Relational Contracts in Strategic Alliances☆

George P. Baker[a,*], Robert Gibbons[b], and Kevin J. Murphy[c,]

[a]*Harvard Business School, Boston, MA 02163*
, Massachusetts Institute of Technology, Cambridge, MA 02142-1347
[c]*University of Southern California, Los Angeles, CA 90089-1427*

Received 30 August 2002

Abstract

Strategic alliances range from short-term cooperative projects, through long-term partnerships and joint ventures, to transactions that permanently restructure firm boundaries and asset ownership. In this paper, we draw on detailed discussions with practitioners to present a rich model of feasible governance structures. Our model focuses on three issues emphasized by practitioners: spillover effects (as opposed to hold-ups motivated by specific investments), contracting problems ex post (as opposed to only ex ante), and relational contracts (as opposed to spot transactions). Using this model, we first identify the managerial challenges presented by each governance structure and then analyze which governance structure is efficient in which environments.

JEL classification: D2; L14; L22

Key Words: Strategic alliances; Theory of the Firm; Relational Contracts

---

☆ We are very grateful for discussions with several practitioners, especially Judy Lewent and Richard Kender of Merck, Mark Edwards of Recombinant Capital, Douglas Birdsall of Northwest Airlines, and Gregory Gardiner of Yale University. We also appreciate discussions with Oliver Hart, Bengt Holmstrom, and Jon Levin, and research support from Harvard Business School (Baker and Gibbons), MIT's Program for Research on Innovation in Markets and Organizations (Gibbons), and USC's Marshall School (Murphy).

*Corresponding author: Tel.: (617) 495-6119; fax (617) 496-4191

E-mail addresses: gbaker@hbs.edu (G. Baker),
rgibbons@mit.edu (R. Gibbons),
kjmurphy@usc.edu (K. Murphy).

Relational Contracts in Strategic Alliances

1. Introduction

Strategic alliances exist in a bewildering variety of forms, ranging from short-term cooperative projects, through long-term partnerships and joint ventures, to transactions that permanently restructure firm boundaries and asset ownership. Even brief inspection of the existing governance structures in industries such as pharmaceuticals, biotechnology, medical devices, airlines, and telecommunications shows that firms have invented far more ways to work together than organizational economics has so far expressed (not to mention evaluated).

To investigate this plethora of observed attempts to coordinate activities across firms, we conducted a series of detailed interviews with practitioners who design, implement, consult to, and negotiate terms for these governance structures. Several important ideas arose during these discussions—some familiar from the organizational-economics literature, but others more novel. Three ideas emerged as especially important factors determining the form and performance of strategic alliances: spillovers (or externalities) from the joint project onto the parents; the need for governance structures to induce efficient behavior ex post, since contracts cannot; and the importance of relationships in the successful implementation of these alliances. Standard ideas, such as hold-ups motivated by specific investments, played markedly smaller roles in what we heard from practitioners.

In this paper, we develop a model that integrates the three factors emphasized by the practitioners – spillovers, contracting problems ex post, and relationships. We use this model to examine a collection of governance structures that our interviewees described, such as alliances (where non-integrated parties coordinate activities without changing firm boundaries or asset ownership), acquisitions (where one parent acquires and controls the joint project), mergers (where all the assets are merged into a single firm), mutual divestitures (where an autonomous entity is created to pursue the joint project without parental ownership or direct control), and joint ventures (where a new entity is created and jointly owned by the parents).[1]

We examine the relative efficiency of alternative governance structures in two cases: a static environment (where the parties engage in a one-shot transaction), and a relational environment (where the parties are engaged in ongoing relationships). Our theoretical analysis of the relational case identifies the managerial challenges associated with each governance structure and provides a simple characterization of the efficient governance structure in a relational environment.

