

Partnership in a Dynamic Production System with Unobservable Actions and Noncontractible Output

Erica L. Plambeck

Graduate School of Business, Stanford University, Stanford, California 94305, elp@stanford.edu

Terry A. Taylor

Graduate School of Business, Columbia University, New York, New York 10027, tat2002@columbia.edu, and
Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755, terry.taylor@dartmouth.edu

This paper considers two firms that engage in joint production. The prospect of repeated interaction introduces dynamics, in that actions that firms take today influence the costliness and effectiveness of actions in the future. Repeated interaction also facilitates the use of informal agreements (relational contracts) that are sustained not by the court system, but by the ongoing value of the relationship. We characterize the optimal relational contract in this dynamic system with double moral hazard. We show that an optimal relational contract has a simple form that does not depend on the past history. The optimal relational contract may require that the firms terminate their relationship with positive probability following poor performance. We show how process visibility, which allows the firms to better assess who is at fault, can substantially improve system performance. The degree to which process visibility eliminates the need for termination depends on the nature of the dynamics: If the buyer's action does not influence the dynamics, the need for termination is eliminated; otherwise, termination may be required.

Key words: relational contracts; Markov decision process; double moral hazard

History: Accepted by Sunil Chopra, operations and supply chain management; received December 31, 2004.

This paper was with the authors 4 months for 2 revisions.

1. Introduction

The success of a buyer-manufacturer relationship in creating a product often depends on the actions of both parties. Each firm has comparative strengths in understanding different aspects of the product's design and the technical aspects related to its production. Accordingly, the buying firm may provide critical technical expertise to assist the supplier in design, engineering, and production. Although a manufacturer typically provides the bulk of the infrastructure for production, a buying firm may provide critical inputs, either in the form of specialized equipment or raw materials. The output of the production process, then, depends on the diligence with which both firms provide the associated physical and managerial inputs.

When firms engage in joint production, each firm observes the output of the production process, but cannot observe the full scope of actions taken by its partner. This causes a free-rider problem: Each firm shirks on costly, unobservable action because the benefits would be shared with its partner (Holmstrom 1982). The free-rider problem is exacerbated by stochasticity in the output (each firm knows that it can assert that a bad outcome is due to the failure

of the other firm or simple bad luck) and by difficulties in contracting on the output. We focus on a setting with unobservable actions and stochastic, noncontractible output.

When a buyer and manufacturer interact, they should evaluate the potential for continuing to do business in the future. The prospect of potential future interaction shapes how firms behave in two ways. First, the prospect of future interaction facilitates the development of trust and cooperation. Firms will be more hesitant to behave opportunistically if they anticipate that doing so will damage their prospects for engaging trading partners in the future. Second, the prospect of future interaction introduces dynamics as conditions and the relationship itself evolves over time. In particular, a firm's actions today impact the costliness and effectiveness of actions in the future. For example, a manufacturer may invest in technologies that reduce its costs of producing certain types of products in the future. Such an investment could impact the effectiveness of the buying firm's action positively or negatively. For example, if the buyer is unfamiliar with the technologies, this may reduce the utility of the buyer's production expertise. Further, exogenous factors, such as general economic condi-

tions, evolve over time, and these also shape the costliness and effectiveness of the firms' actions.

Biopharmaceutical contract manufacturing is our primary motivating example. A drug developer (e.g., Eli Lilly or Genentech), outsources production to a contract manufacturer (e.g., Lonza) with the intention of collaborating over the lifetime of the drug. At the outset, the drug developer may contribute scientists and engineers to help the contract manufacturer develop an efficient manufacturing protocol. Thereafter, each month, the drug developer provides genetically modified mammalian cells and culture medium; the manufacturer cultures (multiplies) the cells in a metal tank and then extracts therapeutic protein. The yield of therapeutic protein from this batch process depends on the quality of the raw materials contributed by the drug developers and also depends on microbial contamination, temperature, pH, pressure, and other process variables controlled by the manufacturer. The drug developer contracts to buy a specific number of batches per unit time at a specified price per batch. However, the yield per batch is stochastic and noncontractible.

Firms can provide stronger incentives for action by developing informal agreements that make payments contingent on noncontractible output (see §2 for an example in the biopharmaceutical context). Because such payments are discretionary, they must be enforced by the value of the ongoing cooperative relationship rather than the court system. Our objective is to characterize how firms should optimally structure informal agreements in the face of dynamics and the temptation to free ride.

We also examine the value of jointly monitoring the production process. For example, in the biopharmaceutical industry, firms can invest in information systems that enable manufacturers and their buyers to observe detailed process data. This helps the firms to identify and disentangle problems with process control (the manufacturer's responsibility) and with raw materials (the buyer's responsibility). In the language of economists, the information systems provide signals of the firms' actions that are commonly observed by both firms. We show how informal agreements should be adapted to use such signals and that doing so can substantially increase the firms' expected profits.

The primary vehicle in economic theory for studying long-term relationships where trust, cooperation, and reputation are important is the repeated game, in which players face the same "stage game" in every time period and each player seeks to maximize the discounted sum of his payoffs. Typically, a repeated game has many possible Nash equilibria, but the players are assumed to coordinate on one that is mutually advantageous. Cooperation is enforced by the threat

of transition to an undesirable Nash equilibrium in the continuation game.

Klein and Leffler (1981) and Taylor and Wiggins (1997) consider settings where product quality is noncontractible and is solely a function of the manufacturer's effort. Klein and Leffler show that in a competitive market, buyers will pay a premium above variable production cost to firms that maintain a reputation for high quality. In Taylor and Wiggins (1997), a buyer inspects every shipment from her manufacturer and rejects faulty items. Taylor and Wiggins show how the buyer can avoid costly inspection by paying a premium for every shipment and threatening to terminate this practice if he later discovers faulty items.

Baker et al. (2001, 2002) emphasize that players may shape their repeated game through transfer payments. They have popularized the term *relational contract* for an informal agreement regarding actions and voluntary payments, enforced by reputational concerns, between parties that interact repeatedly. They study a repeated game with relationship-specific investment by one party ("hold up") and derive insights regarding optimal ownership structure. Levin (2003) examines relational contracting in a principal-agent model with moral hazard or hidden information. Levin proves that if an optimal relational contract exists, then a simple stationary optimal relational contract exists. That is, in searching for an optimal relational contract, one may restrict attention to stationary relational contracts. For the case with moral hazard, under a stationary optimal relational contract, Levin shows that the relationship is never terminated on the equilibrium path, and the voluntary payment to the agent is "one-step:" a bonus if output exceeds a threshold.

Whereas economists have focused on i.i.d. repeated games, sociologists and organizational scholars have observed through field studies and surveys that relational contracts evolve dynamically. Granovetter (1985), Powell (1990), Sako (1992), and Uzzi (1996, 1997) stress the role of personal relationships and social sanctions in generating and sustaining interorganizational agreements, and argue that trust and collaboration between trading partners can only be developed gradually over time. An alternative perspective is that successful collaboration, especially between large firms, requires developing and refining *institutional* routines for partnering and reciprocity (Powell et al. 1996, Zaheer et al. 1998). In a survey of suppliers to automobile manufacturers, Dyer and Chu (2000) find that ongoing automaker assistance with quality, cost reduction, and inventory has the greatest positive effect on supplier trust and cooperation. Frequency of trade, the length of the relationship, and strong potential for future business are also positively associated with trust and cooperation between

industrial buyers and suppliers (Heidi and Miner 1992, Dyer and Chu 2000). For automobile manufacturers and their suppliers, Dyer (1996) concludes that sequential and cumulative relationship-specific investments tend to reduce the new-model-cycle-time and increase profitability. In the biotechnology industry, Robinson and Stuart (2002) observe that through rapidly evolving *networks* of alliances, firms learn about, monitor, and sanction their (prospective) R&D and manufacturing partners. Kenworthy et al. (1996) observe a breakdown in relational contracting in the U.S. automobile industry in the 70s and early 80s, and attribute this to oil-price shocks and increased foreign competition.

This paper extends the theory of relational contracts to the *collaborative* and *dynamic* settings of interest to researchers in supply chain and operations management. Technically, this paper generalizes Levin's (2003) model of relational contracting with i.i.d. one-sided moral hazard by incorporating double moral hazard (unobservable actions by the principal (buyer) as well as the agent (supplier)), and by incorporating Markovian dynamics (the buyer's and supplier's actions stochastically influence output in the current period and the state of the system in the next period). Our model is formulated in §2. In §3 we construct a simple optimal relational contract. In contrast to Levin (2003), we have a discrete state space and a continuous random variance in the output, which together guarantee existence of an optimal relational contract. Our simple optimal relational contract has very different structural properties than in Levin's basic i.i.d. repeated game with one-sided moral hazard. It is not stationary. Actions depend only on the current state, but payments depend upon the observed state transition and output. Furthermore, §4 identifies conditions under which an optimal relational contract may require that the firms terminate their relationship with positive probability following an undesirable transition. Termination may be required even when the firms observe an independent signal for the action of each firm that allows them to assign blame. Interestingly, the need for termination and the dynamic structure of the game are intertwined. If the buyer's action in period t influences only the output in period t , but not the state in subsequent periods, then the relationship is *never* terminated in equilibrium under the optimal relational contract. Section 5 applies the theory both to illustrate how relational contracts can be used when production is a collaborative process and to demonstrate the value of process visibility. Section 6 provides concluding remarks.

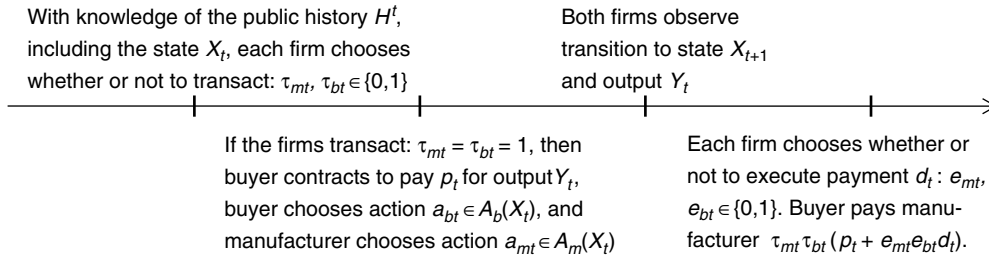
Two recent working papers, developed independently of ours, also consider double moral hazard and relational contracting, but in a stationary environment with common observation of an independent signal

for the effort of each firm. The paper that is closest in spirit to ours, Doornik (2006), shows that the optimal relational contract requires terminating the cooperative relationship when the signal for both firms is low and, if the relationship continues, a one-step payment analogous to Levin (2003). Doornik's formulation is more general in that it allows both firms to receive a portion of the output produced, whereas we consider the case where the buyer receives all the output. Our formulation is more general in that it allows for dynamics and considers the case where independent signals of effort are not available. Rayo (2004) characterizes optimal ownership structure and optimal (within a limited class) relational contracts in repeated team production. In particular, he restricts attention to nonterminating relational contracts. Ownership determines the allocation of profit in the event of a disagreement, i.e., refusal to execute the transfer payments specified in the relational contract. Rayo shows that when the signals are very noisy, ownership of 100% of joint output should be assigned to a single player (as to the buyer in our paper). Although we focus on the case where the buyer receives the output, all theoretical results in §3 and §4, with the exception of Propositions 4 and 5, hold with a general division of ownership of output. Structural differences between the results in Rayo (2004) and Doornik (2006) and this paper show how supply chain partners should adapt their relational contracts to a dynamic business environment.

The aforementioned papers all assume that firms have common knowledge of cost structure and how effort influences the distribution of output and signals. For analysis of collaborative production under asymmetric information, we refer the reader to Iyer et al. (2005) and references therein.

