The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents*

Isaiah Andrews[†] and Daniel Barron[‡]

April 21, 2014

Abstract

Consider a repeated moral hazard problem involving a principal and several agents. If formal contracts are unavailable and agents observe only their own relationships, an optimal relational contract allocates business among agents depending on past performance. If first-best is attainable, the principal favors an agent who performs well with future business, even if he later performs poorly. The agent loses favor only if he cannot produce and a replacement performs well. If first-best is unattainable, some relationships may deteriorate into persistent low effort. In the first best, the principal need not conceal information from agents; otherwise, she optimally conceals information. JEL Codes: C73, D82, L14, L22

*We are tremendously indebted to Bob Gibbons, Juuso Toikka, and Glenn Ellison for advice, insight, and support. Special thanks to Simon Board, Alessandro Bonatti, Gabriel Carroll, Sebastian Di Tella, Sambuddha Ghosh, Marina Halac, Anton Kolotilin, Ruitian Lang, Greg Leiserson, Hongyi Li, James Malcomson, Mike Powell, Luis Rayo, Eric Van den Steen, Sarah Venables, Xiao Yu Wang, Joel Watson, Nils Wernerfelt, Birger Wernerfelt, Alexander Wolitzky, and Luis Zermeno. Thanks to participants in lunches at MIT and seminars at LSE, Harvard Business School, Rochester, Northwestern, UCSD, Washington University St Louis, Notre Dame, and Boston University. Andrews gratefully acknowledges NSF Graduate Research Fellowship support under Grant No. 1122374. Any errors are entirely our own.

†Doctoral Candidate, MIT Economics Department.
Website: economics.mit.edu/grad/iandrews

‡Post-Doc, Yale University Cowles Foundation. Email: daniel.barron@yale.edu
Website: economics.mit.edu/grad/dbarron

1 Introduction

When formal contracts are unavailable, costly, or incomplete, individuals rely on informal relationships to sustain cooperation. The resulting relational contracts are pervasive both within and between firms. For example, informal agreements typically complement formal contracts in the "just-in-time" supply chains used by Chrysler, Toyota, and others. Similarly, within the firm, 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. 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 allocates business among them critically determines the credibility of her promises. In this paper, we explore how a principal optimally allocates business over time in a repeated game. In each period, the principal requires a single product that only a subset of the agents can produce. The principal chooses one agent from this subset to exert effort that stochastically determines the profitability of the final product. We assume that monitoring is imperfect and private. The principal observes output but not effort, and each agent observes only his own output and pay, and so has limited information about the principal's other relationships.

A *favored* agent is allocated business even when other agents are also available. Because the principal interacts with a favored agent more frequently, she can credibly promise stronger relational incentives to that agent. Ideally, the principal would be able to promise enough future business to each agent to motivate him. However, the same business cannot be promised to different agents; unless the parties are very patient, the principal must prudently allocate business over time to make strong relational incentives credible.

We show that the optimal allocation of future business is fundamentally different depending on whether an agent produces high and low output. An agent is motivated by his expected continuation payoff, which is determined by future wages and bonuses. The principal can credibly promise a large payoff to an agent only if that payoff is accompanied by a substantial amount of future business. Therefore, an optimal allocation rule *rewards success* by favoring a high-performing agent, which allows the principal to credibly promise a large reward to that agent. In contrast, the principal can punish a low-performing agent using transfers, regardless of how much future business has been promised to that agent. As a result, an optimal allocation rule *tolerates failure*, as the principal prefers to use short-term

[1] See Sako (2004) and Liker and Choi (2004) on Toyota, Dyer (1996) on Chrysler, and Baker, Gibbons, and Murphy (2002) for many other examples and references.

transfers to punish poor performance rather than withdrawing future business and weakening the long-term relationship. In other words, the allocation of future business determines the *total continuation surplus* produced by each agent, while wages and bonuses determine *the fraction of that surplus received by that agent* at any given time.

More formally, we characterize a dynamic relational contract that rewards success, tolerates failure, and induces first-best effort whenever any equilibrium does. In this relational contract, an agent immediately rises to favored status when he performs well, then is gradually displaced by more recent high performers. If players are too impatient to attain first-best, the principal sometimes lets one agent shirk in order to promise enough future business to *another* agent to credibly reward his good performance. To analyze this case, we consider the game with two agents and make a realistic but substantive equilibrium restriction.[2] Under these restrictions, one optimal relational contract eventually permanently favors one agent, while the relationship with the other agent devolves into persistent perfunctory performance. In other words, the solution to a short-term moral hazard problem entails a permanent inefficiency.

Methodologically, the game we study has imperfect private monitoring—agents observe neither the output produced by nor the payments made to their counterparts—so we cannot rely on standard recursive methods and instead develop alternative tools. This monitoring structure implies that agents cannot coordinate to jointly punish deviations. The principal may also be able to exploit agents' limited information to better motivate them. So long as first-best is attainable, we show that the principal cannot gain by concealing information from the agents. In contrast, the principal typically finds it optimal to keep agents in the dark about the true history if first-best is unattainable.

Broadly, this paper illustrates how a principal solves short-term moral hazard problems by committing to a long-term and potentially inefficient allocation of business. For example, 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, then the firm can allocate resources over time to motivate both divisions, so that one division might be temporarily favored but neither becomes permanently dominant. If resources are instead scarce, the firm might allow a division that performed well in the past to become dominant for a long time or potentially permanently. In the language of Cyert and March (1963), the firm makes a long-term "policy commitment" to permanently favor one division over another in response to a short-run moral hazard problem. Note that the principal can both pay bonuses and make policy commitments in our setting. Policies that favor one agent over another are

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

important motivational tools that complement monetary incentives and make them credible.

Dynamic allocation decisions play a pivotal role in many real-world environments. For example, companies commonly allocate business among their suppliers in response to past performance. Toyota's *keiretsu* is a classic example of relational supply chain management. Asanuma (1989) notes that in the 1980s, Toyota ranked its suppliers in four different tiers. Better-ranked suppliers were given priority on valuable contracts, and suppliers moved between tiers based on past performance.[3] Similarly, Farlow et al (1996) report that Sun Microsystems regularly considers past performance when deciding how to allocate a project among suppliers. In a survey, Krause et al (2000) find that firms motivate their suppliers using "increased volumes of present business and priority consideration for future business." Similarly, managers allocate tasks and promotions to reward successful employees.

This paper is part of the growing literature on relational contracting spurred by Bull (1987), MacLeod and Malcomson (1989), and Levin (2003); Malcomson (2012) has an extensive survey. Levin (2002) considers relational contracts among a principal and several agents. Unlike that paper, our analysis focuses on the allocation of business among agents, and we motivate bilateral relationships as the result of private monitoring. Segal (1999) considers a principal who can write private (formal) contracts with each of several agents, which is related to our own assumption that monitoring is bilateral. Taylor and Wiggins (1997) compare "just-in-time" relational supply chains to competitive bidding, but do not consider allocation dynamics.

Like us, Board (2011) considers how a principal allocates future business among a group of agents. He shows that the principal separates potential trading partners into "insiders" and "outsiders," trades efficiently with the insiders, and is biased against outsiders. However, the principal in Board's paper perfectly observes agents' actions; consequently, his results are driven by neither moral hazard nor imperfect private monitoring, which are central forces in our paper. In the context of repeated procurement auctions, Calzolari and Spagnolo (2010) 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 related research, Barron (2012) shows that by adopting flexible but less efficient production technology, suppliers can strengthen their relational contracts with a downstream

[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 7 that qualitatively similar dynamics hold with when agents can communicate, albeit under stronger assumptions.

buyer.

Our model is also related to a burgeoning applied-theory literature using games with private monitoring. Kandori (2002) provides an 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 (2013) considers enforcement in games where agents are connected by a network and observe their neighbors' actions. 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 (2013) analyze which networks best sustain cooperation in a repeated game, but do not consider the allocation of business within those networks. Harrington and Skrzypacz (2011) discuss collusion when firms privately observe both prices and quantities.

The rest of the paper proceeds as follows. The next section describes the game, particularly the monitoring structure. In Section 3, we show that an 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. The principal's allocation rule determines each agent's production and hence whether these conditions hold. Section 4 fully characterizes a dynamic 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 5 shows that non-stationary allocation rules are typically optimal. Restricting attention to equilibria that provide *ex post* incentives, we characterize an optimal allocation rule in which one relationship eventually becomes "perfunctory:" that agent is allocated production infrequently and perpetually shirks. When one agent performs well, the *other* relationship is likely to become perfunctory. In Section 6, we show that first-best can be attained using *ex post* incentives whenever it can be attained at all. In contrast, if no equilibrium attains first-best, the optimal relational contract typically conceals information from the agents. In Section 7, we relax the stark assumption that agents cannot communicate. Under an additional assumption, results similar to those from the previous sections continue to hold. Section 8 concludes.

Omitted proofs are in Online Appendix A. Additional results referenced in the text are in Supplementary Online Appendices B, C, and D.

2 Model and Assumptions

2.1 Model

Consider a repeated game with $N+1$ players, denoted $\{0, 1, ..., N\}$. Player 0 is the principal ("she"), while players $i \in \{1, ..., N\}$ are agents ("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 interpreted 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 either accepts or rejects. If he accepts, the agent chooses whether or not to work hard, which determines the probability of high output.

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 from distribution $F(\mathcal{P}_t)$.

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

3. For all $i \in \{1, ..., N\}$, 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 rejects or accepts 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, $c > 0$.

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. For all $i \in \{1, ..., N\}$, the principal transfers $\tau_{i,t}^A \geq 0$ to agent i, who simultaneously transfers $\tau_{i,t}^P \geq 0$ to the principal. $\tau_{i,t}^A, \tau_{i,t}^P$ are observed only by the principal and agent i. The *net transfer to agent i* is $\tau_{i,t} = \tau_{i,t}^A - \tau_{i,t}^P$.

[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 similar to those in Sections 3 - 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.

Let $1_{i,t}$ be the indicator function for the event $\{x_t = i\}$. Then stage-game payoffs in round t are

$$u_0^t = d_t y_t - \sum_{i=1}^N (w_{i,t} + \tau_{i,t})$$
$$u_i^t = w_{i,t} + \tau_{i,t} - d_t c e_t 1_{i,t}$$

for the principal and agent $i \in \{1, ..., N\}$, respectively.

Importantly, 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 7, we show that the allocation of business can play an important role even if agents can communicate. 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.[10] 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.[11]

2.2 Histories, Strategies, and Continuation Payoffs

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

$$\mathcal{H}_0^T = \left\{\mathcal{P}_t, x_t, \{w_{i,t}^A\}, \{w_{i,t}^P\}, d_t, e_t, y_t, \{\tau_{i,t}^A\}, \{\tau_{i,t}^P\}\right\}_{t=0}^T .$$

[6] In particular, agents cannot pay one another. In Appendix C, 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 do not directly interact with one another.

[7] An alternative would be to assume public monitoring and restrict attention to "bilateral punishments" as in Levin (2002). In this model, bilateral punishments are difficult to define because relationships are linked through the allocation rule. In addition, the principal could not conceal information from the agents in this alternative.

[8] The results in Sections 3 - 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 some agent i was always available, then the principal would optimally allocate business only to i.

[11] 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.

It will frequently be convenient to refer to histories that occur in the middle of a period. Define $\mathcal{A}$ as the set of all variables observed in a period, and for $a \in \mathcal{A}$ define $\mathcal{H}_a^T$ as the set of histories immediately following the realization of a. For example, $(h^{T-1}, \mathcal{P}_T, x_T, \{w_{i,T}^A\}, \{w_{i,T}^P\}, d_T, e_T, y_T) \in \mathcal{H}_y^T$. The set of all histories is

$$\mathcal{H} = \bigcup_{T=1}^{\infty} \bigcup_{a \in \mathcal{A} \cup \{0\}} \mathcal{H}_a^T.$$

For each player $i \in \{0, 1, ..., N\}$ and $a \in \mathcal{A} \cup \{0\}$, let $\mathcal{H}_{a,i}^T$ be the set of i's private histories following a in period T.[12] Then

$$\mathcal{H}_i = \bigcup_{T=1}^{\infty} \bigcup_{a \in \mathcal{A} \cup \{0\}} \mathcal{H}_{a,i}^T.$$

Histories $h^T, \hat{h}^T \in \mathcal{H}$ satisfying $h_i^T = \hat{h}_i^T$ are *indistinguishable to i.*

Strategies for player i are denoted $\sigma_i \in \Sigma_i$ with strategy profile $\sigma = (\sigma_0, ..., \sigma_N) \in \Sigma = \Sigma_0 \times ... \times \Sigma_N$.

Definition 1 *For all $i \in \{0, ..., N\}$, player i's continuation surplus $U_i : \Sigma \times \mathcal{H} \to \mathbb{R}$ is*

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

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

$$U_0^i(\sigma, h^t) = E_\sigma \left[\sum_{t'=1}^{\infty} (1-\delta)\delta^{t'-1} (1\{x_{t+t'} = i\} d_{t+t'} y_{t+t'} - w_{i,t+t'} - \tau_{i,t+t'}) | h^t \right]. \tag{2}$$

Intuitively, U_0^i is the expected surplus produced by agent i and earned by the principal. Agents do not know the true history, so their beliefs about continuation surplus are expectations conditional on their private history: $E_\sigma [U_i(\sigma, h^t) | h_i^t]$.

