitment *has an identical stage game as the baseline model, but the action in round t is given by $\mathcal{M} : \mathcal{H} \times M \to A$ that implements action $\mathcal{M}(h^{t-1}, m_t)$ at history (h^{t-1}, m_t). Let $\mathcal{A} = \{\mathcal{M} : \mathcal{H} \times M \to A\}$ be the set of all* pure-strategy mechanisms*; an* allocation rule *is a function $\alpha : M \to A$.*

[5] A set $W \subseteq \mathbb{R}^N$ is *smooth* if

1. it is closed and convex;
2. it has a non-empty interior;
3. at each boundary point v, there is a unique tangent hyperplane P_v, which varies continuously with v. Thus, the boundary is a C^2-submanifold of $\mathbb{R}^N$.

The mechanism $\mathcal{M}$ is defined only in terms of pure strategies, and indeed the construction presented in this paper cannot enforce mixed strategies in equilibrium. To be willing to play a mixed strategy, an agent must be indifferent between multiple actions. The type θ_t is informative about continuation play in my construction, so an agent's beliefs about this type affect his continuation surplus and he cannot be made *exactly* indifferent between two strategy profiles.

In this section, I construct a mechanism that targets allocation rule α and gives each player i a payoff close to $E_\pi\left[g_i(\alpha(\theta),\theta)\right]$. Similar to Jackson and Sonnenschein (2007), this mechanism assigns a maximum number of times that each action $\alpha(m)$ can be played in the T period game, thereby creating an intertemporal tradeoff for player 1: she forgoes the option of reporting $m \in M$ in a future period if she reports m today. If the allocation rule $\alpha : M \to A$ is Pareto efficient, player 1's potential gains from misreporting are limited, so her payoff approximates $E_\pi\left[g_1(\alpha(\theta),\theta)\right]$ when T is large.

Definition 13 *In the T-round dynamic game, a mechanism $\mathcal{M}$ is a T-*period quota mech-anism with allocation $\alpha : \Theta \to A$ *if the following hold:*

1. *Let*

$$\hat{Q}(\theta) = \mathit{floor}\{T * \pi(\theta)\}$$

and note that $T - \sum_{\theta\in\Theta}\hat{Q}(\theta) \in [0, M]$. *Define the* quota corresponding to an announce-ment $\theta \in \Theta$ *as*

$$Q(\theta_j) = \begin{cases} \hat{Q}(\theta_j) + 1, & j \leq T - \sum_{\theta\in\Theta}\hat{Q}(\theta) \\ \hat{Q}(\theta_j), & \text{otherwise} \end{cases}.$$

2. *In period t of the T-period block, $\mathcal{M}(h^{t-1}, m_t) = \alpha(m_t)$ if $\sum_{t'\leq t} 1\{m_{t'} = m_t\} \leq Q(m_t)$. Otherwise, $\mathcal{M}(h^{t-1}, m_t) = \alpha(\theta_{\underline{j}})$, where*

$$\underline{j} = \arg\min_j \left\{ j \,|\, \sum_{t'<t} 1\{m_{t'} = \theta_j\} < Q(\theta_j) \right\}.$$

A few features of this mechanism will prove useful in the rest of the paper. First, the sum of the quotas $Q(\theta)$ add up to exactly T: $\sum_{\theta\in\Theta} Q(\theta) = T$. Therefore, every player knows the number of times each action $\alpha(\theta)$ is played; player $i \neq 1$'s payoff approximates the weighted average $\frac{1}{T}\sum_{\theta\in\Theta} Q(\theta)g_i(\alpha(\theta))$ as $\delta \to 1$. Second, the quota is determined by the invariant distribution π and does not depend on the initial distribution ν. I will show that this mechanism approximately implements the desired allocation rule for *any* prior ν. Finally, if player 1 exceeds the quota for a given θ, then that message is treated as if she reported a type with quota remaining. Thus player 1 has no incentive to report any type for which she has already exceeded her quota.

This mechanism differs from the one used by ET, which tracks the empirical distribution of reported *transitions* between types and matches that probability to $\{P(\theta_t|\theta_{t-1})\}_{\theta_t,\theta_{t-1}\in\Theta}$. If multiple players have private information, players' payoffs might differ substantially from the utilities targeted by the mechanism if the joint distribution of reports were to differ dramatically from the type distribution. For example, a subset of players could coordinate their reports to collude and increase their collective payoffs, or tailor their messages to previous reports in the block in order to harm other players. Requiring that reported transitions match $P(\theta_t|\theta_{t-1})$ ensures that different players' reports are approximately independent. In my setting, only player 1 reports a type in each period and so these issues are not a concern. Instead, the game with commitment is a decision problem because only player 1 acts; I use features of decision problems to prove the results that follow.

The T-period quota mechanism only approximates payoffs from the desired allocation rule if players are patient and the time horizon is long. Thus, I define implementability as an asymptotic property when $T \to \infty$ and $\delta \to 1$.

Definition 14 *A class of mechanisms in the T-period dynamic game* implements a set of allocations $\mathcal{A} = \{\alpha : \Theta \to A\}$ *if* $\forall\epsilon > 0$, $\exists T^* < \infty$ *such that* $\forall T \geq T^*$, $\exists\, \delta^* \in (0,1)$ *such that* $\forall\, \delta \geq \delta^*$, $\forall\alpha \in \mathcal{A}$, *and* $\forall$ *prior* ν, $\exists$ *a mechanism in that class with equilibrium* σ *such that*

$$\sum_{t=0}^{T-1}(1-\delta)\delta^t E_\sigma\left[g(\alpha(m_t),\theta_t)\right] \in B\left(E_\pi\left[g(\alpha(\theta),\theta)\right],\epsilon\right).$$

Proposition 8 *Let $\mathcal{A}$ consist of all Pareto efficient allocation rules. Then $\mathcal{A}$ can be implemented using T-Period Quota Mechanisms.*

Proof: See online appendix.

To prove Proposition 8, consider player 1's payoff if she were to truthfully report her type in each period. If T is large, then the empirical distribution of reports closely matches π; as a result, she is unlikely to exceed her quotas until the very end of the T periods. Truthful reporting is a feasible strategy, so player 1's utility is bounded from below by $E_\pi\left[g(\alpha(\theta),\theta)\right]-\epsilon$ for T and δ sufficiently large. Player $i \neq 1$'s payoff approximates $\sum_{\theta\in\Theta}\pi(\theta)g_i(\alpha(\theta)) = E_\pi\left[g_i(\alpha(\theta))\right]$ if T and δ are large, regardless of player 1's strategy. Because α is Pareto efficient, player 1's utility is also bounded above by $E_\pi\left[g(\alpha(\theta),\theta)\right]+\epsilon$; otherwise, there would exist some allocation rule that Pareto dominates $\alpha(\cdot)$. These facts together prove the Proposition.

When α is a min-max or max-max payoff, I implement it using a different class of mechanisms.

Definition 15 *The T-Period Unrestricted Mechanism implementing $\alpha(\theta)$ satisfies:*

1. *In each round $t \in \{1,...,T\}$, player 1 announces $m_t \in \Theta$.*

2. *The implemented allocation rule is $\alpha(m_t)$.*

In the equilibrium constructed in Section 2.4, i's min-max or max-max allocation rules are played exactly when intertemporal incentives cannot be used. When $i=1$, the player must find it optimal to truthfully reveal her type in a min-max or max-max allocation rule to provide viable incentives for the *other* players to play the desired actions. Truth-telling is an equilibrium in an unrestricted mechanism but not necessarily in a quota mechanism, particularly at histories where player 1 is close to exceeding her quota.

Proposition 9 *Let $\mathcal{A} = \{\alpha | \exists i \in \{1,...,N\}$ s.t. $\alpha \in \{\alpha^{m,i},\alpha^{M,i}\}\}$. Then $\mathcal{A}$ can be implemented using an unrestricted mechanism, and truth-telling $m_t = \theta_t$, $\forall h^t \in \mathcal{H}$, is an equilibrium of the unrestricted mechanism.*

Proof: See online appendix.

By definition, truth-telling maximizes player 1's myopic payoff in a min-max or max-max allocation rule, so player 1 is willing to truthfully reveal her type in each period. As a result,

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T}E\left[g_i(\alpha(\theta_t),\theta_t)\right]=E_\pi\left[g_i(\alpha(\theta),\theta)\right]$$

$\forall i$, which proves Proposition 9.

Next, I formalize the intuition that players' beliefs about payoffs in the far future converge. To do so, I introduce a class of *κ-delayed mechanisms* in the $(\kappa+T)$-round dynamic game.

Definition 16 *Let $\alpha:\Theta\to A$ be a Pareto efficient, min-max, or max-max payoff. A mechanism in the $(\kappa+T)$-round dynamic game is* a κ-delayed mechanism with allocation rule α *if:*

1. *For $t\in\{1,...,\kappa\}$, some κ-period mechanism $\hat{\mathcal{M}}$ is played.*
2. *For $t\in\{\kappa+1,...,T\}$, a T-period quota or unrestricted mechanism that implements α is played.*

The mechanism in the first κ rounds of a κ-delayed mechanism is irrelevant; the purpose of these rounds is to ensure that players have similar *ex ante* beliefs about the distribution of types in the final T periods. Because types are approximately distributed according to π in the distant future, (normalized) payoffs in the final T rounds converge to some fixed value as $\kappa\to\infty$ regardless of the prior ν.

For a *fixed* T, the set of payoffs that can be approximated by κ-delayed mechanisms are called *invariant payoffs.* I focus on these invariant payoffs for most of Section 2.4, adapting the proof techniques from FLM to apply to them.

Definition 17 *Fix $T<\infty$ and allocation rule $\alpha:\Theta\to A$, and consider any κ-delayed mechanism with allocation rule α. A vector $v^T(\alpha)\in\mathbb{R}^N$ is an* invariant payoff for α *if $\forall\zeta>0$, $\exists\kappa^*,\delta^*$ such that $\forall\kappa\geq\kappa^*$, $\delta\geq\delta^*$, and $\forall$ prior $\nu\in\Delta(\Theta)$, there exists a PBE σ^* in rounds $\{\kappa,....,\kappa+T-1\}$ such that $E_{\sigma^*}\left[\sum_{t=0}^{T-1}\delta^t(1-\delta)g(\alpha(m_{\kappa+t}),\theta_{\kappa+t})\right]\in B(v^T(\alpha),\zeta)$.*

The set of invariant payoffs *is*

$$V^T = co\{v^T(\alpha)|\alpha \text{ is either Pareto-efficient, or } \alpha \in \{\alpha^{i,m}, \alpha^{i,M}\} \text{for some } i\}.$$

The individually-rational invariant payoffs *are*

$$V^{T*} = \left\{v \in V^T | \forall i, v_i \geq v^T(\alpha^{i,m})\right\}.$$

For a fixed T, the set of invariant payoffs may differ from the set of interest V^*. However, Propositions 8 and 9 show that payoffs in a T-period quota or unrestricted mechanism approximate the corresponding stationary payoffs as $T \to \infty$, and so $V^{T*} \to V^*$ (in the Hausdorff sense) in this limit. I will consider a large but fixed T for most of Section 2.4 and take this limit as the final step in the proof.

Next, I argue that some $v^T(\alpha)$ exists for any Pareto-efficient, min-max, or max-max allocation rule.

Proposition 10 *Fix T and a Pareto-efficient, min-max, or max-max action $\alpha : \Theta \to A$. Then there exists at least one invariant payoff $v^T(\alpha)$. Moreover, if α is a min-max or max-max strategy, then player 1 reports truthfully in the final T rounds of the κ-delayed mechanism in the equilibrium that approximates $v^T(\alpha)$.*

Proof: See online appendix.

The game with commitment is a decision problem, so player 1's optimal payoff in continuous in both the discount factor δ and the prior ν, which in turn ensures that an invariant payoff $v_1^T(\alpha)$ exists for every Pareto efficient, min-max, and max-max allocation rule. For player $i \neq 1$, the action profile $\alpha(\theta)$ will be implemented exactly $Q(\theta)$ times in the quota mechanism, so player i's payoff continuously approximates $\frac{1}{T}\sum_{t=1}^{T} Q(\theta)g_i(\alpha(\theta)) = v_i^T(\alpha)$ as $\delta \to 1$, regardless of player 1's strategy. In an unrestricted mechanism, player i does not precisely know the number of times player 1 will report each type. However, truthful reporting is an equilibrium strategy, so the distribution of reports approaches π as κ increases; hence $v_i^T(\alpha)$ exists for any α implemented by an unrestricted mechanism, as well.

The fact that private information is one-sided is critical for this result. The invariant payoff $v^T(\alpha)$ is defined as the limit of equilibrium payoffs as $\kappa \to \infty$ for *any* initial distribution of types ν. If only player 1 acts in the mechanism, then her strategy in rounds $\{\kappa, ..., \kappa + T - 1\}$ depends only on the realization of θ_κ. In contrast, if multiple players were to act in the mechanism, then equilibrium outcomes would depend on beliefs about θ_κ. The set of equilibrium payoffs is not lower hemicontinuous in prior beliefs, so payoffs do not necessarily converge to a single value for *every* prior ν as $\kappa \to \infty$. This problem—that equilibrium payoffs are not lower hemicontinuous in prior beliefs—is a central difficulty to extending this result to multi-sided private information.

Next, I extend the intuition from Proposition 10 to the infinite-horizon dynamic game by introducing a mechanism that implements a sequence of T-period mechanisms.

Definition 18 *Consider the infinite-horizon dynamic game with commitment. Fix* $T, K \in \mathbb{N}$. $\forall\ j \in \mathbb{N}$, $k \leq K$, block (k, j) *consists of periods*

$$T^{(k,j)} = \{(K(j-1)+k-1)T, (K(j-1)+k-1)T+1, ...(K(j-1)+k)T-1\}$$

with (k, j)-block history

$$h^{(k,j)} = \{\xi_t, m_t, y_t\}_{t\in\{(K(j-1)+k-1)T,(K(j-1)+k-1)T+1,...(K(j-1)+k)T-1\}}\,.$$

A mechanism $\mathcal{M} : \mathcal{H} \to \mathcal{A}$ *is* (T, K)-Recurrent *if:*

1. $\forall j \in \mathbb{N}, k \leq K$, *block* (k, j) *is either a* T*-period quota mechanism implementing Pareto efficient allocation rule* $\alpha^{(k,j)}$, *or an unrestricted mechanism implementing min-max or max-max* $\alpha^{(k,j)}$.

2. *For any two histories* h^t *and* $\tilde{h}^t$ *with* t *a round in block* (k, j), *if* $h^{(k,j')} = \tilde{h}^{(k,j')}\ \forall j' < j$, *then* $\alpha^{(k,j)} = \tilde{\alpha}^{(k,j)}$.

3. *Messages don't affect outcomes:* $\forall$ *strategies* σ, $\tilde{\sigma}$, $\forall k \leq K$, $\forall j \in \mathbb{N}$,

$$Prob_\sigma\left\{\alpha^{(k,j)} = \alpha | \alpha^{(k,1)}, \alpha^{(k,2)}, ..., \alpha^{(k,j-1)}\right\} = Prob_{\tilde{\sigma}}\left\{\alpha^{(k,j)} = \alpha | \alpha^{(k,1)}, \alpha^{(k,2)}, ..., \alpha^{(k,j-1)}\right\}. \quad (2.1)$$

The (T, K)-Recurrent mechanism—which forms a crucial part of this analysis—satisfies the three important properties given in the definition. Property 1 ensures that $\mathcal{M}$ implements a distinct T-period mechanism in each block of the infinite-horizon game, while 2 guarantees that the public randomization devices, messages, and signals (ξ_t, m_t, y_t) observed in round t of block (k, j) only impact future allocation rules $\alpha^{(k,j')}$ in blocks with the same index k. Property 3 guarantees that for a fixed k, the distribution over possible allocation rules $\{\alpha^{(k,j)}\}_{j\in\mathbb{N}}$ evolves independently of the sequence of messages sent by player 1. Together with property 2, this requirement ensures that player 1 cannot manipulate the future allocation rule through her reports, so that she chooses a reporting strategy to maximize her payoff *within* each T-period block. By Proposition 10, one such optimal strategy yields payoffs that approximate the invariant payoff as $\kappa \to \infty$.

Intuitively, a (T, K)-Recurrent mechanism splits the infinite-horizon game into sets of K blocks, each of which consists of T periods. A quota or unrestricted mechanism is played in each T-period block, and the public signals from these rounds determine the mechanism that is played in the blocks that are separated by $K - 1$ blocks from one another. In effect, this construction creates a set of K "auxiliary games" to the baseline dynamic game; each round of an auxiliary game consists of T periods of the baseline game, with a discount rate of $\delta^{T(K-1)}$ between "rounds."

Increasing K in a (T, K)-Recurrent mechanism mimics an increase in κ in a κ-delayed mechanism, delaying the impact that any round t has on payoffs in the continuation game. The next corollary proves that players' expectations in block (k, j) of the payoff in future blocks (k, j'), $j' > j$, approximates the invariant payoff $v^T(\alpha^{(k,j')})$ as $K \to \infty$.

Corollary 2 *Consider a (T, K)-Recurrent mechanism. Then $\forall\ \zeta > 0$, $\exists$ a $K^* < \infty$ and $\delta^* < 1$ such that $\forall\ K \geq K^*$, $\delta \geq \delta^*$, $\exists$ an equilibrium σ such that $\forall$ history h^t with t in block (k, j), and any $j' > j$, expected payoffs in block (k, j') conditional on h^t satisfy*

$$\frac{1}{1-\delta^T} \sum_{t' \in T^{(k,j)}} E_\sigma\left[\delta^{t'-(K(j'-1)+K-1)} u_{t'} | h^t\right] \in B\left(E_\sigma\left[v^T(\alpha^{(k,j')}) | h^t\right], \zeta\right). \qquad (2.2)$$

Proof: See online appendix.

As in Proposition 10, this result relies crucially on the fact that only player 1 has private information. Given the structure of a (T, K) recurrent equilibrium, player 1 considers only the public history and her current type when calculating her optimal action. Therefore, continuation payoffs in block (k, j) depend only on $\alpha^{(k,j)}$ and player 1's type at the *beginning* of that block. In particular, so long as reports in block (k, j) do not affect the expected allocation rule in any future block, player 1 is willing to follow an optimal strategy for the T-period mechanism that implements $\alpha^{(k,j)}$. Properties (2) and (3) of Definition 18 imply that player 1 cannot affect future allocation rule by her report, so she plays an optimal T-period reporting strategy in each block.

By Corollary 2, continuation payoffs in the far future of a (T, K)-Recurrent mechanism are close to invariant payoffs. The allocation rule $\alpha^{(k,j')}$ is independent of messages in block (k, j) for $j' > j$, but it *can* depend on the public signals in block (k, j). In the next section, I consider the game without commitment, fix T, and construct a relationship between $\alpha^{(k,j')}$ and the signals observed in block (k, j) to punish deviations from the desired allocation rule.

2.4 Equilibrium

The goal of this section is to build a conceptual apparatus that mimics the notion of enforceability introduced by Abreu, Pearce, and Stachetti (1990). With this construction in hand, I modify the basic proof technique of FLM to prove that points in the interior of V^* can be approximated by an equilibrium for sufficiently patient players.

Consider a $(T, 2)$-Recurrent mechanism. Intuitively, the odd $(1, j)$ and even $(2, j)$ blocks form two distinct games: outcomes observed in block $(1, 1)$ only affect the targeted allocation rule for blocks $(1, 2)$, $(1, 3)$, and so on, and similarly for blocks $(2, j)$. Players' *beliefs* in blocks $(1, j)$ and $(2, j)$ are still related, so I construct an equilibrium that deters deviations from the equilibrium action for any such beliefs. The key observation is if an action profile a be deterred with *uniformly strict* bonus schemes in V^T, then these bonus schemes can be approximated by continuation payoffs in the far future and so a can be enforced in equilibrium.

Before considering enforceability in this class of games, I define a few features that continuation payoffs in an enforceable profile must satisfy.