While much of the recent literature on vertical integration and the theory of the firm has emphasized hold-ups motivated by specific investments, we utilize the "relational-adaptation theory of the firm" developed in Baker, Gibbons, and Murphy (2002a), which formalizes and extends the theory of the firm begun by Simon (1951) and Williamson (1975). The key ideas in the relational-adaptation theory of the firm are that (1) ongoing relationships can help parties achieve efficient adaptations as states of the world are realized and (2) integration can either enhance or inhibit the parties' effort to use their relationship towards this end. In this paper, we apply this idea to a much richer set of governance structures than the simple "make or buy" problem studied in our 2002a paper.

Our relational-adaptation approach complements two other streams of research – one based on property rights, the other on agency theory – that also emphasize the importance of ongoing relationships in organizational settings. Regarding property rights, Garvey (1995), Baker, Gibbons, and Murphy (1999, 2002b), and Halonen (2002) enriched static models in the tradition of Grossman and Hart (1986), Hart and Moore (1990), and Hart (1995) by adding ongoing relationships to static property-rights models. In the static models, ex post surplus shares create ex ante incentives for non-contractible specific investments. Adding relationships

[1] Note that we use "alliances" both as an umbrella term (as in "strategic alliances exist in many forms") and as a specific governance structure (in which non-integrated parties coordinate activities). We trust our meaning will be clear from context.

to the static property-rights models enriches the feasible set of ex post surplus shares, and so improves ex ante incentives (and often also changes the efficient structure of asset ownership, compared to the static model).

Regarding agency theory, Baker, Gibbons, and Murphy (1994), Che and Yoo (2001), and Rayo (2002) enriched static models in the tradition of Holmstrom (1982), Holmstrom and Milgrom (1991), and Baker (1992), again by adding ongoing relationships to the static models. Adding relationships to the static agency models enriches the set of feasible incentive contracts, this time by allowing relational incentive contracts to link pay to subjective performance measures (i.e., variables that are observable but not verifiable). The feasibility of such relational incentive contracts changes the role, and hence typically also the efficient design, of incentive contracts based on objective performance measures, compared to the static agency models.

In contrast to both these property-rights and agency-theory streams of research, our relational-adaptation approach has no ex ante actions, so relationships have no role to play in improving ex ante incentives. Instead, our focus is on the complementary problem of ex post adaptation. This distinction between ex ante incentives and ex post adaptation has both theoretical and empirical implications, as we discuss below.

The paper is organized as follows. Section 2 develops a simple model with two transferable assets and two firms and then conducts a static analysis of this simple model. Section 3 turns to our main focus, relational governance, first by providing evidence on the potential importance of relationships in strategic alliances, then by analyzing relational governance in the simple model from Section 2, and finally by analyzing relational governance in a much more general model (with an arbitrary number of assets and firms, and very general payoffs and action spaces). Our analysis delivers three primary insights. First, we identify the managerial challenges associated with each governance structure, and show that these challenges differ in static and relational environments. Second, we define and identify the "efficient" governance structure as the structure that minimizes the managerial challenges in each environment. Third, we establish that relationships not only improve the performance of any given governance structure, but also that the efficient governance structure under relational governance differs from the efficient structure under static governance.

Sections 2 and 3 restrict attention to governance structures with "unique control" (i.e., any given asset is owned by exactly one firm). Section 4 sketches a new approach to "joint control," where certain assets are jointly owned by multiple firms. Because our model emphasizes contracting problems ex post, joint control raises theoretical issues that have gone unnoticed in the property-rights approach to joint ownership (e.g., Hart and Moore (1990), Garvey (1995), Rey and Tirole (2001), and Halonen (2002)). These theoretical issues—including how utilization decisions are made for jointly owned assets and how non-contractible profits from these decisions are subsequently divided—correspond to difficulties that practitioners described to us concerning the management of joint ventures.

Section 5 follows Maskin and Tirole (1999), by noting that what we call owning an asset could just as well be owning a decision right, and the latter right could be conveyed by contract. Asset-ownership structures with unique control then correspond to contracts in which every decision right is allocated to one party or another, whereas ownership structures with joint control correspond to contracts in which some decision rights are not allocated to any one contracting party (and so are jointly controlled by the community of contracting parties). We believe that the latter contracts are prominent—both in the world and in recent empirical work (e.g., Lerner and Merges, 1998; Elfenbein and Lerner, 2002; and Ryall and Sampson, 2002) —so we expect that our extension of the Maskin-Tirole logic to joint-control contracts will be useful.

