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

Isaiah Andrews[†] and Daniel Barron[‡]

September 23, 2015

Abstract

We consider how a firm dynamically allocates business among several suppliers to motivate them in a relational contract. The firm chooses one supplier who exerts private effort. Output is non-contractible, and each supplier observes only his own relationship with the principal. In this setting, allocation decisions constrain the transfers that can be promised to suppliers in equilibrium. Consequently, optimal allocation decisions condition on payoff-irrelevant past performance to make strong incentives credible. We construct a dynamic allocation rule that attains first-best whenever any allocation rule does. This allocation rule performs strictly better than any rule that depends only on payoff-relevant information.

JEL Codes: C73, D82, L14, L22

---

1 Introduction

Many firms rely on relational contracts to motivate their suppliers. For example, informal promises are an essential feature of the supply chains used by Toyota, Honda, Chrysler, and firms with "just-in-time" suppliers. In these relational contracts, firms promise monetary compensation to reward or punish their suppliers. They also promise to allocate business among suppliers to strengthen their relationships (Dyer (1994), Liker and Choi (2004), Krause et al (2000)). The firm cannot formally commit to these promises; instead, they are made credible by the understanding that if the firm betrays a supplier, their relationship sours and surplus is lost.

This paper explores how a principal (downstream firm) can dynamically allocate business among her agents (suppliers) to overcome this commitment problem. In our framework, the principal repeatedly interacts with a group of agents whose productivities vary over time. The principal allocates production to a single agent in each period. The chosen agent exerts a binary effort that determines the probability that the principal earns a profit. The parties have deep pockets, so the chosen agent could in principle be motivated using large bonuses and fines. Output is not contractible, however, so players must have the incentive to follow through on these payments. Compounding this commitment problem, agents do not observe the details of one anothers' relationships with the principal, and so are unable to jointly punish the principal if she reneges on one of them.

We show that allocation decisions make payments credible and so play a central role in optimal relational contracts. The principal motivates an agent by promising him bonuses and high future wages following success and demanding fines and low future wages following failure. If the principal or an agent reneges on these payments, that player is punished with the breakdown of the corresponding bilateral relationship. Consequently, an agent's payoff is bounded from below by his outside option—he would renege on any payments that gave him a lower payoff. His payoff is bounded from above by the total surplus he produces in the future—the principal would renege on any payment

that resulted in a higher payoff. Importantly, this upper bound depends on the principal's future allocation decisions.

The principal would like to promise a large payoff to an agent who performs well, so the upper bound on an agent's payoff binds when he produces high output. To relax this binding constraint, an optimal relational contract allocates more production to an agent after he performs well. We construct a dynamic allocation rule—the **Favored Supplier Allocation** (FSA)—that does this, and prove that it attains first-best whenever any allocation rule does if the set of feasible agent productivities is not too dispersed. To attain first-best, the FSA must allocate production to one of the most productive agents in each period. Among those agents, the FSA chooses the one who has produced high output most recently. Therefore, this allocation rule depends both on an agent's current productivity and on his (payoff-irrelevant) past performance. We show that conditioning on past performance is important: the FSA attains first-best for a strictly wider range of discount factors than any allocation rule that depends only on current productivities.

The FSA *favors past success* in that an agent is allocated production more frequently following good performance. It also *tolerates past failures*: the most recent high performer is allocated business, regardless of how many times he has failed since his last success. The probability that an agent is chosen decreases only if another agent has higher productivity, is allocated production, and performs well.

These properties of the FSA reflect how payments and allocation decisions interact in a relational contract. We can construct transfers so that the principal is indifferent between allocation decisions and so is willing to implement any allocation rule, including the FSA. An optimal allocation rule is therefore only constrained by the need to give each agent strong incentives that are credible in equilibrium. The FSA tolerates past failure because an agent can be punished by low transfers following low output, regardless of the allocation rule. In particular, an agent might be "favored" by the FSA but nevertheless earn a low payoff. In contrast, an agent's maximum payoff following high output is constrained by future allocation decisions; the FSA favors past successes

to relax this constraint when it binds.

In many real-world supply chains, a supplier's past performance plays a pivotal role in whether it is allocated production. Asanuma (1989) documents that in the 1980s, Toyota preferentially allocated business to firms that performed well in the past. Farlow et al (1996) notes that Sun Microsystems uses a similar system, and Krause et al (2000) survey manufacturing firms and conclude that many of them allocate business based on past performance. These relationships are also shaped by the possibility of supplier failure. Consistent with our intuition that the FSA "tolerates past failures," Liker and Choi (2004) argue that companies like Honda do not immediately withdraw business following poor performance. Instead, they impose costs on the supplier—intense scrutiny, extra shifts, and expedited delivery—while maintaining the relationship. "Metalcraft," the pseudonymous company studied by Kulp and Narayanan (2004), is similarly hesitant to pull business from a supplier who performs poorly.

The extensive case-study literature on allocation decisions in supply chains has spurred a limited theoretical literature. Board (2011) is among the first to formally model allocation dynamics. He argues that dynamics can arise in a repeated hold-up problem if suppliers are liquidity constrained. In contrast, our model considers a fundamentally different contracting friction: limited commitment by the principal (due to moral hazard with non-contractible output), rather than liquidity constraints. Section 4 has a detailed comparison. Other contributions include Taylor and Wiggins (1997), who compare competitive and relational supply chains but do not consider dynamics, and Li, Zhang, and Fine (2011), who focus on formal cost-plus contracts.

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) is among the first to consider relational contracts with multiple agents. Calzolari and Spagnolo (2010) analyze how the number of bidders in a procurement auction affects relational incentives, but do not consider allocation dynamics. In related research, Barron and Powell (2015) expand the tools developed here to consider inefficient

policies in relational contracts.