Definition 19 *An allocation rule* $\alpha : [0,1]\times\Theta \to A$ *and a finite collection* $\{w(\xi, y, \theta)\}_{\xi\in[0,1],y\in Y,\theta\in\Theta} \subset W$ ***satisfy condition*** 0 ***on*** W *if* $\forall \xi \in [0,1]$, *there exists weights* $\lambda^j(\xi, y, \theta)$ *for a finite set of points* $\{w^1, ..., w^{L(\xi)}\} \subseteq W$ *such that* $w(\xi, y, \theta) = \sum_{l=1}^{L(\xi)} \lambda^l(\xi, y, \theta) w^j$ *and*

$$\sum_{y\in Y} \lambda^l(\xi, y, \theta) F(y|\alpha(\xi, \theta)) = \sum_{y\in Y} \lambda^l(\xi, y, \theta') F(y|\alpha(\xi, \theta')), \qquad \forall \theta, \theta' \in \Theta, \forall l \leq L(\xi). \quad (2.3)$$

α *and* $\{w(y, \theta)\}$ ***satisfy condition*** i ***on*** W *if* $\forall \xi \in [0,1]$,

$$\begin{matrix} \sum_{y\in Y} \lambda^l(\xi, y, \theta) F(y|a_i, \alpha_{-i}(\xi, \theta)) = \\ \sum_{y\in Y} \lambda^l(\xi, y, \theta') F(y|a_i', \alpha_{-i}(\xi, \theta')) \end{matrix}, \qquad \forall \theta, \theta' \in \Theta, \forall a_i, a_i' \in A_i, \forall l \leq L(\xi).$$

Conditions 0 and i play an essential role in my definition of enforceability but have no clear analogue in repeated games. Intuitively, the collection $\{w(\xi, y, \theta)\}_{\xi\in[0,1],y\in Y,\theta\in\Theta}$ are a set of bonus payments for the players; in equilibrium, these bonus payments will be mimicked by continuation play. Condition 0 imposes that conditional on players following the allocation rule α, every report by player 1 induces the *same distribution over* ***extremal*** *bonus payments.* If Condition 0 did not hold, then requirement (3) of Definition 18 would not be satisfied and so player 1 might have an incentive to alter her report in the hopes of securing better continuation play. Likewise, Condition i ensures that player i cannot change the distribution of bonus payments by unilaterally deviating from his action. This assumption deters player i from playing a myopically suboptimal action when he is being min-maxed or max-maxed in order to induce favorable continuation play.

Conditions 0 and i require that the *distribution* over continuation play is unaffected by changes in player 1's report or player i's action, respectively. Continuation payoffs are a function of players' beliefs about the distribution of future types, which depend on player 1's private information and actions. These beliefs will *not* be precisely pinned down in equilibrium; as a result, if player 1 could influence the *distribution* of continuation payoffs by altering her message, she might do so even if both distributions lead to the same expected invariant surplus.

With Definition 19 in hand, I define decomposable payoff vectors and enforceable allocation rules.

Definition 20 *We say that a payoff* v *is* (T,ζ,W,δ)*-decomposable if* $\exists$ *an implementable allocation rules* $\alpha:[0,1]\times\Theta\to A$ *and "bonus payments"* $\{w(\xi,y,\theta)\}_{\xi\in[0,1],y\in Y,\theta\in\Theta}\subseteq W$ *such that the following three properties hold:*

1. α *and* $\{w(\xi,y,\theta)\}$ ***satisfy Condition*** 0 ***on*** W*, with* $E_y[w(\xi,y,\theta)]=\bar{w}(\xi)$, $\forall\theta\in\Theta$.

2. *The* ***adding up constraint*** *holds:*

$$v=E_\xi[(1-\delta)v_I(\alpha(\xi,\cdot))+\delta\bar{w}(\xi)]. \tag{2.4}$$

3. *The* ***enforceability constraint*** *holds:* $\forall i$, $\forall\xi\in[0,1]$, *EITHER*

$$\begin{gathered}(1-\delta)g_i(\alpha(\xi,m),\theta)+\tfrac{\delta}{T}\bar{w}(\xi)-\zeta\geq\\ \max_{a_i\in A_i}\{(1-\delta)g_i(a_i,\alpha_{-i}(\xi,m),\theta)\}+\max_{a_i\in A_i\setminus\alpha_i(\theta)}\left\{\tfrac{\delta}{T}E_y[w_i(\xi,y,m)\mid a_i,\alpha_{-i}(m)]\right\}\end{gathered} \tag{2.5}$$

holds, OR Condition i *on* W *is satisfied and*

$$g_i(\alpha(\xi,\theta),\theta)\geq g_i(a_i,\alpha_{-i}(\xi,m),\theta),\ \forall a_i\in A_i,\forall m\in\Theta. \tag{2.6}$$

If there exist bonus payments such that (3) holds, then $\alpha:[0,1]\times\Theta\to A$ *is* (T,ζ,W,δ)*-strictly enforceable. If every* $v\in W$ *is* (T,ζ,W,δ)*-decomposable, then* W *is* (T,ζ,δ)-self decomposable.

This notion of enforceability is defined in terms of actions within a T*-period block* and differs from the repeated-games definition in several ways. A type-dependent allocation rule $\alpha:\Theta\to A$ is enforceable if one of two conditions holds. First, player i might be *uniformly strictly* motivated to play $\alpha_i(m)$ for *any* beliefs about θ, given that continuation play generates surplus $\{w_i(\xi,y,m)\}$.[6] Second, if i is not given strict intertemporal incentives, then he must play a myopic best-response to the other players. In either case, Condition 0 ensures that the second statement of Definition 18 is satisfied if players conform to the allocation rule $\alpha^{(k,j)}$. If player i is playing a myopic best-response, then Condition i guarantees that

[6] Mixed strategies are ruled out because a player typically cannot have uniformly strict incentives to play multiple actions.

she has no incentive to change her action in order to induce a favorable continuation payoff by requiring that player i cannot affect continuation play when she is min-maxed or max-maxed. The allocation rule $\alpha^{(k,j)}$ must be enforced for *every* round of a T-period block. I use a linear incentive scheme to do so, which is why the continuation payoffs in (2.5) are multiplied by $\frac{1}{T}$.

Following FLM, I use continuation payoffs $\{w(\xi, y, \theta)\}$ that lie on a translate of the hyperplane tangent to a boundary point $v \in W$ to support the payoff v.

Definition 21 *A (translate of a) hyperplane $P \subset \mathbb{R}^N$ is an i^{th}* coordinate *hyperplane if it is parallel to the i^{th} coordinate axis - that is, if $P = \{x \in \mathbb{R}^N | x \cdot \beta = c\}$ for some $c \in \mathbb{R}$, $\beta \in \mathbb{R}^N$ such that $\beta_i \neq 0$ and $\forall j \neq i$, $\beta_j = 0$. P is a* normal *hyperplane if it is not a coordinate hyperplane.*

A bonus scheme on hyperplane P is a linear contract. Any continuation payoffs $\{w(\xi, y, \theta)\}$ that lie on the i^{th} coordinate hyperplane satisfy $w_i(\xi, y, \theta) = \frac{c}{\beta_i}$, so player i cannot be strictly incentivized to take the equilibrium action. Thus, such an allocation can be enforced using payoffs on such a hyperplane only if (1) α is a min-max or max-max payoff for player i, *and* (2) continuation payoffs $\{w(\xi, y, \theta)\}$ satisfy Condition i. On a normal hyperplane, strict incentives can be provided to each player, so the primary problem is to find continuation payoffs that uniformly deter every player's deviations.

Lemma 5 shows that continuation payoffs that satisfy these properties can be found so long as $F(y|\alpha(\theta))$ satisfies pairwise full rank for every θ.

Lemma 5 *Suppose the signals at action profile α satisfy pairwise full rank and* either*:*

1. *P is a normal hyperplane and α is a Pareto-efficient, min-max, or max-max allocation rule;*

2. *$\exists i \in \{1, ..., N\}$ such that P is parallel to the i^{th} coordinate axis and α is a min-max or max-max allocation for player i.*

Then $\forall$ $\delta \in (0,1)$, $\zeta > 0$, $T > 0$, α is (T, ζ, P', δ) strictly enforceable, where P' is any translate of P.

Proof: See online appendix.

This lemma depends critically on F satisfying pairwise full rank at *every* Pareto-efficient, min-max, and max-max allocation rule, but this assumption plays different roles depending on whether P is a normal or coordinate hyperplane. If P is normal, then pairwise full rank ensures that different players' deviations are statistically distinguishable from one another so that there exists a budget-balanced bonus scheme that induces the desired actions. If P is parallel to the i^{th} coordinate hyperplane, then pairwise full rank guarantees that continuation payoffs exist that satisfy Condition i. These bonus payments deter deviations by players $j \neq i$ while remaining unaffected (in expected distribution) by player i's action.

Lemma 6 proves several key properties of enforceability and decomposability.

Lemma 6 *1. If a payoff v is (T,ζ,W,δ)-strictly decomposable, then v is also (T,ζ',W,δ)-strictly decomposable for any $\zeta' < \zeta$.*

2. Let P be any hyperplane (normal or coordinate), and suppose that α is (T,ζ,P,δ)-strictly enforceable. Then:

(a) For every $\delta' > \delta$, there exists a continuous function $\hat{\zeta}(\delta) \leq \zeta$ such that α is $(T,\hat{\zeta}(\delta),P,\delta')$-strictly enforceable. Moreover, $\exists$ $\{w(y,\theta,\delta,\zeta)\}$ that enforce α and $\exists \kappa > 0$ such that

$$||w(y,\theta,\delta',\hat{\zeta}(\delta'))|| = \kappa\frac{1-\delta'}{\delta'}.$$

where the extremal payoffs also satisfy $||w|| = \kappa\frac{1-\delta'}{\delta'}$.

(b) α is (T,ζ,P',δ)-strictly enforceable for any translate P' of P.

Proof: See online appendix

Properties 1 and 2b follow immediately from the definition of enforceability. To show Property 2a, I specify a function $\hat{\zeta}(\delta) > 0$ that decreases linearly in δ; plugging this expression into (2.5) and rearranging yields the desired property.

Next, I use the properties of enforceability and invariant payoffs that I have already proven to show that every payoff in a self-decomposable set W is "close" to an equilibrium payoff when the block length T is large (but finite).

Proposition 11 *$\forall \epsilon > 0$, fix $T > 0$ so $\exists \delta_M < 1$ such that if $\delta \geq \delta_M$ and α is a Pareto-efficient, min-max, or max-max allocation rule, then the T-period quota or unrestricted mechanism implementing α has payoffs $U \in B\left(E_\pi\left[g_i(\alpha(\theta),\theta)\right], \frac{\epsilon}{2}\right)$. For a closed, convex, bounded set W, suppose $\exists\ \hat{\delta} < 1$ such that $\forall \delta \geq \hat{\delta}$, $\exists \zeta(\delta) > 0$ such that W is $(T, \zeta(\delta), \delta)$-strictly self-decomposable. Then $\exists \delta^*$ such that $\forall \delta \geq \delta^*$, $w \in W$, $\exists$ an equilibrium of the infinite-horizon game with payoff $v \in B(w, \epsilon)$.*

Proof: See online appendix.

The proof of Proposition 11 constitutes an important building block in the construction of my main result, and relies critically on the properties of (T, K)-Recurrent mechanisms. Suppose that a set W is self-decomposable and let $w \in W$. By (2.4), there exists an allocation rule α^1 and continuation payoffs w^1 that together with the appropriate discounting generate v. Payoff w^1 can be similarly deconstructed into an allocation rule α^2 and *further* continuation payoffs w^2; continuing this process leads to a sequence of allocation rules that generate v. If players follow these allocation rules and are sufficiently patient, then the resulting payoff v is close to w by the definition of invariant payoffs.

A (T, K)-Recurrent mechanism can be used to enforce these allocation rules. The allocation rule α^1 is played in block $(k, 1)$ $\forall k \leq K$. In block (k, j), continuation play is determined by choosing one round of $(k, j-1)$ at random and targeting the continuation value specified by the report and public signal from that round. If $T = K = 2$, then the allocation rule in rounds $\{5,6\}$, $\{9,10\}$, and so forth targets continuation value $w(\xi_1, y_1, m_1)$ with probability $\frac{1}{2}$ and otherwise targets $w(\xi_1, y_2, m_2)$. By Condition 0, the expected continuation value $\bar{w}(\xi_1)$ can be implemented using the same convex combination of points in W, regardless of the report m; as a result, continuation play can be made independent of player 1's report, so this mechanism is (T, K)-Recurrent. As K grows, the realization of a public signal and the resulting penalties and rewards are separated by many rounds, so these incentives approximate the invariant payoff. Because rewards and penalties are realized in the more distant future as K increases, the *effective between-block discount factor* δ^{TK} decreases; to enforce the equilibrium actions, $K \to \infty$ as $\delta \to 1$.

The rest of the argument presented here is modified from FLM, Lemma 4.2 and Theorem 4.1. Together, the next set of lemmas demonstrate that a set W which satisfies some regularity properties is (T, ζ, δ)-self decomposable if every payoff on the boundary of W is decomposable on a translate of the hyperplane tangent to that point. Combining this result with Lemma 5 yields the desired efficiency result.

The first step in this argument is to replace the notion of decomposability with a local version for compact sets W.

Definition 22 *A set W is T locally self-decomposable if for each $w \in W$, there is a $\delta^*_w < 1$ and an open set U containing v such that, for every $\delta \geq \delta^*_w$, there exists a $\zeta^*_w > 0$ such that, for every $\zeta \leq \zeta^*_w$, $U \cap W$ is (T, ζ, W, δ)-decomposable.*

So long as W is compact, then local self-decomposability implies (T, ζ, δ)-decomposability for δ sufficiently close to 1.

Lemma 7 *If a subset $W \subseteq \mathbb{R}^N$ is compact, convex, and T-strict locally self-decomposable, then there exists a $\delta^* < 1$ such that for every $\delta \geq \delta^*$, there exists an $\zeta^* > 0$ such that for every $\zeta \leq \zeta^*$, W is (T, ζ, δ)-strictly self-decomposable.*

Proof: See online appendix.

The collection of open sets in the definition of T locally self-decomposable together form an open cover of W. If W is compact, then this cover has a finite subcover. If δ is larger than the maximal δ^*_w for all neighborhoods that together form that subcover and ζ is smaller than the minimal ζ^*_w in the subcover, then all points in W can be enforced and so W is (T, ζ, δ)-self decomposable.

Finally, I argue that a set W is locally decomposable if it is *decomposable on tangent hyperplanes* - that is, if every point $w \in \text{bd}(W)$ can be decomposed using some Pareto-efficient, min-max, or max-max allocation rule $\alpha : \Theta \to A$ that is separated from W by the hyperplane tangent to W at w, and payoffs on a translate of that hyperplane.

Definition 23 *A smooth subset $W \subseteq V^*$ is T-*strictly decomposable on tangent hyperplanes *if, for every point v on the boundary of W, there exists a profile $\alpha(\cdot)$ such that:*

1. *There exists an element $g(\alpha) \in V^T$ that is separated from W by the unique $(n-1)$-dimensional hyperplane P_v that is tangent to W at v, and*

2. *There exist $\delta < 1$ and $\zeta > 0$ such that α is (T, ζ, P_v', δ)-enforceable for some translate P_v' of P_v.*

So long as W is smooth, Lemma 8 shows that decomposability on tangent hyperplanes implies local decomposability.

Lemma 8 *Let W be a smooth set. Suppose W is T*-strictly decomposable on tangent hyperplanes. *Then W is locally T-strictly self-decomposable.*

Proof: See online appendix.

For (2.4) to hold, the translate of the tangent hyperplane used to enforce α must cut through the interior of W. Because the boundary of W is a C^2-submanifold, the boundary of W decreases at a square-root rate about w. By Lemma 6, the bonus scheme on the hyperplane tangent to w used to enforce α decreases linearly. Thus, this bonus scheme lies in the *interior* of W for δ sufficiently large. Similarly, a neighborhood about w can be supported by slightly varying the continuation payoffs while keeping α constant.

A few distinctions between this argument and FLM are worth noting. First, *only* min-max, max-max, and Pareto-efficient allocation rules are played in the constructed equilibrium because other allocation rules may be impossible to implement with a quota or unrestricted mechanism. This fact explains why V^* is not the full set of individually-rational payoffs: Theorem 1 holds because V^* contains only convex combinations of implementable payoffs. For a similar reason, mixed-strategy min-max payoffs cannot be sustained in the class of equilibria constructed here, even though they might generate a strictly lower payoff and hence serve as a better punishment for the min-maxed player. Continuation play at the end of block (k, j) randomizes among a small set of extreme continuation payoffs, which ensures that every sequence of messages leads to the same distribution over continuation payoffs and in turn guarantees that player 1's reporting strategy is not influenced by *between*-block incentive concerns. For this reason, the public signal ξ_t observed in each round is absolutely critical in the proof.

To complete the proof, I argue that a smooth set $W \subseteq \text{int}(V^*)$ also satisfies $W \subseteq \text{int}(V^{T*})$ if T is sufficiently large. If $W \subseteq V^{T*}$, then the arguments from the previous section imply that W is self-decomposable for δ sufficiently large, which ensures that an equilibrium payoff approximates every $w \in W$. These steps together yield Theorem 1.

Lemma 9 *Let $W \subseteq V^{T*}$ be a smooth set. Then there exists a $\delta^* < 1$ such that for every $\delta \geq \delta^*$, $\exists\ \zeta > 0$ such that W is (T, ζ, δ)-decomposable.*

Proof: See online appendix.

If $W \subseteq V^{T*}$, then any hyperplane tangent to a point $w \in \text{bd}(W)$ separates the set W from *some* point $v \in \text{bd}(V^{T*})$. If this hyperplane is normal, then v corresponds to some convex combination of Pareto-efficient, min-max, and max-max payoffs. If the hyperplane is instead parallel to i^{th} coordinate axis, then it separates W from either the min-max or max-max invariant payoffs for player i. By Lemma 5, these payoffs are supportable on the i^{th} coordinate hyperplane. Hence, W is (T, ζ, δ)-decomposable.

Lemma 10 *Let $W \subseteq V^*$ be a smooth set, and suppose $\exists\ \epsilon > 0$ such that*

$$\bigcup_{w \in W} B(w, \epsilon) \subseteq V^*.$$

Then $\exists\ T^* < \infty$ such that $\forall\ T \geq T^*$, $\exists\ \delta^* < 1$ such that $\forall \delta \geq \delta^*$,

$$W \subseteq V^{T*}.$$

Proof: See online appendix.

As $T \to \infty$, Propositions 8 and 9 imply that V^{T*} closely approximates V^*, in the sense that every pure-strategy, min-max, and max-max invariant payoff $v^T(\alpha)$ satisfies

$$\lim_{T \to \infty} v^T(\alpha) = E_\pi \left[g(\alpha(\theta), \theta) \right] \in V^*.$$

Hence, for any $W \subseteq \text{int}(V^*)$, $W \subseteq \text{int}(V^{T*})$ for sufficiently large T. Combining the results from the last two sections yields Theorem 1.

2.5 θ Affects Every Player's Payoff

So long as Proposition 10 and Corollary 2 are satisfied, the proofs in Section 2.4 can be used in other settings. To demonstrate this point, I adapt RSV to extend Theorem 1 to some games in which θ affects *every* player's payoff. Formally, suppose θ affects everyone's payoffs but is known only to player 1.

Definition 24 A game with common state of the world *is identical to Section 2.2, except* $\forall i \in \{2, ..., N\}$, *player* i *has payoff* $u_i(a_i, y, \theta)$, *with* $g_i(a, \theta) = E_y[u_i(a_i, y, \theta)|a]$. *Players do not observe their stage-game payoffs.*

There are two substantial difficulties in extending the arguments from Sections 2.3 and 2.4. First, because every player's payoff depends on the type θ, Proposition 8 must be refined to show that player 1 reports truthfully "most of the time." RSV's techniques can be applied to prove this result under additional assumptions about 1's utility. Second, for $i \in \{2, ..., N\}$, player i's optimal action depends on his beliefs about θ. So long as i's actions are enforced on a normal hyperplane, deviations can be deterred for any beliefs. However, player i cannot be incentivized through continuation play if the supporting hyperplane is parallel to the i^{th} coordinate axis. Therefore, i's best response cannot depend on his beliefs in a min-max and max-max allocation rule, which limits the generality of this theorem.

The first additional assumption constrains player 1's utility profile.