2. Model

Joint production is modeled as a dynamic game. The state of the system at the beginning of period t , X_t , takes values in a finite, discrete state space \mathcal{X} . The state reflects both external factors, such as economic conditions, and internal factors, such as the capabilities of the firms. At the beginning of each period t , the buyer (she) and the manufacturer (he) decide whether or not to transact. If both parties agree to transact, they sign a formal contract under which the buyer commits to pay a constant p_t in return for the output Y_t from joint production in period t . (Y is mnemonic for "yield.") Then the manufacturer undertakes a noncontractible, productive action $a_{mt} \in A_m(X_t)$ and incurs cost $c_m(a_{mt}, X_t)$, and the buyer chooses noncontractible action $a_{bt} \in A_b(X_t)$ and incurs cost $c_b(a_{bt}, X_t)$.

Figure 1 Sequence of Events in Period t 

Notes. The formal payment p_t may depend on the public history H^t , but is a known constant at the beginning of period t . The discretionary payment d_t may depend on H^t , X_{t+1} , and Y_t . This payment for performance can induce the manufacturer to choose a costly but productive action.

The feasible action set for each firm is a real interval: for every $x \in \mathcal{X}$

$$\begin{aligned} A_m(x) &= [\underline{a}_m(x), \bar{a}_m(x)] \\ A_b(x) &= [\underline{a}_b(x), \bar{a}_b(x)]. \end{aligned} \quad (1)$$

The initial state X_t and actions (a_{mt}, a_{bt}) determine the distribution of the output Y_t and the probability of transition to state $X_{t+1} = z \in \mathcal{X}$ at the end of period t , according to the transition matrix $P(a_m, a_b)$ with elements

$$P_{xz}(a_m, a_b) = \Pr\{X_{t+1} = z \mid X_t = x; a_m, a_b\}.$$

Given $(X_t, X_{t+1}) = (x, z)$, the distribution of Y_t is independent of the firms' actions; let $Y(x, z)$ denote the expected value of Y_t . If either firm refuses to transact in period t , then both buyer and manufacturer incur zero cost, $Y_t = 0$, and the distribution of X_{t+1} is governed by transition matrix $P(0, 0)$.

Each firm observes the state of the system X_t at the beginning of period t , the transition to state X_{t+1} , and the output Y_t , but cannot observe the other firm's action. The action sets, cost functions, transition matrix, and distribution of output are common information. The firms cannot write a formal contract with payment contingent on actions, state transition, or output. Instead, ongoing interaction allows the firms to promise discretionary payments contingent on observed performance in order to create incentives for costly but productive actions. A *relational contract* consists of four parts, for each period $t = 1, 2, \dots$:

i. A formal (court-enforced) contract. If both firms agree to transact in period t , the buyer contracts to pay a constant p_t to the manufacturer in return for the output Y_t . The payment p_t , a known constant at the beginning of period t , may depend on the public history at the beginning of period t .

ii. A discretionary payment. If both firms agree to transact in period t , the buyer promises to pay d_t to the manufacturer. The payment d_t may depend on the public history at the end of period t , which includes the transition to state X_{t+1} and the output Y_t . The payment d_t may be negative, representing a promise for the manufacturer to pay the buyer.

iii. A strategy for the manufacturer that specifies whether or not to transact with the buyer $\tau_{mt} \in \{0, 1\}$ and, in the event that both parties agree to transact in period t , action a_{mt} and whether or not to execute the discretionary transfer payment $e_{mt} \in \{0, 1\}$. τ_{mt} and a_{mt} may depend on the public history at the beginning of period t ; e_{mt} may depend on the public history at the end of period t .

iv. A strategy for the buyer that specifies whether or not to transact with the manufacturer $\tau_{bt} \in \{0, 1\}$ and, in the event that both parties agree to transact in period t , action a_{bt} and whether or not to execute the discretionary transfer payment $e_{bt} \in \{0, 1\}$. τ_{bt} and a_{bt} may depend on the public history at the beginning of period t ; e_{bt} may depend on the public history at the end of period t .

Formally, the public history at the beginning of period t is $H^t = \{X_1, \dots, X_t; Y_1, \dots, Y_{t-1}; \tau_{m1}, \dots, \tau_{m,t-1}; e_{m1}, \dots, e_{m,t-1}; \tau_{b1}, \dots, \tau_{b,t-1}; e_{b1}, \dots, e_{b,t-1}\}$. The sequence of events in each period t is depicted in Figure 1.

Note that if the firms transact in period t , the buyer formally contracts to pay p_t even if the output Y_t turns out to be low or zero. Formal contracts of this nature are common in the semiconductor industry, where the buyer purchases "wafer starts," but her yield on these wafers is stochastic. Similarly, biopharmaceutical contract manufacturers sell "batch fermentation starts" rather than actual output. An example of a discretionary payment in the biopharmaceutical industry was described to the authors by managers at a large contract manufacturer: The manufacturer agrees informally that if the yield of a batch is low due to some error in its process control, the buyer will not be required to make the full payment (d_t is negative). However, if the manufacturer attributes the low yield to problems with the raw material provided by the buyer, it will not give the discount. Occasionally, disputes occur over who is responsible for a low yield, and such disputes can lead to a breakdown in cooperation.

The manufacturer's discounted profit starting from the beginning of period T is given by

$$\Pi_{mT} = \sum_{t=T}^{\infty} \delta^{t-T} \tau_{bt} \tau_{mt} [p_t + d_t e_{bt} e_{mt} - c_m(a_{mt}, X_t)]. \quad (2)$$

The buyer's discounted profit starting from period T is given by

$$\Pi_{bT} = \sum_{t=T}^{\infty} \delta^{t-T} \tau_{bt} \tau_{mt} [Y_t - p_t - d_t e_{bt} e_{mt} - c_b(a_{bt}, X_t)]. \quad (3)$$

The objective for each firm is to maximize its discounted expected profit.

We say that a *relational contract is self-enforcing* if, given the prices and discretionary transfer payments in (i) and (ii), the firms' strategies constitute a perfect public equilibrium (PPE) with $e_{bt} = e_{mt} = 1$ for all $t = 1, 2, \dots$. That is, the firms are willing to execute the discretionary transfer payment in every period that they transact. As defined in Fudenberg et al. (1994), a profile of strategies $\{(\tau_{mt}, a_{mt}, e_{mt}), (\tau_{bt}, a_{bt}, e_{bt})\}_{t \in 1, 2, \dots, \infty}$ is public if for every period t ; $\tau_{mt}, a_{mt}, \tau_{bt}, a_{bt}$ depend only on the public history H^t at the beginning of period t , and e_{mt}, e_{bt} depend only on H^t, X_{t+1}, Y_t . A PPE is a profile of public strategies that, at each decision point in time and possible history at that time, constitute a Nash equilibrium from that time onward.

In particular, a self-enforcing relational contract must satisfy, for every period t , public history $H^t \in \mathcal{X}^t \times \mathcal{R}_+^t \times \{0, 1\}^{4(t-1)}$, initial state $X_t \in \mathcal{X}$, transition to state $X_{t+1} \in \mathcal{X}$, and output $Y_t \in \mathcal{R}_+$

$$E[\Pi_{mt} | H^t] \geq 0 \quad (4)$$

$$E[\Pi_{bt} | H^t] \geq 0 \quad (5)$$

$$a_{mt} \in \arg \max_{a \in A_m(X_t)} \left\{ -c_m(a, X_t) + \sum_{z \in \mathcal{X}} P_{X_t z}(a, a_{bt}) \cdot E[d_t + \delta \Pi_{m(t+1)} | H^t, X_{t+1} = z] \right\} \quad (6)$$

$$a_{bt} \in \arg \max_{a \in A_b(X_t)} \left\{ -c_b(a, X_t) + \sum_{z \in \mathcal{X}} P_{X_t z}(a_{mt}, a) \cdot (Y(X_t, z) + E[d_t + \delta \Pi_{b(t+1)} | H^t, X_{t+1} = z]) \right\} \quad (7)$$

$$d_t(H^t, X_{t+1}, Y_t) + \delta E[\Pi_{m(t+1)} | H^t, X_{t+1}, Y_t] \geq 0 \quad (8)$$

$$\delta E[\Pi_{b(t+1)} | H^t, X_{t+1}, Y_t] - d_t(H^t, X_{t+1}, Y_t) \geq 0. \quad (9)$$

Because each firm can refuse to transact in period t , each firm is guaranteed positive discounted expected profit ((4) and (5)). The incentives for action in period t depend on the discretionary transfer payment $d_t(H^t, X_{t+1}, Y_t)$, but not the formal price p_t . Equation (6) specifies that the manufacturer's action

maximizes his infinite-horizon discounted expected profit, assuming that the buyer chooses effort a_{bt} in the current period and that both parties adhere to the relational contract in all subsequent periods. Equation (7) plays the analogous role for the buyer. Equations (8) and (9) ensure that both parties prefer to execute the discretionary transfer payment rather than terminate the relationship. Because termination is the most severe credible punishment that can be imposed on a party that fails to execute the discretionary payment, (8) and (9) are necessary conditions for the relational contract to be self-enforcing. Intuitively, if a relational contract is self-enforcing, then neither firm wishes to deviate unilaterally. As observed by Abreu (1988), conditions (4)–(9) are sufficient for a relational contract with “trigger strategies” to be self-enforcing. A trigger strategy is to adhere to the relational contract in every period until the other firm first refuses to execute the discretionary transfer payment, and then to refuse to transact in subsequent periods. In summary, a relational contract is self-enforcing if and only if it satisfies (4)–(9).

As observed in Theorem 1 of Levin (2003), if there exists any self-enforcing relational contract that generates total expected discounted profit $v \geq 0$, then an initial transfer payment can be used to divide this total expected profit in any way that respects the firms' participation constraints. Specifically, for any $\alpha \in [0, 1]$, the firms can allocate expected discounted profit of αv to the manufacturer and $(1 - \alpha)v$ to the buyer. Therefore, it is natural to focus on maximizing total expected discounted profit. We say that a relational contract is optimal if it is self-enforcing and no other self-enforcing relational contract generates strictly greater total expected discounted profit.

In practice, firms often write a detailed long-term contract, containing terms that cannot be enforced by a court. The reader should imagine the buyer and manufacturer writing down the payment terms (i)–(ii) for $t = 1, 2, \dots, \infty$ for the optimal relational contract, with the understanding that these terms will induce the actions (iii)–(iv) in the optimal relational contract. Although the reader may be skeptical that firms would write down such a potentially complex agreement, in the next section we show that the full scope of complexity is unnecessary: An optimal relational contract has significantly simplifying structural properties.

For ease of exposition, we define the payment terms (i)–(ii) separately from the firms' strategies (iii)–(iv). That is, we deliberately restrict the strategy space so that each firm simply chooses whether or not to execute the payment terms specified in the relational contract. Alternatively, one might extend the strategy space to include alternating offers of payment at the beginning and end of each period, building

on Rubinstein's (1982) alternating-offers model of bargaining over the partition of \$1. However, a PPE in the extended state space cannot achieve greater expected discounted profit than the optimal relational contract with the simple structure (i)–(iv). For purposes of maximizing expected discounted profit, our assumption of a simple strategy space is without loss of generality.

Furthermore, in deriving an optimal relational contract, we will without loss of generality, as noted by Abreu (1988), restrict attention to trigger strategies. In a laboratory experiment using a repeated trust game, Schweitzer et al. (2005) observed that when a subject is deceived by its partner (the partner promises to make a payment in return for cooperative action, and breaks that promise), in subsequent periods the subject tends to distrust his partner and to behave noncooperatively. Even when the deceived subject receives a promise, an apology, and a series of cooperative actions from his partner, noncooperative behavior persists. The behavioral observation that deception causes significant and enduring harm to trust provides support for focusing on trigger strategies (refusal to transact after a partner breaks a promise to pay for performance).