Section 6 turns from theory towards testing. We survey the existing empirical literature and relate it to our theoretical approach. We also describe how the static and relational predictions of our model might be tested, and also how one might investigate comparative predictions about how the optimal governance structure differs in static versus relational environments. Finally, Section 7 concludes.

2. A Simple Model of Static Governance

In this section, we first develop a simple model that incorporates two of the factors practitioners emphasized to us: spillovers and contracting problems ex post. We then conduct a static analysis of several governance structures that firms commonly use to coordinate their activities, including mergers, alliances, acquisitions, and divestitures.

2.1 A Simple Model

Suppose there are four assets, {A, a, B, b}, and (initially) two firms, A and B. Firm A owns {A, a} and Firm B {B, b}. Asset A represents the core activity of Firm A, and asset B the core activity of Firm B. Assets {a, b}, on the other hand, are valuable only if they are used together, in coordinated fashion. Formally, assets {a, b} can be used together to produce profits $\pi_a = \pi_b = \pi > 0$, or used separately to produce $\pi_a = \pi_b = 0$. The profits π_a and π_b accrue to the owners of assets a and b.

In addition to producing profits $\pi_a = \pi_b = \pi$, coordinated use of {a, b} can also affect the profits from the core activities of Firms A and B. Let π_A and π_B denote the payoffs from these "spillover effects" on A and B. (That is, any profit from the core activity of Firm A that is independent of the use of assets {a, b} is excluded from π_A and ignored hereafter, and likewise for Firm B.) As in most of the literature on asset ownership, the payoffs π_A,

Figure 1

Economic Environment

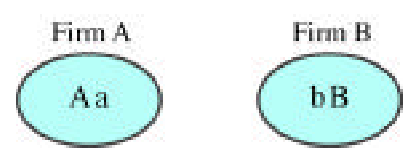

Coordinated use of {a,b} yields:

Direct profits: π_a, π_b (alienable)
Spillovers: π_A, π_B (inalienable)

π_B, π_a, and π_b are observable but not verifiable.[2] All of this is summarized in Figure 1.

The coordinated use of {a, b} could either complement or compete with the core activities of one or both firms. To capture these possibilities, the spillover payoffs depend on a state variable, s, which also is observable but not verifiable. The spillover payoffs $\pi_A(s)$ and $\pi_B(s)$ can each take on three values: high (H), medium (M), or low (L), where H > 0 > M > L. There are thus nine possible states, denoted s_{ij}, where i=H, M, or L is the realization of π_A, and j=H, M, or L is the realization of π_B. Denote the probability of state s_{ij} by p_{ij}. In this simple environment, we assume not only symmetric payoffs (i.e., the realizations of both π_A and π_B are H, M, or L), but also symmetric probabilities (i.e.,$p_{ij} = p_{ji}$).

While our focus on asset ownership is in the Grossman-Hart-Moore (GHM) tradition, we depart from this literature by imposing the following assumption: the right to determine the utilization of an asset is inalienably attached to the asset's owner, and is therefore non-contractible *even ex post*. By making this assumption, we intend to focus on settings where there is an important difference between perfunctory and consummate performance, but the latter is not contractible. For example, Firm A may sign a contract committing it to use {a, b} jointly with Firm B, but Firm A may go through the motions to satisfy the letter while violating the spirit of the contract.[3]

Our assumption that asset utilization is non-contractible ex post rules out bargaining over asset utilization after the state has been observed (because implementing the bargained utilization of an asset would require a contract, unless the bargained utilization coincides with the utilization preferred by the asset's owner). Even without ex post contracts, however, one

---

2 The assumption that the payoffs are not verifiable allows us to omit court-enforceable contracts from our analysis. The usual argument given for this omission is that the asset-ownership analysis pertains to the inevitable gaps in the court-enforceable contracts. But this argument ignores the possibility that the effects of asset ownership may interact with the terms of court-enforceable contracts, as in our (1994) paper. Thus, a superior approach would include the incomplete court-enforceable contracts in the analysis, along with asset ownership and relational contracts, as we began to do in a simple setting in our (2001) paper.