Assumption 7 *Let* $\alpha : \Theta \to A$ *be a Pareto efficient allocation rule. Then*

$$\sum_{\theta \in \Theta} g_1(\alpha(\theta), \theta) \geq \sum_{\theta \in \Theta} g_1(\alpha(\psi(\theta)), \theta) \tag{2.7}$$

for any permutation $\psi : \Theta \to \Theta$, *with equality if and only if* $\psi(\theta) = \theta$, $\forall \theta \in \Theta$.

Condition (2.7) is a strict version of Rochet's (1987) necessary and sufficient condition for the allocation rule $\alpha(\cdot)$ to be implementable in a static mechanism design problem with quasi-linear utility.[7] RSV assume that an equivalent condition holds for the decision rule

[7] Rochet's Theorem 1 requires that $\forall$ finite cycles $\{\theta_0, ..., \theta_Q = \theta_0\} \subseteq \Theta$,

$$\sum_{k=0}^{Q} \{g_1(\alpha(\theta_k), \theta_{k+1}) - g_1(\alpha(\theta_k), \theta_k)\} \leq 0$$

they implement in equilibrium. In the class of games I consider, monitoring is imperfect and thus (2.7) must hold at *every* Pareto-efficient allocation rule so that continuation play can be used to enforce the desired actions.

In the equilibrium I construct, players are required to occasionally "waste" rounds, in the sense that no information is communicated and every player chooses a constant action. These wasted rounds will be useful precisely when one of the players $i \in \{2, ..., N\}$ cannot be incentivized through continuation play. Hence, I assume that there exists a profile $\alpha^{TI,i}$ for each player $i \in \{2, ..., N\}$ such that $\alpha_i^{TI,i}$ is a myopic best-response regardless of i's beliefs about *last period's* θ.

Assumption 8 $\forall i \in \{2, ..., N\}$, $\exists$ *an action profile* $\alpha^{TI,i}$ *that satisfies*

$$\alpha_i^{TI,i} \in \arg\max_{a_i} \sum_{\theta' \in \Theta} g_i(a_i, \alpha_{-i}^{TI,i}, \theta') P(\theta'|\theta)$$

$\forall \theta \in \Theta$.

Assumption 8 states that for every $i \in \{2, .., N\}$, there exists some action profile $\alpha^{TI,i}$ such that i is playing a best response, regardless of his beliefs about *last period's* type. This assumption holds if both players have an "opt out" action in which they choose not to interact with one another, as in Assumption (v) of Athey and Segal (2007). It may also hold in sender-receiver games if the Markov chain $P(\theta'|\theta)$ is not too persistent.

The results in Section 2.4 rely critically on the fact that for $i \neq 1$, player i's best response to an action profile is independent of θ. In particular, player i's min-max and max-max allocation rules must be implementable without knowing the true state of the world, which is not true if player i's payoff depends on θ. Corollary 2 requires that only player 1 acts in the mechanism, so player i's min-max and max-max allocation rules cannot depend on his beliefs, which fundamentally constrains the set of payoffs supported in equilibrium.

Definition 25 *Define the* type-invariant *min-max and max-max payoffs for player* $i \in \{2, ..., N\}$ *as*

$$g_i^{mTI} = \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} E_\pi \left[g_i(a, \theta) \right]$$

I additionally assume that this inequality is strict unless $Q = 0$ and the cycle is trivial.

and

$$g_i^{MTI} = \max_{a \in A} E_\pi \left[g_i(a, \theta) \right]$$

respectively. Let $\alpha^{mTI,i}$ *and* $\alpha^{MTI,i}$ *be corresponding pure-strategy action profiles, and define*

$$V^{mCommon} = \left\{ v | \exists i \text{ s.t. } v = u(\alpha) \text{ for } \alpha \in \{\alpha^{mTI,i}, \alpha^{MTI,i}, \alpha^{TI,i}\} \right\}.$$

Then the type-invariant *payoff set is*

$$V^{**} = \left\{ v \in V | v \in co(V \cup V^{mCommon}), v_1 \in [g_1^m, g_1^M], v_i \in [g_i^{mTI}, g_i^{MTI}], \forall i \in \{2, ..., N\} \right\}.$$

The main theorem of this section states that any $v \in V^{**}$ can be approximated arbitrarily closely by an equilibrium payoff when players are sufficiently patient. Like V^*, V^{**} is a convex hull of implementable payoffs. However, V^{**} also bounds player $i \in \{2, ..., N\}$'s payoff from above by g_i^{MTI}, which in some games might be strict enough to rule out any Pareto efficient payoffs. Intuitively, this restriction is relatively mild if players have substantial capacity to help or harm one another independently of the state θ. For example, if players can pay one another in each period,[8] then g_i^{MTI} may be large because it includes a large transfer to player i (even if players are not risk-neutral or such transfers are inefficient).

Assumption 9 *There exists an* $\omega > 0$ *so that* $\forall\ \pi' \in \Delta(\Theta)$ *such that* $||\pi' - \pi|| < \omega$, $\forall i \in \{2, ..., N\}$,

$$\alpha_i^{mTI,i} \in \arg\max_{a_i \in A_i} E_{\pi'} \left[g_i(a_i, \alpha_{-i}^{mTI,i}, \theta) \right]$$

and

$$\alpha_i^{MTI,i} \in \arg\max_{a_i \in A_i} E_{\pi'} \left[g_i(a_i, \alpha_{-i}^{MTI,i}, \theta) \right].$$

Assumption 9 ensures that player i's best response to his type-independent min-max and max-max allocation rules is robust to small changes in the type distribution. For example, this assumption is satisfied if either $E_\pi \left[g_i(\alpha^{mTI,i}, \theta) \right] > E_\pi \left[g_i(a_i, \alpha_{-i}^{mTI,i}, \theta) \right]\ \forall a_i \neq \alpha_i^{mTI,i}$, or $\alpha_i^{mTI,i} \in \arg\max_{a_i} g_i(a_i, \alpha_{-i}^{mTI,i}, \theta)$, $\forall \theta \in \Theta$.

[8]Formally, the action a_i includes a payment chosen from some large but finite feasible set.

Under these additional assumptions, any $v \in \text{int}(V^{**})$ is arbitrarily close to an equilibrium payoff when players are patient.

Theorem 2 *Consider a game with a common state of the world such that Assumptions 5 - 9 hold, and suppose $W \subseteq int(V^{**})$ is a smooth set. $\forall \epsilon > 0$, $\exists\, \delta^* < 1$ such that if $\delta \geq \delta^*$, $\forall w \in W$, there exists an equilibrium with payoff v such that $v \in B(w, \epsilon)$.*

Proof: See online appendix.

The argument for Theorem 2 is in Supplemental Appendix ??; here, I discuss notable differences between this proof and that of Theorem 1. One key distinction occurs when a type-independent min-max or max-max allocation for player $i \in \{2, ..., N\}$—say $\alpha^{mTI,i}$—is played in block (k, j) and enforced by continuation play on a hyperplane P_i parallel to the i^{th} coordinate axis. In this case, player i will play a myopic best-response to $\alpha_{-i}^{mTI,i}$, but $\alpha_i^{mTI,i} \in \arg\max_{a_i} E_{\pi'}\left[g_i(a_i, \alpha_{-i}^{mTI,i}, \theta)\right]$ only when $||\pi' - \pi|| < \omega$. To circumvent this problem, the allocation rule $\alpha^{TI,i}$ is played at the beginning of the block, while $\alpha^{mTI,i}$ is played for the remaining rounds. Because $\alpha_i^{TI,i} \in \arg\max_{a_i} \sum_{\theta'} g_i(a_i, \alpha_{-i}^{TI,i}, \theta) P(\theta'|\theta)$, $\forall \theta$, player i conforms to $\alpha^{TI,i}$ in the beginning of the block. If $\alpha^{TI,i}$ is played for a sufficiently long time, player i's beliefs about θ is close to the invariant distribution π, so he is willing to conform to $\alpha^{mTI,i}$ by Assumption 9.

A second difference is that player 1 must tell the truth with high probability in order to implement an allocation rule α. As shown by RSV, player 1 cannot systematically lie and improve his payoff in a T-period quota mechanism if Assumption 7. Intuitively, any lie simply reorders the sequence of messages in a quota mechanism, which under this assumption strictly hurts utility.

The final difference in the proof lies in the definition of invariant payoffs $v^T(\alpha)$: I use the fact that a constant Blackwell policy is optimal when player 1 is sufficiently patient to show that *every* player's payoff approximates some *constant* $v^T(\alpha)$ for all δ sufficiently close to 1. Other than these complications (and the changes to definitions required to take them into account), the proof is similar to the argument outlined in Sections 2.3 - 2.4.

2.6 Conclusion and Future Directions

This paper has proven that when players are patient, there exists an equilibrium that attains nearly efficient payoffs in a large class of games with imperfect private monitoring in which one player has a Markov private type. Adapting techniques from both the repeated games and dynamic mechanism design literatures, I demonstrate that an intuitive equilibrium construction can be used to induce players to cooperate, even when they face both moral hazard and adverse selection problems.

The current analysis assumes that every player observes an identical - albeit imperfect - signal in each round. The information structure in many real-world settings is typically more complicated: players might observe different signals about the actions taken in each period, so that the game is one of *imperfect private monitoring.* In Kandori and Matsushima's (1998) canonical Folk Theorem with imperfect private monitoring and public communication, the equilibrium periodically induces a player truthfully reveal her private history, which is used to reward or penalize the other players. This incentive scheme requires that the monitoring player be *exactly indifferent* between telling the truth and lying, which is difficult in games with persistent private information because each player's continuation payoff depends on her beliefs about the current type. One class of private monitoring games that might remain relatively tractable—originally studied by Ben-Porath and Kahneman (1996)—assumes that at least two players observe the same signal in each period. In this setting, those players can be asked to reveal their shared signal and punished if their reports disagree; hence, I conjecture that an efficiency result similar to Theorem 1 holds in such games with a Markov private type.

An important shortcoming of the approach detailed here is that it does not naturally extend to games with multi-sided private information. If several players have private information, then player 1 is not the only decision-maker in the mechanism studied in Section 2.3. As a result, payoffs would not be lower hemicontinuous in the prior ν, so a *single* invariant payoff approximated by equilibrium payoffs for *every* prior ν might not exist. In the current approach, continuation payoffs are only used to enforce the desired actions; in principle, I conjecture that continuation play might also be useful in inducing truth-telling by replicating transfers from static mechanism design, which might lead to a generalization of the results in both Sections 2.4 and 2.5.

Finally, the quota mechanism is a powerful way to mitigate adverse selection problems when transfers are not available. Most of the papers that study these mechanisms focus on their asymptotic properties, which arise from the Law of Large Numbers; however, the intertemporal tradeoffs generated by such mechanisms seem relevant even when players are impatient. Frankel (2011) demonstrates that quota mechanisms can be optimal in some dynamic mechanism design problems. Elucidating the costs and benefits of these mechanisms in a general setting with impatient players appears to be a formidable task, but I believe that a further investigation of these problems might lead to new insights in settings as diverse as public goods problems and optimal contracts in a workplace.

Chapter 3

Putting the Relationship First

3.1 Introduction

Businesses in a relation-based system will expand at those margins where diminishing returns set in most slowly. This will mean preserving the closeness of the relation, even at the cost of undertaking a new activity that is not economically so close—not such a good complement in production or consumption. - Avinash Dixit (2007)

Repeat dealing, cultural homogeneity...and a lack of third-party enforcement...have been typical conditions. Under them transactions costs are low, but because specialization and division of labor is rudimentary, transformation costs are high. - Douglass North (1990)

Individuals and companies rely on informal relationships with one another to encourage cooperation when formal contracts are incomplete or unavailable. Participants in these relationships invest to strengthen their bonds, rather than simply cutting costs or maximizing productive efficiency. For example, suppliers might alter their production process or produce different goods in order to better meet a favored buyer's needs, leading to markets that look structurally different from those with readily available formal contracts.

In this paper, I illustrate one way that relational contracts can affect investment decisions. A single downstream firm requires several inputs from a group of suppliers. Before the relationship begins, each supplier chooses a set of products to manufacture: highly specialized firms are very efficient at manufacturing a small number of products, while generalist suppliers can inefficiently produce many different goods. The game has imperfect private monitoring—each supplier is unable see the details of the downstream firm's relationship with *other* suppliers—which prevents the upstream firms from jointly punishing a deviation by the downstream firm.

The main result of this paper links suppliers' investments to the underlying contractual environment. When formal contracts are available, many upstream firms enter the market, and each specializes in a small set of products in order to minimize manufacturing costs. In contrast, generalist upstream firms have an advantage when enforceable contracts cannot be written. A generalist supplier can meet many of the downstream firm's needs, increasing the future value generated by that relationship. The supplier can then threaten to withhold production if the downstream firm does not adequately compensate it for output, which induces the downstream firm to pay bonuses and so encourages high effort. This leads to a tension between efficiency and adaptability in a relational contracting setting that is absent when formal contracts are available. I develop applications dealing with employment and legal reform that emphasize this connection between *ex ante* investments and the underlying contracting environment.

As implied in the opening quotes by Dixit (2007) and North (1990), real markets are rife with interactions that are tailored to maximize the efficacy of relational contracts. In an attempt to mimic successful Japanese car companies, Chrysler revolutionized its production process in the early 1990s by developing close relationships with a small number of upstream firms.[1] In the context of this model, Chrysler's decision to use informal contracts naturally led to a reduction in the number of regular suppliers; I argue that the remaining suppliers exerted higher effort precisely *because* they dramatically expanded the set of products they manufacture. Along this line, Liker and Choi (2004) point out that both Toyota and Honda ask their top-tier suppliers to "produce subsystems instead of components." Similarly, Nistor (2012) finds that if a restaurant requires customized ingredients, it tends to source from fewer suppliers. In a survey, Guinipero (1990) reports that manufacturers implementing just-in-

[1]See Dyer (1996) for an in-depth analysis of Chrysler's transformation.

time techniques tend to both reduce the size of their supply networks and emphasize quality.

The fundamental trade-off in this model resembles Bernheim and Whinston's (1990) analysis of multimarket contact. If duopolists compete in several different markets, they can threaten to revert to competition in *every* market following a deviation from the collusive price. Bernheim and Whinston show that if the different markets exhibit certain kinds of heterogeneity, then the duopolists can exploit this heterogeneity in order to better sustain collusive outcomes. In my model, a *generalist supplier* has a similar sort of "multimarket contact" with the downstream firm: it produces—and so can threaten to withhold—many different products. In my basic model, trade occurs in only one market in each round, which creates a very stark benefit for multimarket contact: parties who interact in multiple markets trade with one another more frequently. In Section 3.5.1, I consider a setting in which generalist suppliers are sometimes optimal even if trade occurs in every market in every round.

While the seminal papers by Bull (1987) and Levin (2003) have instigated an extensive literature on relational contracts, relatively few papers consider the interaction between these informal arrangements and market structure (see Malcomson (2012) for a review of the large and growing relational contracting literature). Notable exceptions include Board (2011) and Calzolari and Spagnolo (2009), who argue that relational contracts tend to lead to small markets. In Board's model, a downstream firm pledges future surplus to its suppliers in order to induce cooperation. Because the downstream firm must sacrifice some rent every time it contracts with a supplier, it chooses to contract with a strict subset of the available upstream firms. Similarly, Calzolari and Spagnolo argue that restricting entry in a procurement auction can increase bidders' expected future surplus and thus induce higher effort. By considering a game that emphasizes the importance of ex ante investments, I generate insights that are complementary to both of these papers.

Many of the standard tools in game theory cannot be applied in environments with imperfect private monitoring; see Kandori (2002) for an overview. As a result, much of the theoretical literature focuses on either "Folk Theorem"-type results (investigating the set of equilibria among very patient players), or restricts attention to *belief-free equilibria*, which have a simplifying recursive structure.[2] In contrast, I focus on a simple principal-agents game in which output produced by and bonuses paid to an agent are observed by the principal and that agent, but not by other agents. Hence, this model loosely resembles Ellison's (1994) work on communal enforcement and Wolitzky's (2011) analysis of public goods provision, albeit without any contagion-style punishments.

In Section 3.2, I discuss the timing of the model and introduce several important assumptions. Section 3.3 covers three different benchmark solutions that provide useful comparisons to the main results. I explore the central trade-off between specialization and adaptability in Section 3.4, along with a simple example that has a closed-form solution and is used in the extensions. Section 3.5 covers three applications: the first explores what happens if multiple products are required in each period, the second investigates why upstream firms might resist the introduction of formal contracts, and the third considers human capital investments by employees. I conclude with discussion in Section 3.6.

3.2 Model

I propose a model in which a repeated game is preceded by *ex ante* investments. Section 3.2.1 describes the timing and monitoring structure of the game, and Section 3.2.2 gives definitions and assumptions.

3.2.1 Timing

Consider an intermediate goods market with a single principal, who requires one of many different products in each period. At the beginning of the game, agents decide whether to enter the market or not. They also choose a *specialization* $\mathcal{P}_i \subseteq [0,1]$, which determines the products that they can manufacture *and* their efficiency at producing each good. $\mathcal{P}_i$ captures the fundamental trade-off between efficiency and flexibility which lies at the heart of the analysis: if the measure of $\mathcal{P}_i$ is large, then agent i can make many different goods but must pay a high fixed cost to produce. After entry decisions and specializations are observed by everyone, the repeated game begins. In each period, a single good $\phi_t \in [0,1]$ is randomly

[2] See Ely and Valimaki (2002) and Ely, Horner, and Olszewski (2005) for more details.

drawn as required by the principal. Any agent with $\phi_t \in \mathcal{P}_i$ is able to produce this good; the principal pays a wage to each agent and asks one to produce. That agent pays a fixed cost that depends on $\mathcal{P}_i$ and also chooses a private and costly effort e_t that determines output y_t. After observing output, the principal can choose whether to pay a discretionary bonus $\tau_{i,t}$ to agent i. Importantly, output y_t is observable only by the principal and producing agent, while wage $w_{i,t}$ and discretionary bonus $\tau_{i,t}$ are observable to only the principal and recipient of these transfers.

Players share a common discount factor δ in the repeated game, and do not discount between $t = 0$ and the repeated game $t = 1, 2, ...$ Formally, the game has the following timing:

- *At the beginning of the game $t = 0$:*

  1. A countably infinite number of agents simultaneously choose whether to enter or exit the market. Entry costs $F_E > 0$, and agents that do not enter have no additional actions. Let $\{1, ..., M\}$ be the set of agents in the market.
  2. Each agent $i \in \{1, ..., M\}$ publicly chooses a measurable specialization $\mathcal{P}_i \subseteq [0, 1]$. Denote $\mu_i = \mu(\mathcal{P}_i)$ as the Lebesgue measure of this set.
  3. The principal and agent i simultaneously make transfer payments $\tau^A_{i,0} \geq 0$, $\tau^P_{i,0} \geq 0$, respectively, to one another. Let $\tau_{i,0} = \tau^A_{i,0} - \tau^P_{i,0}$ be the net transfer to agent i.[3]

- *In each round $t = 1, 2, ...$:*

  1. A required product $\phi_t \sim U[0, 1]$ is publicly observed.
  2. The principal offers production to one agent $x_t \in \{\emptyset\} \cup \{1, 2, ..., M\}$. This offer is observed only by agent x_t.
  3. $\forall i \in \{1, ..., M\}$, the principal and agent i simultaneously make wage payments $w^A_{i,t} \geq 0$, $w^P_{i,t} \geq 0$ to one another, with net wage to agent i $w_{i,t} = w^A_{i,t} - w^P_{i,t}$. These payments are observed only by the principal and agent i.

[3]For all transfer payments, if $\tau_{i,0} \geq 0$, the convention is that $\tau^P_{i,0} = 0$, and similarly $\tau^A_{i,0} = 0$ if $\tau_{i,0} \leq 0$.

4. Agent x_t accepts or rejects production: $d_t \in \{0, 1\}$. This decision is observed only by the principal.

5. If x_t accepts the contract, then he pays fixed cost $\gamma(\mu_i)$, where $\gamma : [0, 1] \to \mathbb{R}_+$, and privately chooses effort $e_t \in \{0, 1\}$ at cost ce_t.