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

$$S_i(\sigma, h^t) = U_0^i(\sigma, h^t) + U_i(\sigma, h^t).$$

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

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

The principal's history is

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

Dyad surplus plays a critical role in the analysis. Intuitively, S_i is *agent i's contribution to total surplus*—the surplus from those rounds when agent i is allocated production. We will show that S_i is a measure of the relational capital available to make incentive pay for agent i credible. 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 $PBE(\delta) \subseteq \mathbb{R}^{N+1}$ for discount factor $\delta \in [0, 1)$.[13] We focus on *optimal* relational contracts, which maximize *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 for all results unless otherwise noted.

Assumption 1 *Agents are symmetric:* $\forall \mathcal{P}, \mathcal{P}' \in 2^{\{1,...,N\}}$, $F(\mathcal{P}) = F(\mathcal{P}')$ *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:* $y_H p_1 - c > y_H p_0$.

Assumption 1 simplifies the analysis and rules out dynamics that are driven by asymmetries among the players. Assumption 2 implies (1) $\mathcal{P} = \{i\}$ with positive probability, so the principal must sometimes allocate business to agent i, and (2) $|\mathcal{P}| > 1$ occurs, so the allocation decision is non-trivial. Assumption 3 implies $e = 1$ in the first-best.

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. The second gives necessary and sufficient conditions for a relational contract in terms of dyad surpluses.

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

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}.

[13] A Perfect Bayesian Equilibrium is an assessment (σ^*, μ^*) consisting of strategy profile σ^* and beliefs $\mu^* = \{\mu_i^*\}_{i=0}^N$. Player i's beliefs $\mu_i^* : \mathcal{H}_i \to \Delta(\mathcal{H})$ give a distribution over true histories h^t at each of i's private histories h_i^t. Given μ_i^*, σ_i^* is a best response at each $h_i^t \in \mathcal{H}_i$. Beliefs updated according to Bayes Rule whenever possible. Otherwise, μ_i^* assigns positive weight only to histories that are consistent with h_i^t, but is otherwise unconstrained. This definition is adapted from Mailath and Samuelson (2006).

[14] Proof sketch: Consider an optimal PBE σ^*. At $t = 0$, agent i pays a transfer equal to $U_i(\sigma^*, \emptyset)$ to the principal. If this transfer is paid, the equilibrium is σ^*. Otherwise, then $\forall t$ $w_{i,t} = \tau_{i,t} = 0$, agent i chooses $d_t = 0$, and x_t is chosen to maximize the principal's continuation payoff. These strategies form an equilibrium. Agent i pays $U_i(\sigma^*, \emptyset)$ because he earns 0 otherwise.

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 each relationship's strength is determined by the allocation of future business. In this setting, formal contracts must be incomplete for the allocation of business to be important.

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 surpluses $\{S_i\}$. In Appendix C, we show that any PBE payoff can be attained by an equilibrium in which agents do not condition on their past effort choices. Without loss of generality, any deviation that is observed by at least two players is optimally punished as harshly as possible.[15] Lemma 1 proves that if the principal or agent i reneges on a promised payment, the harshest possible punishment is the bilateral breakdown of their relationship.

Lemma 1 *Fix equilibrium σ^* and an on-path history $h^t \in \mathcal{H}$. Consider two histories $h^{t+1}, \tilde{h}^{t+1} \in \mathcal{H}_0^{t+1}$ that are successors to h^t such that $h^{t+1} \in supp\{\sigma|h^t\}$ but $\tilde{h}^{t+1} \notin supp\{\sigma|h^t\}$, 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*

$$E_{\sigma^*}\left[U_i(\tilde{h}^{t+1})|\tilde{h}_i^{t+1}\right] \geq 0, \tag{3}$$

$$U_0(\tilde{h}^{t+1}) \geq \max_{\hat{h}^{t+1}|\forall j\neq i, \hat{h}_j^{t+1} = \tilde{h}_j^{t+1}} \left[\sum_{j\neq i} U_0^j(\hat{h}^{t+1})\right] \geq \sum_{j\neq i} U_0^j(h^{t+1}). \tag{4}$$

Proof:

If (3) were 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.

If (4) is not satisfied, let

$$\hat{h}^{t+1} \in \arg\max_{\hat{h}^{t+1}|\forall j\neq i, \hat{h}_j^{t+1} = \tilde{h}_j^{t+1}} \sum_{j\neq i} U_0^j(h^{t+1})$$

Following $\tilde{h}^t$, recursively define the following strategy for the principal. $\forall t' \geq 1$, 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 $h^{t+t'+1}$ be the observed

[15] Analogous to Abreu (1988).

history at the beginning of round $t+t'+1$. 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 by construction, $\hat{h}_j^{t+t'} = h_j^{t+t'}\ \forall j \neq i$, while round $t+t'$ actions that are observed by any agent $j \neq i$ are identical to actions prescribed by $\sigma_0^*(\hat{h}^{t+t'})$.

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 $\sum_{j\neq i} U_0^j(\hat{h}^{t+1})$. ■

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 those agents did not observe the deviation. In particular, the principal can act as if the true history is any $\hat{h}^t$ that is consistent with the beliefs of agents $j \neq i$. By choosing $\hat{h}^t$ to maximize her surplus, the principal exploits agents' limited information to mitigate her loss from i's refusal to trade.

The next result builds on Lemma 1 to prove that agent i's beliefs about S_i determine whether he is willing to work 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, a relational contract is summarized by a set of intuitive dynamic enforcement constraints that must be satisfied when an agent *works hard* and *produces high output*.

Lemma 2 1. *Let σ^* be a PBE. Then*

$$\begin{gathered}(1-\delta)\tfrac{c}{p_1-p_0} \leq \delta E_{\sigma^*}\left[S_i(h^{t+1}) \mid h_i^t\right], \\ \forall i \in \{1,...,N\}, \forall h^{t-1} \in \mathcal{H}_y^{t-1} \text{ on the equilibrium path} \\ \text{such that } x_t = i,\ e_t = 1,\ y_t = y_H.\end{gathered} \tag{5}$$

2. *Let σ be a strategy that generates total surplus V and satisfies (5). 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 A.

For any t and $h^{t+1} \in \mathcal{H}_0^{t+1}$, define

$$D_i(h^{t+1}) = U_0(h^{t+1}) - \max_{\hat{h}^{t+1} | \forall j \neq i, \hat{h}_j^{t+1} = h_j^{t+1}} \sum_{j \neq i} U_0^j(\hat{h}^{t+1}).$$

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 $\tilde{h}^{t+1}$ 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 two histories $h_H^t, h_L^t \in \mathcal{H}_y^t$ that differ only in that $y_t = y_H$ in h_H^t and $y_t = 0$ in h_L^t. Let $\tau_i(y_H) = E_{\sigma^*}\left[\tau_{i,t} \mid h_H^t\right]$, $U_i(y_H) = E_{\sigma^*}\left[U_i \mid h_H^t\right]$, and similarly for $\tau_i(0)$ and $U_i(0)$. Define $B_i(y) = (1-\delta)\tau_i(y) + \delta U_i(y)$. On the equilibrium path, agent i chooses $e_t = 1$ only if

$$p_1 B_i(y) + (1-p_1)B_i(0) - (1-\delta)c \geq p_0 B_i(y) + (1-p_0)B_i(0) \tag{6}$$

For bonus $\tau_{i,t}$ to be paid in equilibrium, it must be that

$$\begin{aligned} (1-\delta)E_\sigma\left[\tau_{i,t}|h_y^t\right] &\leq \delta E_\sigma\left[D_i(h^t)|h_y^t\right] \\ -(1-\delta)E_\sigma\left[\tau_{i,t}|h_{y,i}^t\right] &\leq \delta E_\sigma\left[U_i(h^t)|h_{y,i}^t\right] \end{aligned} \tag{7}$$

or else either the principal or agent i would choose not to pay $\tau_{i,t}$. Plugging these constraints into (6) and applying the definition of S_i results in (5). ■

Define $\tilde{S} = \frac{1-\delta}{\delta}\frac{c}{p_1 - p_0}$. Then $\tilde{S}$ is the difference between payoffs following high or low output required to motivate an agent to work hard. Statement 1 of Lemma 2 says that an agent will work hard only if he expects to produce dyad-specific surplus of at least $\tilde{S}$ following high output. Following output y, an agent's total compensation $B_i(y)$ consists of his bonus payment today and continuation surplus: $B_i(y) = (1-\delta)\tau_i(y) + \delta U_i(y)$. By Lemma 1, $B_i(0) \geq 0$, while $B_i(y_H) \leq \delta S_i(y_H)$. Agent i is willing to work hard only if $B_i(y_H) - B_i(0) \geq \delta\tilde{S}$ and thus $S_i(y_H) \geq \tilde{S}$ as in (5). This argument is similar to the proofs by MacLeod and Malcomson (1989, Proposition 1) and Levin (2003, Theorem 3) that reneging constraints can be aggregated in bilateral relational contracts.

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 (5) holds. Third, because (5) depends on agent i's beliefs, we must ensure that each agent has the appropriate 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 that ensures the principal's expected payoff is 0 in each period. Because the principal's payoff is strictly increasing in effort, agent x_t infers

d_t^* and e_t^* from this wage and chooses these actions. After output is realized, no transfer is made unless the agent accepted production, was expected to work hard, and produced $y_t = 0$. In that case, agent x_t pays his expected continuation surplus to the principal, which he infers from the initial wage.[16] Crucially, the principal expects to earn 0 in every period, so she is willing to implement the desired allocation rule. Agent x_t earns the expected surplus produced in period t, but then pays the principal if he performs poorly. So long as (5) is satisfied, these payments communicate exactly enough information to ensure that x_t works hard and repays the principal if he produces low output. The principal punishes any (observed) deviation by terminating trade.

Agent i's dyad 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 (5), 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 after the agent both *work hards* and *produces high output.* This result highlights how *wages and bonus payments* interact with the *allocation of business.* Following low output, an agent's continuation payoff can be made 0 using transfers, regardless of how much business he is promised in the future. In contrast, an agent can be promised at most S_i continuation surplus following high output. The agent is only willing to work hard if he believes that the difference between his payoffs following high and low output is at least $\tilde{S}$, which implies condition (5). The agent's beliefs are typically coarser than the true history, so the principal's allocation decision influences agent expectations.[17] Therefore, private monitoring shapes both off-equilibrium punishments (Lemma 1) and on-path dynamic enforcement constraints (Lemma 2).

Because (5) depends only on the allocation rule, accept/reject decisions, and effort choices, subsequent results will define relational contracts in terms of these variables. 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 relational contract.

[16] Note that the principal pays a large efficiency wage to the producing agent and is repaid if output is low. This is essential to make the principal willing to implement *any* allocation rule. Other transfer schemes may also work, but must still satisfy (5).

[17] The principal has substantial control over agent beliefs. In Appendix C, we show that set of equilibrium total surpluses remains the same if the principal could costlessly send private messages to each agent in every period.

4 Allocation Dynamics if First-Best is Attainable

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

By Lemma 2, expected dyad surplus $E\left[S_i|h_i^t\right]$ determines whether agent i works hard. If δ is close to 1, players care tremendously about the future, so any allocation rule leads to a large S_i and high effort. An allocation rule is *stationary* if x_t is independent of history. A relational contract with a stationary allocation rule attains first-best if players are patient.[18]

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

$$(1-\delta)\frac{c}{p_1-p_0} \leq \delta \frac{1}{N} V^{FB}. \tag{8}$$

Proof:

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

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

which proves that (8) is a sufficient condition to induce first-best with stationary contracts.

To prove that (8) is also necessary, note that $\sum_i \text{Prob}\{x_{Stat} = i|\mathcal{P} \neq \emptyset\} \leq 1-F(\emptyset)$. Thus, in any stationary allocation rule there exists some i such that $\text{Prob}\{x_{Stat} = i|\mathcal{P} \neq \emptyset\} \leq \frac{1-F(\emptyset)}{N}$. If (8) 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 (8) does not hold.

[18] Given a stationary allocation rule, the optimal contract in each dyad resembles Levin (2003), Theorem 6.

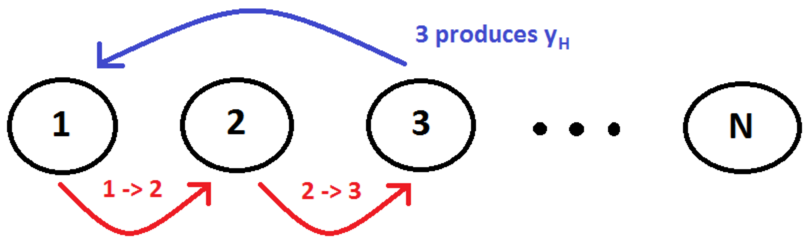

Figure 1: The Favored Producer Allocation. In this picture, the agents ranked 1 and 2 are not available, so production is given to the third-ranked agent. This agent produces y_H and so displaces the 1st and 2nd ranked agents.

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. If he produces high output, then he immediately becomes the most favored agent. 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 Allocation 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.*

4. *On the equilibrium path, $d_t = e_t = 1$ iff (5) holds, and otherwise $d_t = 1$, $e_t = 0$.*

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

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, $\exists\, \underline{\delta} < \bar{\delta}$ such that if $\delta \in [\underline{\delta}, \bar{\delta}]$, the Favored Producer Allocation attains first-best but no stationary equilibrium does.*

Proof:

See Appendix A.

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 if 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 (5) and so make it more difficult to motivate i at these earlier histories. At the same time, the harsher allocation rule would *not* relax this constraint for the other agents to the same extent because (5) depends only on histories following *high output*. In a harsh allocation rule, agents recognize that their rewards for strong performance are fleeting, since they will lose favored status as soon as they perform poorly. Rather than sacrificing her long-term relationship by withdrawing future business from a favored agent, the principal prefers to punish that agent using immediate transfers. In particular, transfers can hold an agent at 0 if $y_t = 0$. The maximum surplus that can be credibly promised to an agent if $y_t = y_H$ is S_i, which is largest in an allocation rule that tolerates failure.