Observe that our model formulation allows for each firm to have a state-dependent outside alternative to joint production in each period. The manufacturer's cost function $c_m(a_m, x)$ represents actual production costs and any forgone profit from working with an alternative partner. (Increasing the value of the manufacturer's outside alternative in state x increases $c_m(a_m, x)$ by a constant for all $a_m \in A_m(x)$.) Similarly, the buyer's cost function $c_b(a_b, x)$ represents actual production costs and any forgone profit from working with an alternative partner. Then, the "profit" functions in (2)–(3) represent discounted profit in excess of the outside alternative. However, to be consistent with our assumption that each firm seeks to maximize this profit, the outside alternative should evolve exogenously rather than be influenced by the firms' actions.

For simplicity, we have focused on the case where the buyer works with a single manufacturer. However, in practice, the buyer might work with a team of manufacturers. For example, in developing a new model, an automobile manufacturer must coordinate and integrate the design efforts of various component suppliers. All of our theoretical results hold in a setting with N manufacturers in which the transition probabilities and output depend on the actions of the manufacturers as well as the buyer (see Plambeck and Taylor 2006).

Finally, for brevity of mathematical expression in §3, we have formulated a rich state space \mathcal{X} in which

the state X_{t+1} at the end of period t includes information about the output Y_t . Conditional on the event that the firms transact in period t and $(X_t, X_{t+1}) = (x, z)$, the output Y_t is a continuous random variable with distribution function F_{xz} and is independent of the firms' actions in period t and all previous states, output, and actions. The distribution function $F_{xz}(y)$ is continuous and has support $[\underline{y}_{xz}, \bar{y}_{xz}]$ where $\underline{y}_{xz} < \bar{y}_{xz}$. Recall that $Y(x, z)$ is the expected value of \tilde{Y}_t conditional on $(X_t, X_{t+1}) = (x, z)$, so

$$Y(x, z) = E[Y_t | (X_t, X_{t+1}) = (x, z)] = \int_{\underline{y}_{xz}}^{\bar{y}_{xz}} y dF_{xz}(y).$$

Although the state space \mathcal{X} is discrete, the output Y_t exhibits continuous random noise. This random noise, commonly observed by both firms, will serve to coordinate the firms' strategies in the optimal relational contract we construct in §3.

3. Derivation of an Optimal Relational Contract

Our main result is that there exists an optimal relational contract with an attractively simple structure, characterized by an unusual sort of dynamic program, in which the cost of action depends on the ongoing value function. Before presenting the main result (Theorem 1), we need to develop some machinery and intuition.

The total expected discounted profit with perfect coordination is given by the dynamic programming recursion

$$\bar{V}(x) = \max \left[\delta \sum_{z \in \mathcal{X}} P_{xz}(0, 0) \bar{V}(z); \max_{a_m \in A_m(x), a_b \in A_b(x)} \left\{ -c_m(a_m, x) - c_b(a_b, x) + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a_b) [Y(x, z) + \delta \bar{V}(z)] \right\} \right]. \quad (10)$$

Let $\bar{\mathcal{X}} \subset \mathcal{X}$ denote the states in which it is optimal to transact, and $\bar{a}_m(x)$, $\bar{a}_b(x)$ denote the optimal actions in state $x \in \bar{\mathcal{X}}$, obtained by solving (10). We will subsequently call these the first-best transaction states and actions. Clearly, the resulting first-best expected discounted profit \bar{V} is an upper bound on the total expected discounted profit that the firms can achieve under any relational contract. If actions were contractible, the buyer and manufacturer could achieve \bar{V} . However, the ability to write a formal long-term contract with state- and output-contingent payments $p_t(Y_1, \dots, Y_t, X_1, \dots, X_t, X_{t+1})$ would *not* necessarily enable the buyer and manufacturer to achieve \bar{V} . A free-rider problem arises when the incentive payments to the buyer and manufacturer add up to zero

in every period (Holmstrom 1982). To create second-best incentives for action in the current period, an undesirable outcome must be followed by punishment through inefficient actions that destroy profit in subsequent periods. We will derive an optimal relational contract in which punishment occurs simply through termination of the relationship, i.e., a refusal by both firms to transact in all subsequent periods.

Now let us develop a dynamic programming recursion for the total expected discounted profit under an optimal relational contract. For each $x \in \mathcal{X}$ and $v: \mathcal{X} \rightarrow R^+$, define the operator

$$T(v)(x) = \max \left[\delta \sum_{z \in \mathcal{X}} P_{xz}(0, 0)v(z); \right. \\ \left. \max_{a_m \in A_m(x), a_b \in A_b(x)} \left\{ -C(a_m, a_b, v, x) \right. \right. \\ \left. \left. + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a_b)[Y(x, z) + \delta v(z)] \right\} \right], \quad (11)$$

where the cost function is given by

$$C(a_m, a_b, v, x) = c_m(a_m, x) + c_b(a_b, x) \\ + \min_{V_m, V_b} \sum_{z \in \mathcal{X}} P_{xz}(a_m, a_b)Q(x, z)\delta v(z) \quad (12)$$

subject to:

$$V_m(x, z) \geq 0, \quad V_b(x, z) \geq 0, \\ V_m(x, z) + V_b(x, z) \leq v(z) \quad \text{for } z \in \mathcal{X} \\ a_m \in \arg \max_{a \in A_m(x)} \left\{ -c_m(a, x) + \sum_{z \in \mathcal{X}} P_{xz}(a, a_b)\delta V_m(x, z) \right\} \\ a_b \in \arg \max_{a \in A_b(x)} \left\{ -c_b(a, x) \right. \\ \left. + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a)[Y(x, z) + \delta V_b(x, z)] \right\} \\ Q(x, z) = [v(z) - V_m(x, z) - V_b(x, z)]/v(z).$$

The operator Tv gives the maximum total discounted expected profit under a self-enforcing relational contract, assuming that if the firms do not terminate in Period 1, the total discounted expected profit at the beginning of Period 2 is given by v . The cost function $C(a_m, a_b, v, x)$ has two components: the direct cost of action (c_m and c_b) and the expected cost associated with possible termination. In the minimization embedded in (12), $V_m(x, z)$ is the portion of the total discounted expected profit from Period 2 allocated to the manufacturer, conditional on $(X_1, X_2) = (x, z)$; $V_b(x, z)$ is the analogous quantity for the retailer. These quantities reflect both the discretionary payment in Period 1 and the ongoing value of the relationship from Period 2. The constraints $V_m(x, z) \geq 0$

and $V_b(x, z) \geq 0$ ensure that it is in the interest of each firm to execute the discretionary payment. The termination function $Q(x, z)$ specifies the probability with which the firms terminate their relationship following state transition $(X_1, X_2) = (x, z)$ (Q is mnemonic for “Quit”). Allowing for termination with positive probability weakly decreases the total cost of any given action (a_m, a_b) by allowing the inequality $V_m(x, z) + V_b(x, z) \leq v(z)$ to hold strictly. Thus, deliberately destroying profit (by termination) following some state transitions may increase total expected discounted profit.

The operator Tv is distinctive in that the cost of an action depends upon the ongoing value function v as well as the state x . The set of actions (a_m, a_b) that are feasible depends on the ongoing value function v . For example, if $v = 0$, then the minimization embedded in (12) must have $V_m = V_b = 0$; consequently, the only feasible action pairs for the manufacturer and buyer in state x are

$$\left\{ (a_m, a_b): a_m \in \arg \max_{a \in A_m(x)} \{-c_m(a, x)\}, \right. \\ \left. a_b \in \arg \max_{a \in A_b(x)} \left\{ -c_b(a, x) + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a)Y(x, z) \right\} \right\}. \quad (13)$$

These are Nash equilibria of the single-period game in which the firms transact without the value of an ongoing relationship (the potential for repeat business) to induce cooperative behavior. For any actions (a_m, a_b) not satisfying (13), $C(a_m, a_b, 0, x) = \infty$. As v increases, the set of feasible action pairs expands because the constraints in the minimization problem in (12) are relaxed.

The operator T has useful structural properties. Because the objective $-C(a_m, a_b, v, x) + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a_b) \cdot [Y(x, z) + \delta v(z)]$ increases with v , the operator Tv is isotone. Then existence of a largest fixed point follows from Tarski’s fixed-point theorem (Tarski 1955).

PROPOSITION 1. *The operator T has a largest fixed point V^* . That is,*

$$V^* = TV^*$$

and for any other fixed point $V = TV$, $V^(x) \geq V(x)$ for all $x \in \mathcal{X}$. Furthermore, $V^*(x) \in [0, \bar{V}(x)]$ for all $x \in \mathcal{X}$.*

All proofs, with the exception of that of Theorem 1, are in the online supplement to this paper on the *Management Science* website at <http://mansci.pubs.informs.org/eocompanion.html>.

To describe the simple optimal relational contract, we now introduce some notation. For each $x \in \mathcal{X}$, $(a_m^*(x), a_b^*(x))$ denotes the actions obtained by solving

(11) with $v = V^*$, $(V_m^*(x, z); V_b^*(x, z))$ denotes the minimizers of $C(a_m^*(x), a_b^*(x), x, V^*)$; and

$$\tau_m^*(x) = \tau_b^*(x) = \begin{cases} 1 & \text{if } V^*(x) > \delta \sum_{z \in \mathcal{X}} P_{xz}(0, 0) V^*(z) \\ 0 & \text{if } V^*(x) = \delta \sum_{z \in \mathcal{X}} P_{xz}(0, 0) V^*(z) \end{cases}$$

indicates whether or not transacting in state x is profitable. Finally, we define the termination probability

$$Q^*(x, z) = [V^*(z) - V_m^*(x, z) - V_b^*(x, z)] / V^*(z)$$

and termination period

$$\mathcal{T} = \inf\{t: Y_t < F_{X_t, X_{t+1}}^{-1}(Q^*(X_t, X_{t+1}))\}.$$

THEOREM 1. *The total expected discounted profit under an optimal relational contract is $V^*(X_1)$, and a simple optimal relational contract is characterized as follows. The firms' strategies for whether or not to transact are*

$$\tau_{mt} = \begin{cases} \tau_m^*(X_t) & \text{if } t \leq \mathcal{T} \text{ and } e_{bs} = e_{ms} = 1 \text{ for all } s < t \\ 0 & \text{if } t > \mathcal{T} \text{ or } e_{ms} e_{bs} = 0 \text{ for some } s < t \end{cases}$$

$$\tau_{bt} = \begin{cases} \tau_b^*(X_t) & \text{if } t \leq \mathcal{T} \text{ and } e_{bs} = e_{ms} = 1 \text{ for } s < t \\ 0 & \text{if } t > \mathcal{T} \text{ or } e_{ms} e_{bs} = 0 \text{ for some } s < t. \end{cases}$$

That is, the firms terminate the relationship at the end of period \mathcal{T} . In each period that the firms transact, the formal price depends only on the current state, and the discretionary transfer payment depends only upon the observed transition and output

$$p_t = \alpha V^*(X_t) + c_m(a_m^*(X_t), X_t) - \sum_{z \in \mathcal{X}} P_{X_t z}(a_m^*(X_t), a_b^*(X_t)) \delta V_m^*(z)$$

$$d_t = \begin{cases} [1 - Q^*(X_t, X_{t+1})]^{-1} \delta V_m^*(X_t, X_{t+1}) - \alpha \delta V^*(X_{t+1}) & \text{if } Y_t \geq F_{X_t, X_{t+1}}^{-1}(Q^*(X_t, X_{t+1})) \\ 0 & \text{if } Y_t < F_{X_t, X_{t+1}}^{-1}(Q^*(X_t, X_{t+1})), \end{cases}$$

where $\alpha \in [0, 1]$ is the fraction of expected total discounted profit allocated to the manufacturer; the action strategies depend only on the current state

$$a_{mt} = a_m^*(X_t), \quad a_{bt} = a_b^*(X_t) \quad \text{for } t = 1, 2, \dots$$

and each firm is willing to execute the discretionary transfer payment

$$e_{bt} = e_{mt} = 1 \quad \text{for } t = 1, 2, \dots$$

The proof of Theorem 1 is in the appendix.