6. Output $y_t \in Y \subseteq \mathbb{R}$ is realized, where $y_t \sim F(y|e_t)$ if both $\phi_t \in \mathcal{P}_i$ and $d_t = 1$, and $y_t = 0$ otherwise. y_t is observed only by the principal and x_t.

7. $\forall i \in \{1, ..., M\}$, the principal and agent i simultaneously make bonus payments $\tau_{i,t}^A \geq 0$, $\tau_{i,t}^P \geq 0$, with net payment $\tau_{i,t} = \tau_{i,t}^A - \tau_{i,t}^P$. This transfer is observed only by i and the principal.

8. Payoffs are realized: agent i earns

$$u_{i,t} = (1 - \delta)\left(\tau_{i,t} + w_{i,t} - 1\{x_t = i\}d_t(\gamma(\mu_i) + ce_t)\right),$$

while the principal earns

$$u_{0,t} = (1 - \delta)\left(1\{x_t \neq \emptyset\}d_t y_t - \sum_{i=1}^{M}(\tau_{i,t} + w_{i,t})\right).$$

Three of the assumptions in this model require special consideration. First, specialization determines the subset of goods $\mathcal{P}_i$ that agent i can produce and his fixed cost of producing $\gamma(\mu_i)$. This assumption is stark but cleanly captures the intuition that specialization increases efficiency at the cost of flexibility. Second, an agent's specialization cannot be changed once it is chosen. This is a notion of *lock-in*: an agent tailors its production process to manufacture certain goods, and changing this role is prohibitively costly. Third, agents have *no means of communicating with one another*. While extreme, this assumption is a tractable way to model bilateral relationships between the principal and each agent.[4]

[4]Why don't specialized agents horizontally integrate in order to share information with one another? One possible answer is that "generalist agents" in this model are in fact horizontally integrated. Under this interpretation, $\gamma(\mu)$ represents the production costs associated with a bloated organizational structure. There might also be legal or financial constraints that prevent horizontal integration, or behavioral restrictions that limit joint punishment by different divisions within a company.

3.2.2 Histories and Equilibrium

Because players observe different outcomes as the game progresses, I separately track each player's private history over time.

Definition 26 *The* set of baseline histories $\mathcal{H}_B^T$ at time T *is*

$$\mathcal{H}_B^T = \left\{M, \{\mathcal{P}_i\}_{i=1}^M, \{\tau_{i,0}^A, \tau_{i,0}^P\}_{i=1}^M, \left\{\phi_t, x_t, d_t, \{w_{i,t}^A, w_{i,t}^P\}_{i=1}^M, e_t, y_t, \{\tau_{i,t}^A, \tau_{i,t}^P\}_{i=1}^M\right\}_{t=1}^T\right\}$$

The principal observes all actions except for effort e_t, so the set of principal's baseline histories at time T *is* $\mathcal{H}_{B,i}^T = \left\{h^T \backslash \{e_t\}_{t=1}^T | h^T \in \mathcal{H}_B^T\right\}$. *The* set of agent i's baseline histories at time T *is*

$$\mathcal{H}_{B,i}^T = \left\{M, \{\mathcal{P}_i\}_{i=1}^M, \tau_{i,0}^A, \tau_{i,0}^P, \{\phi_t, x_t, d_t 1\{x_t = i\}, w_{i,t}, \tau_{i,t}, 1\{x_t = i\}e_t, 1\{x_t = i\}y_t\}_{t=1}^T\right\}$$

Let $\mathcal{N}$ be the nodes of the stage game; then (h^T, n_{T+1}) indicates a history at node $n_{T+1} \in \mathcal{N}$ of round $T+1$. let $\mathcal{I}_i(n_t)$ be agent i's information set at stage-game node n_t, so that a private history is $(h_i^T, \mathcal{I}_i(n_{T+1}))$.

Strategies are denoted σ_i for agent $i \in \{1, ..\}$ and σ_0 for the principal, with profile $\sigma = \{\sigma_0, ...\}$. A relational contract is a Perfect Bayesian Equilibrium (PBE) of the repeated game.[5] A relational contract is *stationary* if on the equilibrium path, actions in period t depend only on variables observe in period t, and is *optimal* if it maximizes total *ex ante* expected surplus. Because monitoring is imperfect and private, I cannot use standard recursive techniques in this analysis and so rely on other methods.

The cost function $\gamma(\mu)$ is constrained so that a meaningful trade-off between specialization and flexibility exists.

Assumption 10 *$\gamma(\mu)$ is differentiable with $\gamma', \gamma'' > 0$ and $\gamma'(0) = 0$, and $\mu(E[y|e = 1] - c - \gamma(\mu))$ is strictly increasing in μ.*

[5] A Perfect Bayesian Equilibrium consists of a strategy profile σ and belief system $\rho = \{\rho_i\}_{i=0}^{\infty}$ over true histories for each player such that (1) given beliefs $\rho_i(h_i^t)$, σ_i maximizes player i's continuation surplus, and (2) ρ_i updates according to Bayes Rule whenever it is well-defined. When Bayes Rule is not well-defined, ρ_i assigns weight only to histories that are consistent with agent i's information but is otherwise unconstrained.

Assumption 10 implies that if agent i is allocated production of every good in $\mathcal{P}_i$ and works hard, then the total surplus produced by i is increasing in μ_i. In addition to driving the tension between adaptability and productive efficiency, this assumption ensures that it is optimal for specializations $\{\mathcal{P}_i\}_{i=1}^M$ to cover the entire interval if the agents are expected to work hard.

Finally, I constrain $F(y|e)$ so that it is efficient for an agent to accept production if and only if he works hard.

Assumption 11 *1. F first-order stochastically increases in effort:* $F(y|e=1) >_{FOSD} F(y|e=0)$.

2. $e = 1$ *is strictly efficient:* $E[y|e=1] - c - \gamma(\mu) > 0 \geq E[y|e=0] - \gamma(\mu)$, $\forall \mu \in [0,1]$.

It will turn out in this analysis that the critical determinant of the strength of agent i's relationship with the principal can be measured by the *total surplus produced by agent* i, which is determined by whether (1) i is allocated production of $\phi_t \in \mathcal{P}_i$, (2) i accepts production if it is offered, and (3) i works hard if he accepts production.

Definition 27 *The* total per-period surplus produced by agent i *in round* t *is*

$$\pi_{i,t}^{TOT} = (1-\delta)1\{x_t = i\}d_t(y_t - ce_t - \gamma(\mu_i)).$$

The principal's per-period surplus from agent i *is*

$$\pi_{0,t}^{i} = (1-\delta)1\{x_t = i\}d_t y_t - w_{i,t} - \tau_{i,t}.$$

Given strategy profile σ and history (h^{t-1}, n_t), the continuation surplus for agent i *is*

$$U_i(h^{t-1}, n_t; \sigma) = E_\sigma\left[\sum_{t'=t+1} \delta^{t'-t-1} u_{i,t} | h^{t-1}, n_t\right]$$

and the continuation surplus for the principal from agent i *is*

$$U_0^i(h^{t-1}, n_t; \sigma) = E_\sigma\left[\sum_{t'=t+1} \delta^{t'-t-1} \pi_{0,t}^{i} | h^{t-1}, n_t\right],$$

with the principal's total continuation payoff equal to $U_0 = \sum_{i=1}^{M} U_0^i$.

When unambiguous, I suppress the notation $(h^{t-1}, n_t; \sigma)$ in U_i, U_0^i, and U_0. By construction, $\pi_{i,t}^{TOT} = \pi_{0,t}^i + u_{i,t}$ and the total surplus produced in a period is $\sum_{i=1}^{M} \pi_i^{TOT}$. Therefore, π_i^{TOT} captures the *contribution of each agent to total surplus* in a period. Intuitively, this total surplus is at stake in each relationship in the sense that it is lost as a punishment if either the principal or agent i reneges on a promised payment. If agent i works hard when he is allocated production, then giving production to agent i more frequently increases π_i^{TOT} and thus increases the size of this punishment. Because $\mu(\mathcal{P}_i)$ determines how frequently agent i is able to produce, the size of $\mu(\mathcal{P}_i)$ determines the maximum value of π_i^{TOT}, which in turn determines the set of credible bonuses that can be supported in a relational contract. This fundamental tension drives my main result.

Following Levin (2003) and much of the subsequent relational contracting literature, I focus on the optimal equilibrium. In this setting, any PBE is payoff equivalent to a PBE in which agents do not condition on their past efforts,[6] so I consider only relational contracts that are independent of past effort decisions.

3.3 Benchmarks - First Best, One-Shot, and Public Monitoring

Three different benchmarks are relevant in this model. The first characterizes the first-best by supposing that output y_t is contractible, the second considers the one-shot game, and the final considers optimal relational contracting if monitoring were public.

3.3.1 Optimal Entry and Specializations with Formal Contracts

Suppose that output y_t, entry, and specializations $\mathcal{P}_i$ are contractible at the beginning of the game. Then the principal can efficiently induce high effort from every agent, since all parties are risk-neutral and have deep pockets. Because each agent exerts high effort regardless of her specialization $\mathcal{P}_i$, these specializations are chosen to balance the fixed costs $\gamma(\mu_i)$ against

[6] The proof of this fact may be found in Andrews and Barron (2012), Appendix B.

the cost of entering the market F_E. Therefore, it is straightforward to calculate the efficient number of entrants and their specializations.

Proposition 12 *Let Assumptions 1 and 2 hold. Every optimal equilibrium entails M^{FB} entrants and takes the following form:*

1. *For any agents i, k in the market, $\mu(\mathcal{P}_i) = \mu(\mathcal{P}_k) = \mu$, where $\mu = \frac{1}{M^{FB}}$.*

2. *$\forall t$, effort is $e_t = 1$.*

The first best number of firms M^{FB} decreases in F_E and is independent of δ.

Proof: See Appendix 3.7.

A principal who has access to formal incentive contracts can always motivate her workers, so the number of firms in the market is determined by setting marginal production costs $\frac{1}{M^2}\gamma'(\frac{1}{M})$ equal to the cost of entry F_E (subject to integer constraints). In Section 3.4, I will show that M^{FB} is large and specializations $\mu = \frac{1}{M^{FB}}$ is specialized relative to the optimal relational contract when players are impatient. Intuitively, when formal contracts are available, market size is limited only by the cost of entry and each firm specializes in a narrow band of products.

The distributions $F(y|e = 1)$ and $F(y|e = 0)$ are statistically distinguishable, so some contract exists that induces high effort. I record this result as a corollary.

Corollary 3 *There exists a bounded set of transfers $\tau(y)$ that induce the agent to exert high effort, and such that $\infty > \sup_y \tau(y) - \inf_y \tau(y)$. For any transfers $\tau(y)$ that induces high effort, $\sup_y \tau(y) - \inf_y \tau(y) > 0$.*

Proof: One such contract with bounded transfers is exhibited in the proof of Proposition 1. Note that the agent's IC constraint is

$$E[\tau(y)|e = 1] - E[\tau(0)|e = 0] \geq c$$

and so $\sup_y \tau(y) - \inf_y \tau(y) \geq c > 0$ as desired. ■

Rather than explicitly calculating the optimal incentive contract in what follows, I will instead use Corollary 3 to argue that an agent's payoff must vary with y in order to motivate him to work hard. In a relational contract, the principal must prefer to pay this bonus rather than renege and face a punishment by the betrayed agent, and the role of $\mathcal{P}_i$ is to ensure that the principal is indeed willing to do so.

3.3.2 Equilibrium in the One-Shot Game

A second important benchmark is the one-shot version of the repeated game. The principal cannot credibly promise to reward the agent for high output, so no surplus is generated in equilibrium. Because entry is costly and low effort is inefficient, no agents enter the market ($M_{SPOT} = 0$).

Proposition 13 *The unique payoff in the spot game is* 0*, with* $M_{SPOT} = 0$.

Proof: Using backwards induction, $\tau_i(y) = 0$, $\forall i, y \in \{0, y_H\}$. Therefore, $e^* = 0$ in equilibrium, and so total surplus generated by any firm in the market is no larger than 0 because $E[y|e = 0] - \gamma(\mu) \leq 0$. But then no firm chooses to enter the market, since they must incur the cost $F_E > 0$ to do so. ■

Proposition 13 demonstrates that repeated interaction between the principal and each agent is required to induce the agents to work hard. While all players earn their min-max payoff in this one-shot equilibrium, Section 4 will prove that punishments following a deviation that is not publicly observed do *not* typically mix-max the deviator.

3.3.3 What Happens if Monitoring is Public?

It is instructive to consider the optimal equilibrium if x_t, y_t, $\{w_{i,t}\}$, and $\{\tau_{i,t}\}$ were publicly observed so that the game was one of imperfect public monitoring. This benchmark will highlight role of private monitoring in the baseline model.

Many of the tools developed in the foundational relational contracting paper by Levin (2003) can be adapted to this setting. As a first step, I show that stationary contracts are optimal in this setting.[7]

Lemma 11 *There exists an optimal relational contract that is stationary in the public monitoring game.*

Proof: See Appendix 3.7.

The proof of Lemma 11 constructs a stationary optimal relational contract and is similar to Levin (2003), Theorem 2. In this equilibrium, the principal earns all of the surplus produced in each period, and is punished by reversion to the static equilibrium following any deviation. As a result, whenever the agents are willing to work hard in any equilibrium, the principal is willing to pay them an incentive scheme that motivates them to work hard in each period of a stationary equilibrium.

The next result demonstrates that the efficient contract either has $M = 0$, or $M \geq M^{FB}$.

Proposition 14 *In an optimal relational contract,* $\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$, $\forall i, k \in \{1, ..., M\}$. $\mu(\mathcal{P}_i) = \frac{1}{M}$, *and either* $M^{Pub} = 0$ *or* $M^{Pub} \geq M^{FB}$ *firms enter the market.*

Proof: See Appendix 3.7.

With public monitoring and transfers between agents, *at least* M^{FB} firms enter the market. As in the first best, additional entry leads to lower production costs $\gamma(\frac{1}{M})$. This reduction in costs has two benefits: it both directly increases surplus *and* indirectly allows the principal to credibly promises larger bonuses in equilibrium, since she would lose her entire continuation surplus—which does not include the cost of entry F_E—were she to renege on this promise. The optimal relational contract weighs both of these benefits against the cost F_E of additional entry to determine the efficient number of entrants, whereas the first-best weighs only the first of these benefits against F_E. Therefore, $M^{Pub} \geq M^{FB}$.

[7] Note that utility is not transferrable between agents in this setting, and so the model differs from Levin (2003). Indeed, Andrews and Barron (2012) show in a related model that if the set of relational contracts were restricted so that the agents earned a positive fraction of the surplus they produced, then a tension similar to what is explored in this paper continues to be relevant.

3.4 Putting the Relationship First

In this section, I consider the optimal relational contract in the baseline game. The optimal equilibrium is substantially different than the first-best benchmark from Section 3.3.1: a relatively *small* number of agents enter the market, each of whom is able to produce a *broad* range of products. I prove the general result in Section 3.4.1, while Section 3.4.2 considers a specific example with a complete closed-form solution.

3.4.1 The Role of Flexibility in Relational Contracts

This section proves that if formal contracts are unavailable, the optimal market structure involves a small number of agents that are not very specialized. Because agents cannot communicate with one another, they are unable to coordinate and jointly punish deviations by the principal. As a result and unlike the case with public monitoring, the future surplus generated by *each* agent determines the largest bonus that can be paid to that agent. *Ex ante* investments that focus on flexibility (i.e., a large μ_i) increase the total surplus generated by agent i, at the cost of reducing the maximum feasible total surplus.

Together, the next lemmas provide the basic ingredients for the main result. First, I argue that a principal that deviates in her relationship with agent i is punished by no worse than the *bilateral breakdown* of trade with i. To ensure that each agent is sufficiently valuable, the equilibrium measure of specialization $\mu(\mathcal{P}_i)$ must be bounded from below by some $\mu^*(\delta) \geq \frac{1}{M^{FB}}$. Moreover, if specializations do not overlap (so that $\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$), then the efficient relational contract can be replicated by a stationary relational contract. Proposition 15 combines these arguments to demonstrate the main result: when relational contracts are used and players are impatient, the number of agents in the market is smaller and each agent specializes in a broader range of tasks than in the first-best equilibrium.

Other agents do not observe when the principal betrays one of their number, so the punishment in each principal-agent dyad depends on the surplus generated by that relationship. More precisely, if an equilibrium transfer is not made to (or by) agent i, then the size of the punishment following this deviation is bounded above by the value of the relationship to the player paying the bonus. Thus, the contribution of *each agent* to total surplus is the key variable that determines what is at stake in each relationship and thus the incentive pay $\tau_{i,t}$ that can be offered to each agent.

Lemma 12 *Let Assumption 11 hold, and suppose σ^* is an equilibrium. Fix a history (h^{t-1}, n_t), where n_t is a node immediately following output y_t. Then*

$$\begin{aligned} (1-\delta)E_\sigma\left[\tau_{i,t}|h_i^{t-1},\mathcal{I}_i(n_t)\right] &\leq \delta E_{\sigma^*}\left[U_0^i(h^t)|h_i^{t-1},\mathcal{I}_i(n_t)\right] \\ -(1-\delta)E_\sigma\left[\tau_{i,t}|h_i^{t-1},\mathcal{I}_i(n_t)\right] &\leq \delta E_{\sigma^*}\left[U_i(h^t)|h_i^{t-1},\mathcal{I}_i(n_t)\right] \end{aligned} \tag{3.1}$$

Proof: See Appendix 3.7.

The proof of Lemma 12 argues that the player responsible for paying $\tau_{i,t}$ always has the option of refusing to pay and suffering a breakdown in the relationship. Because agent i interacts only with the principal, a breakdown in this realtionship would hold i at his outside option 0. On the other hand, the principal can continue to allocate business to the other agents following a breakdown with agent i. Therefore, she can "cut her losses" following a breakdown and so $(1-\delta)\tau_{i,t}$ is bounded by the value of *that agent* to the principal, U_0^i. The bound (3.1) on the principal's punishment is not necessarily tight, since she may be able to do strictly better by surreptitiously altering her allocation rule following a deviation. For instance, if the principal could award product ϕ_t to either firm i or firm j in round t, then she may be able to reward the product to j following a breakdown with i without triggering any futher punishment.

For Lemma 12 to hold, agents must be unable to communicate; otherwise, they could coordinate to jointly punish a deviation by the principal.[8] Hence, there is an incentive to aggregate production in a small group of firms in order to limit reneging temptation. This impulse towards aggregation leads to a deviation from the first best market structure, since a specialized firm produces at a low cost but is unable exact a revenge sufficiently large to deter the principal from deviating.

The next lemma demonstrates that it is in fact easier to induce high effort from an agent with a large $\mathcal{P}_i$, particularly when that agent is the sole producer of every product in $\mathcal{P}_i$.

[8] Alternatively, I could assume that agents can communicate but are behaviorally restricted to "bilateral breakdown" following a deviation by the principal. Such a behavioral restriction is non-trivial to define because agents' beliefs about how the principal allocates business might change when they observe a deviation.

Lemma 13 *Fix the number of entrants M and specializations $\{\mathcal{P}_i\}$. A necessary condition for there to exist a PBE in which agent i chooses $e_t = 1$ on the equilibium path is that there exists an incentive scheme $\tau(y)$ that induces $e_t = 1$ and*

$$\sup_y \tau_i(y) - \inf_y \tau_i(y) \leq \frac{\delta}{1-\delta}\mu_i \left(E[y|e=1] - c - \gamma(\mu_i)\right) \tag{3.2}$$

Proof: See Appendix 3.7.

Agent i can be motivate to work hard using both contemporaneous transfers $\tau_{i,t}$ and continuation payoffs U_i. From Lemma 12, the total variation in these incentives must satisfy $0 \leq (1-\delta)\tau_{i,t} + \delta U_i \leq \delta(U_i^0 + U_i)$. The sum $U_i + U_i^0$ is bounded above by the right-hand side of (3.2), which is the surplus produced by agent i if he is awarded every $\phi \in \mathcal{P}_i$ in each period and always works hard. If agent i were unwilling to choose $e_t = 1$ if she were the sole producer of every $\phi_t \in \mathcal{P}_i$ in each round, then she would also be unwilling to choose $e_t = 1$ in *any* relational contract. Intuitively, allocating production to agent i whenever possible maximizes the size of the punishment following a deviation in two ways. First, if i produces every $\phi \in \mathcal{P}_i$ then she can threaten to withhold production of these goods, which creates a powerful incentive to maintain the relationship. Moreover, all *other* agents also expect agent i to produce every $\phi \in \mathcal{P}_i$, so the principal cannot deviate from this allocation rule without being punished by *every* agent. Thus, if the principal were to renege on agent i, she would be unable to reallocate production without triggering a punishment.