For (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 the Favored Producer Allocation, an agent who produces y_H is more likely to produce in each future period so $S_i > \frac{1}{N}V^{FB}$.

The Favored Producer Allocation is simple enough that we can calculate and perform comparative statics on the set of parameters for which first-best is attainable. If $N = 2$, then first-best is easier to attain if the probability that *both* agents are available $F(\{1,2\})$ is large.[20] Therefore, frequent (independent) disruptions in a supply chain decrease both total surplus and the probability that both agents are available and so make it difficult to sustain strong relationships.

[20] At $F(\{1,2\}) = 1$ (which is ruled out by Assumption 2), the Favored Producer Allocation is stationary.

We turn now to the proof of Proposition 3. Because strategies in a PBE depend on private histories that grow increasingly complex over time, this game is not amenable to a recursive analysis. Instead, we focus 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, then show that the Favored Producer Allocation attains first-best whenever these conditions hold. We focus on intuition for the proof here and relegate details to Appendix A.

The basic idea of the proof is to sum up the *obligation* to each agent—the expected amount of future business that has been promised to that agent. 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. Once $\tilde{S}$ is promised to an agent, it is "paid off" when the principal allocates business to that agent. Therefore, the *expected obligation owed to agent i at time t* (denoted Ω_t^i) is the net *ex ante* expected dyad surplus that must been promised to agent i to motivate first-best effort in rounds $1, ..., t$.

Definition 4 *Given strategy 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*

$$\Omega_t^i \equiv \frac{\Omega_{t-1}^i}{\delta} + \beta_{i,t}^L p_1 \delta \tilde{S} - \beta_{i,t}^H (1-\delta)(y_H p_1 - c) \tag{9}$$

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

Given a strategy in which all agents work hard, Ω_t^i is a state variable that tracks how much dyad surplus is "owed" to agent i in round t. In $t = 1$, 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 (5). Therefore, $\Omega_1^i \equiv \beta_{i,t}^L p_1 \delta \tilde{S}$ equals the expected dyad surplus that must be promised to agent i at the end of the first period. In $t = 2$, agent i is still owed this initial amount, now worth $\frac{\Omega_1^i}{\delta}$ due to discounting. He also 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 $t = 2$. If i produced 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 in that round, $(1-\delta)(y_H p_1 - c)$. Therefore, agent i's expected obligation in round 2 is

$$\Omega_2^i = \frac{\Omega_1^i}{\delta} + \beta_{i,2}^L p_1 \delta \tilde{S} - \beta_{i,2}^H (1-\delta)(y_H p_1 - c).$$

A similar intuition applies in every other period.

If $\Omega_t^i \to \infty$ as $t \to \infty$, then agent i incurs expected obligation faster than it could possibly be paid off by promises of future business. In Appendix A, we show that $\limsup_t \Omega_t^i < \infty$ is a *necessary* condition for a strategy σ to be an equilibrium. Because Ω_t^i is an expectation across all histories, this necessary condition is independent of agent beliefs. While concealing information from agent i might alter the incentive constraint (5), it cannot systematically trick agent i into expecting more future business than could possibly be given to him.

Given this result, it suffices to find conditions under which $\limsup_{t\to\infty} \Omega_t^i = \infty$ for *every* allocation rule. Agents are symmetric, so obligation is minimized when the equilibrium treats all agents symmetrically at the start of the game. 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. These two observations pin down the allocation rules that minimize Ω_t^i for all t.

To complete the proof, we show that if the Favored Producer Allocation does not attain first-best, then $\limsup_{t\to} \Omega_t^i = \infty$ for any allocation rule and hence first-best is unattainable. Intuitively, the Favored Producer Allocation gives an agent who has produced y_H a larger fraction of future surplus than any other allocation rule that satisfies (5) in each period. Given that agents cannot be tricked into believing that total surplus is larger than V^{FB}, the Favored Producer 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 still produces whenever more-favored agents are unavailable.

In one important sense, agents' private information does not play a role in the Favored Producer Allocation: (5) 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] We define a full-disclosure equilibrium by requiring that the dynamic enforcement constraint (5) hold at the true history.

[21] As introduced by Ely and Valimaki (2002) and Ely, Horner, and Olszewski (2005). In Appendix C, we define belief-free equilibrium for our game and show that a strategy profile is payoff-equivalent to a belief-free equilibrium if and only if it satisfies (5) at the true history.

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

$$\begin{gathered}(1-\delta)\tfrac{c}{p_1-p_0} \leq \delta S_i(\sigma^*, h^t) \\ \forall i \in \{1, ..., N\}, \forall h^t \in \mathcal{H}_y^t \text{ on the equilibrium path} \\ \text{such that } x_t = i,\ e_t = 1,\ y_t = y_H.\end{gathered} \tag{10}$$

If (10) 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 (10) and (5) is that (10) 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.

Proposition 3 implies that if first-best is attainable, then there exists a full-disclosure equilibrium that does so.

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

In this section, we have characterized a relational contract that attains first-best whenever any Perfect Bayesian Equilibrium does. Our optimal allocation rule rewards success and tolerates failure: agents gain favor when they perform well and do not lose it when they perform poorly. The dynamics resemble a tournament: agents compete for a temporary favored status that lasts until another agent produces high output. In this way, the principal ensures that she can credibly promise strong incentives to each agent.

5 Unsustainable Networks and Relationship Breakdown

We turn now to the case with two agents and consider relational contracts if first-best is unattainable. In this case, there is not enough future surplus to motivate every agent to work hard, so the principal must sacrifice some relationships to sustain others. As a result, relational contracts that allocate business based on past performance typically dominate stationary contracts. 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 no PBE has total surplus V^{FB} 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 a non-stationary 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 $0 < |\mathcal{M}_{Stat}| < N$. 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 $\mathcal{M}(\mathcal{P}_1) \cap \mathcal{P}_1 \neq \emptyset$ whenever $\mathcal{P}_1 \neq \emptyset$, 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 y_H, the principal must promise that agent a large amount of future business. The principal satisfies this promise by refraining from motivating 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 deteriorating into permanent perfunctory performance.

For the rest of the section, we restrict attention to full-disclosure relational contracts. This substantial restriction allows us to use recursive techniques to characterize equilibrium behavior. Full-disclosure equilibria also have the important property that they do not rely on subtle features of an agent's beliefs. In complex environments, it might be difficult to control each agent's knowledge of past play; full-disclosure relational contracts are robust to agents learning additional information about the true history.[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 princi-

[22] In a collusion model with adverse selection, Miller (2012) argues that *ex post* incentives are natural for this reason.

pal's relationship with each 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 defined by:*

1. *Begin in state G_1. In state G_i,* $Prob_\sigma \{x_t = i | i \in \mathcal{P}_t\} = Prob_\sigma \{x_t = -i | \mathcal{P}_t = \{-i\}\} = 1$*, 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 might break down. Before breakdown occurs, the allocation rule is the same as the Favored Producer Allocation and both agents are expected to work hard. When agent i produces y_H, he enters exclusive dealing with the principal with probability q_i. Once this happens, agent $-i$ stops working hard. 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 earn exclusive dealing.

Proposition 5 shows that for an appropriately-chosen $q^* \in [0, 1]$, (q^*, q^*)-Exclusive Dealing is optimal among full-disclosure equilibria. A critical caveat is that it is not uniquely optimal: other allocation rules do equally well, including some that do not involve the 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.*

[23] For instance, an allocation rule that grants *temporary* exclusive dealing to a high performer immediately following y_H is also optimal.

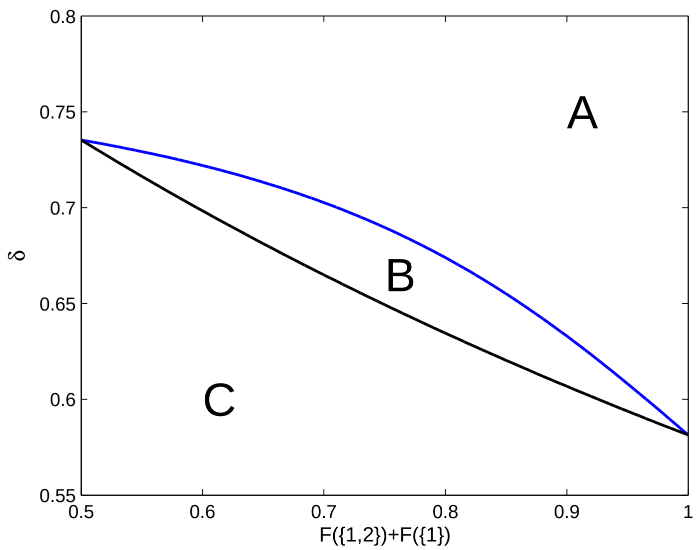

Figure 2: Optimal equilibria for $y_H = 10$, $p_1 = 0.9$, $p_0 = 0$, $c = 5$, $F(\emptyset) = 0$. (A) the Favored Producer Allocation can attain first-best; (B) (q^*, q^*)-Exclusive Dealing is an optimal FDE, and non-stationary equilibria strictly dominate stationary equilibria; (C) no effort can be supported in equilibrium.

Proof:

See Appendix A.

The proof of Proposition 5 is lengthy but the intuition is straightforward. Suppose agent 1 works hard and produces $y_t = y_H$. In subsequent rounds, 1 is allocated production when possible. If only agent 2 can produce, then the principal can either ask 2 to work hard, or not. Asking 2 to work hard has the benefit of generating more surplus. However, it requires that agent 2 be given a lot of future business if he produces y_H, which makes it difficult to fulfill the earlier promise to agent 1. It turns out that this benefit and cost scale at the same rate over time. The *timing* of when agent 2 works hard is irrelevant; only the total discounted probability that 2 works hard matters. In particular, there exists a $q^* \in [0, 1]$ such that (q^*, q^*)-Exclusive Dealing maximizes this discounted probability subject to fulfilling the earlier promises to agent 1.

Exclusive dealing is not uniquely optimal, but it is interesting that an optimal relational contract might entail the permanent break-down of one relationship. 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 inefficient long-term policy in order to solve a short-term 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. Suppose instead that agent 1 is more productive than agent 2, so that high output for agent 1 is $y_H + \Delta_y > y_H$. Then long-term profitability is determined by which relationship breaks down. In Appendix D, 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, 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 observer would see persistent differences in productivity, even though each group was playing the same optimal relational contract.[24]

6 When is it Optimal to Conceal Information?

We have shown that concealing information is unnecessary if first-best can be attained, while (q^*, q^*)-Exclusive Dealing is an optimal full-disclosure equilibrium otherwise. In this section, we prove that full-disclosure equilibria are typically inefficient if first-best is unattainable: the principal might do even better by not providing *ex post* incentives and keeping agents in the dark.

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 relationships—Farlow et al (1996) report that Sun Microsystems used to withhold the details of its relational scorecards from suppliers.[25]

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 $N = 2$. Suppose that first-best cannot be attained in a PBE but that*

$$\frac{c}{p_1 - p_0} < \frac{\delta}{1-\delta}\left(F(\{1\}) + F(\{1,2\})\right)(y_H p_1 - c). \tag{11}$$

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

[24]Persistent performance differences among seemingly similar companies are discussed in Gibbons and Henderson (2013).

[25]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.

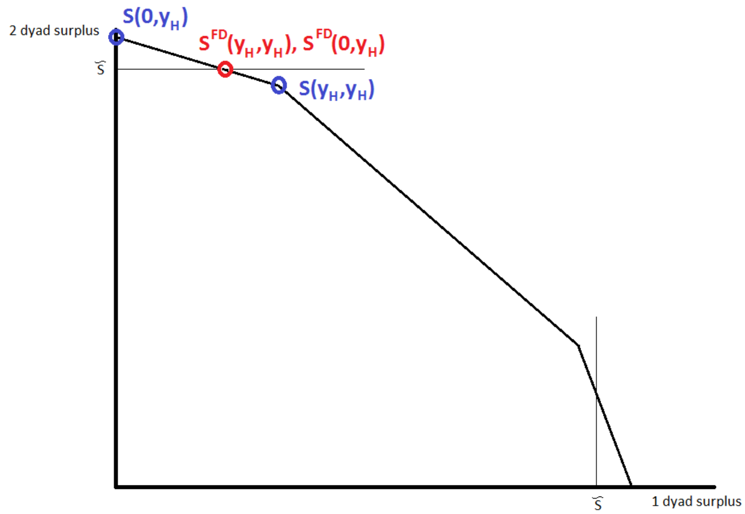

Figure 3: Set of full-disclosure dyad-specific surpluses

Proof:

See Appendix A.

The proof of Proposition 6 constructs an equilibrium that dominates (q^*, q^*)-Exclusive Dealing by concealing information. Exclusive dealing with one agent is inefficient; by concealing information from the agents, the principal can mitigate these inefficiencies.

More precisely, Figure 3 illustrates the set of dyad-specific surpluses that can be sustained in a full-disclosure equilibrium. Consider an optimal full-disclosure equilibrium. 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 a larger $S_2^{FD}(y_H, y_H)$ tightens agent 1's dynamic enforcement constraint (5) in $t = 1$, since $S_1^{FD}(y_H, y_H)$ enters that constraint and is negatively related to $S_2^{FD}(y_H, y_H)$. In contrast, $S_1^{FD}(0, y_H)$ does not enter agent 1's constraint in $t = 1$. Now, suppose that 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 surplus in this alternative. 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 (5) for agent 1 in $t = 1$. This slack can then be used to implement more efficient continuation play 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 a long literature on correlated equilibria including 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 our setting can influence the information learned by an agent and thereby use slack constraints at one history 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.