Having derived an optimal relational contract, we now provide a road map for the remainder of the section. First, we will describe the optimal relational

contract and discuss its implementation. Second, we will discuss how the inability to break off cooperation changes the formulation and results. Third, we will discuss how the optimal relational contract can be computed.

The optimal relational contract in Theorem 1 involves (probabilistic) termination following periods with undesirable performance as reflected in an undesirable state transition and low output. (This is formalized in the next section's Proposition 3.) This termination could be interpreted as resulting from a dispute over who is responsible for poor performance. However, termination occurs despite the fact that in every period in which trade occurs the buyer and supplier take the agreed-upon action. Thus, the firms are not penalizing one another for presumed shirking. Rather, the purpose of termination is to provide stronger incentives for action by jointly punishing the firms for unfavorable stochastic outcomes. These stronger incentives lead to greater expected profit in the periods in which the firms transact; however, profit is, of course, reduced in periods following termination. The optimal relational contract balances the near-term gain from stronger incentives for action against the eventual loss resulting from termination.

That output Y_t is a *continuous* random variable eliminates the need for complex history dependence in the optimal relational contract. In the special case that the termination probability $Q^*(x, z) \in \{0, 1\}$ for all $(x, z) \in \mathcal{X}^2$, the observed output Y_t is irrelevant. However, for $Q^*(x, z) \in (0, 1)$, the observed output Y_t serves to coordinate the firms' strategies for transaction, to destroy exactly the right amount of value $Q^*(x, z) \delta V^*(z)$. Given that the strategy of one firm is to refuse to transact in all subsequent periods, an optimal strategy for the other firm is to refuse to transact in all subsequent periods. If the support of the output Y_t were discrete, the firms might need to transact for several more periods with suboptimal, history-dependent actions, in order to destroy exactly the right amount of value $Q^*(x, z) \delta V^*(z)$. For a class of systems (including the symmetric version of the joint production system in §5), the continuous output Y_t is not needed: If for every (x, z) such that $Q^*(x, z) \in (0, 1)$ there exists a solution to (12) with either $V_m^*(x, z) = 0$ or $V_b^*(x, z) = 0$, then following such a state transition, one firm is indifferent between executing the discretionary payment and terminating the relationship. When this indifferent firm plays a mixed strategy in which with probability $Q^*(x, z)$ it refuses to execute the discretionary payment and thus terminates the relationship, the firms do not need to coordinate their termination decision. Finally, in §4, we assume that the first order conditions for the optimal actions are sufficient, and prove in Proposition 3 that $Q^*(x, z) \in \{0, 1\}$ for all $(x, z) \in \mathcal{X}^2$, except for (at most)

one unique $z \in \mathcal{Z}$. We conclude that probabilistic termination ($Q^*(x, z) \in (0, 1)$) coordinated by the continuous random variable Y_t plays an important role in the proof of Theorem 1, but in many cases the firms can do very well without this.

In practice, one might expect that following poor outcomes, firms would break off cooperation for a limited period of time, rather than forever. One might argue that the firms cannot credibly refuse to transact; in the event of termination, they would renegotiate the relational contract to generate some ongoing profit. The economics literature on repeated games with imperfect monitoring is subject to the same criticism that in a punishment phase, the players have an incentive to renegotiate to a more favorable continuation equilibrium; see, for example, Abreu et al. (1986, 1991). Laboratory experiments in a repeated trust game conducted by Schweitzer et al. (2005) suggest that the degree to which cooperation can be restored following a breakdown in the relationship depends on whether the breakdown is accompanied by deliberate, visible deception. Their observations suggest the firms would break off cooperation following the failure to make a discretionary payment because this involves deliberate, visible deception. However, when a firm observes a low yield, there could be ambiguity as to whether he was deceived by a shirking partner or whether the outcome was due to chance, so it is less clear that the firms would break off cooperation. To the extent that the threat of terminating the relationship because of a low yield is not credible, it is desirable to identify conditions under which the optimal relational contract does not require such termination; §4 explores this issue and shows that process visibility can eliminate the need for termination.

If the optimal relational contract characterized by Theorem 1 requires termination with positive probability, the firms should anticipate the potential for renegotiation of the relational contract. Several papers explore the renegotiation of formal contracts in dynamic games; see Laffont and Tirole (1990), Rey and Salanie (1996), and references therein. They observe that allowing renegotiation is equivalent to restricting attention to long-term contracts that are *immune to renegotiation* (i.e., in every period, the players cannot achieve greater profit by substituting an alternative continuation contract). Suppose that we impose the analogous constraint that the relational contract be immune to renegotiation on the equilibrium path. This simply involves modifying the operator T so that

$$V_m(x, z) + V_b(x, z) \leq v(z) \quad \text{for } z \in \mathcal{Z}$$

in (12) is replaced by

$$V_m(x, z) + V_b(x, z) = v(z) \quad \text{for } z \in \mathcal{Z}.$$

Allowing renegotiation means that $Q^* = 0$. By extension of the proof of Theorem 2 in Levin (2003), if an optimal relational contract exists, then expected discounted profit under this optimal relational contract is the largest fixed point of the modified operator T in $[0, \bar{V}]$; Theorem 1 holds with the modified operator T and $Q^* = 0$. Allowing renegotiation reduces expected discounted profit at time zero (strictly reduces expected discounted profit if and only if the optimal relational contract of Theorem 1 requires termination with strictly positive probability). Therefore, insofar as process visibility eliminates the need for termination in the optimal relational contract of Theorem 1, the potential for renegotiation increases the value of process visibility. In the applications in §5 and the online supplement, for a wide range of parameters, the optimal contract has $Q^* = 0$, i.e., allowing renegotiation does not reduce expected discounted profit.

Other researchers have considered punishment schemes where the magnitude of the punishment is limited, for example, by a finite limit on the duration of punishment (e.g., Atkins et al. 2005). The case in which the firms are willing to destroy at most a fraction $\theta(z)$ of the value that can be created in state z is captured in our framework by adding the constraint

$$V_m(x, z) + V_b(x, z) \geq [1 - \theta(z)]v(z) \quad \text{for } z \in \mathcal{Z}$$

to the minimization problem embedded in (12). If the firms can refuse to transact for at most l periods and then must resume cooperation, we instead have the constraint

$$\begin{aligned} &V_m(x, z) + V_b(x, z) \\ &\geq \delta^l \sum_{x_1 \in \mathcal{X}} P_{zx_1}(0, 0) \sum_{x_2 \in \mathcal{X}} P_{x_1x_2}(0, 0) \cdots \sum_{x_l \in \mathcal{X}} P_{x_{l-1}x_l}(0, 0)v(x_l) \end{aligned}$$

for $z \in \mathcal{Z}$,

which significantly complicates the recursion.

As a practical matter, for any given problem parameters, specifying the optimal relational contract requires calculating the largest fixed point of T , V^* , which from Theorem 1 is also the optimal value function. Proposition 2 provides a theoretical basis for and guidance as to how to employ value iteration to compute the optimal value function. Define $T^0V \equiv V$ and for $n \geq 1$, $T^nV \equiv T(T^{n-1}V)$. Value iteration involves computing T^nV for successively larger values of n , starting with a given value function V .

PROPOSITION 2. *Value iteration converges to the optimal value function V^* when one begins with the first-best value function \bar{V} :*

$$V^* = \lim_{n \rightarrow \infty} T^n \bar{V}.$$

Furthermore, the value function after a finite number of iterations provides an upper bound: $T^n \bar{V} \geq V^*$. Value iteration converges geometrically to the optimal value function at the rate of the discount factor:

$$\sup_{x \in \mathcal{X}} \{T^n \bar{V}(x) - V^*(x)\} \leq \delta^n \sup_{x \in \mathcal{X}} \{\bar{V}(x) - V^*(x)\}. \quad (14)$$

Observe that the convergence result is dependent on the initial value function. In standard dynamic programming analyses, where the cost function does not depend on the value function, convergence is often obtained regardless of the initial value function. The usual approach is to show that the optimal value operator is a contraction and then to appeal to the Banach Fixed-Point Theorem to establish convergence. In our case, because the cost function $C(a_b, a_m, v, x)$ depends on the value function, the optimal value operator T need not be a contraction. However, beginning value iteration with the first-best value function \bar{V} ensures that the resulting value function in each iteration is decreasing. Using this property in conjunction with the definition of T establishes the convergence result.

An immediate implication of (14) is that value iteration yields a value function that is within ε of optimal (i.e., $\sup_{x \in \mathcal{X}} \{T^n \bar{V}(x) - V^*(x)\} \leq \varepsilon$) after $n = \lceil \log(\varepsilon / \sup_{x \in \mathcal{X}} \{\bar{V}(x) - V^*(x)\}) / \log(\delta) \rceil$ iterations. Thus, if the discount factor is small, with relatively few iterations, value iteration will yield a value function that is near optimal. This bound on the number of iterations is large when δ is close to unity. However, when δ is large, the second term on the right-hand side of (14) will tend to be small (because V^* is increasing in δ), which favors more rapid convergence.

4. Structural Properties of an Optimal Relational Contract

We begin with a brief overview of the results in this section. First, the optimal relational contract has a simple threshold rule for termination. That is, with the state space ordered according to expected discounted profit under the optimal relational contract, termination occurs upon transition to a low state, below a specified threshold. Termination is needed because of the free-rider problem that occurs in team production; with a single agent taking action, an optimal relational contract never requires termination on the equilibrium path (Levin 2003).

Second, we give the firms information technology that provides a signal for the manufacturer's action that is independent of the buyer's action. We identify conditions under which this signal enables the firms to avoid termination and achieve the first-best total expected discounted profit with a self-enforcing relational contract. If the signal is contractible, then

the firms achieve the first best even in a single-period interaction. We focus on the case that the signal is not contractible. If the ongoing relationship is very valuable, the buyer can credibly commit to pay the manufacturer for a high signal, and thus the first best is achieved as though the signal were contractible; termination is not required. If the buyer's action in the current period influences the yield in the current period but not the ongoing value of the relationship, then termination is never required under an optimal relational contract. This surprising result relies on the buyer capturing the output in the current period. Doornik (2006) presents a simple repeated game with an independent signal for each firm's action; the firms share the output so the optimal relational contract requires termination. We give examples in which the buyer's current action influences the future value of the relationship, and the optimal relational contract requires termination.

Third, we show how the Markovian dynamics influence the structure of the payment terms in an optimal relational contract. A prominent conclusion in the relational contracts literature is that the optimal discretionary payment has a simple "one-step" structure, i.e., the buyer pays a fixed bonus if the signal is above a specified threshold. This result is proven for i.i.d. repeated games with one-sided moral hazard (Theorem 6 by Levin 2003) and proven for i.i.d. repeated games with double moral hazard and independent signals (Propositions 3 and 4 by Doornik 2006). For dynamic games with double moral hazard, we identify conditions under which a one-step discretionary payment remains optimal and conditions under which a richer payment structure is required.