The next corollary shows that the incentive condition (3.2) is also sufficient if the agents have specializations that do not overlap.

Corollary 4 *If $\mu(\mathcal{P}_i \cap \mathcal{P}_j) = 0\ \forall i,j \in \{1,...,M\}$, then there exists a stationary optimal relational contract. In this equilibrium, agent i picks $e_t = 1\ \forall t$ on the equilibrium path iff $\exists$ transfers $\tau_i(y)$ that induce high effort and satisfy (3.2).*

Proof: See Appendix 3.7.

Consider an equilibrium such that $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$ for *every* agent. In round t, the principal optimally allocates production of $\phi_t \in [0,1]$ to the sole agent who can produce, so the maximum feasible continuation surplus in each relationship does not depend on the principal's allocation rule. Hence, a stationary relational contract is optimal for similar reasons to Levin (2003), and in particular there exists an optimal stationary equilibrium in which the principal earns all of the surplus in each period. In such a stationary contract, the left-hand side of (3.2) captures the temptation to renege on a promised bonus payment, whereas the right-hand side is the amount of surplus lost by the principal if she does not pay the equilibrium bonus τ_i.

So long as $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$, (3.2) implies that the specialization $\mathcal{P}_i$ is the sole determinant of whether agent i works hard in the optimal relational contract. In general, equilibrium behavior in this game can be quite complicated, but much of this complexity stems from how the principal allocates business over time. If this allocation rule is independent of the history—which is natural when specializations don't overlap—then simple contracts are optimal.

The incentive constraint (3.2) suggests an important definition: the minimum specialization required to induce high effort in equilibrium, $\mu^*(\delta)$.

Corollary 5 *There exists a continuous function $\mu^*(\delta)$ that is defined on $\delta \in [\hat{\delta}, 1)$ for some $\hat{\delta} < 1$ such that (1) $\mu^*(\delta)$ is decreasing in δ, (2) $\mu^*(\hat{\delta}) = 1$, and (3) $\lim_{\delta \to 1} \mu^*(\delta) = 0$, and $\exists$ a PBE in which agent i chooses $e_t = 1$ at some history only if $\mu_i \geq \mu^*(\delta)$. Moreover, if $\mu_i \geq \mu^*(\delta)$, $\forall i$, and $\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$, $\forall i, k$, then all agents choose $e_t = 1$, $\forall t$ along the equilibrium path.*

Proof: See Appendix 3.7.

The critical threshold $\mu^*(\delta)$ is the smallest interval of specialization such that inequality (3.2) holds for some incentive scheme that induces high effort. In an optimal equilibrium, every agent that enters the market must satisfy $\mu_i \geq \mu^*(\delta)$; otherwise, that firm would never exert high effort, and so should instead stay out of the market to save the entry cost F_E. Moreover, if $\mathcal{P}_i \cap \mathcal{P}_k = \emptyset$, $\forall i, k$, then $\mu_i \geq \mu^*(\delta)$ is sufficient to induce high effort from i.

Corollary 5 illustrates the central intuition of the model: if δ is far from 1, $\mu^*(\delta) > \frac{1}{M^{FB}}$, so firms must choose a broader specialization than in the first-best. In other words, each firm "puts the relationship first:" rather than specializing to minimize manufacturing costs, as they would if formal contracts were available, firms in the market instead inefficiently produce a broad array of different products. Broad specializations—$\mu(\mathcal{P}_i) > \frac{1}{M^{FB}}$—lead to lower total surplus given $e = 1$ but also increase *dyad-specific surplus* $U_0^i + U_i$, so that better relational incentive contracts can be implemented in equilibrium.

Proposition 15 puts the preceding steps together to show that a market reliant on relational contracts typically has fewer entrants and broader specializations than a market in which formal contracts are available. If participants are more impatient, then the number of agents in the market is smaller and each entrant is responsible for a broader set of products.

Proposition 15 *Suppose Assumptions 1 and 2 hold. Then:*

1. *For every $M \leq M^{FB}$, there exists $\bar{\delta}(M)$ such that for all $\delta \geq \bar{\delta}(M)$, at least M firms will enter the market in any optimal equilibrium.*

2. *For every $M \leq M^{FB}$, there exists an open interval $\Delta(M)$ that satisfies $\sup \Delta(M) \leq \inf \Delta(M+1)$, such that for every optimal equilibrium with $\delta \in \Delta(M)$:*

   (a) *M firms are optimal;*

   (b) *$\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$, $\mu(\mathcal{P}_i) = \frac{1}{M}$, and the efficient equilibrium can be replicated by a stationary equilibrium.*

Proof: See Appendix 3.7.

The proof of Proposition 15 relies on the fact that $\mu(\mathcal{P}_i) \geq \mu^*(\delta)$ $\forall i$ in the efficient equilibrium. So long as $\mu^*(\delta)M \leq 1$, there is no reason for specializations to overlap; because $\gamma(\mu)$ is increasing and strictly convex, it is instead optimal for each firm to specialize in a disjoint subset of the same size, $\mu(\mathcal{P}_i) = \frac{1}{M}$. In this situation, $\frac{1}{M} \geq \mu^*(\delta)$ and so the agents exert high effort in the optimal equilibrium and together produce the entire interval of goods. For $M < M^{FB}$, surplus from this equilibrium is strictly increasing in M conditional on high effort, so any equilibrium with $M < M^{FB}$ firms and $(M+1)\mu^*(\delta) \leq 1$ is dominated by an equilibrium with $M+1$ firms with disjoint specializations of measure $\frac{1}{M+1}$. This bounds the number of entrants M from below. Moreover, if $\mu^*(\delta)M = 1$, then it is uniquely optimal for exactly M firms to enter and choose disjoint specializations with $\mu(\mathcal{P}_i) = \mu^*(\delta)$. This equilibrium both induces high effort and minimizes entry and fixed costs. Because $F_E > 0$, it remains optimal for M firms with disjoint specializations to enter the market on an open interval $\Delta(M)$ about the δ for which $\mu^*(\delta)M = 1$.

Together, Proposition 12 and Proposition 15 illustrate the central trade-off between flexiblity and efficiency. When formal contracts are available, a large number of agents enter the market and each specializes in a relatively small subset of products to minimize the fixed cost $\gamma(\mu)$. In contrast, when high effort can only be induced through a relational contract and players are impatient, the optimal equilibrium often involves a small number of relatively inefficient firms, each of whom is responsible for producing a wide variety of products and generating substantial surplus. Because the set of products is fixed and each agent specializes in a large subset of that set, fewer agents to needed to satisfy the principal's needs.

On the intervals $\Delta(M)$, $\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$ and there is no overlap between specializations. In other words, each product is "single-sourced:" agents never compete to produce the same set of products. Single-sourcing is not uncommon within supplier networks, particularly when the firms rely on relational contracting—for instance, the promise to "carry out business with...suppliers without switching to others" is enshrined in Toyota's 1939 Purchasing Rules.[9] In this model, multi-sourcing—in which several agents can produce the same inputs—increases the principal's outside option and thus her reneging temptation, which makes it more difficult to sustain high effort within a relationship. Now, suppose that multi-sourcing did occur, perhaps for some unmodeled reason. Even in this case, the tension between adaptability and efficiency presented in Proposition 15 would be unlikely to disappear. Indeed, multi-sourcing makes it harder to induce high effort; because flexibility is a way to increase lock-in and effort, flexible *ex ante* investments are one way to *mitigate* the deleterious effects of multi-sourcing on relationships. In short, the trade-off between efficiency and flexibility seems likely to hold, even if multi-sourcing were optimal for other (unmodeled) reasons.

Outside of the intervals $\Delta(M)$ in Proposition 15, specializations might overlap $\mathcal{P}_i \cap \mathcal{P}_j \neq \emptyset$. Andrews and Barron (2012) explore the non-stationary allocation rules that are optimal in such a setting.

[9] As referenced in Sako (2004).

3.4.2 A Simple Framework for Applications

This subsection presents an example for which a simple closed-form optimal relational contract exists. While this example is a special case, it starkly illustrates the broader trade-off between flexibility and specialization.

Consider the following binary output: if $e_t = 0$ or $\phi_t \notin \mathcal{P}_i$, then $y_t = 0$; if $e_t = 1$ and $\phi_t \in \mathcal{P}_i$, then $y_t = y_H > 0$ with probability p and otherwise $y_t = 0$. Suppose the cost function is

$$\gamma(\mu) = \left\{ \begin{array}{ll} 0, & \mu \leq \frac{1}{M} \\ \gamma, & \mu > \frac{1}{M} \end{array} \right\} \tag{3.3}$$

where $M \geq 2$ is an integer and γ satisfies $y_H p - c - \gamma > \frac{1}{M}(y_H p - c) > F_E$. In this market, each agent faces a very simple choice: they either specialize in a subset $\frac{1}{M}$ of the market—what I'll call a "specialist"—or they choose to become a "generalist," able to inefficiently produce whatever product is required. For the purposes of this analysis, assume $MF_E < \gamma$, so that it is optimal to have M specialists enter the market if output is contractible.

Proposition 16 *Define $\bar{\delta}$ by $\frac{c}{p} = \frac{\bar{\delta}}{1-\bar{\delta}}\frac{1}{M}(y_H p - c)$ and $\underline{\delta}$ by $\frac{c}{p} = \frac{\underline{\delta}}{1-\underline{\delta}}(y_H p - c - \gamma)$. In this example with the assuptions given above, any optimal equilibrium satisfies:*

1. *If $\delta \geq \bar{\delta}$, then M firms enter the market and specialize in subsets $\mathcal{P}_i \subseteq [0,1]$ with $\mu(\mathcal{P}_i) = \frac{1}{M}$ and $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$. Firm i produces every $j \in \mathcal{P}_i$.*
2. *If $\delta \in [\underline{\delta}, \bar{\delta})$, then a single firm enters the market, specializes in $\mathcal{P}_1 = [0,1]$, and manufactures every $j \in [0,1]$.*
3. *If $\delta < \underline{\delta}$, then no firms enter the market.*

Proof:

First, consider optimal entry and specialization supposing that $e_t = 1$ in every period. Because $MF_E < \gamma$, the maximum surplus if there are $S \leq M$ specialists and $G \leq 1$ generalists is

$$\frac{S}{M}(y_H p - c) + G\frac{M-S}{M}(y_H p - c - \gamma) - (S+G)F_E.$$

The derivative of this expression with respect to S is $\frac{1}{M}(y_H p - c) - \frac{G}{M}(y_H p - c - \gamma) - F_E$. If $G = 1$, then this derivative is positive if $\gamma > MF_E$, and if $G = 0$, it is strictly positive if $\frac{1}{M}(y_H p - c) > F_E$. These inequalities hold by assumption, so $S = M$ and $G = 0$ is optimal if each specialist chooses $e_t = 1, \forall t$.

Now, consider the optimal relational contract, and define τ_H and τ_0 to be a bonus scheme with minimal $|\tau_H - \tau_0|$ such that

$$p\tau_H + (1-p)\tau_0 - c \geq \tau_0.$$

It is clear that one such bonus scheme is $\tau_H = \frac{c}{p}$, $\tau_0 = 0$. If $\delta \geq \bar{\delta}$, then by Corollary 4 first-best effort can be induced in a relational contract with M entrants and $\mathcal{P}_i \cap \mathcal{P}_k = \emptyset$, and so first-best can be attained. If instead $\delta < \bar{\delta}$, Corollary 5 implies that in order for agent i to choose $e = 1$ in equilibrium, it must be that $\mu_i > \frac{1}{M}$. Because all $\mu_i > \frac{1}{M}$ have the same fixed cost γ of production and $y_H p - c - \gamma > F_E$, it is optimal for a single agent to enter with $\mathcal{P}_1 = [0, 1]$. By Corollary 4 and using the bonus scheme τ_0, τ_H, that agent chooses $e_t = 1 \; \forall t$ so long as

$$\frac{c}{p} \leq \frac{\delta}{1-\delta}(y_H p - c - \gamma)$$

By assumption, $y_H p - c - \gamma > \frac{1}{M}(y_H p - c)$, so there exists a range $\delta \in (\underline{\delta}, \bar{\delta})$ in which a single agent enters the market and chooses $\mathcal{P}_1 = [0, 1]$. ■

This example provides a very sharp result. For sufficiently patient firms, the first best market structure can be achieved: a large number of agents enter the market, and each specializes in a small subset of products. As δ decreases, however, the market abruptly collapses to a *single* agent.[10] This lone remaining agent instead prioritizes his relationship with the principal in order to preserve its own incentive to exert high effort. Because the first-best market structure involves M entrants specializing in $\mu(\mathcal{P}_i) = \frac{1}{M}$, Proposition 16 reiterates that relational contracts tend to involve *fewer agents* and *more flexible investments* than formal contracts.

[10]Notice that this discontinuity is the result of the discountinuous cost function γ and is different than Proposition 4, which proves that for continuous γ *any* $M \leq M^{FB}$ is optimal for an open set of discount factors.

3.5 Applications

The availability and quality of formal contracts differ dramatically between countries, industries, and jobs. In this section, I extend the basic model to explore several implications for production and investment. Extension 3.5.1 shows that the basic trade-off between flexibility and specialization may persist even if every product in required in every round. Extension 3.5.2 illustrates that entrenched firms might agitate *against* socially beneficial legal change because they fear that it would render their investment in flexibility obsolete. Extension 3.5.3 considers the employment relationship and suggests important differences between the skill sets and assigned tasks of employees and independent contractors.

3.5.1 Extension 1 - Multiple Required Products in Each Round

In the baseline model, I assume that the principal requires a single product in each round, which implies that agents with broader specializations interact *more frequently* with the principal. The purpose of this extension is to demonstrate that the trade-off between flexibility and specialization may be relevant even when every product is required in each round. This result follows because the principal can tailor the relational contract to provide stronger incentives to a generalist agent without increasing the maximal reneging temptation.

To make this point, I consider a very simple model that departs from the baseline model in two key ways. First, the products required by the principal are drawn from a *finite set* $\{1, ..., M\}$, so that agents specialize in a subset $\mathcal{P}_i \subseteq \{1, ..., M\}$. This assumption ensures that there is aggregate uncertainty in the market conditional on effort, which is required for the result to hold. Second, *all products are required in every period*, so the principal chooses one agent $x_{\phi,t}$ to produce each good $\phi \in \{1, ...M\}$ in each round t. An agent i assigned the set of products $\chi_{i,t} \subseteq \{1, ..., M\}$ exerts effort $e_{\phi,t}$ on each $\phi \in \chi_{i,t}$ at cost $c\sum_{\phi\in\chi_{i,t}} e_{\phi,t}$. The principal earns surplus $\sum_{\phi=1}^{M} y_{\phi,t}$, $y_{\phi,t} \in \{0, y_H\}$, where

$$Prob\{y_{\phi,t} = y_H|e_\phi\} = \left\{ \begin{array}{ll} p & e_\phi = 1 \\ 0 & e_\phi = 0 \end{array} \right\}$$

if $\phi \in \mathcal{P}_{x_{\phi,t}}$, and $y_{\phi,t} = 0$ if $\phi \notin \mathcal{P}_{x_{\phi,t}}$. As in Section 3.4.2, I assume the specialization function

$\gamma(\cdot)$ takes the simple form that $\gamma(\mathcal{P}_i) = 0$ if $\mathcal{P}_i$ is a singleton and otherwise $\gamma(\mathcal{P}_i) = \gamma$. I also assume that $F_E = 0$.

Proposition 17 illustrates that a single generalist agent can sometimes be optimal in this game.

Proposition 17 *Consider the model in this section, and suppose that $p < \frac{M-1}{M}$ and $\frac{M-1}{M}(y_H p - c) > \gamma$. In an optimal equilibrium with the minimal number of agents in the market, there exists a γ^* such that if $\gamma \leq \gamma^*$, there exist cutoffs $\bar{\delta}_S > \underline{\delta}_S$ such that:*

1. *If $\delta > \bar{\delta}_S$, then M firms enter the market and specialize in a single product;*
2. *If $\delta \in (\underline{\delta}_S, \bar{\delta}_S)$, then a single firm enters the market, specializing in M products;*
3. *Otherwise, no firms enter the market.*

Proof: See Appendix 3.7.

The extent of cooperation in a relational contract depends on the shape of the bonus scheme τ: as in Lemma 13, agent i works hard only if there exists some bonus scheme τ defined on the possible outcomes of products in $\chi_{i,t}$ that induces high effort and satisfies

$$\sup \tau - \inf \tau \leq \frac{\delta}{1-\delta}\left(U_i + U_i^0\right)$$

Suppose agent i manufactures every product $\{1, ..., M\}$; then he must only be motivated to work hard on each $\phi \in \{1, ..., M\}$ based on the *expected* realizations of $(y_1, ..., y_M)$. Because M is finite, agent i faces aggregate uncertainty about the vector of realized outputs. If it is unlikely that $y_j = y_H$ for *every* product j, the optimal contract can offer a relatively low payment for this outcome while still providing incentives for effort. The reneging temptation depends only on the *largest* and *smallest* bonuses paid, so it increases *less rapidly* than the continuation value produced by a single agent as the number of products made that agent increases. Hence, a generalist agent can be induced to work hard when there is aggregate uncertainty, even if a specialist agent would be unwilling to do so.

3.5.2 Extension 2 - Resistance to Contractual Innovation

Formal contracts unambiguously increase total surplus in this setting, since they lead to smaller specializations and higher effort. Suppose now that the agents have an unanticipated opportunity to make formal contracts available *after* they choose their specializations, perhaps by supporting legal reform or codifying the production process. Although the principal would benefit from having access to formal contracts, the existing agents have already chosen to be generalists and so might resist any reform that leads to the entry of specialist competitors.[11]

To make this argument formal, I make several restrictive assumptions. Suppose that *after* the agents have entered the market and chosen their specializations, they have an unanticipated opportunity to make y_t contractible for the principal. For instance, the agents might be able to codify their knowledge of the production process, develop an internal auditing scheme that could be appropriated by the principal, or push for legal reform to eliminate corruption in the courts. Importantly, once this contracting technology is generated, it is a public good in that the principal can write formal contracts with any agent.

In order for the agents to be motivated to resist legal reform, they must have some stake in the relationship in the sense that each expects to earn positive profits whenever he produces. Therefore, I restrict attention to relational contracts in which each agent earns a fraction $\alpha \in (0, 1)$ of the total expected surplus if he is called upon to produce, and earns 0 otherwise. To simplify the argument, I assume that $F_E = 0$ and consider the optimal relational contract with the minimal number of entrants.

Formally, consider the following additions to the timing at the end of $t = 0$:

1. Legal reform becomes available: unless the current agents pay $F_R > 0$, output y_t, entry, and specializations become contractible. Assume that *no players anticipate this stage* before it occurs.

2. If legal reform occurs, the principal writes a long-term formal contract with existing and potential new entrants, who then choose whether or not to enter at cost $F_E = 0$.

[11]Unlike Baker, Gibbons, and Murphy (1994), the introduction of formal contracts in this extension has an unambiguously positive effect on social welfare.

3. If legal reform does not occur, players continue according to an optimal relational contract.[12]

These modifications to the game are *ad hoc*—they are meant to represent a market that is ripe for legal reform after operating for a long time. To make this point in the starkest way, assume the specialization cost $\gamma(\mu)$ is as in Section 3.4.2.

Proposition 18 *Suppose that*

$$\gamma(\mu) = \left\{ \begin{array}{ll} 0, & \mu \leq \frac{1}{M} \\ \gamma, & \mu > \frac{1}{M} \end{array} \right\}$$

Assume $\delta \in [\underline{\delta}, \bar{\delta})$, *where* $\bar{\delta}$ *and* $\underline{\delta}$ *are defined in Proposition 16, and consider only relational contracts of the following form: there exists an* $\alpha > 0$ *such that* $\forall t, h^{t-1}$, *agent i's payoff is at least* $\alpha E\left[\pi_{i,t}^{TOT} | h_i^{t-1}, \phi_t, 1\{x_t = i\}\right] \geq 0$.[13] *In the optimal relational contract in this class, legal reform does not occur if* $F_R \leq \alpha(y_{HP} - c - \gamma)$.