7 Communication

Our 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,[26] agents within an organization or in particularly close-knit supply chains 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.[27] To consider how communication between agents affects our results, we define an augmented game.

Definition 7 *The* augmented game with communication *is identical to the baseline model, except that each player simultaneously chooses a publicly-observed message $m_i \in M$ at the beginning of each period, where $M = \mathbb{R}$.*

Agents can use the message space M to share information with one another and coordinate to punish the principal. If the principal earned all the rents from production, then the allocation rule would typically be irrelevant in the game with communication. On the other hand, the allocation rule remains an important motivational tool if agents retain some of the profits they produce, as formalized in Assumption 4.

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

$$E\left[u_{x_t}^t | h^t\right] = \gamma \sum_{i=0}^{N} E\left[u_i^t | h^t\right],$$

the principal earns $(1-\gamma)\sum_{i=0}^{N} E\left[u_i^t | h^t\right]$, and all $i \notin \{0, x_t\}$ earn 0.

[26] See Segal (1999) for a discussion on bilateral relationships.

[27] 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 increases in the amount of surplus produced by that dyad, the allocation of business will matter for incentives.

Assumption 4 is related to a rent-sharing assumption made in Halac (2012), which 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^**:*

$$\begin{gathered}\tilde{S} \leq E_{\sigma^*}\left[U_0(h^{t+1}) + U_i(h^{t+1})|h_i^t\right] \\ \forall i \in \{1,...,N\}, \forall h^t \in \mathcal{H}_y^t \text{ on the equilibrium path} \\ \text{such that } x_t = i,\ e_t = 1,\ y_t = y_H.\end{gathered} \tag{12}$$

Proof:

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

$$\begin{array}{c}\tau_i(y_t) \leq \frac{\delta}{1-\delta} E_\sigma\left[U_0(h^{t+1})|h^t\right] \\ -\tau_i(y_t) \leq \frac{\delta}{1-\delta} E_\sigma\left[U_i(h^{t+1})|h_i^t\right]\end{array}.$$

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 (6) yields (12). ■

Condition (12) is similar to (5), 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 (12) for i. Unlike Lemma 2, the principal may have an incentive to deviate from an optimal allocation rule, so (12) 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 (12) is satisfied under the Favored Producer Allocation when* $d_t = e_t = 1$, $\forall t$, *then there exists an equilibrium that uses the Favored Producer Allocation and generates* V^{FB} *total surplus.*

Proof:

See Supplementary Appendix B.

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 game one of imperfect public monitoring. The Favored Producer Allocation is a full-disclosure equilibrium, so it remains a relational contract under public monitoring.

Allocating business to i relaxes his constraint (12) 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 $\Delta \subseteq [0,1]$ such that if $\delta \in \Delta$ this equilibrium attains V^{FB} but stationary equilibria do not.*

Proof:

See Supplementary Appendix B

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 (12), 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:

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

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 substantial surplus and thus (12) 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 only affected by how rent is shared in *other* dyads.[28]

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

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. If first-best is attainable, an optimal relational contract resembles a tournament. The prize is a (temporary) larger share of future business and an agent "wins" if he is allocated business and performs well. If first-best cannot be attained, the optimal full-disclosure equilibrium may involve exclusive dealing between the principal and a high-performing agent, while other relationships deteriorate. The principal can mitigate these inefficiencies by concealing information about the history of play from the agents. 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 information she reveals.

Along with the allocation of business, a principal makes many other decisions that affect her relationships. For example, Klein and Leffler (1981) point out that specific investments can improve relational contracts by increasing the surplus at stake in a relationship. In the context of our model, the optimal allocation rule would take these investments into account. Thus, investments by one agent create an externality on the principal's other relationships. Allocation dynamics might also change over time as the principal learns the ability of her agents.

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 continuation surplus. Nevertheless, those companies with very close supplier relations tend to source from a small number of suppliers, who earn substantial rents. Toyota even went so far as to enshrine "[carrying] out business....without switching to others" in their 1939 Purchasing Rules (as noted by Sako (2004)).

This paper also illustrates a broader point: a principal might choose seemingly inefficient actions (e.g., exclusive dealing with one agent) in order to cultivate close relationships. 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, while seemingly welfare-improving changes undermine existing agreements and inhibit cooperation.

References

[1] Abreu, Dilip. 1988. “On the Theory of Infinitely Repeated Games with Discounting.” *Econometrica* 56 (2): 383-396.

[2] Ali, Nageeb, and David Miller. 2013. “Enforcing Cooperation in Networked Societies.” Working Paper.

[3] Asanuma, Banri. 1989. "Manufacturer-Supplier Relationships in Japan and the Concept of Relation-Specific Skill." Journal of the Japanese and International Economies 3 (1): 1 - 30.

[4] Aumann, Robert. 1974. “Subjectivity and Correlation in Randomized Strategies.” *Journal of Mathematical Economics* 1 (1): 67-96.

[5] Baker, George, Robert Gibbons, and Kevin Murphy. 2002. “Relational Contracts and the Theory of the Firm.” *Quarterly Journal of Economics* 117 (1): 39-83.

[6] Barron, Daniel. 2012. “Putting the Relationship First: Relational Contracts and Market Structure.” Working Paper.

[7] Board, Simon. 2011. “Relational Contracts and the Value of Loyalty.” *American Economic Review* 101 (7): 3349-3367.

[8] Bull, Clive. 1987. “The Existence of Self-Enforcing Implicit Contracts.” *Quarterly Journal of Economics* 102 (1): 147-159.

[9] Calzolari, Giacomo and Giancarlo Spagnolo. 2010. “Relational Contracts and Competitive Screening.” CEPR Discussion Paper No. 7434.

[10] Cyert, Richard, and James March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.

[11] Dyer, Jeffrey. 1996. “How Chrysler Created an American Keiretsu.” *Harvard Business Review*, July.

[12] Ely, Jeffrey, Johannes Horner, and Wojciech Olszewski. 2005. “Belief-Free Equilibria in Repeated Games.” *Econometrica* 73 (2): 377-415.

[13] Ely, Jeffrey, and Juuso Valimaki. 2002. “A Robust Folk Theorem for the Prisoner’s Dilemma.” *Journal of Economic Theory* 102 (1), 84-105.

[14] Farlow, David, Charles A. Holloway, Glen Schmidt, and Andy Tsay. 1996. "Supplier Management at Sun Microsystems (A)." Stanford University GSB, Case Study OIT-16 A.

[15] Fong, Yuk-Fai, and Jin Li. 2010. "Information Revelation in Relational Contracts." Working Paper.

[16] Fuchs, William. 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *The American Economic Review* 97 (4): 1432-1448.

[17] Gibbons, Robert, and Rebecca Henderson. 2013. "What Do Managers Do?" In *Handbook of Organizational Economics,* eds. Robert Gibbons and John Roberts, 680-731. Princeton, NJ: Princeton University Press.

[18] Halac, Marina. 2012. "Relational Contracts and the Value of Relationships." *The American Economic Review* 102 (2): 750-779.

[19] Harrington, Joseph, and Andrzej Skrzypacz. 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *American Economic Review* 101 (3): 1-25.

[20] Jerez, F. Asis Martinez, Robert S. Kaplan, and Katherine Miller. 2009. "Infosys's Relationship Scorecard: Measuring Transformational Partnerships." Harvard Business School, Case Study 9-109-006.

[21] Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (3): 2590-2615.

[22] Kandori, Michihiro. 2002. "Introduction to Repeated Games with Private Monitoring." *Journal of Economic Theory* 102 (1): 1-15.

[23] Klein, Benjamin, and Keith Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." *The Journal of Political Economy* 89 (4): 615-641.

[24] Krause, Daniel, Thomas Scannell, and Roger Calantone. 2000. "A Structural Analysis of the Effectiveness of Buying Firms' Strategies to Improve Supplier Performance." *Decision Sciences* 29 (1): 33-55.

[25] Kulp, Susan, and V.G. Narayanan. 2004. "Metalcraft Supplier Scorecard." Harvard Business School, Case Study 9-102-047.

[26] Levin, Jonathan. 2002. "Multilateral Contracting and the Employment Relationship." *Quarterly Journal of Economics,* 117 (3): 1075-1103.

[27] Levin, Jonathan. 2003. "Relational Incentive Contracts." *American Economic Review* 93 (3): 835-857.

[28] Li, Hongmin, Hao Zhang, and Charles Fine. 2011. "Dynamic Performance-Based Business Share Allocation in a Supply Chain with Substitutable Suppliers." Working Paper.

[29] Liker, Jeffrey K., and Thomas Y. Choi. 2004. "Building Deep Supplier Relationships." *Harvard Business Review,* reprint R0412G.

[30] MacLeod, Bentley, and James Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57 (2): 447-480.

[31] Mailath, George, and Larry Samuelson. 2006. Repeated Games and Reputations: Long Run Relationships. New York, NY: Oxford University Press.

[32] Malcomson, James. 2012. "Relational Incentive Contracts." In *Handbook of Organizational Economics,* eds. Robert Gibbons and John Roberts, 1014-1065. Princeton, NJ: Princeton University Press.

[33] Miller, David. 2012. "Robust Collusion with Private Information." *Review of Economic Studies* 79 (2): 778-811.

[34] Myerson, Roger. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323-358.

[35] Rayo, Luis, and Ilya Segal. 2010. "Optimal Information Disclosure." *Journal of Political Economy* 118 (5): 949-987.

[36] Sako, Mari. 2004. "Supplier Development at Honda, Nissan, and Toyota: Comparative Case Studies of Organizational Capability Enhancement." *Industrial and Corporate Change* 13 (2): 281-308.

[37] Segal, Ilya. 1999. "Contracting with Externalities." *Quarterly Journal of Economics* 114 (2): 337-388.

[38] Taylor, Curtis, and Steven Wiggins. 1997. "Competition or Compensation: Supplier Incentives Under the American and Japanese Subcontracting Systems." *American Economic Review* 87 (4): 598-618.

[39] Wolitzky, Alexander. 2013. "Cooperation with Network Monitoring." *Review of Economic Studies* 80 (1): 395-427.

A For Online Publication: Proofs[29]

Lemma 2

Statement 2: Sufficiency

Given a strategy σ satisfying (5), we construct an equilibrium σ^*.

For any strategy profile $\hat{\sigma}$, let $\mathcal{H}(\hat{\sigma}) \subseteq \mathcal{H}$ be the set of on-path histories. Define an *augmented history* as an element $(h^t, \tilde{h}^t) \in \mathcal{H} \times \mathcal{H}$. We will define a set of augmented histories, denoted $\mathcal{H}^{Aug} \subseteq \mathcal{H} \times \mathcal{H}$, to relate on-path histories in σ^* to histories in σ.

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

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

1. If $t = 0$, then begin with $(\emptyset, \emptyset) \in \mathcal{H}^{Aug}$. Otherwise, let $\left(h_0^t, h_0^{t,*}\right) \in \mathcal{H}^{Aug}$ with $h_0^t, h_0^{t,*} \in \mathcal{H}_0^t$.

2. For every history $h_e^t \in \mathcal{H}_e^t$ such that h_e^t follows h_0^t and $h_e^t \in \mathcal{H}(\sigma)$, define the augmented history $(h_e^t, h_e^{t,*}) \in \mathcal{H}^{Aug}$. Define $h_e^{t,*} \in \mathcal{H}_e^t$ from $h_0^{t,*}$ and $h_e^t \in \mathcal{H}_e^t$ from h_0^t by the following: (starred and unstarred variables are part of $h_e^{t,*}$ and h_e^t, respectively)

   (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 $e_t^* = 0$.[30]

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

   $$w_{x_t,t}^* = d_t \left[(1-e_t) y_H p_0 + e_t \left(y_H p_1 + (1-p_1)\frac{\delta}{1-\delta} E_\sigma \left[S_i(\sigma, h_0^{t+1}) | h_{e,i}^t, y_t = 0\right]\right)\right]$$

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

3. Let $h_0^{t+1} \in \mathcal{H}_0^{t+1}$ be an on-path successor to h_e^t. Define the augmented history $\left(h_0^{t+1}; h_0^{t+1,*}\right) \in \mathcal{H}^{Aug}$, where $h_0^{t+1,*} \in \mathcal{H}_0^{t+1}$ follows from $h_e^{t,*}$ by setting $y_t = y_t^*$, $\tau_{i,t}^* = 0$ $\forall i \neq x_t$, and

$$\tau_{x_t,t}^* = -1\{w_{x_t,t}^* \geq y_H p_1, y_t^* = 0\} \frac{w_{x_t,t}^* - y_H p_1}{1 - p_1}$$

[29] 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."

[30] 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 $(h^t, h^{t,*})$ may have rounds in which $d_{t'}^* = 1$ and $e_{t'}^* = 0$ but $d_{t'} = 0$; however, σ^* is still an equilibrium with identical total surplus and distribution over $\{\mathcal{P}_t, x_t, d_t e_t, y_t\}_{t=1}^T$ $\forall T < \infty$ as σ.

Starting with $(\emptyset, \emptyset) \in \mathcal{H}^{Aug}$, this construction associates a unique $(h^t, h^{t,*}) \in \mathcal{H}^{Aug}$ with each $h^t \in \mathcal{H}(\sigma)$.

Next, we recursively define the candidate equilibrium σ^*. The principal's play in σ^* will be defined in terms of an augmented history.

1. At $t = 0$, choose augmented history $(\emptyset, \emptyset) \in \mathcal{H}^{Aug}$. Otherwise, let $\left(h_0^t, h_0^{t,*}\right) \in \mathcal{H}^{Aug}$ satisfy $h_0^t, h_0^{t,*} \in \mathcal{H}_0^t$.