To state and prove these results, we impose additional assumptions about the action sets and how actions influence the transition probabilities. Specifically, Rogerson (1985) proposed sufficient conditions to justify the "first-order approach" (relaxing the constraint that the agent chooses an action that maximizes his utility to a first order necessary condition) in a static principal-agent problem. We extend these conditions to a system with Markovian dynamics. First, for every $x \in \mathcal{X}$ the cost functions $c_m(a_m, x)$ and $c_b(a_b, x)$ are increasing and continuously differentiable in the actions a_m and a_b , respectively. For fixed state $x \in \mathcal{X}$, we can order the states $\mathcal{X} = \{z_1, z_2, \dots, z_N\}$ such that

$$\begin{aligned} Y(x, z_1) + \delta V^*(z_1) &\leq Y(x, z_2) + \delta V^*(z_2) \\ &\leq \dots \leq Y(x, z_N) + \delta V^*(z_N). \end{aligned}$$

Under the optimal relational contract, starting from state x , a transition to state z_{i+1} yields greater expected total discounted profit than a transition to state z_i , for each $i = 1, 2, \dots, N - 1$. The second

assumption is that for any $a_b \in A_b(x)$, $\{a_m, a_m^1, a_m^2\} \in A_m(x)$ and $\beta \in [0, 1]$ such that $c_m(a_m, x) = \beta c_m(a_m^1, x) + (1 - \beta)c_m(a_m^2, x)$, and for each $n \in \{1, 2, \dots, N\}$,

$$\begin{aligned} & \sum_{i=n}^N P_{xz_i}(a_m, a_b) \\ & \geq \beta \sum_{i=n}^N P_{xz_i}(a_m^1, a_b) + (1 - \beta) \sum_{i=n}^N P_{xz_i}(a_m^2, a_b); \end{aligned} \quad (15)$$

for any $a_m \in A_m(x)$, $\{a_b, a_b^1, a_b^2\} \in A_b(x)$ and $\beta \in [0, 1]$ such that $c_b(a_b, x) = \beta c_b(a_b^1, x) + (1 - \beta)c_b(a_b^2, x)$, and for each $n \in \{1, 2, \dots, N\}$,

$$\begin{aligned} & \sum_{i=n}^N P_{xz_i}(a_m, a_b) \\ & \geq \beta \sum_{i=n}^N P_{xz_i}(a_m, a_b^1) + (1 - \beta) \sum_{i=n}^N P_{xz_i}(a_m, a_b^2). \end{aligned} \quad (16)$$

Intuitively, this second assumption implies a decreasing marginal expected discounted profit for each additional dollar's worth of action. The third assumption is that $P_{xz}(a_m, a_b)$ is strictly positive and continuously differentiable in (a_m, a_b) , and for any $a_m \in A_m(x)$, $a_b \in A_b(x)$

$$\frac{(\partial/\partial a_m)P_{xz_i}(a_m, a_b)}{P_{xz_i}(a_m, a_b)} \quad \text{and} \quad \frac{(\partial/\partial a_b)P_{xz_i}(a_m, a_b)}{P_{xz_i}(a_m, a_b)}$$

increase with i . (17)

Rogerson (1985) points out that this assumption is equivalent to the following statistical property. If one is given a prior over a firm's action choice, observes the transition (x, z) , and then calculates a posterior cumulative distribution $G(a | (x, z))$ for the action choice, then for every a and $i = 1, 2, \dots, N - 1$,

$$G(a | (x, z_{i+1})) \leq G(a | (x, z_i)). \quad (18)$$

That is, observing a more desirable transition allows one to infer that the firm took greater action, in the sense of stochastic dominance. Together, these three assumptions guarantee that if a firm's ongoing expected discounted profit (including current-period output for the buyer) contingent on the transition (x, z_i) increases with i , then that firm's objective is a concave function of its action. This allows us to substitute the first order condition for each firm's incentive compatibility constraint in (6)–(7) and (12).

Proposition 3 establishes that there exists an optimal relational contract that requires probabilistic termination (coordinated by the observed yield Y_t) in at most one threshold state z_n . The firms continue to cooperate if $X_{t+1} > z_n$, termination occurs with probability 1 if $X_{t+1} < z_n$, and termination occurs with probability $Q^* \in [0, 1]$ in the threshold state $X_{t+1} = z_n$.

PROPOSITION 3. *There exists an optimal relational contract with the following termination threshold property. For each $x \in \mathcal{X}$ with $V^*(x) > \delta \sum_{z \in \mathcal{X}} P_{xz}(0, 0)V^*(z)$, i.e., for each state x in which the firms transact, there exists a threshold state z_n such that*

$$Q^*(x, z_i) = \begin{cases} 1 & \text{for } z_i < z_n \\ 0 & \text{for } z_i > z_n. \end{cases} \quad (19)$$

In the optimal relational contract in Proposition 3, if termination occurs at all, it occurs in the event of a transition to an undesirable state that, intuitively, allows Bayesian inference that the firms took little action (in the sense of stochastic dominance in (18)). However, although the behavior is consistent with the idea that information is being extracted from the observed state, the firms are not, in fact, making statistical inferences. In each period, both firms take the actions specified in the relational contract. The purpose of this form of termination function is to provide incentives for those actions. Finally, it is straightforward to extend the proof of Proposition 3 to show that *any* optimal relational contract must have the threshold property (19) if the following two conditions are satisfied: (17) holds in the strict sense, and the optimal actions $\{a_m^*(x), a_b^*(x)\}_{x \in \mathcal{X}}$ are unique.

Independent Signals: The Value of Assigning Blame

Now we assume that in each period t that the firms transact, they observe a signal s_{mt} that conveys information about the manufacturer's action a_{mt} and is invariant with respect to the buyer's action. The signal takes values in an ordered set $s_{mt} \in \{s^1, s^2, \dots, s^N\}$, where $s^i \geq s^{i-1}$ for $i = 1, 2, \dots, N$. With a slight adaptation, let $P_{x(z, s^i)}(a_m, a_b)$ denote the probability of observing signal s^i and a transition to state z , given actions (a_m, a_b) and initial state x . Assume that (15)–(17) continue to hold with the substitution of $P_{x(z, s^i)}(a_m, a_b)$ for $P_{xz}(a_m, a_b)$. This generalization of (17) implies existence of a signal threshold $\hat{s}_m(a_m, x)$ such that

$$\frac{\partial}{\partial a_m} P_{x(z, s^i)}(a_m, a_b) \geq 0 \quad \text{if and only if } s^i \geq \hat{s}_m(a_m, x).$$

Firms engaged in team production cannot, in general, use a court-enforced contract that divides the output to create incentives for the first-best actions (Holmstrom 1982). A free-rider problem arises because incentive payments must sum to zero; some form of punishment is needed to break this "budget balance constraint" (take profit from both firms following a negative outcome) in order to induce the first-best actions. However, for our setting, Proposition 4A establishes that if the firms can write a court-enforced contract contingent on the signal of

the manufacturer’s action, then a simple one-step contract achieves the first best. The dynamics of the system become irrelevant. The buyer has optimal incentives because she receives the output; and the payment to the supplier, contingent on the signal, is not influenced by her actions. (In general, with multiple risk-neutral firms and shared output, a contractible, independent signal for every firm’s action would be necessary and sufficient to achieve the first best.)

Unfortunately, in many cases, signals are not contractible. In our biopharmaceutical example, detailed process control data observed by the firms is both complex and subject to manipulation, making contracting on this information difficult. Therefore, Proposition 4B and the remainder of this section turn to the case where the signal is not contractible. When the signal is not contractible, the dynamic structure of the game drives the results. Specifically, Proposition 4B establishes that if the ongoing value of the relationship is sufficiently high, then the noncontractible signal for the manufacturer’s action enables the firms to achieve the first best. Section 3 established that without the signal, termination plays the role of breaking the budget balance constraint. Qualitatively, observing an independent signal of the manufacturer’s action allows the firms to assign blame for a negative outcome in an unbiased manner. Therefore, it seems natural that the independent signal for the manufacturer’s action, although noncontractible, will allow the firms to avoid terminating the relationship. Proposition 4B dashes that hopeful conjecture. If the buyer’s action in the current period influences the ongoing value of the relationship in subsequent periods, then optimal relational contracts continue to require termination. In fact, termination can occur in the optimal relational contract even with independent signals for both the buyer and the manufacturer; for an example, see the online supplement.

Recall that $\{\bar{a}_m(x), \bar{a}_b(x) : x \in \mathcal{X}\}$ denote the first-best actions, obtained by solving (10).

PROPOSITION 4A. *Suppose that the signal s_{mt} is contractible. The first-best actions are induced by a one-step formal contract: for $X_t = x \in \bar{\mathcal{X}}$*

$$p_t = \begin{cases} p + B & \text{if } s_{mt} \geq \hat{s}_m(\bar{a}_m(x), x) \\ p & \text{if } s_{mt} < \hat{s}_m(\bar{a}_m(x), x), \end{cases} \quad (20)$$

with bonus for a high signal of

$$B = \frac{(\partial/\partial a_m)c_m(\bar{a}_m(x), x)}{\sum_{s \geq \hat{s}_m(\bar{a}_m(x), x)} (\partial/\partial a_m)P_{x(z, s)}(\bar{a}_m(x), \bar{a}_b(x))}$$

and base pay given by

$$p = c_m(\bar{a}_m(x), x) - B \sum_{s \geq \hat{s}_m(\bar{a}_m(x), x)} P_{x(z, s)}(\bar{a}_m(x), \bar{a}_b(x)).$$

PROPOSITION 4B. *Suppose that the signal s_{mt} is not contractible. If*

$$\begin{aligned} & \sum_{s \geq \hat{s}_m(\bar{a}_m(x), x)} \sum_{z \in \mathcal{Z}} \frac{\partial}{\partial a_m} P_{x(z, s)}(\bar{a}_m(x), \bar{a}_b(x)) \min_{z \in \bar{\mathcal{Z}}} \bar{V}(z) \\ & \geq \frac{\partial}{\partial a_m} c_m(\bar{a}_m(x), x) \quad \text{for all } x \in \bar{\mathcal{X}}, \end{aligned} \quad (21)$$

then the one-step contract (20), where the bonus is discretionary ($d_t = B$ if $s_{mt} \geq \hat{s}_m(\bar{a}_m(x), x)$ and $d_t = 0$ otherwise) and first best actions are self-enforcing. If (21) is not satisfied and the buyer’s action a_{bt} influences the action set, cost structure, or transition probability matrix for period $t + 1$, then the optimal relational contract may require termination.

Proposition 4B establishes the sufficient condition (21) for an optimal relational contract to have a one-step payment structure when the signal is not contractible. To understand why (21) is sufficient, observe that with the contract (20), the buyer captures the first-best expected discounted profit \bar{V} . Therefore, (21) guarantees that the buyer prefers to pay the discretionary bonus $d_t = B$ when $s_{mt} \geq \hat{s}_m$ rather than terminate the relationship and lose \bar{V} . The relational contract with the first-best actions and one-step payment of type (20), where the bonus B is discretionary, is self-enforcing, and therefore optimal.