Proof: As in Proposition 16, one agent enters the market at the beginning of the game and specializes in $\mathcal{P}_1 = [0, 1]$. Consider the continuation equilibrium if legal reform occurs. Because $(y_{HP} - c - \gamma) < (y_{HP} - c)$ and $F_E = 0$, the principal will write a long-term contract inducing high effort in each period with M new entrants following legal reform. Those M agents enter the market and specialize in disjoint subsets of measure $\mu_i = \frac{1}{M}$. Call agent 1 the original entrant in the market, while agents $2, ..., M + 1$ are the M new entrants following legal reform.

Following legal reform, agent 1 is never allocated production and the principal has no incentive to pay him, so agent 1 earns 0. If legal reform does not occur, the continuation relational contract induces high effort from agent 1, who earns $\alpha(y_{HP} - c - \gamma)$ by the assumption that each agent earns at least $\alpha E[\pi_{i,t}^{TOT} | h_i^{t-1}, \phi_t, 1\{x_t = i\}]$. Therefore, agent 1 is willing to pay any $F_R \leq \alpha(y_{HP} - c - \gamma)$ in order to prevent legal reform. ■

[12]In particular, this assumptions implies that agents are not punished for resisting legal reform. One justification for this assumption is that by successfully resisting legal reform, and agent might ensure that the principal was never even aware of the opportunity for that reform.

[13]That is, if $x_t = i$, then i earns a percent α of the total surplus in the period, and if $x_t \neq i$, then i earns 0.

This example highlights a tension inherent in developing markets that might prevent or delay the introduction of better contracts. The agents in a relational market prize flexibility, which is only valuable if relational contracts are required—if legal reform succeeds, these agents are displaced by a larger group of highly specialized agents. Importantly, the principal *cannot* credibly commit to share the efficiency gains from legal reform. Any such promise would have to satisfy the principal's reneging constraint to be credible, and if formal contracts are available, then the principal's payoff following reneging is too attractive to support these promises.

3.5.3 Extension 3 - The Employment Relationship

In the context of the employment relationship, Proposition 15 suggests that employees and independent contractors (such as consultants) may make systematically different investments in human capital. To make this point precise, I modify the example from Section 3.4.2 so that the "principal" actually represents *two different employers*. At the beginning of the game, agents choose whether to enter at cost $F_E > 0$ and, if they do enter, choose specializations $\mathcal{P}_i$. The fixed cost of a specialization $\mathcal{P}_i$ is given by (3.3), and specializing in a product implies that the agent can make that product for *either* employer.

In each period, *two* products $\phi_t^1, \phi_t^2 \sim U[0,1]$ are independently drawn as required for that period, and for each product ϕ_t^k, the principal asks one agent x_t^k to produce. Critically, each agent can only manufacture a *single* product in each period, so $x_t^1 \neq x_t^2$ unless $x_t^1 = x_t^2 = \emptyset$. For product k, agent x_t^k accepts or rejects production, chooses effort $e_t^k \in \{0,1\}$ at cost ce_t^k, and produces binary output $y_t^k \in \{0, y_H\}$, where the probability of $y_H > 0$ is p if $e_t^k = 1$ and $\phi^k \in \mathcal{P}_{x_t^k}$, and 0 otherwise. The monitoring structure for all variables is just like the baseline game, and transfers between the principal and each agent $w_{i,t}$, $\tau_{i,t}$ are the same as that model.

One interpretation of this set-up is that the principal actually represents two different employers, each of whom has a single task ϕ_t^k that must be accomplished in each period. This interpretation is unusual because these two (presumably independent) employers act as a single player in the game. By modeling them as a single player, I ensure that the employers will always act in their joint best interest in order to highlight an interaction between employment and human capital investments. If I instead modeled the two employers as two separate players, then the optimal equilibrium might entail contagion-style punishments in which one employes serves to disseminate information to the agents when the other employer reneges on a bonus.[14] In order to completely rule out such contagion punishments (which may or may not be realistic in a given setting), I model the employers as a single player.

Agent i is a *contractor* if $\text{Prob}\{x_t^1 = i\}, \text{Prob}\{x_t^2 = i\} \in (0,1)$ $\forall t$ on the equilibrium path, so that i produces either ϕ_t^1 or ϕ_t^2 with positive probability in each period. An agent is an *employee* if he exclusively produces one of ϕ_t^1 or ϕ_t^2, so that $\exists k \in \{1,2\}$ with $\text{Prob}\{x_t^k = i\} = 0$ $\forall t$.

Formalizing this logic, Proposition 19 proves that for some parameters, employees—who are generalists and produce only one of the products $k \in \{1,2\}$—are (non-uniquely) optimal if output is non-contractible, while contractors—who work for both firms and specialize in a small subset of products—are optimal if output is contractible.

Proposition 19 *Consider the example presented in this section. Let*

$$\frac{\delta}{1-\delta}(y_H p - c - \gamma) \geq \frac{c}{p} > \frac{\delta}{1-\delta}(1-(1-\frac{1}{M})^2)(y_H p - c)$$

and suppose that $\max\left\{M\tilde{F}_E,\ y_H p - c - M^2 \tilde{F}_E\right\} < \gamma < \frac{M-1}{M}(y_H p - c)$*. Then:*

1. *If output is contractible, M workers enter the market in any equilibrium. Each specializes in $\frac{1}{M}$ of the interval $[0,1]$ and manufactures products for both principals.*

2. *If output is not contractible, then 2 workers enter in any equilibrium. Each specializes in the interval $[0,1]$.*

Moreover, if output is not contractible, there exists an optimal equilibrium in which each worker works exclusively for a single principal.

Proof: See Appendix 3.7.

[14] For instance, the following chain of events might occur: if principal A reneges on a bonus for agent i, then agent i reneges on principal B. When principal B observes this, she reneges on every other employee, who renege on employer A in turn.

If agent i is a specialist with $\mu(\mathcal{P}_i) = \frac{1}{M}$, then the principal can use that agent to produce either ϕ_t^1 or ϕ_t^2 when one of those products happens to lie in $\mathcal{P}_i$. if M is large, then the probability of *both* $\phi_t^1 \in \mathcal{P}_i$ *and* $\phi_t^2 \in \mathcal{P}_i$ is very small, and so the principal only needs a single agent to specialize in the region $\mathcal{P}_i$. In contrast, if i is a generalist with $\mu(\mathcal{P}_i) = 1$, then they will always be able and needed to produce one of the products, so the principal might assign i to be an *employee* for a single task $k \in \{1, 2\}$. Generalists tend to be optimal when formal contracts are unavailable by a logic similar to the previous sections.

Proposition 19 can be interpreted to suggest a difference between *employees* and *contractors*. *Contractors* (1) produce contractible output, (2) are likely to be highly specialized and very good at their chosen tasks, and (3) work for multiple firms. In contrast, employees (1) produce non-contractible output, (2) are flexible, though perhaps not very efficient at any given task, and (3) are locked into a bilateral relationship with a single firm. In this model the decision whether to hire long-term or short-term workers is driven by the human capital investments required to sustain high effort. Employees and contractors are distinguished both by the differences in the contractibility of their jobs, and by their different *ex ante* specializations. Put another way, the contractibility of output drives the *ex ante* human capital acquired by a worker, and consequently whether that worker operates as a contractor or a employee. If formal labor contracts are not available, employees—who are locked into a bilateral relationship and can flexibly produce whatever product is required of them—might be favored over contractors.

3.6 Conclusion

In closing, I informally discuss a few more applications in the context of this model. First, markets sometimes interact with cultural norms to inhibit trade. For instance, Dixit (2011) points out that in institution-poor environments, firms from developing countries tend to have an advantage over multinationals that are based in developed nations. His explanation is that firms from developing countries understand how to navigate inefficient or corrupt institutions—they know whom to influence to get things done. I present a complementary story: a firm from a developed country might also be *too specialized* to work well in a relational setting. Suppose that a "downstream" domestic firm (a retailer, perhaps) would like to contract with a multinational "upstream" firm. Such a multinational might be a very cost-effective producer of a few inputs, but it lacks the flexibility to be a close partner if formal contracts are not available. As a result, foreign firms face multiple barriers to entry in developing markets: not only do they need to learn how to work around inefficient or corrupt formal institutions, they must also alter their production process to prize adaptability over efficiency.

Second, relational contracts sometimes build upon existing social ties—for instance, a CEO might hire a close relative to supply a service, rather than choosing the most cost-effective producer. In a relational contract, such behavior might be eminently justified, since ostracism and social sanctions are powerful forces that imply the parties can punish one another very harshly if one of them deviates from their agreement, which in turn prevents reneging. In Section 3.5.2, close social ties can also be used to induce an otherwise unwilling agent to accept legal reform rather than resist it. Unlike informal contracts, familial relationships exist independently of business ties. Therefore, social ties provide a "stick" that helps the principal commit to rewarding its suppliers for acquiescing to changes, even when those changes weaker their market relationship. While an outsider might observe inefficient production methods and attribute them to nepotism, closely intertwined business and social networks could also serve a valuable purpose by inducing higher effort and mitigating resistance to efficiency-enhancing institutional changes.

Finally, this model presents a simplistic view of firm specialization and capital acquisition, and it would be worthwhile to expand on the notion of *ex ante* investments and precisely illustrate how they interact with relational contracts. For example, suppliers often choose whether to invest in general or relationship-specific capital. In a static setting, relationship-specific investments can be appropriated ex post; in a repeated game, however, these investments increase lock-in and hence effort provision in a relationship (a point made by Klein and Leffler (1981)). Discussing the hold-up problem in this context requires an assumption about how surplus is split in a relational contract, and hence requires a theory of bargaining in repeated games.

3.7 Appendix: Omitted Proofs

Proof of Proposition 12:

This proof proceeds in four steps. A long-term contract can be written before entry and utility is transferrable, so the optimal contract will set M, $\{\mathcal{P}_i\}$, and $x_t, e_t, \forall t$ to maximize total surplus.

(1) Optimal one-shot contract. Consider the following formal contract $\tau(y)$: let $\bar{y}$ be such that $F(\bar{y}|e=0)-F(\bar{y}|e=1) \in (0,1)$, which must exist because $F(\cdot|e=1) >_{FOSD} F(\cdot|e=0)$. Then $\tau(y) = \frac{c}{F(\bar{y}|e=0)-F(\bar{y}|e=1)}$ if $y \geq \bar{y}$, and otherwise $\tau(y) = 0$. Under this contract, an agent chooses $e = 1$ since

$$\frac{c}{F(\bar{y}|e=0)-F(\bar{y}|e=1)}(1-F(\bar{y}|e=1)) - c \geq \frac{c}{F(\bar{y}|e=0)-F(\bar{y}|e=1)}(1-F(\bar{y}|e=0)).$$

It is an optimal static equilibrium for the principal of offer this formal contract, and so $e_t = 1$ in any round of the repeated game for which $\phi_t \in \mathcal{P}_i$.

(2) If agents choose $e_t = 1$, $\forall t$, then specializations do not overlap. Suppose that $\mu(\mathcal{P}_i \cap \mathcal{P}_k) > 0$, and consider the alternative with $e_t = 1$, $\forall t$, $\hat{\mathcal{P}}_j = \mathcal{P}_j$ $\forall j \neq k$, and $\hat{\mathcal{P}}_k = \mathcal{P}_k \backslash \mathcal{P}_i$. By claim 1, each agent chooses $e_t = 1$ in the static equilibrium with these alternative specializations. If $\hat{\mathcal{P}}_k = \emptyset$, then k need not enter the market, which improves total surplus by $F_E > 0$; If $\hat{\mathcal{P}}_k \neq \emptyset$, then firm k's cost is $\gamma(\hat{\mu}_k) < \gamma(\mu_k)$, also increasing total surplus.

(3) $\mu_i = \mu_j$, $\forall i, j$. Suppose M firms enter the market and choose $\{\mathcal{P}_i\}_{i=1}^{M}$. By claim 1, $e_t = 1$ in the efficient equilibrium. By claim 2, there is no overlap in specializations and so the problem is

$$\min_{\{\mu_i\}} \sum_{i=1}^{M} \mu_i \gamma(\mu_i)$$

subject to

$$\sum_{i=1}^{M} \mu_i \leq 1$$

$\gamma(\cdot)$ is convex and increasing, so $\mu\gamma(\mu)$ is convex. Hence, the solutions to this problem is $\mu_i = \mu$ $\forall i$ and some $\mu \geq 0$.

(4) Optimal entry. The measure μ is chosen to solve

$$\max_{\mu} M\mu(E[y|e=1]-c-\gamma(\mu))$$

subject to the constraint $M\mu \leq 1$. But $M\mu < 1$ cannot be a solution, since $\mu(E[y|e=1]-c-\gamma(\mu))$ is strictly increasing in μ. Hence, the optimal equilibrium sets $\mu = \frac{1}{M}$ and $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$, $\forall i, j$.

The efficient number of entrants M solves

$$\max_{M}(E[y|e=1]-c-\gamma(\frac{1}{M}))-F_E M$$

with first-order condition $\gamma'(\frac{1}{M}) = M^2 F_E$.

The left- and right-hand sides of this expression are strictly decreasing and increasing in M, respectively, so some (non-integer) M^* equates the two sides. Total surplus is continuous in M and M^{FB} must be a natural number, so M^{FB} is either the floor or the ceiling of M^*. ■

Proof of Lemma 11:

A deviation in any variable other than e_t in round t is publicly observed and hence can be punished by breakdown in the market, which leaves every player at the min-max payoff 0. Let σ be an optimal equilibrium. If $e_t = 0$ at every history of σ, then it can be trivially replicated by a stationary contract. Otherwise, let (h^{t-1}, n_t) be a history on the equilibrium path immediately following $e_t = 1$ and output y_t. Let $\tilde{\tau}(y_t) = E_\sigma\left[\tau_{x_t}|h^{t-1}, n_t\right]$ be the expected transfer to the producing agent in this history, and $U_i(y_t) = E_\sigma\left[U_i|h^{t-1}, n_t\right]$ $\forall i \in \{0, ..., M\}$ the expected continuation surplus (where $i = 0$ is the principal).

For $e_t = 1$ in equilibrium,

$$(1-\delta)\tilde{\tau}(y_t) + \delta U(y_t) - c \geq (1-\delta)\tilde{\tau}(y_t) + \delta U(y_t). \tag{3.4}$$

For transfer $\tilde{\tau}(y_t)$ to be supported in equilibrium,

$$\begin{aligned} (1-\delta)\tilde{\tau}(y_t) &\leq \delta U_0(y_t) \\ -(1-\delta)\tilde{\tau}(y_t) &\leq \delta U_{x_t}(y_t) \end{aligned} \qquad \forall y_t \tag{3.5}$$

since otherwise at least one player would prefer to deviate, not pay a bonus, and earn min-max payoff 0 in the rest of the game.

Define $\tau(y_t) = \tilde{\tau}(y_t) + \frac{\delta}{1-\delta} U_{x_t}(y_t)$, and note that $\tau(y_t) \geq 0$ by (3.5). Consider the following stationary equilibrium: At $t = 0$, M firms enter and specialize as in σ, and $w_{i,0} = F_E$. $\forall t > 0$ on the equilibrium path, x_t is the agent with the *smallest specialization* μ_{x_t} such that $\phi_t \in \mathcal{P}_{x_t}$, $w_{x_t,t} = c + \gamma(\mu_{x_t}) - E\left[\tau(y_t)\right]$, and $w_{i,t} = 0\ \forall i \neq x_t$. Agent x_t accepts ($d_t = 1$) and chooses $e_t = 1$. Following output y_t, $\tau_{x_t} = \tau(y_t) \geq 0$ and $\tau_i = 0\ \forall i \neq x_t$. Any deviation is immediately punished by reversion to the static equilibrium.

At $t = 0$, the agent is indifferent between entering or not. The principal is willing to pay $w_{i,0} = F_E$ because MF_E is smaller than the principal's total continuation surplus in an optimal equilibrium. $\forall t > 0$, each agent earns 0. For each $\phi_t \in [0,1]$, let $\mu^m(\phi_t)$ be the smallest specialization $\mu(\mathcal{P}_i)$ such that $\phi_t \in \mathcal{P}_i$. Then the principal earns

$$U_0 = \int_0^1 1\{\exists i \in \{1,...,M\} \text{ s.t. } \phi_t \in \mathcal{P}_i\} E[y_t - c - \gamma(\mu^m(\phi_t))|e_t = 1] d\phi_t$$

from every on-path history in this relational contract.

The principal or agent x_t is willing to pay $w_{x_t,t}$ because doing so earns that player at least 0, while failing to do so earns them 0. $\forall t$, agent x_t is willing to work hard because (3.4) holds for $\tau(y_t)$. The principal is willing to pay $\tau(y_t) \geq 0$ because

$$(1-\delta)\tau(y_t) = (1-\delta)\tilde{\tau}(y_t) + \delta U_{x_t}(y_t) \leq \delta(U_0(y_t) + U_{x_t}(y_t)) \leq \delta U_0 \tag{3.6}$$

where the first inequality follows from (3.5) and the second inequality is because U_0 is the maximum expected surplus attainable given entry M and specializations $\{\mathcal{P}_i\}$. Thus, we have found a stationary equilibrium that produces at least as much total surplus as the posited original equilibrium. ■

Proof of Proposition 14:

Let $V(M, \{\mathcal{P}_i\})$ be the expected total surplus in the optimal stationary equilibrium with M firms and specializations $\{\mathcal{P}_i\}$. By emma 11, there exists a stationary optimal continuation equilibrium in which the principal earns total surplus:

$$V(M, \{\mathcal{P}_i\}) = \begin{cases} \int_0^1 1\{\exists i \text{ s.t. } \phi_t \in \mathcal{P}_i\} E[y_t - c - \gamma(\mu^m(\phi_t)) | e_t = 1] d\phi_t - MF_E & \text{if (3.6) holds} \\ -MF_E & \text{otherwise} \end{cases}.$$

Suppose that at the beginning of the game, the total-surplus maximizing number of agents M^{Pub} enter and choose optimal specializations $\{\mathcal{P}_i\}_{i=1}^M$. The principal pays F_E to each entrant and play continues according to the optimal stationary equilibrium, while any deviation is punished by a breakdown of the market. Since $V(M, \{\mathcal{P}_i\}) \geq 0$, the principal would rather pay MF_E than suffer relational breakdown, so firms are willing to enter because they earn 0 and hence this is an equilibrium. Therefore, it suffices to find the M and $\{\mathcal{P}_i\}$ that maximize V.

Suppose that $\mu(\mathcal{P}_i \cap \mathcal{P}_k) \neq 0$, and consider the alternative equilibrium in which agent k does not specialize in $\mathcal{P}_i \cap \mathcal{P}_k$. If $\mathcal{P}_k \backslash \mathcal{P}_i \subseteq \cup_{l \neq k} \mathcal{P}_l$, then agent k can exit the market without affecting total surplus; otherwise, the fixed cost $\gamma(\mu_k)$ of agent k strictly decreases. In either case, $V(M, \{\mathcal{P}_i\})$ strictly increases, so it must be that $\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$ in the optimal equilibrium.

If $\nexists$ $(M, \{\mathcal{P}_i\})$ such that $V(M, \{\mathcal{P}_i\}) > 0$, then $M = 0$ is optimal in equilibrium. Otherwise, $\mu_i = \frac{1}{M}$ because $\gamma(\mu)$ is convex, as in Proposition 12, so the optimal equilibrium maximizes total surplus subject to satisfying (3.6). Ignoring this constraint, M^{FB} maximizes surplus by definition. The constraint (3.6) is relaxed as M increases, so the optimal number of entrants is either $M = 0$ or $M \geq M^{FB}$. ■

Proof of Lemma 12:

Towards contradiction. Suppose that

$$-(1-\delta) E_\sigma \left[\tau_{i,t} | h_i^{t-1}, \mathcal{I}_i(n_t)\right] > \delta E_{\sigma^*} \left[U_i(h^t) | h_i^{t-1}, \mathcal{I}_i(n_t)\right].$$

Consider the following deviation for agent i: do not pay this transfer, and in future rounds pay no transfers, never accept production, and never exert effort. This deviation yields payoff 0, so is profitable given agent i's information set $\left(h_i^{t-1}, \mathcal{I}_i(n_t)\right)$.