2. Given $h_{\mathcal{P}}^{t,*} \in \mathcal{H}_{\mathcal{P}}^t$ following $h_0^{t,*}$, let $h_{\mathcal{P}}^t \in \mathcal{H}_{\mathcal{P}}^t$ be the unique successor to h_0^t with the same realization of $\mathcal{P}_t$. The principal chooses $h_e^t \in \mathcal{H}_e^t$ according to the conditional distribution on $\mathcal{H}_e^t$ given by $\sigma | h_{\mathcal{P}}^t$. There exists a unique augmented history $(h_e^t, h_e^{t,*}) \in \mathcal{H}^{Aug}$.

3. The principal chooses x_t^* and $\{w_{i,t}^*\}$ according to $h_e^{t,*}$.

4. Agent x_t chooses $d_t = 1$ iff $w_{x_t,t}^* \geq y_H p_0$ and chooses $e_t = 1$ iff $w_{x_t,t}^* \geq y_H p_1$. If $e_t = 1$ and $y_t = 0$, then $\tau_{x_t,t}^* = -\frac{w_{x_t,t}^* - y_H p_1}{1-p_1}$. Otherwise, $\tau_{x_t,t}^* = 0$. $\forall i \neq x_t$, $\tau_{i,t}^* = 0$.

5. Let $h_0^{t+1,*} \in \mathcal{H}_0^{t+1}$ be the history realized by following steps (1) - (4). Let $h_y^t \in \mathcal{H}_y^t$ be the successor to h_e^t with the same y_t as $h_0^{t+1,*}$. The principal chooses history $h_0^{t+1} \in \mathcal{H}_0^{t+1}$ according to the conditional distribution $\sigma | h_y^t$. Note that $(h_0^{t+1}, h_0^{t+1,*}) \in \mathcal{H}^{Aug}$ by construction. Repeat step (1) using $(h_0^{t+1}, h_0^{t+1,*})$.

6. If a deviation occurs in any variable except 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 $d_t = e_t = 0$ and $w_{i,t}^P = \tau_{i,t}^P = 0$.

Payoff-Equivalence of Old and New Strategies: We claim:

1. $\forall\, t \geq 0$, σ and σ^* generate the same distribution over $\{\mathcal{P}_s, x_s, d_s e_s, y_s\}_{s=1}^t, \forall t$. If $p_0 > 0$, then 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 continuation surplus from history $h^t \in \mathcal{H}(\sigma)$ as σ^* does from $(h^t, h^{t,*}) \in \mathcal{H}^{Aug}$.

We prove the first statement by induction. The result is immediate for $t = 0$. Suppose the statement holds up to t, and let $(h_0^t, h_0^{t,*}) \in \mathcal{H}^{Aug}$ with $h_0^t \in \mathcal{H}_0^t$. The distribution $F(\mathcal{P}_t)$ is exogenous. Conditional on $\mathcal{P}_t$, for $p_0 > 0$ actions (x_t^*, d_t^*, e_t^*) are drawn according to $\sigma | h_{\mathcal{P}}^t$, while for $p_0 = 0$ the analogous statement holds for $(x_t^*, d_t^* e_t^*)$. The distribution over output y_t depends only on d_t and e_t, and the conditional distribution over $(h_0^{t+1}, h_0^{t+1,*}) \in$

$\mathcal{H}^{Aug}$ given $(h_y^t, h_y^{t,*}) \in \mathcal{H}^{Aug}$ is given by the distribution $\sigma|h_y^t$ over h_0^{t+1}. Therefore, the conditional distribution given by $\sigma^*|(h_0^t, h_0^{t,*})$ over period $t+1$ augmented histories is the same as the conditional distribution $\sigma|h_0^t$ over $\mathcal{H}_0^{t+1}$. $\forall T$, h_0^T and $h_0^{T,*}$ have the same sequence $\{\mathcal{P}_s, x_s, d_s e_s, y_s\}_{s=1}^T$ if $p_0 = 0$ and $\{\mathcal{P}_s, x_s, d_s, e_s, y_s\}_{s=1}^T$ if $p_0 > 0$, so this proves the claim.

The second statement follows from this proof, since total surplus is determined by $\{d_t, e_t, y_t\}_{t=1}^{\infty}$ if $p_0 > 0$ and $\{d_t e_t, y_t\}_{t=1}^{\infty}$ otherwise.

Optimality of Principal's Actions: Next, we show that the principal has no profitable deviation. Under σ^*, the principal earns 0 in each period, so it suffices to show that the principal cannot earn strictly positive surplus by deviating. At 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 in all future periods. 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 from $w_{i,t}^* = 0$ for $i \neq x_t$ would be unprofitable for the same reasons. Following a deviation from $w_{x_t,t}^*$, the principal would earn $d_t y_t - w_{x_t,t}$. If $w_{x_t,t} < y_H p_0$ then $d_t^* = 0$ and the deviation is unprofitable. If $w_{x_t,t} \geq y_H p_0$ but $w_{x_t} < y_H p_1$, then either $w_{x_t,t} = y_H p_0$, in which case agent x_t chooses $d_t^* = e_t^* = 0$, or w_{x_t} is off-path for agent x_t. In either case, the principal earns no more than $y_H p_0 - w_{x_t,t} \leq 0$. If $w_{x_t,t} \geq y_H p_1$, then either $d_t^* = e_t^* = 1$ and agent x_t pays $\tau_{x_t,t}^*$ given above, or $w_{x_t,t}$ is off-path for agent x_t. Thus, an upper bound on the principal's payoff is

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

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

Claim 1: Suppose $h_i^t = \hat{h}_i^t$. Then $h_i^{t,*} = \hat{h}_i^{t,*}$. Proof by induction. For $t = 0$, $h^t = \emptyset$ is the only history. Suppose the result holds for all histories in period $t-1$, and consider $h^t, \hat{h}^t \in \mathcal{H}(\sigma)$ with unique augmented histories $(h^t, h^{t,*}), (\hat{h}^t, \hat{h}^{t,*}) \in \mathcal{H}^{Aug}$, respectively. Suppose towards contradiction that $h_i^{t,*} \neq \hat{h}_i^{t,*}$. By induction, there $\exists\, a \in \mathcal{A}$ in period t that differs between $h_i^{t,*}$ and $\hat{h}_i^{t,*}$. Suppose first that $x_t^* = i$. Because $h_i^t = \hat{h}_i^t$, $\mathcal{P}_t^*$, x_t^*, d_t^*, e_t^*, and y_t^* are the same in $h^{t,*}$ and $\hat{h}^{t,*}$. Transfer $\tau_{i,t}^*$ is determined by $w_{i,t}^*$, and $w_{i,t}^*$ depends only on d_t^*, e_t^*, and $h_{e,i}^t$, where $h_{e,i}^t = \hat{h}_{e,i}^t$ by assumption. If $x_t^* \neq i$, the same argument proves that $\mathcal{P}_t^*$, x_t^*, $\tau_{i,t}^*$ and $w_{i,t}^*$ are identical in period t. Thus, $h_i^{t,*} = \hat{h}_i^{t,*}$. Contradiction.

Claim 2: Agent i has no profitable deviation from $\tau_{i,t}$. Agent i earns 0 continuation surplus following a deviation. Let $(h_y^t, h_y^{t,*}) \in \mathcal{H}^{Aug}$, and let $(h_0^{t+1}, h_0^{t+1,*}) \in \mathcal{H}^{Aug}$ be a successor history with $h_0^{t+1}, h_0^{t+1,*} \in \mathcal{H}_0^{t+1}$. By the proof that σ and σ^* are payoff equivalent,

$$S_i(\sigma, h_0^{t+1}) = E_{\sigma^*}\left[\sum_{t'=t+1}^{\infty} \delta^{t'}(1-\delta)1_{i,t}(y-c)|(h_0^{t+1}, h_0^{t+1,*})\right]. \tag{13}$$

Note that the right-hand side of (13) is equal to i-dyad surplus if the *augmented* history is $(h_0^{t+1}, h_0^{t+1,*})$. Hence, by construction of σ^*,

$$E_{\sigma^*}\left[\sum_{t'=t+1}^{\infty} \delta^{t'}(1-\delta)u_i^{t'}|(h_0^{t+1}, h_0^{t+1,*})\right] = S_i(\sigma, h_0^{t+1}). \tag{14}$$

Claim 1 implies that σ^* induces a coarser partition over augmented histories than σ. Consider a history $h_y^{t,*}$ such that $e_t^* = 1$ and $y_t^* = 0$. Then

$$\begin{array}{c} E_{\sigma^*}\left[U_i(\sigma^*, h_0^{t+1,*})|h_{y,i}^{t,*}\right] = \\ E_{\sigma^*}\left[S_i(\sigma, \hat{h}_0^{t+1})|(\hat{h}_y^t, \hat{h}_y^{t,*}) \in \mathcal{H}^{Aug} \text{ s.t. } \hat{h}_{y,i}^{t,*} = h_{y,i}^{t,*}\right] = \\ E_{\sigma^*}\left[E_{\sigma}\left[S_i(\sigma, \hat{h}_0^{t+1})|\hat{h}_{y,i}^t\right](\hat{h}_y^t, \hat{h}_y^{t,*}) \in \mathcal{H}^{Aug} \text{ s.t. } \hat{h}_{y,i}^{t,*} = h_{y,i}^{t,*}\right] = \\ E_{\sigma}\left[S_i(\sigma, h_0^{t+1})|h_{y,i}^t\right] \end{array}. \tag{15}$$

The first equality holds by (14). The second equality holds for two reasons: (a) σ^* induces a coarser partition over augmented histories than σ, and (b) by construction, the distribution induced by σ^* over augmented histories $(h^t, h^{t,*}) \in \mathcal{H}^{Aug}$ is the same as the distribution induced by σ over $h^t \in \mathcal{H}$.[31] The final equality holds because if $\hat{h}_{y,i}^{t,*} = h_{y,i}^{t,*}$, then $\hat{w}_{i,t}^* = w_{i,t}^*$. But then $E_{\sigma}\left[S_i(\sigma, h_0^{t+1})|h_{e,i}^t, y_t = 0\right] = E_{\sigma}\left[S_i(\sigma, \hat{h}_0^{t+1})|\hat{h}_{e,i}^t, y_t = 0\right]$ by definition of $w_{i,t}^*$.

Following a deviation, $U_i = 0$. Therefore, agent i has no profitable deviation if

$$-(1-\delta)\tau_{i,t}^* \leq \delta E_{\sigma}\left[S_i(\sigma, h_0^{t+1})|h_{e,i}^t, y_t = 0\right]$$

which holds (with equality) by construction.

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 deviates. Similarly, agent i has no profitable deviation from $e_t = 0$. Agent i has no profitable

[31] Recall that every $h^t \in \mathcal{H}(\sigma)$ corresponds to a unique $(h^t, h^{t,*}) \in \mathcal{H}^{Aug}$.

deviation from $e_t = 1$ if his IC constraint (6) holds. Given $\tau^*_{i,t}$, this constraint may be written

$$\delta E\left[U_i(\sigma^*, h_0^{t+1,*})|h_{e,i}^{t,*}, y_t^* = y_H\right] \geq (1-\delta)\frac{c}{p_1 - p_0}.$$

This inequality holds by an argument similar to (15). Indeed, at every $\left(\hat{h}^t_{y,i}, \hat{h}^{t,*}_{y,i}\right)$ such that $\hat{h}^{t,*}_{y,i} = h^{t,*}_{y,i}$, $\hat{w}^*_{i,t} = w^*_{i,t} \geq y_H p_1$. Hence, $E_\sigma\left[S_i(\sigma, \hat{h}_0^{t+1})|\hat{h}^t_{e,i}, y_t = y_H\right] \geq \frac{1-\delta}{\delta}\frac{c}{p_1-p_0}$ because σ satisfies (5).

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

Proposition 3

We prove this proposition using two lemmas. It will be useful to rewrite obligation Ω^i_t non-recursively:

$$\Omega^i_t = \sum_{t'=1}^{t} \frac{1}{\delta^{t-t'}}\beta^L_{t'} p_1 \delta \tilde{S} - (1-\delta)V^{FB}\sum_{t'=1}^{t}\frac{1}{\delta^{t-t'}}\beta^H_{t'}. \tag{16}$$

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

$$\limsup_{t\to\infty} \Omega^i_t = \infty.$$

Then σ is not an equilibrium.

Proof of Lemma A1

Let $\sigma \in \Sigma$ attain first-best with $\limsup_{t\to\infty} \Omega^i_t = \infty$ for some $i \in \{1, ..., N\}$. Then σ specifies $d_t = e_t = 1$ whenever $\mathcal{P}_t \neq \emptyset$. Towards contradiction, suppose that σ is an equilibrium.

For $h^t \in \mathcal{H}^t_0$, let

$$b_i(h^t_i) = 1\{x_t = i, y_t = y_H\} * 1\{\forall t' < t, x_{t'} = i \Rightarrow y_{t'} = 0\} \tag{17}$$

equal 1 if agent i produces y_H for the first time in period t. Lemma 2 implies that (5) must be satisfied at any on-path history with $b_i(h^t_i) = 1$, so *a fortiori* holds in expectation:

$$E_\sigma\left[b_i(h^t_i)\right] * \delta\tilde{S} \leq E_\sigma\left[(1-\delta)b_i(h^t_i)E_\sigma\left[\sum_{t'=1}^{\infty}\delta^{t'}1_{i,t+t'}V^{FB} \mid h^t_i\right]\right] \tag{18}$$

where $V^{FB} = y_H p_1 - c$ and $1_{i,t} = 1\{x_t = i\}$. Because σ attains first-best, $E_\sigma\left[b_i(h^t_i)\right] = p_1\beta^L_t$.