3 We were led to this assumption of non-contractability ex post by our discussions with practitioners. Hart and Holmstrom (2002) and Aghion, Dewatripont, Rey (2002) explore related static models.

could imagine ex post bargaining over asset ownership. That is, because it is not possible to enforce a contract that directly influences the utilization decision taken by an asset's current owner, there may be gains created by selling the asset to a new owner whose self-interested utilization decision would be closer to the desired decision that could not be achieved by contract. But we rule out such ex post renegotiation of asset ownership, for example because the opportunity to use {a, b} jointly is fleeting, and transferring ownership would take some time. We comment further on this issue below.

To summarize, the timing in the simple environment is as follows: (1) the parties choose a governance structure (e.g., an asset ownership structure; see below), perhaps accompanied by side payments; (2) the state is publicly revealed; (3) the decision to use {a, b} jointly or separately is made (details depend on what governance structure was chosen; see below); and (4) payoffs from asset-utilization decisions are received.

We assume that the parties are risk-neutral and so seek to maximize the expected value of the sum of any side payments and utilization payoffs they may receive. The model is interesting only if the asset-utilization decision differs across states and governance structures, so we assume not only that $H > 0 > M > L$, but also that $M + 2\pi < 0$ (so that an individual firm would not like to see {a, b} used jointly if its spillover payoff were M, even if it owned both assets) and that $H + M + 2\pi > 0 > H + L + 2\pi$ (so that it is efficient to use {a, b} jointly if the spillover payoffs are H and M but not if the spillover payoffs are H and L). Panel A of Figure 2 shows the aggregate payoffs (to Firm A plus Firm B) from joint utilization of {a, b} in each state s_{ij}. The shaded cells of Panel A indicate the states in which it is efficient to use {a, b} jointly: s_{HH}, s_{HM}, and s_{MH}. We will call these three states the efficient states, and the remaining six states the inefficient states.

2.2 *Static Governance under Unique Control*

We now analyze static governance in this simple model, but we restrict attention to governance structures with “unique control” (i.e., any given asset is owned by exactly one firm); governance structures involving joint control are analyzed in Section 4 below. Since our main purpose in this section is to build intuition, we focus on mergers, alliances, acquisitions, and divestitures, discussing other feasible unique-control governance structures only briefly.

Figure 3a
Merger

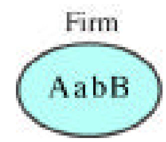

Coordinated use of {a,b} if efficient

MERGERS. We first consider (and then assume away) mergers, where all the assets of Firm A are merged with all the assets of Firm B to form a single firm owning assets {A, a, B, b}, as shown in Figure 3a. The owner of this merged firm will use {a, b} jointly whenever $\pi_A + \pi_B + 2\pi > 0$, which occurs in the three efficient states – s_{HH}, s_{HM}, and s_{MH} (see Panel A of Figure 2). In the six inefficient states, a merged firm will forego the profit of 2π from joint use of {a, b} because the negative spillover payoffs more than outweigh the positive joint-use payoffs. The expected surplus produced by a merger (V^M) therefore coincides with the first-best social surplus (V^{FB}),

$$(1) \qquad V^M = V^{FB} = p_{HH}2(H+\pi) + 2p_{HM}(H+M+2\pi).$$

Since such mergers would completely internalize the externality problems that animate our model, we assume that there are factors outside our model (such as antitrust considerations, or efficiency considerations in the style of Grossman-Hart-Moore) that make a merger either infeasible or woefully inefficient, so that we need not consider it further.[4]

[4] Another interpretation of our model is that there are only two assets, a and b, and the payoffs π_A and π_B are private benefits to parties A and B. Under this interpretation, there are no physical assets A and B, so no single party can receive all four payoffs π_A, π_B, π_a, and π_b. That is, mergers could be assumed to be contractually infeasible, rather than politically infeasible or economically sub-optimal.