Our model has private monitoring because each agent observes only his own relationship with the principal; Segal (1999) makes a similar assumption in his static analysis of formal contracts. This assumption implies that agents cannot jointly punish the principal and so plays an essential role in our dynamics. Games with private monitoring are difficult to analyze because players condition their actions on different variables (see Kandori (2002) for an overview). Nevertheless, a number of applied papers have investigated settings with private monitoring (e.g., Fuchs (2007), Wolitzky (2013), Ali and Miller (2013), Harrington and Skrzypacz (2011)). In principle, optimal equilibria in our game could be non-recursive and hence very complicated. However, we prove that a relatively simple, recursive allocation rule—the FSA—attains first-best whenever any equilibrium does. To prove this result, we develop tools to handle the complexities of private monitoring in our setting.

2   Model and Assumptions

2.1   Timing

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 period the principal requires a single good that can be produced by any one of the agents. Each agent's productivity is drawn from a finite set and publicly observed. After observing productivities, the principal allocates production to one agent, who either accepts or rejects. If he accepts, that agent exerts binary effort that determines the probability of high output, the value of which depends on his productivity. Utility is transferable between the principal and each agent. At the beginning of the game, the principal and each agent can "settle up" by transferring money to one another. Payments are observed only by the two parties involved.

Formally, we consider the infinite repetition $t = 1, 2, ...$ of the following stage game with common discount factor δ:

1. Productivities $v_t = (v_{1,t}, ..., v_{N,t})$ for agents $i \in \{1, ..., N\}$ are publicly drawn from distribution $F(v)$, with each $v_{i,t} \in \{0, v^1, ..., v^K\} \subseteq \mathbb{R}_+$ for $K < \infty$.

2. The principal publicly chooses agent $x_t \in \{1, ..., N\}$ as the *producer*.

3. The principal pays each agent $i \in \{1, ..., N\}$, who simultaneously pays the principal. Define $w_{i,t} \in \mathbb{R}$ as the resulting *net wage to agent i*. Only the principal and agent i observe $w_{i,t}$.

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 $e_t = y_t = 0$.

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

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

7. The principal pays each agent $i \in \{1, ..., N\}$, who simultaneously pays the principal. Define $\tau_{i,t} \in \mathbb{R}$ as the resulting *net bonus* to agent i. Only the principal and agent i observe $\tau_{i,t}$.

Let $1_{i,t}$ be the indicator function for the event that agent i is allocated production, $1_{i,t} = 1\{x_t = i\}$. Then stage-game payoffs in period t are

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

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

Two features of this model warrant further discussion. First, each agent observes only his own output and pay and cannot communicate with other agents.[2] As a result, the principal can renege on one agent without facing

---

[1]Requiring $w_{i,t} = \tau_{i,t} = 0$ for all $i \neq x_t$ would not change any of our results, though it would alter our discussion of transfers in Section 4.

[2]If agents could directly pay one another, then they could use these payments to communicate. In that case, optimal allocation rules would typically be history-independent (see Andrews and Barron (2013) for a proof).

punishment from the other agents. This assumption drives equilibrium dynamics. We explore it further and compare our results to a game with public monitoring in Section 4. Second, the principal pays a wage before the chosen agent accepts or rejects production. This assumption simplifies punishment payoffs by ensuring that an agent can punish the principal if he receives an out-of-equilibrium wage. Adding a further round of transfers after the agent accepts production but before he exerts effort would not change any of our results.

2.2 Histories, Strategies, and Continuation Payoffs

The *set of histories at the start of period T* is $\mathcal{H}_0^T = \{v_t, x_t, \{w_{i,t}\}, d_t, e_t, y_t, \{\tau_{i,t}\}\}_{t=1}^{T-1}$. 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, $(h_0^T, v_T, x_T, \{w_{i,T}\}_i, d_T, e_T, y_T) \in \mathcal{H}_y^T$. The set of histories is $\mathcal{H} = \bigcup_{T=1}^{\infty} \bigcup_{a \in \mathcal{A} \cup \{0\}} \mathcal{H}_a^T$.

For each agent $i \in \{1, ..., N\}$, let $I_i : \mathcal{H} \to 2^{\mathcal{H}}$ be agent i's information partition over histories. That is, $I_i(h_a^t)$ equals the set of histories that agent i cannot distinguish from $h_a^t \in \mathcal{H}_a^t$.[3] A strategy is a mapping from each player's beliefs at each history to feasible actions at that history, and is denoted $\sigma = (\sigma_0, ..., \sigma_N) \in \Sigma = \Sigma_0 \times ... \times \Sigma_N$.

Definition 1 *For $i \in \{0, ..., N\}$, player i's* continuation surplus *equals $U_{i,t} = \sum_{t'=0}^{\infty} (1 - \delta)\delta^{t'} u_i^{t+t'}$. Define i-dyad surplus as the total surplus produced by agent i:*

$$S_{i,t} = \sum_{t'=0}^{\infty} (1 - \delta)\delta^{t'} 1_{i,t+t'}(y_{t+t'} - ce_{t+t'}).$$

Intuitively, i-dyad surplus is agent i's contribution to total surplus—the surplus from those periods in which agent i is allocated production. We will show that dyad surplus constrains the incentives that can be promised to agent i, and so plays a critical role in our analysis. Note that an agent's beliefs

---

[3]For example, $\forall i \in \{1, ..., N\}$, $I_i(h_0^T) = \{v_t, x_t, w_{i,t}, 1_{i,t}d_t, 1_{i,t}e_t, 1_{i,t}y_t, \tau_{i,t}\}_{t=1}^{T-1}$.

about his continuation surplus condition on his information, rather than the true history: $E_\sigma [U_{i,t} | I_i (h^t)]$.

A *relational contract* is a Perfect Bayesian Equilibrium (PBE) of this game, with equilibrium set Σ^*.[4] We focus on *optimal* relational contracts, which maximize *ex ante total* surplus among all relational contracts: $\max_{\sigma^* \in \Sigma^*} \sum_{i=0}^N E_{\sigma^*}[U_{i,1}]$. This is equivalent to maximizing the principal's surplus because utility is transferable between the principal and each agent.[5] An *allocation rule* is a mapping from the histories immediately following the realization of $\{v_{i,t}\}_{i=1}^N$, $\cup_{t=1}^\infty \mathcal{H}_v^t$, to the agent x_t who is allocated production at that history. Without loss of generality, we restrict attention to the set of relational contracts in which players do not condition on past effort choices.[6]

Let $v_{\max,t} = \max_j v_{j,t}$ be the maximum productivity in period t, and define the set of *most productive agents* as $M_t = \{i | v_{i,t} = v_{\max,t}\}$. A *first-best relational contract* is a relational contract that (i) always allocates production to an agent in M_t, who (ii) chooses $e_t = 1$. Define the resulting *first-best total surplus* $V^{FB} = E[v_{\max,t} p_1 - c]$. Without loss of generality, assume $v^1 < v^2 < ... < v^K$. We maintain the following three assumptions for the entire analysis.