Dividing by δ^{K-t} and summing across $t = 1, ..., K$ yields

$$\sum_{k=1}^{K} \frac{p_1 \beta_k^L}{\delta^{K-k}} \delta \tilde{S} \leq (1-\delta) V^{FB} \sum_{k=1}^{K} E_\sigma \left[\sum_{t'=1}^{\infty} \frac{1}{\delta^{K-k-t'}} 1_{i,k+t'} \mid b_i(h_i^{k+t'}) = 1 \right] \text{Prob}_\sigma \{ b_i(h_i^k) = 1 \}. \quad (19)$$

This infinite sum is dominated by $\sum_{t'=1}^{\infty} \delta^{t'} (y_H p_1 - c)$ and so converges absolutely. Switch the order of summation on the right-hand side of (19) and re-index using $\psi = k + t'$. Because the events of (17) are disjoint, $\sum_{k=1}^{K} \text{Prob}\{b_i(h_i^k) = 1\} = \text{Prob}\{\exists k \leq K \ s.t. \ x_k = i, y_k = y_H\}$. Therefore, (19) can be written

$$\begin{array}{c} \sum_{\psi=1}^{K} \frac{p_1 \beta_\psi^L}{\delta^{K-\psi}} \delta \tilde{S} \leq \\ \sum_{\psi=1}^{\infty} \frac{(1-\delta) V^{FB}}{\delta^{K-\psi}} \text{Prob}_\sigma \left\{ x_\psi = i, \exists t' \leq \min\{K, \psi - 1\} \ s.t. \ x_{t'} = i, y_{t'} = y_H \right\}. \end{array} \quad (20)$$

Now, the tail for $\psi > K$ on the right-hand side of (20) is dominated by $\sum_{k=K}^{\infty} \frac{(1-\delta) V^{FB}}{\delta^{K-k}} = V^{FB}$ and so converges absolutely. Hence, (20) implies that $\exists \ C \in \mathbb{R}$ such that $\forall K \in \mathbb{N}$,

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

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

$$\sum_{\psi=1}^{K} \frac{p_1 \beta_\psi^L}{\delta^{K-\psi}} \delta \tilde{S} - (1-\delta) V^{FB} \sum_{\psi=1}^{K} \frac{1}{\delta^{K-\psi}} \beta_\psi^H \leq C \quad (21)$$

By (16), the left-hand side of (21) is Ω_K^i. So Ω_K^i is bounded above by some C for all K. Contradiction; σ cannot be an equilibrium. ■

Statement of Lemma A2

Suppose $\exists$ strategies σ that attain the first-best total surplus V^{FB} such that $\forall i \in \{1, ..., N\}$, $\limsup_{t \to \infty} \Omega_t^i < \infty$. $\forall h^t \in \mathcal{H}_0^t$, let

$$E(h^t) = \left\{ i \in \{1, ..., n\} \mid h^t, \exists t' < t \text{ s.t. } x_{t'} = i, y_{t'} = y_H \right\}$$

and $\forall t$,

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

with $(E_i^t)^C$ the complement of E_i^t. Then $\exists$ strategies $\hat{\sigma}$ attaining V^{FB} with obligation $\{\hat{\Omega}_t^i\}_{i,t}$ satisfying:

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: $\forall t \geq 0, i \in \{1, ..., N\}$,

$$\begin{array}{c} \text{Prob}_{\hat{\sigma}}\left\{x_t = i | E_i^t\right\} = \frac{1}{N}\sum_{j=1}^N \text{Prob}_{\hat{\sigma}}\left\{x_t = j | E_j^t\right\} \\ \text{Prob}_{\hat{\sigma}}\left\{x_t = i | \left(E_i^t\right)^C\right\} = \frac{1}{N}\sum_{j=1}^N \text{Prob}_{\hat{\sigma}}\left\{x_t = j | \left(E_j^t\right)^C\right\}. \end{array}$$

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

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 symmetric, $\tilde{\sigma}$ generates first-best surplus. If $\limsup_{t\to\infty} \max_i \Omega_t^i < \infty$ under σ, then

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

$\forall i \in \{1, .., N\}$ under $\tilde{\sigma}$. $\forall i \in \{1, ..., N\}$,

$$\text{Prob}_{\tilde{\sigma}}\{x_t = i | E_i^t\} = \frac{1}{N}\sum_{j=1}^N \text{Prob}_{\sigma}\{x_t = j | E_j^t\}$$

and a similar statement holds for $(E_i^t)^C$.

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:

$$\hat{x}_t = \begin{cases} \arg\min_{i\in\mathcal{P}_t} \rho(i) & \text{if } \mathcal{P}_t \cap E(h^t) = \emptyset \\ \arg\min_{i\in\mathcal{P}_t\cap E(h^t)} \rho(i) & \text{otherwise} \end{cases}. \tag{22}$$

Agent x_t chooses $d_t = e_t = 1$.

If $\limsup_{t\to\infty} \max_i \tilde{\Omega}_t^i < \infty$ under $\tilde{\sigma}$, then we claim that $\limsup_{t\to\infty} \max_i \hat{\Omega}_t^i < \infty$ in $\hat{\sigma}$.

As with $\tilde{\sigma}$, $\hat{\sigma}$ satisfies

$$\text{Prob}_{\hat{\sigma}}\{x_t = i | E_i^t\} = \frac{1}{N}\sum_{j=1}^{N} \text{Prob}_{\sigma}\{x_t = j | E_j^t\}$$

$\forall i \in \{1, ..., N\}$, so $\exists\, \tilde{\Omega}_t$, $\hat{\Omega}_t$ such that $\forall i \in \{1, ..., N\}$, $\tilde{\Omega}_t^i = \tilde{\Omega}_t$ and $\hat{\Omega}_t^i = \hat{\Omega}_t$.

For all $t \in \mathbb{Z}_+, m \in \{1, ..., N\}$, define $Y_m^t = \{h^t |\ |E(h^t)| = m\} \subseteq \mathcal{H}_0^t$ and $\phi_m^t = \frac{1}{m}\text{Prob}_{\tilde{\sigma}}\{x_t \in E(h^t)|Y_m^t\}$, and note that

$$\phi_m^t \leq \frac{1}{m}\text{Prob}_{\hat{\sigma}}\{x_t \in E(h^t)|Y_m^t\}. \tag{23}$$

Because $\hat{\sigma}$ and $\tilde{\sigma}$ are symmetric, $\text{Prob}\{i \in E_t(h^t)|Y_m^t\} = \frac{m}{N}$ for both strategies. Then for $\tilde{\sigma}$,

$$\begin{array}{c} \beta_t^H = \sum_{m=1}^{N} \frac{m}{N}\text{Prob}_{\tilde{\sigma}}\{Y_m^t\}\phi_m^t, \\ \beta_t^L = \sum_{m=0}^{N-1} \text{Prob}_{\tilde{\sigma}}\{Y_m^t\}\left(\frac{1-F(\emptyset)}{N} - \frac{m}{N}\phi_m^t\right). \end{array}$$

Because $\tilde{\sigma}$ attains first-best,

$$\text{Prob}_{\tilde{\sigma}}\{Y_0^t\} = [F(\emptyset) + (1 - F(\emptyset))(1 - p_1)]\,\text{Prob}_{\tilde{\sigma}}\{Y_0^{t-1}\}$$

and

$$\text{Prob}_{\tilde{\sigma}}\{Y_m^t\} = \begin{array}{c} \phi_m^{t-1}\text{Prob}_{\tilde{\sigma}}\{Y_m^{t-1}\} + (1 - F(\emptyset) - \phi_m^{t-1})(1-p_1)\text{Prob}_{\tilde{\sigma}}\{Y_m^{t-1}\} + \\ (1 - F(\emptyset) - \phi_{m-1}^{t-1})p_1\text{Prob}_{\tilde{\sigma}}\{Y_{m-1}^{t-1}\} \end{array}.$$

To show that $\limsup_{t\to\infty} \tilde{\Omega}_t > \limsup_{t\to\infty} \hat{\Omega}_t$, by (23) we need only show that obligation decreases in ϕ_m^t. In (16), obligation is a smooth function of β_t^L and β_t^H, and hence of $\{\phi_m^t\}_{m,t}$, so it suffices to show $\frac{\partial \Omega_t}{\partial \phi_m^{t'}} \leq 0\ \forall t, t' \in \mathbb{N}, m \in \{1, ..., N\}$.

Using (9) and repeatedly substituting for Ω_{t-1} yields

$$\frac{\partial \Omega_t}{\partial \phi_m^{t'}} = \frac{1}{\delta^{t-t'-1}}\frac{\partial}{\partial \phi_m^{t'}}\Omega_{t'+1} + \sum_{s=1}^{t-t'-1}\frac{1}{\delta^s}\left(\frac{\partial}{\partial \phi_m^{t'}}\beta_{t-s}^L p\delta\tilde{S} - \frac{\partial}{\partial \phi_m^{t'}}\beta_{t-s}^H(1-\delta)V^{FB}\right). \tag{24}$$

We can calculate

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

$$\frac{\partial \beta_t^H}{\partial \phi_m^{t'}} = \begin{cases} \sum_{k=1}^{N} \frac{k}{N} \phi_k^t \frac{\partial \text{Prob}_{\tilde{\sigma}}\{Y_k^t\}}{\partial \phi_m^{t'}} & \text{if } t' < t \\ \frac{m}{N} \text{Prob}_{\tilde{\sigma}}\{E_m^{t'}\} & \text{if } t' = t \\ 0 & t' > t \end{cases}.$$

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

$$\frac{\partial}{\partial \phi_m^{t'}} \beta_t^L \equiv -\frac{\partial}{\partial \phi_m^{t'}} \beta_t^H \tag{25}$$

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

$$\frac{\partial}{\partial \phi_m^{t-1}} \Omega_t = \left(p_1 \delta \tilde{V} + (1-\delta) V^{FB}\right) \frac{\partial}{\partial \phi_m^{t-1}} \beta_{t-1}^L. \tag{26}$$

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

$$\frac{\partial}{\partial \phi_m^{t'}} \beta_{t'}^L < -\sum_{s=t'+1}^{\infty} \delta^{s-t'} \frac{\partial}{\partial \phi_m^{t'}} \beta_s^L < -\sum_{s=t'+1}^{t} \delta^{s-t'} \frac{\partial}{\partial \phi_m^{t'}} \beta_s^L \tag{27}$$

where we have used the fact that $\frac{\partial}{\partial \phi_m^t} \beta_s^L > 0$ for $s > t$. Together, (25) and (27) imply that

$$\begin{gathered} \sum_{s=1}^{t-t'-1} \frac{1}{\delta^s} \left(\frac{\partial}{\partial \phi_m^{t'}} \beta_{t-s}^L p_1 \delta \tilde{S} - \frac{\partial}{\partial \phi_m^{t'}} \beta_{t-s}^H (1-\delta) V^{FB} \right) = \\ \left(p_1 \delta \tilde{S} + (1-\delta) V^{FB}\right) \sum_{s=1}^{t-t'-1} \frac{1}{\delta^s} \frac{\partial}{\partial \phi_m^{t'}} \beta_{t-s}^L < -\left(p_1 \delta \tilde{S} + (1-\delta) V^{FB}\right) \frac{\partial}{\partial \phi_m^{t'}} \beta_{t'}^L. \end{gathered} \tag{28}$$

If we plug this expression into (24) and apply (26), we finally conclude

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

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 Allocation 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. Let σ_{FPA} be a strategy profile using the Favored Producer Allocation.

In the allocation rule (22), the probability that an agent is given production if he and

$m-1$ other agents have produced y_H in the past is $\frac{F^m}{m}$. Plugging β_t^L and β_t^H into (21) yields

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

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 Allocation allocates production to agents who have already produced, and so generates the same obligation as $\hat{\sigma}$ from Lemma A2. Let S_G be agent i's surplus immediately after he produces y_H in the Favored Producer Allocation. If obligation converges when $\tilde{S} = S_G$ and production is allocated as in the Favored Producer Allocation, then it diverges for $\tilde{S} > S_G$ and the proposition is proven.

In the Favored Producer Allocation, let $h_x^t \in \mathcal{H}_x^t$. Define
$b_t^H(h_x^t) = 1\{x_t = i, i \in E(h^t)\}$ and $b_t^L(h_x^t) = 1\{x_t = i, i \notin E(h^t)\}$. Given $h_y^t \in \mathcal{H}_y^t$ such that $x_t = i$ and $y_t = y_H$, agent i's future surplus is

$$\delta S_G = E_{\sigma^{FPA}} \left[(1-\delta) V^{FB} \sum_{t'=1}^{\infty} \delta^{t'} b_{t+t'}^H (h_x^{t+t'}) | h_y^t \right]. \tag{29}$$

Because (29) holds at *every* history h_y^t where $x_t = i$ and $y_t = y_H$,

$$\delta S_G = E_{\sigma^{FPA}} \left[(1-\delta) V^{FB} \sum_{t'=1}^{\infty} \delta^{t'} b_{t+t'}^H (h_x^{t+t'}) | h_y^t \text{ s.t. } y_t = y_H, b_t^L(h_x^t) = 1 \right]. \tag{30}$$

Because $e=1$ at every on-path history, $\text{Prob}_{\sigma^{FPA}}\{y_t = y_H | b_t^L(h_x^t) = 1\} = p_1$. The event $\{b_t^L(h_x^t) = 1, y_t = y_H\}$ can occur only once, so we can multiply both sides of (30) by $\delta^t 1\{b_t^L(h_x^t) = 1, y_t = y_H\}$, sum across t, and take expectations. The resulting expression is

$$\begin{gathered} \delta S_G \textstyle\sum_{t=1}^{\infty} \delta^t \text{Prob}_{\sigma^*} \left\{ b_t^L(h_x^t) = 1, y_t = y_H \right\} = \\ E_{\sigma^*} \left[\textstyle\sum_{t=1}^{\infty} \delta^t \left(\begin{array}{c} 1\left\{ b_t^L(h_x^t) = 1, y_t = y_H \right\} * \\ E_{\sigma^*} \left[(1-\delta) V^{FB} \sum_{t'=1}^{\infty} \delta^{t'} b_{t+t'} | h^t \text{ s.t. } b_t^L(h_x^t) = 1, y_t = y_H \right] \end{array} \right) \right]. \end{gathered} \tag{31}$$

Note $\text{Prob}_{\sigma^*} \left\{ b_t^L(h_x^t) = 1, y_t = y_H \right\} = p_1 \beta_t^L$. We can simplify the right-hand side of (31) using Iterated Expectations, then use dominated convergence to switch the order of summation.