Proposition 5 establishes the converse to the latter part of Proposition 4B: If the buyer’s action in period t does not influence the effective state of the system in period $t + 1$, then there exists a nonterminating optimal contract ($Q^* = 0$). Proposition 5 also identifies the conditions under which the first best is achieved with the optimal relational contract. For ease in presenting previous results, we adopted a state space formulation in which the distribution of output Y_t is determined by the transition (X_t, X_{t+1}) . Hence, the state X_{t+1} contains information about the output in the previous period Y_t . To state Proposition 5, we must associate states that are equivalent except for information about the output in the previous period. For each $x \in \mathcal{X}$, define

$$\begin{aligned} D(x) = \{x' \in \mathcal{X} : & A_m(x) = A_m(x'), A_b(x) = A_b(x'), \\ & c_m(\cdot; x) = c_m(\cdot; x'), c_b(\cdot; x) = c_b(\cdot; x'), \\ & P_x(\cdot, \cdot) = P_{x'}(\cdot, \cdot)\}. \end{aligned}$$

Intuitively, $D(x)$ is the set of all states $x' \in \mathcal{X}$ with the same action set, cost structure, and transition probability matrix as state x . Proposition 5 applies to the biopharmaceutical production example given in the introduction because the quality of raw materials input by the buyer to the current batch influences the yield for the current batch but does not influence the yield of subsequent batches.

PROPOSITION 5. *Suppose that the buyer's action a_{bt} influences only the output Y_t ; that is, $D(X_{t+1})$ is invariant with respect to the buyer's action a_{bt} . Then there exists a nonterminating optimal relational contract ($Q^* = 0$) with discretionary transfer payment*

$$d_t = \begin{cases} \beta(X_t)V^*(X_{t+1}) & \text{if } s_{mt} \geq \hat{s}_m(a_m^*(X_t), X_t) \\ 0 & \text{if } s_{mt} < \hat{s}_m(a_m^*(X_t), X_t), \end{cases} \quad (22)$$

where $\beta(x) \in [0, 1]$, and formal payment

$$p_t = c_m(a_m^*(X_t), X_t) - \sum_{s \geq \hat{s}_m(a_m^*(X_t), X_t)} \sum_{z \in \mathcal{Z}} P_{X_t(z, s)}(a_m^*(X_t), a_b^*(X_t))d(X_t, z).$$

The first best is achieved if and only if, for every state $x \in \mathcal{X}$,

$$\begin{aligned} & \sum_{s \geq \hat{s}_m(\bar{a}_m(x), x)} \sum_{z \in \mathcal{Z}} \frac{\partial}{\partial a_m} P_{x(z, s)}(\bar{a}_m(x), \bar{a}_b(x)) \delta \bar{V}(z) \\ & \geq \frac{\partial}{\partial a_m} c_m(\bar{a}_m(x), x). \end{aligned} \quad (23)$$

Proposition 5 establishes that the optimal relational contract has a simple form. In the special case of a stationary repeated game, $V^*(X_{t+1})$ is constant, so the discretionary payment is one step as in Levin's (2003) stationary game with one-sided moral hazard. With dynamics, the size of the bonus depends on the ongoing value of the relationship. In particular, in a state x where the optimal relational contract has strict underinvestment by the manufacturer:

$$\begin{aligned} & \frac{\partial}{\partial a_m} \left[\sum_{z \in \mathcal{Z}} P_{xz}(a_m^*(x), a_b^*(x)) [Y(x, z) + \delta V^*(z)] \right. \\ & \left. - c_m(a_m, x) \right] > 0, \end{aligned}$$

the optimal relational contract has parameter $\beta(x) = 1$, which means that the manufacturer receives the maximum bonus, all of the relational capital $V^*(X_{t+1})$, when his signal exceeds the threshold. Finally, Proposition 5 provides a necessary and sufficient condition for the first best to be achieved. This condition holds when the discounted expected value under perfect coordination, \bar{V} , is sufficiently large.

Together, Propositions 4B and 5 demonstrate one of our main insights: The presence of dynamics and the need for termination are tightly interconnected. If the system does not exhibit dynamics (or more precisely, if the buyer does not influence the dynamics of the system), then an optimal relational contract does not require termination. If dynamics are present and are influenced by the buyer's action, termination may be required.

The intuition behind these diverging results is the following. When the buyer's action a_{bt} does not influence the ongoing expected discounted profit in the next period $V^*(X_{t+1})$ and does not influence the discretionary payment d_t , then the buyer has optimal incentives for action in period t because she captures the output Y_t . The discretionary payment d_t in (22) is constructed so that, indeed, it only depends on the manufacturer's action, as reflected in the signal s_{mt} and the continuation value $V^*(X_{t+1})$, which is not influenced by the buyer's action a_{bt} . Providing incentives for the manufacturer to take the optimal action does not require destroying value, because any value not allocated to the manufacturer can be transferred as a windfall gain to the buyer, without distorting the buyer's incentives. When the buyer's action a_{bt} does influence the ongoing expected discounted profit in the next period $V^*(X_{t+1})$, this logic breaks down. Providing incentives for optimal actions requires that the discretionary transfer payment depends on both firms' actions. This introduces the free-rider problem, which can be addressed by the joint punishment of termination.

Recall that increasing the value of the manufacturer's outside alternative in state z increases $c_m(a_m, z)$ by a constant for all $a_m \in A_m(z)$. Similarly, increasing the value of the buyer's outside alternative in state z increases $c_b(a_b, z)$ by a constant for all $a_b \in A_b(z)$. This reduces the first-best value of the relationship $\bar{V}(z)$ and, by violating (23), may prevent the firms from achieving the first best.

5. Application

To illustrate how relational contracts can be used and to demonstrate the value of process visibility, this section provides numerical results for a stylized model of the biopharmaceutical contract manufacturing introduced in §1. In each period t , the buyer and manufacturer jointly produce a batch of the therapeutic protein. The yield Y_t depends on both the action of the buyer $a_{bt} \in [0, 1]$ to improve the quality of the raw material input and the action of manufacturer $a_{mt} \in [0, 1]$ to control the batch process. The manufacturer is successful in process control $s_{mt} = 1$ with probability a_{mt} , and otherwise is unsuccessful $s_{mt} = 0$. Similarly, the buyer's raw material input is good $s_{bt} = 1$ with probability a_{bt} , and otherwise is bad $s_{bt} = 0$. The batch is successful if and only if both firms are successful: $X_{t+1} = s_{mt}s_{bt} = 1$. A successful production process is associated with a higher expected yield. Specifically, when $X_{t+1} = s_{mt}s_{bt} = 1$, the yield random variable is $Y_t \sim \text{Uniform}(1.0, 1.1)$. When $X_{t+1} = s_{mt}s_{bt} = 0$, the yield random variable is $Y_t \sim \text{Uniform}(0.0, 0.1)$. We assume that the firms have the same cost of action in every period. The manufacturer's cost of action is

$c_m(a_m) = \gamma_M - c_M \log(1 - a_m)$, where γ_M represents the value of the manufacturer’s outside alternative and the second term represents the direct cost of action. Similarly, the buyer’s cost of action is $c_b(a_b) = \gamma_B - c_B \log(1 - a_b)$. We present results for the case that the cost functions are symmetric $\gamma_M = \gamma_B = \gamma$ and $c_M = c_B = c$, and then explain the effects of relaxing this symmetry. The discount factor is $\delta = 0.9$, the value of each firm’s outside alternative is $\gamma = 0.025$, and the contract allocates $\alpha = 0.7$ of the system profit to the manufacturer. The parameters are chosen to reflect that the yield depends substantially on the firms’ actions, that the firms interact reasonably frequently, and that the bargaining positions may be asymmetric.

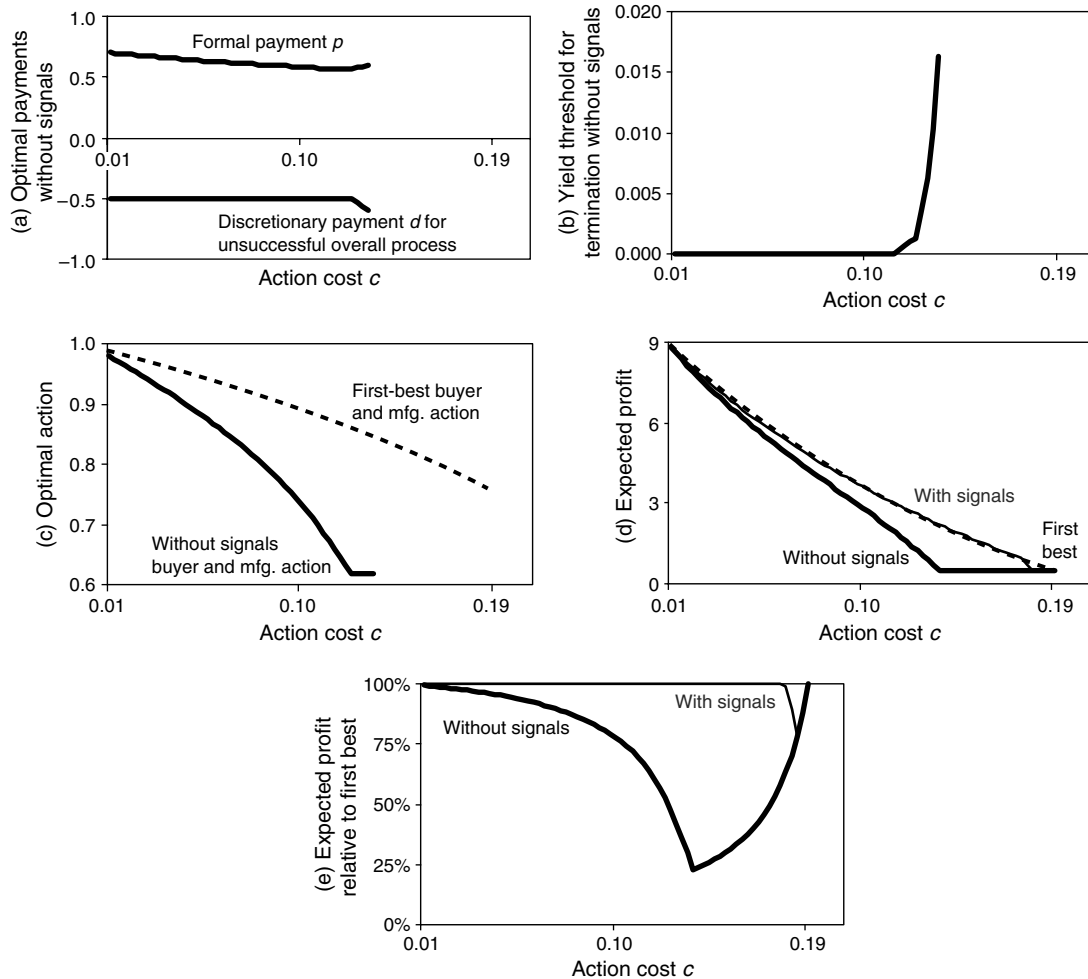
In many cases, the firms may only observe the yield and overall system state, but will not be able to ascertain who is at fault when the overall production process is unsuccessful. However, by investing in information systems, the firms may obtain access to real-time, detailed process control data, which will

help managers to identify problems with process control or faulty materials; in this case, the signals (s_{mt}, s_{bt}) are observable to the firms. Figure 2 depicts the optimal relational contract in the setting without signals and demonstrates the impact of having signal information on system performance. The figure depicts the optimal relational contract and its performance as a function of the action cost parameter c , which reflects the costliness of effort.

To understand Figure 2 it is easiest to begin by focusing on the setting without signals, which corresponds to the thick lines in Figures 2(a)–(e) (subsequently we will describe how observing signals changes the optimal relational contract). Theorem 1 characterizes the optimal relational contract in the setting without signals, and Figures 2(a), (b), and (c) illustrates this result by depicting the corresponding payments, termination threshold, and actions.