Assumption 1 *F is exchangeable: for any permutation ϕ, $F(v) = F(\phi(v))$.*

Assumption 2 *If $v_{i,t} > 0$ then $v_{i,t} p_1 - c > v_{i,t} p_0$, with $Pr\{v_{\max,t} > 0\} = 1$.*

Assumption 3 *$Pr\{|M_t| > 1\} > 0$ and $Pr\{|M_t| = 1\} > 0$.*

Assumption 1 implies that agents are symmetric, greatly simplifying our analysis. By Assumption 2, $e_t = 1$ is efficient in each period unless $v_{x_t,t} = 0$, and it is always efficient for an agent in M_t to exert effort. Assumption 3 has

---

[4]Adapted from Mailath and Samuelson (2006): a Perfect Bayesian Equilibrium is an assessment (σ^*, μ^*) consisting of strategy profile σ^* and beliefs about the true history $\mu^* = \{\mu_i^*\}_{i=0}^N$. σ_i^* is a best response given μ_i^*, which is updated by Bayes Rule whenever possible. Otherwise, μ_i^* assigns positive weight only to histories that are consistent with $I_i (h^t)$.

[5]Proof: for $\sigma^* \in \Sigma^*$, agent i pays $E_{\sigma^*}[U_{i,1}]$ to the principal at the start of the game, or else is punished with $d_t = 0$ and $w_{i,t} = \tau_{i,t} = 0$ $\forall t \geq 1$. The principal's payoff then equals ex ante total surplus $\sum_{i=0}^N E_{\sigma^*}[U_{i,1}]$.

[6]The proof that this is without loss may be found in Andrews and Barron (2013).

two implications that warrant further discussion: (i) multiple agents sometimes have the same productivity in one period, and (ii) each agent is the unique most productive agent with positive probability.

Consider implication (i). To attain first-best, an equilibrium must allocate production in each period to one of the most productive agents in that period (i.e. one of the agents in M_t). If M_t is always a singleton, then the first-best allocation rule is uniquely determined by exogenous productivity draws. Thus, ties are required for non-trivial allocation dynamics to arise in a first-best relational contract. We discuss to what extent our intuition applies in settings without ties in Section 4.

The assumption that productivities are sometimes tied seems reasonable in real-world supply chains. For example, in their study of "Metalcraft," Kulp and Narayanan (2004) highlight an allocation decision in which at least three suppliers have "comparable" productivities. Two suppliers have similar costs and qualities, while the third has higher cost but higher quality. The case emphasizes that "the costs associated with managing upstream suppliers [are] hard to quantify objectively," implying that it is difficult to strictly rank suppliers. More generally, a supplier's "productivity" includes many components—manufacturing costs, quality, and expertise—that are either difficult to measure or naturally discrete (see LaLonde and Pohlen (1996) for further discussion). Hence, a firm might naturally view several of its suppliers as equally efficient when it allocates production.

Implication (ii) of Assumption 3 is also crucial for equilibrium dynamics, since otherwise the principal could attain first-best by allocating business among a strict subset of agents. In practice, the relative productivities of different suppliers vary over time with changes in the downstream firm's needs and suppliers' capabilities. For example, Metalcraft sometimes chooses "non-preferred" suppliers because preferred suppliers are unable to efficiently manufacture a given part. Similarly, Toyota asks several suppliers to submit proposals for a component, then chooses one based on the quality of its proposal ($v_{i,t}$) and its past performance. A supplier chosen in one model-cycle might be less productive than other suppliers in the next. As a result, Toyota's

suppliers have product portfolios that vary substantially over time (Asanuma 1989).

3   The Favored Supplier Allocation

This section gives conditions under which a simple dynamic allocation rule—the Favored Supplier Allocation—is part of a first-best relational contract whenever any relational contract attains first-best. The definition of the Favored Supplier Allocation, statement of the result, and intuition are given in Section 3.1, with a formal proof in Section 3.2. Section 3.3 explores the incentives of the Favored Supplier Allocation in more detail. We focus on strategies yielding first-best total surplus, so our discussion assumes $e_t = 1$ in each period unless otherwise noted.

3.1   Statement of the Main Result

Suppose the principal could commit to a transfer scheme as a function of output. Then regardless of future allocation decisions, the principal could write a short-term formal contract to hold an agent's payoff at 0 and motivate him to work hard.[7] As a result, an allocation rule that depends only on current productivities and ignores past performance would be optimal.

In contrast, allocation dynamics play an important role in the game without commitment. Our main result concerns what we call the Favored Supplier Allocation, which allocates production among agents according to a simple history-dependent rule.

Definition 2 *For each $t \geq 1$, history h_0^t, and agent i, define $T_i(h_0^t) = \max\{t' < t | x_{t'} = i, y_{t'} > 0\}$ as the most recent time agent i produced high output. The Favored Supplier Allocation (FSA) allocates production in each period to the efficient agent who has most recently produced high output: $x_t \in \arg\max_{i \in M_t} T_i(h_0^t)$. If $T_i(h_0^t) = -\infty$ for all agents in M_t, then $x_t \in M_t$ is chosen at random.*

---

[7]For instance, the principal could set agent x_t's wage equal to $c - p_1 \frac{c}{p_1 - p_0}$, pay a bonus of $\frac{c}{p_1 - p_0}$ if $y_t = v_{x_t, t}$, and pay no bonus if $y_t = 0$.