The result is

$$\delta S_G \sum_{t=1}^{\infty} \delta^t \beta_t^L = (1-\delta) V^{FB} E_{\sigma^{FPA}} \left[\sum_{t=1}^{\infty} \delta^t b_t^H (h_x^t) \right]. \tag{32}$$

Apply dominated convergence to interchange expectation and summation in (32). Noting that $E[b_t^H(h_x^t)] = \beta_t^H$, we yield

$$p_1 \delta S_G \sum_{k=1}^{\infty} \delta^k \beta_k^L - (1-\delta) V^{FB} \sum_{k=1}^{\infty} \delta^k \beta_k^H = 0. \tag{33}$$

Suppose $\tilde{S} = S_G$, and consider any allocation rule that minimizes obligation as in Lemma A2. By (16), note that the first t terms in the sum in (33) equal Ω_t. Therefore,

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

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

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

Thus, obligation converges for $\tilde{S} = S_G$, and so $\lim_{t\to\infty} \Omega_t = \infty$ for any allocation rule whenever $\tilde{S} > S_G$. By Lemma A1, this implies that first-best cannot be attained for $\tilde{S} > S_G$. For $\tilde{S} \leq S_G$, the Favored Producer Allocation 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 while no stationary equilibrium does. 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 $\delta > \delta_{Stat} > 0$. Since both sides of (5) are continuous in δ, by Lemma 2 it suffices to show that for $\delta = \delta_{Stat}$, the Favored Producer Allocation satisfies $E_{\sigma^*}\left[S_i(h^{t+1}) | h_{y,i}^t \right] > \frac{1}{N} V^{FB}$ at any on-path $h_y^t \in \mathcal{H}_y^t$ such that $x_t = i$, $e_t = 1$, and $y_t = y_H$. Let $r_{i,t}$ denote the rank of player i at the start of round t in the Favored Producer Allocation. For any on-path $h_0^t \in \mathcal{H}_0^t$, $E_{\sigma^*}[S_i(h^t) | h_0^t] = E_{\sigma^*}[S_i(h^t) | r_{i,t}]$. By Assumptions 1 and 2, $E_{\sigma^*}[S_i(h^t) | r_{i,t}]$ is strictly decreasing in $r_{i,t}$ since $r_{i,t} < r_{j,t}$ implies both that Prob$\{x_t = i\} >$ Prob$\{x_t = j\}$ and that Prob$\{r_{i,t+t'} < r_{j,t+t'}\} > \frac{1}{2}$ for all $t' \geq 0$. Since $\sum_{i=1}^{N} E_{\sigma^*}[S_i(h^{t+1}) | r_{i,t+1}] = V^{FB}$, $E_{\sigma^*}[S_i(h^{t+1}) | r_{i,t+1} = 1] > \frac{1}{N} V^{FB}$, so (5) is slack in the Favored Producer Allocation for $\delta = \delta_{Stat}$. ■

Proposition 5[32]

In Appendix C, we prove that the payoff set $FD(\delta)$ is recursive. Condition (10) constrains the continuation equilibrium at histories in which an agent worked hard and produced y_H. This proof proceeds by considering the periods *after* agent i chooses $e = 1$ and produces $y = y_H$.

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

Let $h_0^t \in \mathcal{H}_0^t$ be in the support of σ^* and satisfy $e_{t'}y_{t'} = 0$ for all $t' < t$. The continuation equilibrium is unconstrained at h^t. It is straightforward to show that both agents choose $e_t = 1$ in period t following history h_0^t. Agents are symmetric, so we conclude that at h_0^t, any allocation rule satisfying $x_t = \emptyset$ iff $\mathcal{P}_t = \emptyset$ is optimal. This condition is clearly satisfied by (q^*, q^*)-Exclusive Dealing.

Next, consider on-path history h_0^t satisfying $x_{t-1} = i$ and $e_{t-1}y_{t-1} = y_H$. Suppose $S_j(\sigma^*, h_0^t) = \hat{S}_j$, $j \in \{1,2\}$, and note that $\hat{S}_i \geq \tilde{S}$ by (10). We will show that these continuation payoffs can be attained using (q^*, q^*)-Exclusive Dealing.

Recall $1_{i,t} = 1\{x_t = i\}$, and define $s_i^t = 1_{i,t}d_t\left(y_H p_0 + e_t\left(y_H(p_1 - p_0) - c\right)\right)$ as the total surplus produced by agent i at time t. At history h_0^t, any optimal continuation equilibrium that gives each agent j at least $\hat{S}_j$ continuation surplus solves:

$$\max_{\{s_i^{t+t'}\}_{i,t'}} \sum_{t'=0}^{\infty} \delta^{t'}(1-\delta)E\left[s_1^{t+t'} + s_2^{t+t'}|h_0^t\right] \tag{34}$$

subject to the constraints that $\{s_j^{t+t'}\}_{j,t'}$ be feasible and $\forall j \in \{1,2\}$,

$$\sum_{t'=0}^{\infty} \delta^{t'}(1-\delta)E\left[s_j^{t+t'}|h_0^t\right] \geq \hat{S}_j, \tag{35}$$

$$\begin{array}{c} \sum_{s=\hat{t}}^{\infty} \delta^{s-\hat{t}}(1-\delta)E\left[s_{x_{t+s}}^{t+s}\,\middle|\,h_y^{t+\hat{t}}\right] \geq \tilde{S} \\ \forall \hat{t} \in \mathbb{Z}_+, \forall \text{ on-path } h_y^{t+\hat{t}} \in \mathcal{H}_y^t \text{ s.t. } e_{t+\hat{t}} = 1, y_{t+\hat{t}} = y_H. \end{array} \tag{36}$$

Constraint (36) follows immediately from (10).

Define

$$L_i^{t+t'} = \left\{h_0^{t+t'} \in \mathcal{H}_0^{t+t'}|h_0^{t+t'} \text{ succeeds } h_0^t \text{ and } \forall 0 \leq \hat{t} \leq t', \text{ if } x_{t+\hat{t}} = i \text{ then } y_{t+\hat{t}}e_{t+\hat{t}} = 0\right\}.$$

[32] For brevity, we omit much of the algebra in this proof. Full details are available upon request.

The set $L_i^{t+t'} \subseteq \mathcal{H}_0^{t+t'}$ contains all histories following h_0^t such that agent $-i$ has not chosen $e = 1$ and produced y_H since before period t. Define $L_{-i}^{t+t'} = \left\{h_0^{t+t'} \in \mathcal{H}_0^{t+t'} | h_0^{t+t'} \text{ succeeds } h_0^t \text{and } h_0^{t+t'} \notin L_i^{t+t'}\right\}$. Let
$\phi^{t+t'} = \text{Prob}_{\sigma^*}\left\{x_{t+t'} = i | L_i^{t+t'}\right\}$ be the probability that i is awarded production in period $t+t'$ conditional on $L_i^{t+t'}$. Let
$\lambda^{t+t'} = \text{Prob}_{\sigma^*}\left\{e_{t+t'} = 1 | x_{t+t'} = -i, L_i^{t+t'}\right\}$ be the probability that agent $-i$ chooses $e = 1$ in round $t+t'$ conditional on being chosen in $t+t'$ and the event $L_i^{t+t'}$. Let $z^{t+t'} = \left\{h_0^{t+t'} \in L_i^{t+t'}, x_{t+t'} = -i, e_{t+t'} = 1\right\} \in \mathcal{H}_e^{t+t'}$ be the event that $h_0^{t+t'} \in L_i^{t+t'}$ and agent $-i$ is awarded production and works hard in period $t+t'$.

Define the vector $S^{t+t'} = (S_1^{t+t'}, S_2^{t+t'})$ by

$$S_i^{t+t'} = \sum_{s=t'}^{\infty} \delta^{s-t'}(1-\delta)E\left[s_i^{t+s} | z^{t+t'-1}, y_{t+t'-1} = y_H\right].$$

$S_j^{t+t'}$ is agent j's expected dyad surplus, conditional on the event that $x_{t+t'-1} = -i$, $e_{t+t'-1}y_{t+t'-1} = y_H$, and that this event occurs in no other period between t and $t+t'-1$.

Consider relaxing the problem (34)-(36) by ignoring constraint (36) for agent i at histories in $L_i^{t+t'}$. For agent $-i$, replace (36) with the weaker set of constraints

$$\begin{array}{c} \forall t' \in \mathbb{Z}_+,\ S_{-i}^{t+t'} \geq \tilde{S} \\ \forall t' \in \mathbb{Z}_+,\ S^{t+t'} \text{ feasible in an FDE} \end{array}.$$

The first of these relaxed constraints follows by taking an expectation across histories in period t at which (36) holds for agent $-i$. The second is implied by the definition of an FDE.

Manipulating this problem results in the following relaxed problem:

$$\max_{\{s_j^{t+t'}\}_{j,s}, \{S_j^{t+t'}\}_{j,s}} \sum_{t'=0}^{\infty} \delta^{t'} \left(\begin{array}{c} (1-\delta)E\left[s_1^{t+t'} + s_2^{t+t'} | L_i^{t+t'}\right] \text{Prob}\left\{L_i^{t+t'}\right\} + \\ \delta\left(S_1^{t+t'+1} + S_2^{t+t'+1}\right) \text{Prob}\left\{z^{t+t'}, y_{t+t'} = y_H\right\} \end{array} \right) \tag{37}$$

subject to the promise-keeping constraint: $\forall j \in \{1,2\}$,

$$\sum_{t'=0}^{\infty} \delta^{t'} \left(\begin{array}{c} (1-\delta)E\left[s_j^{t+t'} | L_i^{t+t'}\right] \text{Prob}_{\sigma^*}\left\{L_i^{t+t'}\right\} + \\ \delta S_j^{t+t'+1} \text{Prob}_{\sigma^*}\left\{z^{t+t'}, y_{t+t'} = y_H\right\} \end{array} \right) \geq \hat{S}_j, \tag{38}$$

the relaxed dynamic enforcement constraint

$$\forall t' \in \mathbb{N},\ S_{-i}^{t+t'} \geq \tilde{S}, \tag{39}$$

and the requirement that $\{s_j^{t+t'}\}_{j,t'}$ be feasible and $\{S^{t+t'}\}_{t'}$ be attainable in an FDE,

$$\begin{array}{c} s_j^{t+t'} = 1_{i,t} d_t \left(y_H p_0 + e_t \left(y_H (p_1 - p_0) - c \right) \right) \\ \forall t' \in \mathbb{N},\ S^{t+t'} \text{ feasible in an FDE.} \end{array} \tag{40}$$

Claim 1: $\exists$ a solution $S = (S_1, S_2) \in \mathbb{R}^2$ such that $\forall t' \in \mathbb{N}$, $S^{t+t'} = S$.

Let $\{s_j^{t+t'}\}_{j,t'}$, $\{S^{t+t'}\}_{t' \in \mathbb{N}}$ be a solution to (37)-(40). For $j \in \{1, 2\}$, define

$$S_j = \frac{1}{\sum_{t'=0}^{\infty} \delta^{t'} \text{Prob}_\sigma\{z^{t+t'}, y_{t+t'} = y_H\}} \sum_{t'=0}^{\infty} \delta^t \text{Prob}_\sigma\{z^{t+t'}, y_{t+t'} = y_H\} S_j^{t+t'+1}$$

with $S = (S_1, S_2)$. Keep $\{s_j^{t+t'}\}_{j,t'}$ the same, but set $S^{t+t'} = S$ $\forall t' \in \mathbb{N}$. S is the limit of a convex combination of dyad surpluses attainable in an FDE, and $FDE(\delta)$ is a bounded, closed, convex set. Thus, S is feasible in an FDE by a straightforward application of the equilibrium construction used in Lemma 2. By construction, S satisfies (38) and (39) and leads to an identical value for (37). Therefore, the alternative with $S^{t+t'} = S$ is also a solution to the relaxed problem.

Returning to Derivations

We further relax the problem by ignoring the constraint (38) for agent $-i$. Next, we rewrite and simplify the relaxed problem (37)-(40).