STATIC ALLIANCES. In our definition of an alliance, there is no change in asset ownership: Firm A owns {A, a} and Firm B owns {B, b}, and the firms engage in arms-length transactions involving the joint use of {a, b}, as shown in Figure 3b. In this static analysis, an arm's-length transaction amounts to unilateral decision-making: Firm i will be willing to participate in coordinated use of {a, b} only if $\pi_i(s)+\pi > 0$. Since $H+\pi > 0 > M+\pi > L+\pi$, the only state in which both firms will be willing to use {a, b} jointly is s_{HH} (see Panel B of Figure 2). Thus, the expected social surplus from an alliance is

Figure 3b
Static Alliance

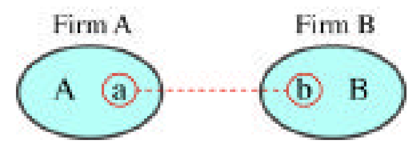

Coordinated use of {a,b} if profitable for *both* Firm A and Firm B

(2) $$V^{AL} = 2p_{HH}(H+\pi).$$

STATIC ACQUISITIONS. In an acquisition (by Firm A), Firm A owns {A, a, b} and Firm B owns {B}, as shown in Figure 3c. Given our assumption that $M + 2\pi < 0$, Firm A will choose to use {a, b} jointly if and only if $\pi_A(s) = H$ (see Panel C of Figure 2). Alternatively, Firm A could own {A} while Firm B owns {B, b, a}, in which case Firm B would choose to coordinate if and only if $\pi_B(s) = H$. Given our symmetry assumptions in this simple model, the expected social surplus is the same in either case:

Figure 3c
Static Acquisition (by A)

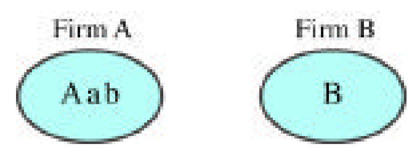

Coordinated use of {a,b} if profitable for Firm A

(3) $$V^{AC} = p_{HH}(2H+2\pi) + p_{HM}(H+M+2\pi) + p_{HL}(H+L+2\pi).$$

STATIC DIVESTITURES. In a mutual divestiture, a third party (call it Firm C) purchases asset {a} from Firm A and asset {b} from Firm B, as shown in Figure 3d. Since Firm C owns {a, b} and joint use of {a, b} delivers profits of $2\pi > 0$ independent of the state, Firm C will use the assets jointly in all states (see Panel D of Figure 2). The total expected surplus (now received by three parties) from divestitures is then

(4) $$V^{DV} = 2\pi + \sum_{i,j=H,M,L} p_{ij}(i + j).$$

Figure 3d

Static Divestiture

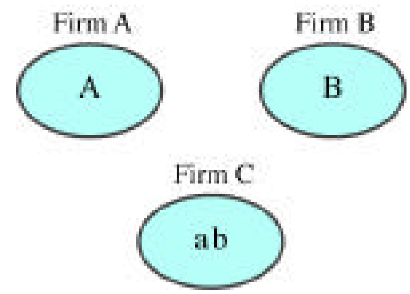

Coordinated use of {a,b} if profitable for Firm C

The efficient static governance structure depends on both the magnitude of the inefficiencies (that is, the cost of joint use in inefficient states, and the opportunity cost of foregoing joint use in efficient states) and the state probabilities. Equations (2') through (4') restate the social surpluses for alliances, acquisitions and mutual divestitures relative to the first-best surplus, V^{FB}:

(2') $$V^{AL} = V^{FB} - 2p_{HM}(H + M + 2\pi),$$

(3') $$V^{AC} = V^{FB} - p_{HM}(H + M + 2\pi) - p_{HL}|H + L + 2\pi|,$$

(4') $$V^{MD} = V^{FB} - 2p_{HL}|H + L + 2\pi| - 2p_{ML}|M + L + 2\pi| - p_{MM}|2M + 2\pi| - p_{LL}|2L + 2\pi|,$$

Comparing (2') through (4') shows that each static governance structure is second-best for some parameter values. For example, if p_{HM} is small, then alliances tend to be second-best (because they do not induce joint use in any inefficient states, and the efficient states in which they fail to achieve joint use have small probability). Similarly, if p_{HM} is large, then mutual divestitures tend to be second-best (because they induce joint use in all the efficient states, and the probabilities of the inefficient states are low). Finally, if p_{HM} is moderate, then acquisitions tend to be second-best (because they induce joint use in most of the efficient states, but avoid joint use in most of the inefficient states).