The Favored Supplier Allocation assigns production to the agent who has produced high output most recently among those agents in M_t. If no agent in M_t has produced high output, then the producer is chosen at random from that set. This allocation rule favors past successes by choosing agents who have performed well over those who have not, and choosing more recent high performers over those who performed well in the distant past. It tolerates past failures by not tracking past low output: the most recent high performer is chosen regardless of his performance in other periods.

Consider a first-best relational contract that uses the Favored Supplier Allocation. Because $e_t = 1$ in each period of this relational contract, an agent's dyad-surplus at the start of a period depends only on the number of agents who have produced high output more recently than him. Define $S_{(j)}^{FSA}$ as the expected dyad-surplus of the jth most recently productive agent—that is, an agent for whom $(j-1)$ other agents have have produced high output more recently. For $j < N$, $S_{(j)}^{FSA}$ can be defined recursively in terms of $S_{(1)}^{FSA}$ and $S_{(j+1)}^{FSA}$, while $S_{(N)}^{FSA}$ can be defined in terms of $S_{(1)}^{FSA}$. The resulting expressions are given in the technical supplement, and may be explicitly obtained as the solution to a system of linear equations.

Our main result finds the lowest discount factor such that a relational contract using the FSA attains first-best. If a further parameter condition is satisfied, then no relational contract attains first-best for smaller discount factors, so the Favored Supplier Allocation attains first-best whenever any allocation rule does.

Proposition 1 *Let $\delta^{FSA} \in (0,1)$ be the smallest δ such that $\frac{\delta}{1-\delta} S_{(1)}^{FSA} \geq \frac{c}{p_1 - p_0}$.[8] Then:*

1. *The FSA is part of a relational contract that attains first best iff $\delta \geq \delta^{FSA}$.*

2. *If $\delta^{FSA} v^K - v^1 \leq (1 - \delta^{FSA}) c \frac{p_0/p_1}{p_1 - p_0}$, then no relational contract attains first-best if $\delta < \delta^{FSA}$. So a relational contract using the FSA attains first-best whenever any relational contract does.*

---

[8]Note that $\frac{1}{1-\delta} S_{(1)}^{FSA}$ is strictly increasing in δ.

Proof: See Section 3.2.

The threshold $\delta^{FSA} \in (0,1)$ may be calculated using the recursive expression for $S_{(1)}^{FSA}$. Note that the parameter constraint in part 2 is always satisfied if there is only one productive type $(K = 1)$.

We focus on developing intuition for Proposition 1 here, deferring the proof to the next subsection. First, we consider how future allocation decisions constrain an agent's equilibrium incentives. As in Levin (2003) an agent's payoff after he produces output is constrained by his expected future production. Agent i's payoff can be no worse than his outside option, 0, since he can always deviate by rejecting future production. Agent i cannot earn more than his expected i-dyad surplus, since the principal would renege on any larger payoff, forfeit that agent's future production, and continue trading with the other agents. Private monitoring leads to this upper bound: following a deviation observed only by agent i, the principal stands to lose only the future surplus produced by *that* agent. Therefore, the upper bound on an agent's reward for high output—and hence his incentive to work hard—depends on his expected dyad surplus following that output.

We show that these payoff bounds are the key constraints imposed by the relational contract. In particular, for any allocation rule, we can construct a relational contract in which (i) an agent earns the lower or upper bounds after producing low or high output, respectively, and (ii) the principal earns payoff 0 at the start of each period and so is willing to implement the allocation rule.

Next, we consider optimal allocation dynamics. In any first-best relational contract the principal chooses exactly one agent in M_t in each period t. Consider a period-t history such that agents 1 and 2 are both in M_t and last produced high output in periods T_1 and T_2, respectively. Assume agent 1 has produced high output more recently: $T_1 > T_2$. As argued above, agent i can be provided stronger incentives to work hard in period T_i if his expected dyad surplus following high output in that period is large. This i-dyad surplus is larger if the principal allocates business to him in period t, but how much it increases depends on two factors.

First, agent i's dyad surplus increases more if $t - T_i$ is small since, from

the perspective of period T_i, production in period t is discounted by δ^{t-T_i}. Agent i's expected dyad surplus also increases more if agent i believes that the period-t history in question is probable, given his information in period T_i. If this probability is larger for agent 1, then choosing agent 1 increases expected 1-dyad surplus in T_1 more than choosing agent 2 would increase expected 2-dyad surplus in T_2. In that case, an optimal allocation rule should choose the agent in M_t who produced high output most recently, which is exactly what the FSA does.[9]

However, the informal logic outlined above is incomplete. In particular, agent 1 in period T_1 might assign a lower probability to a given period-t history than agent 2 in T_2. Agents have different beliefs about the true history because they observe different variables. For example, agent 2 might believe a history in period t is probable given his information in T_2, while agent 1 already knew this same history was impossible when he produced in period T_1. To overcome this complication and prove Proposition 1, we derive a necessary condition for any relational contract to attain first-best. Under the parameter restriction in part 2 of Proposition 1, this necessary condition holds exactly when $\delta \geq \delta^{FSA}$, proving the result.

3.2   Proof of Proposition 1

Let σ^{FSA} be a strategy profile that uses the Favored Supplier Allocation. Dependence on strategies σ is suppressed wherever possible without loss of clarity. Our proof consists of four lemmas. Lemma 1 gives necessary and sufficient conditions for an allocation rule and effort choices to be part of a relational contract. Lemma 2 then establishes a necessary condition for strategies to be part of a first-best relational contract. Under a parameter condition, Lemma 3 shows that this necessary condition holds for any first-best strategy profile σ only if it holds for σ^{FSA}. Finally, Lemma 4 shows that the necessary condition established by Lemma 2 holds only if a first best relational contract using the FSA exists. Hence we establish that, under a parameter condition,

---

[9]We thank an anonymous referee for suggesting an intuition along these lines.

no relational contract may attain first best unless one that uses the FSA does so.

The first step of the proof is to identify necessary and sufficient conditions for an allocation rule and sequence of effort choices to be part of a relational contract.