Figure 2(a) shows that the optimal contract calls for a nonzero discretionary payment only when the

Figure 2 In (a)–(c), the Optimal Relational Contract Without Signals; in (d)–(e), System Performance with and Without Signals



Notes. Thick lines represent the case without signals, thin lines the case with signals, and dashed lines the first best. In (a), the optimal payments; in (b), the yield threshold for termination; the firms terminate upon observing a yield below this threshold; and in (c), the optimal actions for each firm. In (d), discounted expected system profit; and in (e), discounted expected profit as a percentage of the first-best profit.

overall process is unsuccessful, and in this case, calls for a “negative” payment, i.e., a payment from the manufacturer to the buyer. Thus, the optimal payments have a very natural interpretation: The buyer is obligated to pay the full purchase price, but the manufacturer voluntarily refunds a portion of the purchase price to the buyer when the overall process is unsuccessful.

It is natural that the firms might break off cooperation after observing poor performance. Imposing the joint punishment of termination following a low yield strengthens the incentive for the firms to exert effort in order to reduce the likelihood of a low yield. The optimal relational contract calls for the firms to terminate their relationship if the yield falls below the threshold in Figure 2(b). The figure shows that the optimal relational contract requires termination if and only if c is sufficiently large. When the cost of exerting effort is very high ($c > 0.13$ in Figure 2), cooperation is impossible to sustain: The optimal relational contract calls for the firms to pursue their outside options instead of transacting.

Figure 2(c) shows that the optimal relational contract calls for the firms to exert less effort as the cost of effort increases, which is intuitively appealing. More importantly, the figure demonstrates that the firms underinvest in effort relative to the first best. Even with the optimal relational contract, the firms cannot overcome the free-rider problem: Each firm shirks in exerting effort because it anticipates having to share the benefits of its effort with its partner.

Figures 2(d) and (e) report the discounted expected system profit, where the thick line again corresponds to the setting without signals. Figure 2(d) reports the profit under an optimal relational contract and the first-best profit, and Figure 2(e) reports profit under an optimal relational contract, as a percentage of the first-best profit, V^*/\bar{V} . When the firms only observe the overall system’s success, the loss in system efficiency can be substantial. When the cost of exerting effort is very high ($c > 0.13$ in Figure 2), the optimal relational contract without signals calls for the firms to pursue their outside options, and so the resulting profit is insensitive to c . When the cost of exerting effort is lower, cooperation can be sustained. Here, there is a loss in system efficiency due both to underinvestment in action and to endogenous termination. When c is moderately large, providing incentives for effort requires imposing the joint punishment of termination, and it is here that the loss of system efficiency is largest.

Second, we describe how observing the signals (s_{mt}, s_{bt}) changes the optimal relational contract. Proposition 5 characterizes the optimal relational contract when the firms observe a signal of the manufacturer’s action, and this relational contract

continues to be optimal when the firms also observe a signal of the buyer’s action. As described in Proposition 5, because the buyer’s actions do not influence the dynamics of the system, when the firms observe signals, the need for termination is eliminated. Further, when the firms observe a signal of the manufacturer’s action, the discretionary payment depends solely on that signal. Thus, in the optimal relational contract, the manufacturer compensates the buyer whenever $s_{mt} = 0$. Because the signal provides the means to accurately assign blame, the magnitude of the discretionary payment is larger, and this provides stronger incentives for both firms to exert greater effort. In fact, when the cost of exerting effort is not extremely high ($c < 0.18$ in Figure 2), the optimal relational contract with signals induces the first-best action from the buyer and manufacturer. However, when the cost of exerting effort is extremely high ($c > 0.18$ in Figure 2), the ongoing value of the relationship is low, which constrains the discretionary payment and thus causes underinvestment in effort. This is reflected in the thin line, which corresponds to the case with signals, in Figures 2(d) and (e). In summary, the optimal relational contract with signals achieves the first-best profit except when the cost of exerting effort is extremely high. The intuition is that observing signals allows the firms to mitigate the free-rider problem by assigning blame accurately. If the signal were contractible, observing it would allow the firms to eliminate the free-rider problem entirely (Proposition 4A). However, with noncontractible signals, the magnitude of the discretionary payment that is credible is limited, and consequently so is the strength of incentives that can be provided, which explains why underinvestment persists at extremely high effort costs.

A central purpose of this example is to demonstrate the power of process visibility: Even though the signal information is noncontractible, using it appropriately substantially increases system profit and, for a wide range of parameters, completely eliminates inefficiency. Even when inefficiency persists, its magnitude is relatively small. On both an absolute and a relative basis, system inefficiency and the gain from observing signals is increasing and then decreasing in the cost parameter c . Thus, firms that face moderate costs in creating a successful production process gain the most from obtaining information about the success of each firm’s contribution to the process.

Next, we discuss how the results change when the parameters differ from those in this example. The total system profit under an optimal relational contract depends on the total value of the outside alternative (with profit decreasing in the total value), but not on how this is allocated between the two firms. Thus, Figure 2 represents optimal relational

contracts for any (γ_m, γ_b) satisfying $\gamma_m + \gamma_b = 0.050$. The only impact of asymmetric outside alternatives is to change the minimum amount that each firm can be allocated in an optimal relational contract. In contrast, asymmetry in the direct cost of action portion of the cost function introduces asymmetry in the optimal actions. When the discount factor decreases or the yield distributions become less favorable, the range of production costs over which cooperation can be sustained shrinks, and the degree of inefficiency grows. Changing the split in profits α changes the magnitude of the formal and discretionary payments, but it does not change the optimal actions or system profit.

We conclude this section by describing how the biopharmaceutical production model described above can be adapted to reflect important decisions that impact the dynamics of the system. In biopharmaceutical manufacturing, before commencing full-scale production, the buyer's and manufacturer's scientists and engineers work together to develop a manufacturing protocol. This protocol describes the details of the manufacturing process that the firms are to follow, and once the FDA has approved the protocol, the protocol rarely changes. Thus, the degree of success in the up-front protocol design effort, which depends stochastically on the first-period actions of both the buyer and supplier (a_{b1}, a_{m1}) to develop an efficient manufacturing protocol, determines the production technology for the full-scale batch production that follows in all subsequent periods: the distribution function of the output Y_t conditional on the actions (a_{mt}, a_{bt}) , as well as the cost functions $\{c_m(a_{mt}), c_b(a_{bt})\}$ for $t = 2, 3, \dots$.

The online supplement describes a model of manufacturing protocol design in which the firms observe signals that are sufficient statistics for each firm's actions. The contribution is to demonstrate that the optimal relational contract may call for termination with strictly positive probability, despite the fact that the firms observe informative signals that allow them to assign blame for failure in the protocol design process. This contrasts sharply with the theoretical result illustrated in the joint production example above, that in systems that do not exhibit dynamics, observing signals eliminates the need for termination. Finally, although the FDA's regulatory requirements for pharmaceutical manufacturing lend themselves to a setting where the firms make a one-time, up-front effort to develop a manufacturing process, in many other contexts, the manufacturing process progressively evolves as the manufacturer and buyer work together to improve its efficiency. For example, Dyer and Chu (2000) document the important role that buyers' progressive efforts to improve manufacturing efficiency play in sustaining trust and informal agreements in the auto industry. It is straightforward

to apply our framework to capture such progressive process improvement dynamics.

6. Discussion

Managing a long-term supply chain partnership is inherently a dynamic process. The costliness and effectiveness of the actions that the firms take depends on both their previous actions (e.g., investments or divestments of human and physical assets) and the ever-evolving external business environment. Further, trust, once destroyed, may be difficult or impossible to restore. This paper shows how firms engaged in joint production should structure informal agreements in the face of dynamics and the temptation to free ride. We demonstrate that an optimal relational contract has a simple, memoryless form. To address dynamics, it is sufficient that the agreed-upon actions depend only on the current state, and payments depend only on the observed transition. To discourage free riding, the optimal relational contract may impose termination of the relationship to jointly punish the firms following a transition associated with low effort. We show how process visibility can improve system performance by eliminating the need for termination. A key question is to what extent optimal relational contracts nonetheless require termination. One of our main insights is that dynamics and the need for termination are tightly interlinked. If the buyer does not influence the dynamics, then the optimal relational contract does not require termination. If the buyer does influence the dynamics, the optimal relational contract may require termination.

The framework we provide for addressing dynamic, joint production is quite general. For example, the framework applies to progressive investment by a manufacturer in technology or capacity and progressive investment by a retailer in marketing or branding that increases the value of the manufactured good. We are optimistic that the framework can be applied to a number of specific problem contexts (e.g., in operations or at the operations/marketing interface) to obtain sharper insights about dynamics and relationships, and we hope that future work will follow.

An online supplement to this paper is available on the *Management Science* website (<http://mansci.pubs.informs.org/ecompanion.html>).

Acknowledgments

The authors are grateful to the associate editor, referees, Gerard Cachon, Robert Gibbons, Panos Kouvelis, and Matthew Sobel, as well as seminar participants at Bell Labs, Case Western Reserve University, Duke University, Massachusetts Institute of Technology, Purdue University, University of North Carolina at Chapel Hill, University of Pennsylvania, University of Texas at Dallas, Washington University, and the Northwestern Multiechelon Inventory

Conference, for helpful comments. This research was supported by the National Science Foundation under Grant 0239840. An earlier version of this paper was titled “Partnership in a Dynamic Production System.”

Appendix

PROOF OF THEOREM 1. The proof proceeds in three steps. Any optimal relational contract must have certain basic properties. The first step is to describe these properties. The second step demonstrates that for any relational contract with these basic properties, there exists a simple self-enforcing relational contract that achieves the same expected total discounted profit. In this simple relational contract, the firms’ actions depend only on the current state of the system; the discretionary transfer payment and termination probability depend only on the observed transition. We conclude from the second step that in searching for an optimal relational contract, we can restrict attention to self-enforcing relational contracts with this simple structure. The third step constructs the simple optimal relational contract by solving the dynamic program (11).

Step 1. Properties of any Optimal Relational Contract. Consider a relational contract o with the following terms for the first period: formal payment $p^o(x)$, discretionary transfer payment $d^o(x, z, y)$, strategy for the manufacturer of $\{\tau_m^o(x), a_m^o(x)\}$ and strategy for the buyer of $\{\tau_b^o(x), a_b^o(x)\}$, conditional on $(X_1, X_2, Y_1) = (x, z, y)$. Let $V_1^o(x)$ denote the total expected discounted profit, conditional on $X_1 = x$.

$$\begin{aligned} V_1^o(x) &= E^o[\Pi_{m1}^o + \Pi_{b1}^o \mid X_1 = x] \\ &= E^o\left[\sum_{t=1}^{\infty} \delta^{t-1} \tau_{bt}^o \tau_{mt}^o [Y_t - c_m(a_{mt}^o, X_t) - c_b(a_{bt}^o, X_t)] \mid X_1 = x\right], \end{aligned}$$

where the superscript o indicates that the expectation E^o is taken with respect to the distribution induced by the relational contract. Similarly, let $V_2^o(x, z)$ denote the total expected discounted profit under the optimal relational contract starting from Period 2, conditional on $(X_1, X_2) = (x, z)$.

$$V_2^o(x, z) = E^o\left[\sum_{t=2}^{\infty} \delta^{t-2} \tau_{bt}^o \tau_{mt}^o [Y_t - c_m(a_{mt}^o, X_t) - c_b(a_{bt}^o, X_t)] \mid (X_1, X_2) = (x, z)\right].$$

To be considered as a candidate for optimality, the relational contract o must satisfy

$$V_2^o(x, z) \leq V_1^o(z) \quad \text{for every } x, z \in \mathcal{X}. \quad (24)$$