Suppose now that

$$(1-\delta)E_\sigma\left[\tau_{i,t}|h_i^{t-1}, \mathcal{I}_i(n_t)\right] > \delta E_{\sigma^*}\left[U_0^i(h^t)|h_i^{t-1}, \mathcal{I}_i(n_t)\right].$$

Then it must be that $\tau_{i,t} > U_0^i(h^t)$ for at some history in this information set. Consider the following deviation for the principal: do not pay $\tau_{i,t}$ in round t. In the continuation game, replicate equilibrium play following history (h^{t-1}, n_t), with the exception that $w_{i,t'} = \tau_{i,t'} = 0$ in every future period whenever $x_t = i$ is specified by the equilibrium, instead play $x_t = \emptyset$. Because $w_{i,t'}$, $\tau_{i,t'}$, and $x_t = \emptyset$ are not observed by any agents $j \neq i$, all other agents do not detect a deviation. The principal earns no less than 0 from agent i in each future period, so her continutation surplus is bounded below by

$$\delta E\left[\sum_{k\neq i} U_0^i(h^t)|h^{t-1}, n_t\right] > -(1-\delta)E_\sigma\left[\tau_i|h^{t-1}, n_t\right] + \delta E\left[\sum_{i=1}^{N} U_0^i(h^t)|h^{t-1}, n_t\right]$$

and hence this deviation is profitable. ■

Proof of Lemma 13:

Suppose there exists an equilibrium σ and an on-path history (h_i^{t-1}, n_t) immediately preceding effort e_t such that $x_t = i$ and $e_t = 1$. Then

$$\begin{array}{c} E_\sigma\left[(1-\delta)\tau_{x_t,t} + \delta U_{x_t}(h^t) \mid h_{x_t}^{t-1}, \mathcal{I}_{x_t}(n_t), e_t = 1\right] - c \geq \\ E\left[(1-\delta)\tau_{x_t,t} + \delta U_{x_t}(h^t) \mid h_{x_t}^{t-1}, \mathcal{I}_{x_t}(n_t), e_t = 0\right] \end{array} \tag{3.7}$$

for $e_t = 1$ to be incentive compatible. Define $\tilde{\tau}(y_t) = E\left[\tau_{x_t} + \frac{\delta}{1-\delta}U_{x_t}|h_{x_t}^{t-1}, \mathcal{I}_i(n_t), y_t\right]$.

By Lemma 12, the following inequalities must hold for τ_i to be incentive compatible $\forall y_t$:

$$(1-\delta)E_\sigma\left[\tilde{\tau}|h_i^{t-1}, \mathcal{I}_i(n_t, y_t)\right] \leq \delta E_{\sigma^*}\left[U_0^i(h^t) + U_i(h^t)|h_i^{t-1}, \mathcal{I}_i(n_t, y_t)\right]$$

$$-(1-\delta)E_\sigma\left[\tilde{\tau}|h_i^{t-1},\mathcal{I}_i(n_t,y_t)\right]\leq 0$$

By definition,

$$U_0^i(h^{t-1},n_t)+U_i(h^{t-1},n_t)=E_\sigma\left[\sum_{t'=t+1}^{\infty}\delta^{t'-t-1}(u_i+\pi_0^i)|h^{t-1},n_t\right]=$$
$$\leq\mu_i\left(E\left[y_t|e=1\right]-c-\gamma(\mu_i)\right)$$

because $\mu_i\left(E\left[y_t|e=1\right]-c-\gamma(\mu_i)\right)$ is the maximum surplus produced by agent i given $\mathcal{P}_i$. $\tilde{\tau}(y)$ satisfies (3.7) and so is an incentive scheme that induces high effort, and moreover

$$(1-\delta)\left(\sup_y\tilde{\tau}(y)-\inf_y\tilde{\tau}(y)\right)\leq\mu_i\left(E\left[y-c-\gamma(\mu_i)|e=1\right]\right).$$

Hence, whenever agent i chooses $e_t=1$ in equilibrium, there exists an incentive scheme $\tilde{\tau}(y)$ that satisfies (3.7) and (3.2). ■

Proof of Corollary 4:

Suppose that $\mu(\mathcal{P}_i\cap\mathcal{P}_k)=0$ for every $i,k\in\{1,...,M\}$. By Lemma 13, agent i exerts high effort in equilibrium only if (3.2) holds; I construct a stationary equilibrium in which this condition is also sufficient.

Fix some $\tau(y)$ that minimizes $\sup_y\tau(y)-\inf_y\tau(y)$ among all incentive schemes such that (3.7) holds, and assume without loss that $\inf_y\tau_i(y)=0$. Consider the following strategy profile: the principal allocates ϕ_t to i such that $\phi_t\in\mathcal{P}_i$ (if multiple agents are available—which occurs with probability 0—then the principal chooses the lowest-numbered agent that has the smallest μ_i among those for which (3.2) holds). If (3.2) holds for x_t, then $w_{x_t,t}=c+\gamma(\mu_i)-E\left[\tau_i(y)|e=1\right]$, $d_t=e_t=1$, and $\tau_{x_t,t}=\tau(y_t)\geq 0$ following output y_t. If (3.2) does not hold, then the principal pays $w_{x_t,t}=\tau_{x_t,t}=0$ and the agent rejects production. In either case, $w_{i,t}=\tau_{i,t}=0$ $\forall i\neq x_t$. After a commonly observed deviation, the entire market breaks down. After a deviation observed by agent i and the principal, agent i thereafter rejects production, $w_{i,t}=\tau_{i,t}=0$ in each period, and $x_t=\emptyset$ whenever i is the only agent able to produce. The principal plays the on-path actions for all agents $j\neq i$.

I claim that this strategy profile is an equilibrium. Let $\mathcal{U}$ be the set of agents for which

(3.2) holds. In each period on the equilibrium path, every agent earns 0 and so the principal earns the total expected surplus for that period. The principal has no profitable deviation from the allocation rule on the equilibrium path, since either there is only one available agent or the principal is allocating production to the agent that maximizes expected surplus in that period. If (3.2) does not hold, then play is a mutual best response both on and off the equilibrium path. If (3.2) holds, then the agent is indifferent between accepting and rejecting production and willing to choose $e = 1$ by construction of $\tau_i(y)$. Whoever is responsible for paying $w_{x_t,t}$ is willing to do so because paying $w_{x_t,t}$ yields higher continuation surplus by construction. If the principal does not pay $\tau(y_t)$, then she no longer earns any surplus from agent i. Because $\mu(\mathcal{P}_i \cap \mathcal{P}_k) = 0$ for every $i, k \in \{1, ..., M\}$, the total loss from this deviation is

$$\delta\mu_i \left(E[y|e = 1] - c - \gamma(\mu_i)\right)$$

and the total gain from the deviation is $\tau_i(y)$. This is not a profitable deviation because (3.2) holds.

Off the equilibrium path, strategies among those that have observed the deviation are a mutual best response by construction. Pay among those who have not observed are likewise a mutual best-response. Therefore, this strategy profile is an equilibrium that induces high effort from every agent i whose $\mathcal{P}_i$ satisfies (3.2). It is optimal because (3.2) is a necessary condition for high effort. ■

Proof of Corollary 5:

From Lemma 13, agent i is only willing to choose $e = 1$ in an equilibrium if (3.2) holds. The same set of transfer schemes induce high effort regardless of μ_i or δ. Define $\mu^*(\delta)$ to solve the following

$$\inf_{\{\{\tau\}|\{\tau\} \text{ induces } e = 1\}} \left\{\sup_y \tau(y) - \inf_y \tau(y)\right\} = \frac{\delta}{1-\delta}\mu^*(\delta)\left(E[y|e = 1] - c - \gamma(\mu^*(\delta)\right)$$

whenever this is well-defined. Because $\sup \tau_i - \inf \tau_i \geq c > 0$, $\mu^*(\delta) > 0$ defines the *minimum* measure of specialization required to induce an agent to choose $e = 1$. Because $\gamma(\mu)$ is continuous and $\mu(E[y|e=1]-c-\gamma(\mu))$ is increasing in μ for $\mu \in [0,1]$, $\mu^*(\delta)$ exists, is continuous, and satisfies $\lim_{\delta\to 1}\mu^*(\delta)=0$ and $\lim_{\delta\to\hat{\delta}}\mu^*(\delta)=1$ for some $\hat{\delta}<1$.

If $\mu(\mathcal{S}_i\cap\mathcal{S}_k)=0$, then a stationary contract is optimal and this condition is sufficient to induce high effort by Corollary 4, so $e=1$ at every history in the optimal equilibrium iff $\mu_i\geq\mu^*(\delta)$, $\forall i$. ■

Proof of Proposition 15:

To prove this result, I introduce an equilibrium with the desired properties and show that it is optimal.

First, fix the number of entrants $M\leq M^{FB}$, and consider the continuation game. Suppose that $\mu^*(\delta)M\leq 1$. I first claim that the optimal equilibrium with M firms satisfies $\mu(\mathcal{P}_i\cap\mathcal{P}_k)=0$ and $\mu(\mathcal{P}_i)=\mu_M$, $\forall i,k$ and some $\mu_M\in[0,1]$. Suppose not, and assume that the optimal equilibrium generates total surplus v^*.

Let $\tilde{v}$ is the surplus generated under the same allocation rule if $e_t=1$ in each period. Then $v^*\leq\tilde{v}$, where this inequality hold strictly if v^* ever has $e_t=0$ on the equilibrium path.

$$\tilde{v}=E_\sigma\left[\sum_{t=0}^{\infty}\delta^t(1-\delta)\sum_{i=1}^{M}\int_{\mathcal{P}_i}1\{x_t=i\}\left\{E[y|e=1]-c-\gamma(\mu(\mathcal{P}_i))\right\}d\phi_t\right].$$

In turn, $\tilde{v}$ is dominated by the strategy profile in which (1) $e_t=1$ in each period, (2) specializations do not overlap, and (3) the principal allocates to the unique producer in each period. Define $\{\mathcal{P}_i^{NO}\}$ as the set of disjoint specializations created by removing $\mathcal{P}_i\cap\mathcal{P}_k$ from one of i,k's specializations. Let v_{NO} be the resulting surplus, so $v_{NO}\geq\tilde{v}$ with strict inequality if $\mu(\mathcal{P}_i\cap\mathcal{P}_k)>0$ for some $i\neq k$.

$$v_{NO}=E_\sigma\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t\left\{\mu(\mathcal{P}_1^{NO}\cup\ldots\cup\mathcal{P}_M^{NO})\{E[y|e=1]-c\}-\sum_{i=1}^{M}\mu(\mathcal{P}_i^{NO})\gamma(\mu(\mathcal{P}_i^{NO}))\right\}\right].$$

Because $\mu\gamma(\mu)$ is strictly convex, v_{NO} is dominated by the surplus generated if all agents choose $e_t=1$ and have equally-sized, non-overlapping specializations $\mu^{ES}=\frac{1}{M}\mu(\mathcal{P}_1^{NO}\cup\ldots\cup\mathcal{P}_M^{NO})$. This alternative strategy profile generates total surplus v_{ES}, where $v_{ES}\geq v_{NO}$ and this inequality holds strictly if specializations $\{\mathcal{P}_i^{NO}\}$ are not of equal size.

$$v_{ES} = M\mu^{ES}\left(E[y|e=1] - c - \gamma(\mu^{ES})\right).$$

Finally, recall that $\mu(E[y|e=1]-c-\gamma(\mu))$ is increasing in μ. Therefore, v_{ES} is dominated by the surplus generated if each agent picks $e=1$, has an equally-sized specialization, and specializations collectively cover $[0,1]$. Define

$$v_{OPT} \equiv \sum_{t=0}^{\infty}\delta^t\left(E[y|e=1] - c - \gamma(\frac{1}{M})\right)$$

as the surplus in this alternative.

By Lemma 13 and Corollary 4, because $\frac{1}{M} \geq \mu^*(\delta)$, v_{OPT} can be generated in an equilibrium of the repeated game, provided that (1) $\mu_i = \frac{1}{M}$ with $\mathcal{P}_i \cap \mathcal{P}_k = \emptyset$, and (2) $M \leq M^{FB}$ firms enter the market. It remains to show that such an equilibrium exists.

Consider the following strategies:

1. M agents enter the market, labelled $\{1,...,M\}$.
2. agent $i \in \{1,...,M\}$ specializes in the interval $\mathcal{P}_i = [\frac{i-1}{M}, \frac{i}{M}]$.
3. $i \in \{1,...,M\}$, $w_{i,0} = \max\{0, F_E - U_i\}$, where U_i is the surplus earned by agent i in the continuation game.
4. Play continues as in the optimal stationary contract with agents $\{1,...,M\}$.
5. If either fewer or more than M firms enter, or specializations differ from those specified in step 2, then continuation play specifies $w_i = \tau_i = 0$ and $e = 0$, $\forall i$, at every history.

Because both specializations and the set of entrants are public knowledge, the punishment strategy specified in step 5 is feasible. Hence, the number of entrants will not exceed M, and entrants will specialize in the specified interval. If agent i expects to be paid $w_{i,0} = \max\{0, F_E - U_i\}$, then she weakly prefers to enter the market, since the surplus from entering is $w_i + U_i - F_E \geq 0$. Moreover, the principal is willing to pay w_i: if she does not, her loss is $\frac{1}{M}\left(E[y|e=1] - c - \gamma(\frac{1}{M})\right) - U_i$, which is larger than w_i because $F_E \leq \frac{1}{M}(E[y|e=1] - c - \gamma(\frac{1}{M}))$ for $M \leq M^{FB}$.

If $\mu^*(\delta) \leq \frac{1}{M}$, the total surplus generated by the efficient equilibrium is

$$E[y|e=1] - c - \gamma(\frac{1}{M}) - MF_E$$

This expression is increasing in M for $M < M^{FB}$ by definition of M^{FB}. Thus, the number of entrants in the optimal equilibrium is bounded below by $M(\delta) = \min\left\{M^{FB}, \text{floor}\{\frac{1}{\mu^*(\delta)}\}\right\}$.

To prove statements 2 and 3 of the Proposition, suppose that $M\mu^*(\delta) = 1$. For $M \leq M^{FB}$, such a δ exists because $\mu^*(\delta)$ varies continuously on $[0,1]$. By the argument above, at least M firms enter the market. If exactly M firms enter the market and specialize in disjoint subsets of length $\mu^*(\delta)$, then total surplus is

$$\sum_{t=0}^{\infty} \delta^t \{y_H p - c - \gamma(\frac{1}{M})\} - MF_E \tag{3.8}$$

For $M < M^{FB}$, (3.8) is increasing in M. $\mu^*(\delta) = \frac{1}{M}$, so $\gamma(\mu_i) \geq \frac{1}{M}$ by Corollary 5 and hence (3.8) is the *largest surplus* that can be generated in an equilibrium with discount rate δ. Therefore, every efficient equilibrium has M firms enter the market and specialize in subsets $\frac{1}{M}$ of the unit interval. Hence, when $M\mu^*(\delta) = 1$, exactly M firms enter.

Now, consider a small increase in δ. At least M agents will enter the market by statement 1, so suppose that $M + K$ agents enter. Then total surplus is no more than

$$\sum_{t=0}^{\infty} (1-\delta)\delta^t \{y_H p - c - \gamma(\mu^*(\delta))\} - (M+K)F_E \tag{3.9}$$

But μ^* and γ are both continuous, and $F_E > 0$. Hence, there exists some open set $\Delta(M)$ such that for $\delta \in \delta(M)$, (3.8) is strictly larger than (3.9) and so the optimal equilibrium entails M agents.■

Proof of Proposition 17:

I first claim that first-best can be attained if and only if

$$\frac{c}{p} \leq \frac{\delta}{1-\delta}(y_H p - c).$$

Suppose this inequality holds, and consider the following strategy profile:

1. M agents enter the market, each specializing in a single product with $\mathcal{P}_i \cap \mathcal{P}_k = \emptyset \; \forall i, k$. The principal pays each entrant $w_{i,0} = F_E$.

2. In each period, the principal allocates production to the unique i with $\phi_t \in \mathcal{P}_i$.

3. The principal offers wage $w_{i,t} = 0$; the agent accepts and chooses $e_t = 1$.

4. If $y_t = 0$, $\tau_{x_t,t} = 0$; if $y_t = y_H$, $\tau_{x_t,t} = \frac{c}{p}$. $\tau_{i,t} = 0$ for all $i \neq x_t$.

5. If the principal deviates on x_t, then every agent rejects and chooses $e_t = 0$ in each future round, and $w_{i,t} = \tau_{i,t} = 0 \; \forall i$. If either principal or agent deviates on $w_{i,t}$ or $\tau_{i,t}$, then whenever $x_t = i$, the agent rejects and chooses $e_t = 0$ in every future round, and $w_{i,t} = \tau_{i,t} = 0$ for that agent, but otherwise play continues as on the equilibrium path.

The principal earns $M(y_H p - c)$ in continuation surplus following any on-path history, and the agent earns 0. The principal loses $(y_H p - c)$ following a deviation in $w_{i,t}$ or $\tau_{i,t}$, and earns 0 following a deviation in x_t. Each agent is indifferent between entering the market or not, and the principal is willing to pay $w_{i,0} = F_E$ because $y_H p - c - F_E > 0$. The principal and agent are trivially willing to pay the wage, and the principal is willing to pay $\tau_{x_t,t}$ so long as

$$\frac{c}{p} \leq \frac{\delta}{1-\delta}(y_H p - c)$$

which holds by assumption. The agent is willing to work hard because $p\frac{c}{p} - c = 0$. Players mutually best-respond off the equilibrium path, so this describes an equilibrium that attains first-best.

Suppose instead that $\frac{c}{p} > \frac{\delta}{1-\delta}(y_H p - c)$, and fix a strategy profile σ that attains first-best. By an argument similar to Lemmas 12 and 13, it must be that

$$(1-\delta)\left(\sup_y E_\sigma\left[\tau_i | h_i^{t-1}, \mathcal{I}_i(n_t, y_t)\right] - \inf_y E_\sigma\left[\tau_i | h_i^{t-1}, \mathcal{I}_i(n_t, y_t)\right]\right) \leq \delta\left(y_H p - c\right).$$

Plugging in the optimal contract, $\frac{c}{p} \leq \frac{\delta}{1-\delta}(y_H p - c)$, which contradicts the assumption.

Next, I claim that there exists a set of δ for which $\frac{c}{p} > \frac{\delta}{1-\delta}(y_H p - c)$ but a single agent with $\mathcal{P}_1$ can be motivated to work hard. By the argument above, any firm specializing in one product will never exert high effort under this condition. Therefore, every firm entering must specialize in at least two products, and so the maximum surplus in the market is $M(y_H p - c - \gamma)$. Using an argument similar to Lemmas 12 and 13. A stationary relational contract can be constructed similarly to above, but with a single entrant who specializes in $\mathcal{P}_1 = \{1, ..., M\}$ and $w_{1,t} = \gamma$ in each period.

Define an incentive scheme as $\{\tau^M, \tau^{M-1}, ..., \tau^0\}$, where τ^k is the bonus payment if y_H is observed for k products. This incentive scheme can be implemented in the stationary relational contract if and only if

$$\sup_k \tau^k - \inf_k \tau^k \leq \frac{\delta}{1-\delta} M(y_H p - c - \gamma)$$

I claim that there exists a threshold γ^* such that if $\gamma \leq \gamma^*$, high effort can be supported with a single firm but not with M firms. It suffices to show that for some δ such that $\frac{c}{p} > \frac{\delta}{1-\delta}(y_H p - c)$, there exists a bonus scheme $\{\tau^i\}$ that induces an agent to work hard on all M products and satisfies $\sup_i \tau^i - \inf_i \tau^i \leq \frac{\delta}{1-\delta} M(y_H p - c - \gamma)$.