Lemma 1 1. Let σ^* be a PBE in which $e_t = 0$ whenever $v_{x_t,t} = 0$. Then

$$(1-\delta)\frac{c}{p_1-p_0} \leq \delta E_{\sigma^*}\left[S_{x_t,t+1} \mid I_{x_t}\left(h_y^t\right)\right], \qquad (1)$$
$$\forall \text{ on-path } h_y^t \in \mathcal{H}_y^t \text{ such that } e_t = 1, \ y_t > 0$$

2. Let strategy σ be such that $e_t = 0$ if $v_{x_t,t} = 0$, total surplus equals V, and (1) holds. There exists a relational contract σ^* such that (i) total surplus equals V, and (ii) the joint distribution over $\{v_t, x_t, y_t\}_{t=1}^T$ is identical to that implied by σ for all $T \geq 1$.

Proof: See the supplement.□

Constraint (1) requires that a producing agent x_t expect at least $\tilde{S} = \frac{1-\delta}{\delta}\frac{c}{p_1-p_0}$ dyad surplus whenever he exerts effort ($e_t = 1$) and produces high output ($y_t > 0$). The right-hand side of this constraint equals the largest payoff an agent can receive following high output. Since the lowest payoff an agent can receive following low output equals 0, (1) is a necessary condition for agent x_t's incentive constraint for high effort to hold.

Lemma 1 implies that (i) each agent's equilibrium incentives can be summarized by his dyad surplus following high effort *and* high output, and (ii) we can induce the principal to follow any allocation rule in equilibrium. To prove this lemma, we construct transfers so that each agent i earns $E\left[S_{i,t}|h_0^t\right]$ continuation surplus at the start of each period. Then the principal earns 0 and so is willing to follow the equilibrium allocation rule. Agent x_t is willing to pay his expected continuation surplus as a fine following low output because he would lose that surplus following a deviation. So agent x_t earns $\delta E[S_{x_t,t+1}|I_{x_t}(h_y^t)]$ if $y_t > 0$ and 0 if $y_t = 0$, which can be plugged into his effort incentive constraint to yield (1). In this construction, the principal is tempted to renege on the

14

wage she pays the producing agent, while that agent is tempted to renege on the fine he pays following low output.

Ceteris paribus, an allocation rule and efforts that satisfy (1) at δ also satisfy (1) for $\delta' > \delta$. Hence, Lemma 1 implies part 1 of Proposition 1.

Condition (1) depends on agent beliefs and so is challenging to check for an arbitrary strategy profile. We handle this complication by finding necessary conditions for (1) to be satisfied at each history, which is required in any first-best relational contract. To satisfy (1), an allocation rule must ensure that each agent expects at least $\delta \tilde{S}$ dyad-suplus whenever that agent produces high output.

Let $H\left(h^t\right) = \{i : \text{ there exists } s < t \text{ such that } x_s = i, y_s > 0\}$ denote the set of agents who have produced high output before period t. For any first-best strategy σ, define $\beta_{i,s}^L = \Pr_\sigma \{x_t = i, i \notin H(h^t)\}$ as the probability that agent i is allocated production and has never produced high output in the past, and $\beta_{i,s}^H(k) = \Pr_\sigma \left\{v_{\max,t} = v^k, x_t = i, i \in H(h^t)\right\}$ as the probability that $v_{\max,t} = v^k$, agent i is allocated production, and i has produced high output in the past. Let $V_k^{FB} = v^k p_1 - c$.

Definition 3 *For any first-best strategy σ, define the* expected obligation *owed to agent i at time t, $\Omega_{i,t}$, as $\Omega_{i,1} = 0$ and*

$$\Omega_{i,t} \equiv \beta_{i,t}^L p_1 \delta \tilde{S} - (1 - \delta) \sum_{k=1}^{K} \beta_{i,t}^H(k) V_k^{FB} + \frac{\Omega_{i,t-1}}{\delta}. \tag{2}$$

For each $t \geq 1$, one can show that $E_\sigma[S_{i,t}] \geq \Omega_{i,t}$ in any first-best relational contract, where the expectation is taken with respect to beliefs at the start of the game. The first term in $\Omega_{i,t}$ is a lower bound on the additional expected i-dyad surplus in equilibrium based on agent i's production in period t: he produces high output for the first time in that period with probability $\beta_{i,t}^L p_1$, in which case he must produce $\delta \tilde{S}$ expected dyad surplus in the continuation game by (1). In subsequent periods, agent i produces some of this dyad surplus whenever he is allocated production. The second term captures this fact: the probability that agent i has both previously produced high output and is

allocated production in period t with $v_{\max,t} = v^k$ is $\beta_{i,t}^H(k)$, and in this case i produces $(1-\delta)V_k^{FB}$ surplus. The final term, $\frac{\Omega_{i,t-1}}{\delta}$, captures any expected dyad surplus from previous periods that agent i has not yet produced.

If $\Omega_{i,t}$ diverges to ∞, then (1) does not hold in expectation across histories and so must be violated at some history on the equilibrium path. Thus, $\Omega_{i,t}$ must be bounded from above in any first-best relational contract.

Lemma 2 *In any first-best relational contract* σ, $\limsup_{t\to\infty} \Omega_{i,t} < \infty$ *for all* $i \in \{1, ..., N\}$.

Proof: Let $b_i^L(h^t) = 1\{x_t = i\}\,1\{i \notin H(h^t)\}$ indicate the event that i is allocated production in period t and has never previously produced high output. Define $1_{i,t}(k) = 1\{x_t = i, v_{\max,t} = v^k\}$, and let $b_i^H(h^t, k) = 1_{i,t}(k)\,1\{i \in H(h^t)\}$ indicate the event that i is allocated production in period t, has previously produced high output, and the maximal productivity is v^k.