If $V_2^o(x, z) > V_1^o(z)$ the firms could achieve strictly greater expected total discounted profit by starting with the continuation contract from Period 2, rather than the initial contract for state z . Note that (24) may be a strict inequality in order to create incentives for action in Period 1. However, if $\tau_m^o(x) \cdot \tau_b^o(x) = 0$, then

$$V_2^o(x, z) = V_1^o(z) \quad \text{for every } z \in \mathcal{X}. \quad (25)$$

To be considered as a candidate for optimality, the relational contract o must also be self-enforcing in the first period, which implies that

$$E^o[\Pi_{m1}^o \mid X_1 = x] \geq 0 \quad \text{and} \quad E^o[\Pi_{b1}^o \mid X_1 = x] \geq 0 \quad \text{for every } x \in \mathcal{X}, \quad (26)$$

and for every x such that $\tau_m^o(x) = \tau_b^o(x) = 1$ so that the firms transact in the first period:

$$a_m^o(x) = \arg \max_{a \in A_m(x)} \left\{ -c_m(a, x) + \sum_{z \in \mathcal{Z}} P_{xz}(a, a_b^o) \cdot E^o[d^o + \delta \Pi_{m2}^o \mid X_1 = x, X_2 = z] \right\} \quad (27)$$

$$a_b^o(x) = \arg \max_{a \in A_b(x)} \left\{ -c_b(a, x) + \sum_{z \in \mathcal{Z}} P_{xz}(a_m^o, a) (Y(x, z) + E^o[-d^o + \delta \Pi_{b2}^o \mid X_1 = x, X_2 = z]) \right\} \quad (28)$$

$$\delta E^o[\Pi_{m2}^o \mid (X_1, X_2, Y_1) = (x, z, y)] \geq -d^o(x, z, y) \quad (29)$$

$$\delta E^o[\Pi_{b2}^o \mid (X_1, X_2, Y_1) = (x, z, y)] \geq d^o(x, z, y). \quad (30)$$

Step 2. The Equivalent Simple Relational Contract. Now we will construct a simple, self-enforcing relational contract with the same expected total discounted profit as the relational contract o . Define

$$\begin{aligned} Q(x, z) &= [V_1^o(z) - V_2^o(x, z)] / V_1^o(z) \\ \mathfrak{T} &= \inf\{t: Y_t < F_{X_t, X_{t+1}}^{-1}(Q(X_t, X_{t+1}))\}. \end{aligned}$$

The firms’ strategies for whether or not to transact are, for $t = 1, 2, \dots$,

$$\tau_{mt} = \begin{cases} 1 & \text{if } t \leq \mathfrak{T}, e_{bs} = e_{ms} = 1 \text{ for } s < t \text{ and} \\ & X_t \in \{x: \tau_m^o(x) = \tau_b^o(x) = 1\} \\ 0 & \text{otherwise} \end{cases}$$

$$\tau_{bt} = \begin{cases} 1 & \text{if } t \leq \mathfrak{T}, e_{bs} = e_{ms} = 1 \text{ for } s < t \text{ and} \\ & X_t \in \{x: \tau_m^o(x) = \tau_b^o(x) = 1\} \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, action strategies are, for $t = 1, 2, \dots$,

$$\begin{aligned} a_{mt} &= a_m^o(X_t) \\ a_{bt} &= a_b^o(X_t). \end{aligned}$$

The formal price is $p^o(X_t)$ and the discretionary transfer payment is $d_t = 0$ if $Y_t < F_{X_t, X_{t+1}}^{-1}(Q(X_t, X_{t+1}))$, and otherwise is

$$d_t = d(X_t, X_{t+1}),$$

where, for each $(x, z) \in \mathcal{X} \times \mathcal{X}$,

$$\begin{aligned} d(x, z) &= [1 - Q(x, z)]^{-1} E^o[d^o + \delta \Pi_{m2}^o \mid X_1 = x, X_2 = z] \\ &\quad - \delta E^o[\Pi_{m1}^o \mid X_1 = z]. \end{aligned}$$

From the definition of the simple relational contract,

$$\begin{aligned} [1 - Q(x, z)] \{d(x, z) + \delta E^o[\Pi_{m1}^o \mid X_1 = z]\} \\ = E^o[d^o + \delta \Pi_{m2}^o \mid X_1 = x, X_2 = z] \end{aligned} \quad (31)$$

$$\begin{aligned} [1 - Q(x, z)] \{-d(x, z) + \delta E^o[\Pi_{b1}^o \mid X_1 = z]\} \\ = E^o[-d^o + \delta \Pi_{b2}^o \mid X_1 = x, X_2 = z]. \end{aligned} \quad (32)$$

(31) implies that the simple relational contract achieves the same expected discounted profit for the buyer for the

manufacturer as relational contract o , for each initial state $X_1 \in \mathcal{X}$. (32) implies the analogous result for the buyer. It remains to show that the simple relational contract is self-enforcing. (26) implies that the simple relational contract satisfies (4)–(5). Together, (27) and (31) imply that the simple relational contract satisfies (6). Similarly, (28) and (32) imply that the simple relational contract satisfies (7). (29) implies

$$\delta E^o[\Pi_{m1}^o | X_1 = z] \geq -d(x, z),$$

which implies that the simple relational contract satisfies (8). Similarly, (30) implies

$$\delta E^o[\Pi_{m1}^o | X_1 = z] \geq -d(x, z),$$

which implies that the simple relational contract satisfies (9). We conclude that the simple relational contract is self-enforcing.

Step 3. The Simple Optimal Relational Contract. In searching for an optimal relational contract, we can restrict attention to self-enforcing relational contracts with the simple structure described in Step 2. We can also assume without loss of generality that the manufacturer is allocated a fraction $\alpha \in [0, 1]$ of the total discounted expected profit. Let $V(z)$ denote the maximum total discounted profit that can be achieved with such a relational contract, starting in state z . Suppose that the firms will adopt this relational contract in the second period, and would like to develop a discretionary transfer payment, action strategies for the two firms, and a termination function for the first period, that are self-enforcing and maximize expected total discounted profit. Given that the system is initially in state x , this must result in expected total discounted profit of $V(x)$.

$$V(x) = \max \left[\delta \sum_{z \in \mathcal{X}} P_{xz}(0, 0) v(z); \max_{a, Q, a_m, a_b} \left\{ -c_m(a_m, x) + c_b(a_b, x) + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a_b) [Y(x, z) + \delta [1 - Q(x, z)] V(z)] \right\} \right]$$

subject to:

$$a_m \in \max_{a \in A_m(x)} \left\{ -c_m(a, x) + \sum_{z \in \mathcal{X}} P_{xz}(a, a_b) [1 - Q(x, z)] \cdot [d(x, z) + \delta \alpha V(z)] \right\}$$

$$a_b \in \max_{a \in A_b(x)} \left\{ -c_b(a, x) + \sum_{z \in \mathcal{X}} P_{xz}(a_m, a) \{ Y(x, z) + [1 - Q(x, z)] \} \cdot [-d(x, z) + \delta(1 - \alpha)V(z)] \right\}$$

$$\delta \alpha V(z) \geq -d(x, z)$$

$$\delta(1 - \alpha)V(z) \geq d(x, z)$$

$$0 \leq Q(x, z) \leq 1.$$

This is equivalent to

$$TV = V.$$

From Proposition 1, we know that T has a largest fixed point V^* , and therefore $V = V^*$. Thus, the optimal terms are as given in the statement of the theorem. \square

References

- Abreu, D. 1988. On the theory of infinitely repeated games with discounting. *Econometrica* 56(3) 386–396.
- Abreu, D., P. Milgrom, D. Pearce. 1991. Information and timing in repeated partnerships. *Econometrica* 59(6) 1713–1733.
- Abreu, D., D. Pearce, E. Stacchetti. 1986. Optimal cartel equilibria with imperfect monitoring. *J. Econom. Theory* 39(1) 251–269.
- Atkins, D., H. Krishnan, X. Zhao. 2005. Relational price-only contract in supply chains. Working paper, Sauder School of Business, University of British Columbia, Vancouver, Canada.
- Baker, G., R. Gibbons, K. J. Murphy. 2001. Bringing the market inside the firm. *Amer. Econom. Rev.* 91(2) 212–218.
- Baker, G., R. Gibbons, K. J. Murphy. 2002. Relational contracts and the theory of the firm. *Quart. J. Econom.* 117(1) 39–83.
- Doornik, K. 2006. Relational contracting in partnerships. *J. Econom. Management Strategy*. Forthcoming.
- Dyer, J. H. 1996. Specialized supplier networks as a source of competitive advantage: Evidence from the automobile industry. *Strategic Management J.* 17(4) 271–291.
- Dyer, J. H., W. Chu. 2000. The determinants of trust in supplier–automaker relationships in the U.S., Japan and Korea. *J. Internat. Bus. Stud.* 31(2) 259–285.
- Fudenberg, D., D. Levine, E. Maskin. 1994. The folk theorem with imperfect public information. *Econometrica* 62(5) 997–1039.
- Granovetter, M. 1985. Economic action and social structure: The problem of embeddedness. *Amer. J. Sociol.* 91(3) 481–510.
- Heidi, J. B., A. S. Miner. 1992. The shadow of the future: Effects of anticipated interaction and frequency of contact on buyer–seller cooperation. *Acad. Management J.* 35(2) 265–291.
- Holmstrom, B. 1982. Moral hazard in teams. *Bell J. Econom.* 13(2) 324–340.
- Iyer, A. V., L. B. Schwarz, S. A. Zenios. 2005. A principal–agent model for product specification and production. *Management Sci.* 51(1) 106–119.
- Kenworthy, L., S. Macaulay, J. Rogers. 1996. “The more things change...”: Business litigation and governance in the American automobile industry. *Law Soc. Inquiry* 21(3) 631–678.
- Klein, B., K. Leffler. 1981. The role of market forces in assuring contractual performance. *J. Political Econom.* 89(4) 615–641.
- Laffont, J. J., J. Tirole. 1990. Adverse selection and renegotiation in procurement. *Rev. Econom. Stud.* 57(4) 597–625.
- Levin, J. 2003. Relational incentive contracts. *Amer. Econom. Rev.* 93(3) 835–847.
- Powell, W. W. 1990. Neither market nor hierarchy: Network forms of organization. *Res. Organ. Behav.* 12 295–336.
- Powell, W. W., K. W. Koput, L. Smith-Doerr. 1996. Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Admin. Sci. Quart.* 41(1) 116–145.
- Rayo, L. 2004. Relational team incentives and ownership. Working paper, University of Chicago, Chicago, IL.
- Rey, P., B. Salanie. 1996. On the value of commitment with asymmetric information. *Econometrica* 64(6) 1395–1414.
- Robinson, D., T. Stuart. 2002. Just how incomplete are incomplete contracts? Evidence from biotech strategic alliances. *J. Financial Econom.* Forthcoming.
- Rogerson, W. P. 1985. The first-order approach to principal–agent problems. *Econometrica* 53(6) 1357–1368.
- Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50(2) 97–108.
- Sako, M. 1992. *Prices, Quality, and Trust*. Cambridge University Press, Cambridge, UK.

- Schweitzer, M. E., J. C. Hershey, E. T. Bradlow. 2005. Promises and lies: Restoring violated trust. Working paper, The Wharton School, University of Pennsylvania, Philadelphia, PA.
- Tarski, A. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* **5** 285–309.
- Taylor, C. R., S. N. Wiggins. 1997. Competition or compensation: Supplier incentives under the American and Japanese subcontracting systems. *Amer. Econom. Rev.* **87**(4) 598–618.
- Uzzi, B. 1996. The sources and consequences of embeddedness for the economic performance of organizations: The network effect. *Amer. Sociol. Rev.* **61**(4) 674–698.
- Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. *Admin. Sci. Quart.* **42** 35–67.
- Zaheer, A., B. McEvily, V. Perrone. 1998. Does trust matter? Exploring the effects of interorganizational and interpersonal trust on performance. *Organ. Sci.* **9**(2) 141–159.