To begin, consider the incentive scheme $\tilde{\tau}^k = k\frac{c}{p}$. Define Ψ_K as the event that $e_{k,t} = 1$ for K of the M products. Then under the incentive scheme $\{\tilde{\tau}^k\}$, I first argue that

$$\frac{c}{p} E[\#\{y = y_H\} \mid \Psi_K] - Kc$$

is weakly increasing in K. Fix every outcome except y_k, and denote this vector of $M-1$ outputs y_{-k}. For *any* fixed y_{-k}, the value of choosing $e_k = 1$ is $\frac{c}{p}p - c \geq 0$. Therefore, *regardless of the outcome of the other tasks*, the agent always weakly prefers to pick $e_k = 1$ under the specified contract, so the agent's payoff is weakly increasing in K.

Now, consider the IC constraints for an arbitrary incentive scheme $\{\tau^k\}$. The agent is willing to choose $e_{k,t} = 1$, $\forall k, t$, if

$$E_k[\tau^k \mid \Psi_M] - Mc \geq E_k[\tau^k \mid \Psi_K] - Kc, \forall K \in \{0, ..., M\}$$

Consider the incentive scheme $\tau^k = \frac{c}{p}k$, $\forall\ k < M-1$, $\tau^{M-1} > \frac{c}{p}(M-1)$, and $\tau^M < \frac{c}{p}M$, where τ^{M-1} and τ^M are chosen so that

$$E_k[\tau^k \mid \Psi_M] \geq \frac{c}{p}E[\#\{y = y_H\} \mid \Psi_M]$$

Such an incentive scheme exists so long as $p < 1$. I claim that this alternative contract continues to satisfy the IC constraints. Since $\tau^k = \tilde{\tau}^k$ for all $k < M-1$ and $\tau^{M-1} > \tilde{\tau}^{M-1}$, $E_k[\tau^k \mid \Psi_K] - Kc$ is increasing for all $K \leq M-1$. Therefore, it suffices to show that

$$E_k[\tau^k \mid \Psi_M] - Mc \geq E_k[\tau^k \mid \Psi_{M-1}] - (M-1)c$$

or

$$\sum_{k=0}^{M} \begin{pmatrix} M \\ k \end{pmatrix} p^k(1-p)^{M-k}\tau^k \geq \sum_{k=0}^{M-1} \begin{pmatrix} M-1 \\ k \end{pmatrix} p^k(1-p)^{M-1-k}\tau^k + c$$

strictly slackens when τ^{M-1} increases, since this slack can then be used to decrease τ^M.

The coefficient on τ^{M-1} is $\frac{M!}{(M-1)!}p^{M-1}(1-p)$ on the left-hand side, and $\frac{(M-1)!}{(M-1)!}p^{M-1} = p^{M-1}$ on the right-hand side. Thus, increasing τ^{M-1} strictly relaxes the IC constraint if $p < \frac{M-1}{M}$. Under this parameter restriction, $\sup_k \tau^k < M\frac{c}{p}$ when a single agent enters the market and manufactures every product. Hence, for $\gamma > 0$ sufficiently small, there exists an open interval of δ for which $M\frac{c}{p} > \frac{\delta}{1-\delta}(y_H p - c)$ but

$$\sup_i \tau^i - \inf_i \tau^i \leq \frac{\delta}{1-\delta}(y_H p - c - \gamma)$$

On this interval, a single generalist firm enters the market in the optimal equilibrium. ∎

Proof of Proposition 19:

First, suppose that output is not contractible. I claim that an agent with specialization $\mu_i \leq \frac{1}{M}$ will never choose $e_t = 1$ in a relational contract. As in the baseline model, following any deviation in $\tau_{i,t}$ with agent i, the principal can always allocate business as if no deviation has occurred, but set $\tau_{i,t} = w_{i,t} = 0$ and $x_t^k = \emptyset$ whenever he would have chosen $x_t^k = i$ on the equilibrium path. Similarly, agent i can always set $\tau_{i,t} = w_{i,t} = 0$ and $d_t = 0$. As a result, in equilibrium it must be that

$$(1-\delta)\tau_{i,t} \leq \delta E\left[U_0^i(h^t)|h^{t-1}, n_t\right]$$

$$-(1-\delta)\tau_{i,t} \geq \delta E\left[U_i(h^t)|h_i^{t-1}, \mathcal{I}_i(n_t)\right]$$

for any (h^{t-1}, n_t) immediately following the realization of output y_t. Worker i can produce only once per round and only if $\phi_t^k \in \mathcal{P}_i$ for some $k \in \{1,2\}$. Therefore, agent i is able to product at least one of the products in a period with probability $(1-(1-\mu_i)^2)$. By an analogous argument to Lemma 13, in any PBE in which agent i chooses $e_t = 1$, it must be that

$$\frac{c}{p} \leq \frac{\delta}{1-\delta}(1-(1-\mu_i)^2)(y_H p - c). \tag{3.10}$$

By assumption, this inequality does not hold for $\mu_i \leq \frac{1}{M}$ because the right-hand side of (3.10) is increasing in μ_i.

Hence, in any optimal equilibrium every entrant specializes in $\mu_i > \frac{1}{M}$. Then the maximum feasible surplus is $2(y_H p - c - \gamma) - 2\tilde{F}_E$, which is attained in the following stationary equilibrium:

1. Two workers enter and specialize in $\mu(\mathcal{P}_1) = [0,1]$, $\mu(\mathcal{P}_2) = [0,1]$. $w_{i,0} = \tilde{F}_E$

2. ϕ_t^1 is always assigned to $i=1$, ϕ_t^2 is always assigned to $i=2$.

3. $w_{i,t} = \gamma$ for both agents. Agents accept and choose $e=1$.

4. If $y_t = y_H$, $\tau_i = \frac{c}{p}$; otherwise, $\tau_i = 0$.

The agents are willing to follow the equilibrium entry and effort decisions by construction. The principal is willing to pay $w_{i,t}$ for reasons similar to those argued in Proposition 17, and willing to pay τ_i so long as

$$\frac{c}{p} \leq \frac{\delta}{1-\delta}(y_H p - c - \gamma)$$

which holds by assumption. Note that each agent is an employee in this relational contract.

Suppose instead that output is contractible, so that every agent chooses $e_t = 1$ in every period. Note that it is weakly optimal if $\forall$ agent i, $\mu_i \in \{\frac{1}{M}, 1\}$. Call i a "specialist" if $\mu_i = \frac{1}{M}$ and a "generalist" if $\mu_i = 1$. Because only two products are required in equilibrium, there will be at most two generalists in the market. Fix the number of generalists in the market and consider the number of specialists. Suppose it is optimal for at least one specialist to enter the market, which is implied by the condition $\gamma > M\tilde{F}_E$. Then it is optimal for at least M specialists to enter the market, because these specialists can choose disjoint specializations and each generate the same additional surplus as the first specialist. If M specialists enter the market, then it is not optimal to have two generalists in the market.

More than $2M$ specialists will never enter the market in the optimal equilibrium. Fixing the number of generalists $G \in \{0, 1, 2\}$, suppose that $2M \geq K > M$ specialists enter the market. Because only two products are required in each period, it is never optimal for more than two specialists to produce the same product. The total measure of production is $2 \geq \frac{K}{M} > 1$. Let μ^1 be the measure of products that have exactly one producer and μ^2 the measure of products that have two producers. Then surplus can be written

$$(1 - \mu^1 - \mu^2)G(y_H p - c - \gamma) + \mu^1(y_H p - c + 1\{G \neq 0\}(y_H p - c - \gamma)) + \mu^2 2(y_H p - c)$$

subject to the constraint that measures of specialization are between 0 and 1 and sum to $\frac{K}{M}$:

$$\mu^1 + 2\mu^2 = \frac{K}{M}, \qquad \mu^k \in [0, 1], k \in \{1, 2\}.$$

Plugging in summing-up constraint yields

$$(1-\mu^1-\frac{1}{2}(\frac{K}{M}-\mu_1))G(y_H p-c-\gamma)+\mu^1(y_H p-c+1\{G \neq 0\}(y_H p-c-\gamma))+\left(\frac{K}{M} - \mu^1\right)(y_H p-c)$$

or

$$(1 - \frac{1}{2}\mu^1 - \frac{1}{2}\frac{K}{M})G(y_H p - c - \gamma) + \mu^1 1\{G \neq 0\}(y_H p - c - \gamma) + \frac{K}{M}(y_H p - c).$$

Notice that total surplus increases linearly in K, so if $K > M$ specialists are optimal then it is optimal to have $K = 2M$. We have already shown 2 generalists are dominated by M

specialists and one generalist, so we need only consider whether M specialists, M specialists and 1 generalist, or $2M$ specialists are optimal.

If $2M$ specialists enter the market, then total surplus is $2(y_H p - c) - 2M\tilde{F}_E$. If M specialists enter the market, then the same specialist is required for both products with probability $\frac{1}{M^2}$. Hence, total surplus is

$$2(y_H p - c) - \frac{1}{M^2}(y_H p - c) - M\tilde{F}_E$$

Finally, if M specialists and a generalist enter the market, then total surplus is

$$2(y_H p - c) - \frac{1}{M^2}\gamma - (M+1)\tilde{F}_E$$

So long as $M\tilde{F}_E > \frac{1}{M^2}(y_H p - c)$, M specialists strictly dominates $2M$ specialists. So long as $\frac{1}{M^2}(y_H p - c - \gamma) < \tilde{F}_E$, M specialists strictly dominates M specialists and a generalist. Rearranging these conditions, we have that M specialists are optimal so long as

$$M^2\tilde{F}_E > \max\{\frac{1}{M}(y_H p - c), (y_H p - c - \gamma)\}$$

By assumption, we know that $\frac{1}{M}(y_H p - c) < (y_H p - c - \gamma)$, so we require only that $M^2\tilde{F}_E > y_H p - c - \gamma$ which holds by assumption. Therefore, under the given conditions, it is optimal for M specialists to enter the market when formal contracts are available. Each specialist is allocated either x_t^1 or x_t^2 with positive probability in each period, proving the claim. ■

Bibliography

[1] Abreu, Dilip. "On the Theory of Infinitely Repeated Games with Discounting." *Econometrica* 56, 383-396; March 1988.

[2] Ali, Nageeb, and David Miller. "Enforcing Cooperation in Networked Societies." Working Paper; October 2012.

[3] Athey, Susan, and Kyle Bagwell. "Collusion with Persistent Cost Shocks." *Econometrica* 76 (3), 493-540; May 2008.

[4] Athey, Susan, and Ilya Segal. "An Efficient Dynamic Mechanism." Working Paper; September 2012.

[5] Aumann, Robert. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1, 67-96; 1974.

[6] Baker, George, Robert Gibbons, and Kevin Murphy. "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics* 109 (4), 1125-1156; November 1994.

[7] Baker, George, Robert Gibbons, and Kevin Murphy. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics* 117, 39-83; 2002.

[8] Baker, George, Michael Jensen, and Kevin Murphy. "Compensation and Incentives: Practice vs. Theory. *The Journal of Finance* 43 (3), 593-616; July 1988.

[9] Ben-Porath, Elchanan, and Michael Kahneman. "Communication in Repeated Games with Private Monitoring." *Journal of Economic Theory* 70, 281-297; 1996.

[10] Bernheim, Douglas, and Michael Whinston. "Multimarket Contact and Collusive Behavior." *RAND Journal of Economics* 21 (1), 1-26; Spring 1990.

[11] Blackwell, David. "Discrete Dynamic Programming." *Annals of Mathematical Statistics* 60, 719-726; 1962.

[12] Board, Simon. "Relational Contracts and the Value of Loyalty." *American Economic Review* 101 (4), 3349-3367; December 2011.

[13] Bull, Clive. "The Existence of Self-Enforcing Implicit Contracts." *The Quarterly Journal of Economics* 102 (1), 147-159; 1987.

[14] Calzolari, Giacomo, and Giancarlo Spagnolo. "Relational Contracts and Competitive Screening." CEPR Working Paper DP7434; August 2010.

[15] Crawford, Vincent, and Joel Sobel. "Strategic Information Transmission." *Econometrica* 50 (6), 1431-1451; November 1982.

[16] Cyert, Richard, and James March. A Behavioral Theory of the Firm. Prentice-Hall, Englewood Cliffs, NJ; 1963.

[17] Dixit, Avinash. "International Trade, Foreign Direct Investment, and Security." *Annual Reviews of Economics,* forthcoming August 2011.

[18] Dixit, Avinash. Lawlessness and Economics: Alternative Modes of Governance. Princeton University Press; 2007.

[19] Dutta, Prajit. "A Folk Theorem for Stochastic Games." *Journal of Economic Theory* 66, 1 - 32; 1995.

[20] Dyer, Jeffrey. "How Chrysler Created an American Keiretsu." *Harvard Business Review*; July-August 1996.

[21] Ellison, Glenn. "Cooperation in the Prisoner's Dilemma with Annonymous Random Matching." *Review of Economic Studies* 61 (3), 567-588; July 1994.

[22] Ely, Jeffrey, Johannes Horner, and Wojciech Olszewski. "Belief-Free Equilibria in Repeated Games." *Econometrica* 73 (2), 377-415; March 2005.

[23] Ely, Jeffrey, and Juuso Valimaki. "A Robust Folk Theorem for the Prisoner's Dilemma." *Journal of Economic Theory* 102 (1), 84-105; January 2002.

[24] Escobar, Juan, and Juuso Toikka. "Efficiency in Games with Markovian Private Information." Working Paper; April 2012.

[25] Farlow, David, Charles A. Holloway, Glen Schmidt, and Andy Tsay. "Supplier Management at Sun Microsystems (A)." Stanford University Graduate School of Business, Case Study OIT-16 A; March 1996.

[26] Feinberg, Eugene, and Adam Shwartz, editors. Handbook of Markov Decision Processes: Methods and Applications. Kluwer Academic Publishers: Norwell, MA; 2002.;

[27] Fong, Yuk-Fai, and Jin Li. "Information Revelation in Relational Contracts." Working Paper; 2010.

[28] Fuchs, William. "Contracting with Repeated Moral Hazard and Private Evaluations." *The American Economic Review* 97 (4), 1432-1448; September 2007.

[29] Fudenberg, Drew, and David Levine. "Efficiency and Observability in Games with Long-Run and Short-Run Players." *Journal of Economic Theory* 62, 103-135; 1994.

[30] Fudenberg, Drew, and Eric Maskin. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." *Econometrica* 54 (3), 533-554; 1986.

[31] Fudenberg, Drew, David Levine, and Eric Maskin. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62 (5), 997-1039; September 1994.

[32] Fudenberg, Drew, and Jean Tirole. Game Theory. The MIT Press: Cambridge, MA; 1991.

[33] Gibbons, Robert, and Rebecca Henderson. "Relational Contracts and Organizational Capabilities." *Organizational Science* 23 (5), 1350-1364; September 2012.

[34] Guinipero, Larry. "Motivating and Monitoring JIT Supplier Performance." *Journal of Purchasing and Materials Management* 26 (3), 19-24; Summer 1990

[35] Halac, Marina. "Relational Contracts and the Value of Relationships." *The American Economic Review* 102 (2), 750-779; April 2012.

[36] Harrington, Joseph, and Andrzej Skrzypacz. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *American Economic Reivew* 101 (3), 1-25; October 2011.

[37] Hohn, Michaela Isabel. Relational Supply Contracts: Optimal Concessions in Return Policies for Continuous Quality Improvements. Springer Heidelberg Dordrecht London New York; 2010.

[38] Holmstrom, Bengt, and Paul Milgrom. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55 (2), 303-328; March 1987.

[39] Jackson, Matthew, and Hugo Sonnenschein. "Overcoming Incentive Constraints by Linking Decisions." *Econometrica* 75 (1), 241-257; January 2007.

[40] Jerez, F. Asis Martinez, Robert S. Kaplan, and Katherine Miller. "Infosys's Relationship Scorecard: Measuring Transformational Partnerships." Harvard Business School, Case Study 9-109-006; December 2009.

[41] Kamenica, Emir, and Matthew Gentzkow. "Bayesian Persuasion." *American Economic Review* 101 (3), 2590-2615; October 2011.

[42] Kandori, Michihiro. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory* 102, 1-15; 2002.

[43] Kandori, Michihiro. "Weakly Belief-Free Equilibria in Repeated Games with Private Monitoring." *Econometrica* 79 (3), 877-892; May 2011.

[44] Kandori, Michihiro, and Hitoshi Matsushima. "Private Observation, Communication, and Collusion." *Econometrica* 66 (3), 627-652; May 1998.

[45] Klein, Benjamin, and Keith Leffler. "The Role of Market Forces in Assuring Contractual Performance." *The Journal of Political Economy* 89 (4); August 1981.

[46] Krause, Daniel, Thomas Scannell, and Roger Calantone. "A Structural Analysis of the Effectiveness of Buying Firms' Strategies to Improve Supplier Performance." *Decision Sciences* 29 (1), 33-55; Winter 2000.

[47] Kulp, Susan, and V.G. Narayanan. "Metalcraft Supplier Scorecard." Harvard Business School, Case Study 9-102-047; March 2004.

[48] Levin, Jonathan. "Multilateral Contracting and the Employment Relationship." *Quarterly Journal of Economics,* 117 (3), 1075-1103; August 2002.

[49] Levin, Jonathan. "Relational Incentive Contracts." *The American Economic Review* 93 (3), 835-857; June 2003.

[50] Li, Hongmin, Hao Zhang, and Charles Fine. "Dynamic Performance-Based Business Share Allocation in a Supply Chain with Substitutable Suppliers." Working Paper; 2011.

[51] Li, Lode. "Information Sharing in a Supply Chain with Horizontal Competition." *Management Science* 48 (9), 1196-1212; September 2002.

[52] Liker, Jeffrey K., and Thomas Y. Choi. "Building Deep Supplier Relationships." *Harvard Business Review,* reprint R0412G; December 2004.

[53] MacLeod, Bentley, and James Malcomson. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57 (2), 447-480; March 1989.

[54] Mailath, George, and Larry Samuelson. Repeated Games and Reputations: Long Run Relationships. Oxford University Press: New York, NY; 2006.

[55] Malcomson, James. "Relational Incentive Contracts." Oxford Discussion Paper Series; October 2012.

[56] Matsushima, Hitoshi. "Repeated Games with Private Monitoring: Two Players." *Econometrica* 72 (3), 823-852; May 2004.

[57] McAdams, David. "Performance and Turnoever in a Stochastic Partnership." *American Economic Journal: Micro* 3 (4), 107-142; November 2011.

[58] Miller, David. "Robust Collusion with Private Information." *Review of Economic Studies* 79, 778-811; 2012.

[59] Myerson, Roger. "Multistage Games with Communication." *Econometrica* 54 (2), 323-358; March 1986.

[60] Nistor, Cristina. "Pricing and Quality Provision in a Channel: A Model of Efficient Relational Contracts." Working Paper; 2012.

[61] North, Douglass. Institutions, Institutional Change and Economic Performance. Cambridge University Press, Cambridge; 1990.

[62] Ostrom, Elinor. Governing the Commons. Cambridge University Press: Cambridge, United Kingdom; 1990.

[63] Radner, Roy. "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship." *Econometrica* 49 (5), 1127-1148; September 1981.

[64] Radner, Roy. "Repeated Principal-Agent Games with Discounting." *Econometrica* 53 (5), 1173-1198; September 1985.

[65] Rayo, Luis, and Ilya Segal. "Optimal Information Disclosure." *Journal of Political Economy* 118 (5), 949-987; October 2010.

[66] Renault, Jerome, Eilon Solan, and Nicholas Vieille. "Dynamic Sender-Receiver Games." Working Paper; March 2012.

[67] Rochet, Jean-Charles. "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context." *Journal of Mathematical Economics* 16 (2), 191-200; 1987.

[68] Sako, Mari. "Supplier Development at Honda, Nissan, and Toyota: Comparative Case Studies of Organizational Capability Enhancement." *Industrial and Corporate Change* 13 (2), 281-308; April 2004.

[69] Sugaya, Takuo. "The Folk Theorem in Repeated Games with Imperfect Private Monitoring." Working Paper; June 2012.

[70] Wolitzky, Alexander. "Communication and Money in Repeated Games on Networks." Working Paper; June 2012b.

[71] Wolitzky, Alexander. "Cooperation with Network Monitoring." Forthcoming, *Review of Economic Studies*; February 2012a.

[72] Zhang, Hongtao. "Vertical Information Exchange in a Supply Chain with Duopoly Retailers." *Production and Operations Management* 11 (4), 531-546; Winter 2002.