In a first-best relational contract, (1) must hold the first time an agent produces $y_t > 0$. Multiplying both sides of (1) by $b_i^L(h^t)\,1\{y_t > 0\}\,\delta^t$, summing across $s = 1, ..., t$, and taking expectations in $t = 0$ yields

$$\sum_{s=1}^{t} \delta^s \beta_{i,s}^L p_1 \delta \tilde{S} \le E_\sigma \left[\sum_{s=1}^{t} \sum_{s'=1}^{\infty} \sum_{k=1}^{K} (1-\delta)\delta^{s+s'} b_i^L(h^s)\,1\{y_s > 0\}\,b_i^H\left(h^{s+s'}, k\right) V_k^{FB} \right].$$
(3)

We can isolate expectations that depend on events in periods $s > t$ by rewriting the right-hand side

$$R(i,t) + E_\sigma \left[\sum_{s=1}^{t} 1\{i \in H(h^s)\} \sum_{k=1}^{K} (1-\delta)\delta^s b_i^H(h^s, k) V_k^{FB} \right], \qquad (4)$$

where $R(i,t) = (1-\delta)E_\sigma \left[1\{i \in H(h^{t+1})\} \sum_{s=t+1}^{\infty} \sum_{k=1}^{K} \delta^s V_k^{FB} b_i^H(h^s, k) \right]$.

Since $E_\sigma \left[b_i^H(h^s, k)\,1\{i \in H(h^s)\} \right] = \beta_{i,s}^H(k)$, we can re-arrange (3) to obtain $\Omega_{i,t} \le \delta^{-t} R(i,t)$. But $\delta^{-t} R(i,t) \le \delta V_K^{FB}$, so $\limsup_{t\to\infty} \Omega_{i,t} \le \delta V_K^{FB}$ in any first-best relational contract.$\square$

The FSA treats agents symmetrically ex ante, so its implied obligation

is symmetric as well, $\Omega_{i,t}^{FSA} = \Omega_t^{FSA}$. We next show that if Ω_t^{FSA} diverges, then under a parameter restriction obligation must diverge for any first-best strategy.

Lemma 3 *Suppose* $\delta V_K^{FB} \leq V_1^{FB} + p_1 \delta \tilde{S}$. *If* $\limsup_{t \to \infty} \Omega_t^{FSA} = \infty$, *then for any first-best strategy* σ, $\exists i \in \{1, ..., N\}$ *such that* $\lim_{t \to \infty} \Omega_{i,t} = \infty$.

Proof: Given any strategy σ, define a symmetric strategy profile by randomizing agent identities in $t = 1$ and then playing as in σ. Obligation diverges in the symmetric strategy profile only if it diverges for at least one agent in σ. So we can restrict attention to symmetric strategy profiles.

Given a symmetric first-best strategy $\tilde{\sigma}$ with obligation $\tilde{\Omega}_t$, Lemma A.2 proves that $\Omega_t^{FSA} \leq \tilde{\Omega}_t$ for all t if $\delta V_K^{FB} \leq V_1^{FB} + p_1 \delta \tilde{S}$. The proof constructs a decreasing sequence of obligations Ω_t^n such that $\Omega_t^0 = \tilde{\Omega}_t$ and $\lim_{n \to \infty} \Omega_t^n = \Omega_t^{FSA}$. We use the restriction $\delta V_K^{FB} \leq V_1^{FB} + p_1 \delta \tilde{S}$ to prove that obligation is minimized by an allocation rule that chooses agents who have already produced high output whenever at least one such agent is in M_t. Then $\Omega_t^{FSA} \leq \tilde{\Omega}_t$ for any symmetric $\tilde{\sigma}$, implying Lemma 3.□

Lemma 3 shows that, under a parameter condition, divergent FSA obligation implies divergent obligation for all first-best strategies. Thus, Lemmas 2 and 3 together imply that first-best is attainable only if Ω_t^{FSA} is bounded above as $t \to \infty$. The final step of our proof shows that Ω_t^{FSA} diverges exactly when the Favored Supplier Allocation is not part of a relational contract that attains first-best.

Lemma 4 *If* $\delta < \delta^{FSA}$, *then* $\limsup_{t \to \infty} \Omega_t^{FSA} = \infty$.

Proof: We need only show that $\limsup_{t \to \infty} \Omega_t^{FSA} = \infty$ if $S_{(1)}^{FSA} < \tilde{S}$, since this condition is equivalent to $\delta < \delta^{FSA}$. Using the definition of $S_{(1)}^{FSA}$ and repeating the derivation of (3)-(4) with $S_{(1)}^{FSA}$ in place of $\tilde{S}$ yields

$$\sum_{s=1}^{t} \delta^s \beta_s^{L,FSA} p_1 \delta S_{(1)}^{FSA} - \sum_{s=1}^{t} \delta^s \sum_{k=1}^{K} \beta_s^{H,FSA}(k) V_k^{FB} = R^{FSA}(t) \qquad (5)$$

17

where $\beta_s^{L,FSA}$, $\beta_s^{H,FSA}(k)$, and $R^{FSA}(t)$ are the analogues for $\beta_{i,s}^L$, $\beta_{i,s}^H$, and $R(i,t)$ if $\sigma = \sigma^{FSA}$, where we drop the i index because the FSA is symmetric.

The argument in Lemma 2 implies that $0 \leq \delta^{-t}R^{FSA}(t) \leq \delta V_K^{FB}$ for all t. The left-hand side of (5) equals $\delta^t \Omega_t^{FSA}$ if $\tilde{S} = S_{(1)}^{FSA}$. It follows that Ω_t^{FSA} is asymptotically bounded for $\tilde{S} = S_{(1)}^{FSA}$, from which it is straightforward to see that it diverges for $\tilde{S} > S_{(1)}^{FSA}$. $\square$

Finally, we relax the condition $\delta V_K^{FB} \leq V_1^{FB} + p_1 \delta \tilde{S}$ from Lemma 3. If this condition holds at δ^{FSA}, then it holds *a fortiori* for $\delta < \delta^{FSA}$. So Lemmas 2-4 imply that no relational contract attains first-best for $\delta < \delta^{FSA}$. The parameter constraint at δ^{FSA} may be written $\delta^{FSA}v^K - v^1 \leq (1 - \delta^{FSA})c\frac{p_0/p_1}{p_1 - p_0}$. $\blacksquare$

3.3 Implications of Proposition 1

This section presents two corollaries of Proposition 1 that explore the incentives provided by the Favored Supplier Allocation. The first proves that the FSA performs strictly better than a natural benchmark: an allocation rule that does not condition on past performance. The second considers the role of information in the FSA.