Clearly, in the relaxed problem $e_{t+t'} = 1$ whenever $x_{t+t'} = i$. Therefore,

$$\begin{array}{c} E\left[s_1^{t+t'} + s_2^{t+t'} | L_i^{t+t'} \right] = \\ \phi^{t+t'}(y_H p_1 - c) + (1 - F(\emptyset) - \phi^{t+t'}) \left(y_H p_0 + \lambda^{t+t'} (y_H (p_1 - p_0) - c) \right). \end{array}$$

The events $\left\{ z^{t+k}, y_{t+k} = y_H \right\}_{k=0}^{t'-1}$ collectively partition $L_{-i}^{t+t'}$ and are mutually exclusive. Since $\text{Prob}\{y_{t+k} = y_H | z^{t+k}\} = p_1$, we may write $\text{Prob}\left\{ L_i^{t+t'} \right\} = 1 - \left(\sum_{k=0}^{t'-1} \text{Prob}_\sigma\{z^{t+k}\} p_1 \right)$. Given these two expressions, we can write total surplus as the sum of four terms:

$$\sum_{t'=0}^{\infty} \delta^{t'} \left(\begin{array}{c} (1-\delta)(y_H p_1 - c)\phi^{t+t'} \left(1 - \left(\sum_{k=0}^{t'-1} \text{Prob}_\sigma\{z^{t+k}\} p_1 \right) \right) + \\ (1-\delta) y_H p_0 (1 - F(\emptyset) - \phi^{t+t'}) \left(1 - \left(\sum_{k=0}^{t'-1} \text{Prob}_\sigma\{z^{t+k}\} p_1 \right) \right) + \\ (1-\delta)(y_H(p_1 - p_0) - c)(1 - F(\emptyset) - \phi^{t+t'}) \lambda^{t+t'} \text{Prob}\left\{ L^{t+t'} \right\} + \\ \delta (S_1 + S_2) \text{Prob}\left\{ z^{t+t'} \right\} p_1 \end{array} \right)$$

Two further notes are required. First, $\lambda^{t+t'} \left(1 - F(\emptyset) - \phi^{t+t'} \right) \text{Prob}_\sigma \left\{ L_i^{t+t'} \right\} = \text{Prob}_\sigma \left\{ z^{t+t'} \right\}$

follows immediately from the definitions of these variables. Second,

$$\sum_{t'=0}^{\infty}\delta^{t'}\left(\sum_{k=0}^{t'-1}\text{Prob}_{\sigma}\{z^{t+k}\}\right)=\sum_{k=0}^{\infty}\sum_{t'=k+1}^{\infty}\delta^{t'}\text{Prob}_{\sigma}\{z^{t+k}\}=\sum_{k=0}^{\infty}\text{Prob}\{z^{t+k}\}\frac{\delta^{k+1}}{1-\delta}$$

by the Dominated Convergence Theorem. These expressions may be used to simplify (37) and (38). We omit the details.

We further relax this problem in two ways. First, we ignore constraint (38) for agent $-i$. Second, we assume that $\{\phi^{t+t'}\}_{t'}$ and $\left\{\text{Prob}_{\sigma}\{z^{t+t'}\}\right\}_{t'}$ may be chosen independently of one another.[33] It is straightforward to show that (37) is increasing and (38) is relaxed when $\phi^{t+t'}$ increases, $\forall t'$. Therefore, any solution to the program has $\phi^{t+t'}$ equal to its maximum, $\phi^{t+t'}=\alpha$.

With these relaxations, we can substantially simplify the problem (37)-(40). Following these simplifications, the new problem may be written

$$\max_{\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}\left\{z^{t+t'}\right\},S}\left(\begin{array}{c}\delta\left(\alpha(y_H(p_1-p_0)-c)+y_Hp_0\right)+\\ \sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}\left\{z^{t+t'}\right\}*\left(\begin{array}{c}\delta p_1\left(S_1+S_2-\alpha(y_Hp_1-c)\right)-\\ \delta p_1(1-\alpha-F(\emptyset))y_Hp_0+\\ (1-\delta)(y_H(p_1-p_0)-c)\end{array}\right)\end{array}\right) \tag{41}$$

subject to

$$(1-\delta)\alpha(y_Hp_1-c)+\delta p_1\left(S_i-\alpha(y_Hp_1-c)\right)\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}\left\{z^{t+t'}\right\}\geq\hat{S}_i, \tag{42}$$

(39), and (40).

The optimal S is Pareto efficient. Any Pareto efficient equilibrium has total surplus at least equal to the surplus from exclusive dealing, so $S_1+S_2\geq\alpha(y_Hp_1-c)-(1-\alpha-F(\emptyset))y_Hp_0$. Moreover, $S_i\leq\alpha(y_Hp_1-c)$. Hence, (41) is increasing in the sum $\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}\{z^{t+t'}\}$, while (42) is decreasing in this sum. Fixing S, any strategy that maximizes $\sum_{t'=1}^{\infty}\delta^{t'}\text{Prob}\{z^{t+t'}\}$ subject to the promise-keeping constraint is optimal.

Claim 2: without loss of generality, S satisfies $S_{-i}=\tilde{S}$.

By (39), $S_{-i}\geq\tilde{S}$. Suppose $S_{-i}>\tilde{S}$. If $S_1+S_2=V^{FB}$, then total surplus is V^{FB}; contradiction because first-best is unattainable. So $S_1+S_2<V^{FB}$.

The payoff frontier is convex, so S_1+S_2 is decreasing in S_{-i}. Define $S^*=(S_1^*,S_2^*)$

[33] In reality, these variables are related through the equilibrium.

as the vector of dyad-surpluses in a Pareto efficient FDE that satisfy $S^*_{-i} = \tilde{S}$. Then $S_1^* + S_2^* \geq S_1 + S_2$.

Define $Q = \sum_{t'=1}^{\infty} \delta^{t'} \text{Prob}\left\{z^{t+t'}\right\}$ as the solution to the relaxed problem if continuation dyad surpluses equals S. Let Q^* be the value of this variable such that (42) binds when continuation dyad surpluses equal S^*. $S_i^* > S_i$ because $S^*_{-i} = \tilde{S} < S_{-i}$ and S^* is Pareto efficient. Therefore, $Q^* > Q$, since (42) is decreasing in Q. But $S_1^* + S_2^* \geq S_1 + S_2$ and (41) is increasing in Q, so (41) is weakly greater under S^* and Q^* relative to S and Q. So S^*, Q^* is optimal.

Returning to Derivations

Finally, we prove that (q^*, q^*)-Exclusive Dealing solves the relaxed problem. In this equilibrium, $\phi^{t+t'} = \alpha$, while $\sum_{t'=1}^{\infty} \delta^{t'} \text{Prob}\left\{z^{t+t'}\right\}$ is continuous in q^*. At $q^* = 1$, $\sum_{t'=1}^{\infty} \delta^{t'} \text{Prob}\left\{z^{t+t'}\right\} = 0$ and so (42) is satisfied whenever high effort is sustainable in any equilibrium. At $q^* = 0$, we assume $S_{-i} = \tilde{S}$ by claim 2. Then (42) cannot be satisfied, since if it were then the Favored Producer Allocation would attain first-best. Thus, there exists $q^* \in (0, 1]$ such that (42) binds. This q^* solves the constrained problem and gives agent i exactly $\hat{S}_i$ surplus.

To conclude the proof, note that both agents work hard at the beginning of the game. After one agent produces y_H, Claim 2 implies that they are given exactly $\tilde{S}$ continuation surplus. They can be optimally given this surplus using (q^*, q^*)-Exclusive Dealing. Hence, (q^*, q^*)-Exclusive Dealing is an optimal FDE. ■

Proposition 6

We construct a relational contract that is not a full-disclosure equilibrium and strictly dominates (q^*, q^*)-Exclusive Dealing. Recall $\alpha \equiv F(\{1\}) + F(\{1, 2\})$.

By (11), (5) would be slack if $q^* = 1$. Therefore, $q^* \in (0, 1)$. For any $\epsilon > 0$ such that $q^* - \epsilon \in (0, 1)$, define

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

For sufficiently small $\epsilon > 0$, there exist probabilities $\bar{q}(\epsilon), \underline{q}(\epsilon)$ such that $\bar{q}(\epsilon) > q^* > \underline{q}(\epsilon)$ and $\psi(\epsilon)\underline{q}(\epsilon) + (1 - \psi(\epsilon))\bar{q}(\epsilon) = q^*$.

Restrict attention to periods 1 and 2 and consider the following 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 $t = 2$, if $1 \in \mathcal{P}_2$, then $x_2 = 1$. If $\mathcal{P}_2 = \{2\}$, then $x_2 = 2$. If $x_2 = 1$, then $d_2 = e_2 = 1$. If $x_2 = 2$ but $y_1 = 0$, then $d_2 = e_2 = 1$. If $x_2 = 2$ and $y_1 = y_H$, then with probability $1 - q^* + \epsilon$, agent 2 is asked to choose $d_2 = e_2 = 1$.

Otherwise, agent 2 is asked to choose $d_2 = 1$, $e_2 = 0$. Importantly, if agent 2 is asked to choose $e_2 = 1$ in $t = 2$, then he does not know whether $y_1 = 0$ or $y_1 = y_H$. Therefore, this strategy profile conceals information.

From $t = 3$ onwards, continuation play is chosen from the full-disclosure Pareto frontier. We describe continuation play in terms of dyad surpluses. Among all optimal FDE, let $\hat{S}^i \in \mathbb{R}^2$ be the vector of dyad surpluses from the FDE that maximizes agent i's dyad surplus. Let $\hat{S}^{Ex,i} \in \mathbb{R}^2$ be the dyad surpluses that result when agent i is given exclusive dealing: $\hat{S}_i^{Ex,i} = \alpha(y_H p_1 - c)$, $\hat{S}_{-i}^{Ex,i} = (1-\alpha) y_H p_0$.

If $x_2 = 1$, then continuation play is as in (q^*, q^*)-Exclusive Dealing: $\hat{S}_1^1$ with probability $1 - q^*$ and $\hat{S}_1^{Ex,1}$ otherwise. If $x_2 = 2$ and $e_2 = 0$, then continuation dyad surpluses are $S^{2,L} = \hat{S}^{Ex,1}$. Finally, suppose $x_2 = 2$, $e_2 = 1$, and outputs in periods 1 and 2 are y_1 and y_2, respectively. Then continuation dyad surpluses are $S^{2,H}(y_1, y_2)$, where $S^{2,H}(0,0) = S^{2,H}(y_H, 0) = \hat{S}^1$, $S^{2,H}(y_H, y_H) = \underline{q}(\epsilon)\hat{S}^{Ex,2} + (1 - \underline{q}(\epsilon))\hat{S}^2$, and $S^{2,H}(0, y_H) = \bar{q}(\epsilon)\hat{S}^{Ex,2} + (1 - \bar{q}(\epsilon))\hat{S}^2$.

It remains to show that this strategy is an equilibrium that strictly dominates (q^*, q^*)-Exclusive Dealing. For $t \geq 3$, the continuation equilibrium is full-disclosure and so (5) holds. For $t = 2$ and $x_2 = 1$, this constraint may be written

$$\tilde{S} \leq (q^* \hat{S}_1^{Ex,1} + (1 - q^*)\hat{S}_1^1), \tag{43}$$

which holds because (q^*, q^*)-Exclusive Dealing is an equilibrium. If $t = 2$, $x_2 = 2$, and $e_2 = 1$, note that $\text{Prob}\left\{y_1 = y_H | h_{0,2}^2\right\} = \psi(\epsilon)$ at the beginning of period 2. Therefore, agent 2's expected continuation surplus following $y_2 = y_H$ is

$$\psi(\epsilon) S^{2,H}(y_H, y_H) + (1 - \psi(\epsilon)) S^{2,L}(0, y_H) = q^* \hat{S}_2^{Ex,2} + (1 - q^*)\hat{S}_2^2 \geq \tilde{S}$$

by (43).

Consider $t = 1$ and suppose $y_1 = y_H$. Then agent 1's dyad surplus is

$$S_1 = \begin{array}{c} \alpha(1-\delta)(y_H p_1 - c) + \delta(1 - q^* + \epsilon)\alpha \left\{ p_1 q^* \hat{S}_1^{Ex,1} + (1 - p_1 q^*)\hat{S}_1^1 \right\} + \\ \delta(1 - q^* + \epsilon)(1-\alpha) \left[p_1(\underline{q}(\epsilon)\hat{S}_1^{Ex,2} + (1 - \underline{q}(\epsilon))\hat{S}_1^2) + (1 - p_1)\hat{S}_1^1 \right] + \\ \delta(q^* - \epsilon)\hat{S}_1^{Ex,1}. \end{array}$$

Recall that $\underline{q}(\epsilon) < q^*$ and $\hat{S}_1^{Ex,2} < \hat{S}_1^2$. Hence, if $\epsilon = 0$, then $S_1 > q^* \hat{S}_1^{Ex,1} + (1 - q^*)\hat{S}_1^1$ and so $S_1 > \tilde{S}$ by (43). Therefore, for small $\epsilon > 0$, $S_1(y_H, \epsilon) \geq \tilde{S}$.

Finally, we show that this strategy is strictly better than (q^*, q^*)-Exclusive Dealing. In $t = 1$, total surplus is the same for both. In $t = 2$, the equilibrium that conceals information generates strictly higher surplus because $\epsilon > 0$. At $t = 3$, recall that players are symmetric. Hence, continuation play yields total surplus equal to either $\hat{S}_1^{Ex,1} + \hat{S}_2^{Ex,1}$ or $\hat{S}_1^1 + \hat{S}_2^1$. Under (q^*, q^*)-Exclusive dealing, the probability that continuation surplus at $t = 3$ is $\hat{S}_1^{Ex,1} + \hat{S}_2^{Ex,1}$ is

$$p_1 q^* + (1 - p_1 q^*) p_1 q^*.$$

Consider the probability of exclusive dealing at $t = 3$ in the proposed strategy. Using the definitions of $\psi(\epsilon)$, $\underline{q}(\epsilon)$, and $\bar{q}(\epsilon)$, this probability equals

$$p_1(q^* - \epsilon) + [1 - p_1 q^* + p_1 \epsilon]\, p_1 q^*.$$

Since $\epsilon > 0$ and $p_1 q^* < 1$, the candidate equilibrium is less likely to lead to exclusive dealing than (q^*, q^*)-Exclusive Dealing. Thus, the candidate equilibrium dominates an optimal FDE.
■