A relational contract is *stationary* if on the equilibrium path, actions in one period are independent of previous periods and of t. Corollary 1 proves that a relational contract using the Favored Supplier Allocation attains first-best for a strictly wider parameter range than any stationary relational contract.

Corollary 1 *Let $\delta^{Stat} \in (0,1)$ satisfy $\frac{1 - \delta^{Stat}}{\delta^{Stat}}\frac{c}{p_1 - p_0} = \frac{1}{N}V^{FB}$. Then (i) a stationary relational contract attains first-best if and only if $\delta \geq \delta^{Stat}$, and (ii) $\delta^{Stat} > \delta^{FSA}$.*

Proof: See the supplement. $\square$

A stationary allocation rule does not respond to past performance to make large rewards to high performers credible, and so cannot perform as well as a relational contract that uses the FSA. To illustrate this point, Figure 1 shows how δ^{FSA} and δ^{Stat} change with the parameters p_1 and $Pr\{|M_t| = 2\}$ in a two-agent example, holding the other parameters constant.
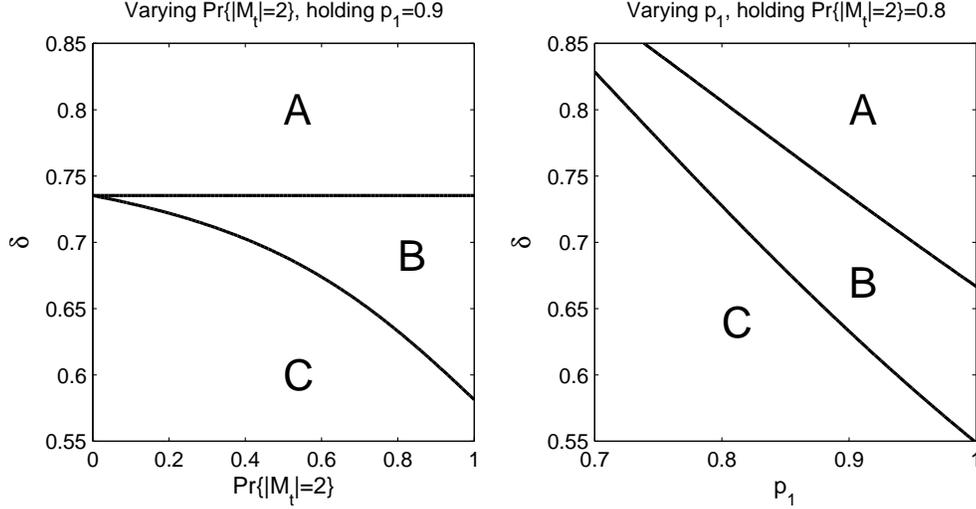
18

Figure 1: Comparing δ^{FSA} and δ^{Stat}. Both panels assume $N = 2$, $p_0 = 0$, $c = 5$, $K = 1$, and $v^1 = 10$. The left panel holds $p_1 = 0.9$ and varies δ and $Pr\{|M_t| = 2\}$. The right panel holds $Pr\{|M_t| = 2\} = 0.8$ and varies δ and p_1. In region A, both stationary and non-stationary relational contracts can attain first-best. In region B, no stationary relational contract attains first-best, but one that uses the FSA does. In region C no first-best relational contract exists.

Finally, we consider the role of information in the Favored Supplier Allocation. The constraint (1) must hold only with respect to agent x_t's information, $I_i(h_y^t)$. In particular, the allocation rule might satisfy (1) but "conceal information" by setting $\delta E[S_{x_t,t+1}|h_y^t] < (1-\delta)\frac{c}{p_1-p_0}$ for some h_y^t, so that agent x_t would be unwilling to exert effort if he learned the true history. The Favored Supplier Allocation does not conceal information in this way.

Corollary 2 *In any first-best relational contract σ^{FSA} that uses the Favored Supplier Allocation, (1) holds with respect to the true history:*

$$(1-\delta)\frac{c}{p_1-p_0} \leq \delta E_{\sigma^{FSA}}\left[S_{x_t,t+1} \mid h_y^t\right], \tag{6}$$
$$\forall \text{ on-path } h_y^t \in \mathcal{H}_y^t \text{ such that } e_t = 1, y_t > 0.$$

Proof: $E_{\sigma^{FSA}}\left[S_{x_t,t+1}|h_y^t\right] = S_{(1)}^{FSA}$ if $y_t > 0$, so (6) holds for $\delta \geq \delta^{FSA}$. ∎

It can be shown that (6) is a necessary and sufficient condition (in the sense

of Lemma 1) for an allocation rule to be part of a belief-free equilibrium.[10]
Hence, Corollary 2 implies that under the conditions of Proposition 1, a belief-free equilibrium that uses the Favored Supplier Allocation attains first-best whenever any Perfect Bayesian Equilibrium does.

If first-best is not attainable, then the optimal equilibrium does not necessarily satisfy (6). Andrews and Barron (2013) considers an example in which first-best is unattainable and constructs a relational contract that strictly dominates any relational contract satisfying (6). Hence, when first best is unattainable, the principal can use the fact that agents do not observe the true history to induce more effort, but (under the conditions of Proposition 1) concealing information cannot expand the set of discount factors for which first-best can be attained.

4   Discussion

Alternative Transfer Schemes -   In the proof of Lemma 1, transfers ensure that each agent's continuation payoff equals his dyad surplus. Therefore, the principal's continuation payoff equals 0 in every period except the very start of the game. This section discusses alternative transfers schemes that also implement the FSA, including schemes in which the principal earns a strictly positive payoff in each period.

Consider the following equilibrium. Every agent pays a fixed "participation fee" to the principal in each period. The principal allocates production as in the FSA, pays the chosen agent a bonus if he produces high output, and otherwise demands a fine from that agent. Participation fees are bounded below by 0 and bounded above by the amount an agent currently in the last rank of the FSA would be willing to pay, $\delta S_{(N)}^{FSA}$. Hence, the principal's continuation payoff in each period can be anywhere between 0 and $N\delta S_{(N)}^{FSA}$ in these relational contracts.

In Appendix B, we show that a relational contract of this form attains first-best whenever $\delta \geq \delta^{FSA}$. Each agent is induced to exert effort by a

---

[10]See Ely, Hörner, and Olszewski (2005) for further discussion of belief-free equilibria.

combination of bonuses and fines. The principal's payoff is independent of the allocation decision and so she is willing to follow the FSA. Agent i stops paying his participation fee if he observes a deviation, so the principal has an incentive to pay a bonus if an agent produces high output. As a result, the principal can earn a strictly positive continuation payoff in a relational contract that implements the FSA. Furthermore, if $\delta > \delta^{FSA}$, then we can reduce the participation fees paid by non-producers without decreasing the surplus earned by the principal. Indeed, for sufficiently patient players the principal can earn the entire expected surplus in each period and motivate agents through bonuses.

In practice, agent participation fees might take the form of "pay-to-play" contracts, in which a supplier makes an explicit monetary payment to secure a spot in the downstream firm's supply chain. More generally, suppliers might be "awarded" a routine contract that gives them negative economic profit in each period. A supplier accepts this negative-profit contract in the hopes of being allocated more lucrative contracts in the future. In that case, the negative-profit contract serves as a participation fee paid by all suppliers and entails no non-contractible effort, while the allocation decision determines which supplier receives a contract that requires non-contractible effort in each period.

The Role of Private Monitoring - In our model, allocation dynamics arise because each agent observes only his own relationship with the principal. Indeed, it can be shown that a stationary allocation rule would be optimal under public monitoring. Suppose that all variables except effort e_t are publicly observed. Agents can still earn no less than 0 in equilibrium. However, the agents can now jointly punish the principal if she betrays any one of them. So the principal is willing to pay the total continuation surplus produced by *every* agent, $\delta E\left[\sum_{i=1}^{N} S_{i,t+1} | h_y^t\right]$, following high output. Every agent can be motivated equally well regardless of the allocation rule because these bounds are the same for each agent. In particular, a stationary allocation rule would be optimal. See Barron and Powell (2015) for details.

Relationship to Board (2011) - In Board (2011), the principal chooses to invest in a single agent in each period at a cost that varies across agents. The chosen agent earns a payoff from which he can choose to repay the principal. In a principal-optimal equilibrium, the principal promises future rent to an agent to induce him to repay today. These rents serve as "switching costs" that encourage the principal to bias trade towards a group of "insider" agents.

Our paper differs from Board's analysis in two crucial ways. First, the allocation rule matters for fundamentally different reasons in the two papers. In Board, suppliers are liquidity-constrained. As a result, the principal must promise her agents rents to solve the hold-up problem. The principal then prefers to trade with agents who have already been promised rents, since she would have to promise additional rents to trade with an agent who has not already been chosen. Board shows that these biases arise if the principal can commit to an allocation scheme, then proves that identical biases are optimal without commitment if players are patient. In contrast, our allocation dynamics arise precisely because the principal **cannot** commit. We do not have liquidity constraints, so the principal need not promise her agents rents to motivate them; indeed, our equilibrium construction in Proposition 1 requires the agents to pay the principal. In our setting, the allocation rule induces the principal not to renege on her promised compensation to the agents.

Second, we consider a setting with moral hazard. As noted in the introduction, occasional failures are an inescapable feature of many business relationships. This focus leads to new equilibrium dynamics. Our results apply to the setting without moral hazard by setting $p_1 = 1$ and $p_0 = 0$. While Proposition 1 applies in this case, the FSA would make no distinction between past allocation decisions and past performance on the equilibrium path because high output would occur with probability one. The fact that the FSA tracks the past successes of each agent, but ignores the past failures, is a potentially surprising result of allowing moral hazard.

Conclusion - The Favored Supplier Allocation entails meaningful dynamics only if the set of most productive agents, M_t, contains multiple agents with

positive probability in each period. If productivities $\{v_{i,t}\}_i$ were instead drawn from a continuous distribution, then M_t would be a singleton in each period and hence every first-best allocation rule would be stationary.

However, non-stationary allocation rules can still be optimal if first-best is unattainable. A suitable analogue of Lemma 1 continues to hold in this setting, so non-stationarity arises from a straightforward generalization of (1). Dynamics might be required for the principal to credibly promise sufficiently strong incentives to motivate high effort, even though they would imply that less productive agents are sometimes allocated production. However, constructing optimal equilibria if first-best is unattainable is intractable with a continuum of productivities. Andrews and Barron (2013) analyze optimal relational contracts if first-best is unattainable using an example with two feasible productivities, but a full characterization is difficult even in this simple setting.

While we have focused on supplier relationships, allocation rules matter in many settings. Managers allocate tasks and promotions among their employees. CEOs allocate scarce time and attention among their divisions. Our analysis suggests that these policies may be fundamentally shaped by commitment problems in long-term relationships.

References

[1] Ali, Nageeb, and David Miller. 2013. "Enforcing Cooperation in Networked Societies." Working Paper.

[2] Andrews, Isaiah, and Daniel Barron. 2013. "The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents." Working Paper. Available at: https://sites.google.com/site/danielbarronecon/

[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] Barron, Daniel, and Michael Powell. 2015. "Policy Commitments in Relational Contracts." Working Paper.

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

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

[7] Calzolari, Giacomo and Giancarlo Spagnolo. 2010. "Relational Contracts and Competitive Screening." CEPR Discussion Paper No. 7434.

[8] Dyer, Jeffrey. 1996. "How Chrysler Created an American Keiretsu." *Harvard Business Review*, July.

[9] Ely, Jeffrey, Johannes Hörner, and Wojciech Olszewski. 2005. "Belief-Free Equilibria in Repeated Games." *Econometrica* 73 (2): 377-415.

[10] 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.

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

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

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

[14] 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.

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

[16] Lalond, Bernard, and Terrance Pohlen. 1996. "Issues in supply chain costing." *The International Journal of Logistics Management* 7 (1): 1-12.

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

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

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

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

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

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

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

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

[25] 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.

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