Even in this simple environment, there are other possible unique-control governance structures. Some are limiting cases of what we have just described, such as disaggregating the four assets into four separate firms. Others are hybridizations of what we have just described, such as divesting only asset b to Firm C, so that Firm A owns {A, a} and Firm B owns {B}. Since our main purpose in analyzing this simple model is to build intuition, we restrict attention to the three unique-control governance structures analyzed above: alliances, acquisitions, and mutual divestitures. From these three options, the second-best static governance structure is the one that maximizes expected surplus,

(5) $V^{ST} \equiv \max[V^{AL},V^{AC},V^{DV}]$.

More generally, V^{ST} is the maximum across all feasible static governance structures, including the other unique-control governance structures mentioned above and the joint-control governance structures described in Section 4.

3. Relational Governance

In this section, we begin by presenting evidence that ongoing relationships between firms are potentially important. Turning to theory, we then analyze relational governance of alliances, acquisitions, and divestitures in the simple model defined in Section 2.1. Finally, we generalize this relational-governance analysis of unique-control structures to allow many assets, many firms, and very general payoffs and action spaces.

3.1 Relationships in Strategic Alliances

There are several ways in which the shadow of the future can loom large for alliance partners. First, alliances are often long-lived and involve continuing interactions between the parties over an extended period. For example, the Fuji-Xerox relationship lasted for decades and included several important restructurings at key junctures (McQuade and Gomes-Casseres, 1992). Second, firms often engage in repeat alliances with the same partners (Gulati, 1995a). In both of these settings, each partner may choose its current actions with an eye on the likely future responses of the other party.

A third possible way that the future may loom large is through indirect ties. For example, if Firms A and B have one alliance, and Firms B and C have another, then A's current actions with B may be influenced by A's potential future dealings with C. More generally, a network of indirect ties can facilitate information flows between firms that have not yet been alliance partners (Gulati, 1995b).

To begin to document the potential importance of relationships in strategic alliances, we analyzed data collected by industry specialists at Recombinant Capital on nearly 12,500 publicly disclosed alliances in the pharmaceutical and biotechnology industries from 1973 to 2001. We use these data to explore each of the three forms of relationships just described: long-lived contracts, repeated contracting, and indirect ties.

Regarding long-lived contracts, the database does not offer complete information on the longevity of individual alliances, but we can nonetheless provide some suggestive evidence. First, of the 12,5000 alliances in the data, only 372 are listed as formally terminated between 1973 and 2001. Second, even for those that were terminated, the median time between the initial contract and the termination was 33 months.[5] Third, 1,548 alliances were formally revised (but not terminated) during the sample period, and the median time from the initial contract to the revision was 21 months (constituting a lower bound on alliance longevity for these contracts). Finally, for over 10,000 alliance contracts, there is no evidence that the contract was not open-ended. In sum, these data suggest that alliances are often not one-shot transactions, but instead hold the prospect of continuing interactions.

Regarding repeat contracting, Table 1 presents evidence on repeat alliances between the same partners. In the Recombinant Capital database, most pairs of firms (9,462) do only one

---

5 These data exclude 12 proposed mergers or acquisitions that were terminated prior to completion.

deal with each other, but over a thousand pairs of firms do more than one deal together; 57 pairs do five or more deals together. Thus, the prospect of doing another deal is not negligible.

Finally, regarding indirect ties, Table 2 shows the number of alliances (and the number of partners) for the top 12 pharmaceuticals and top 12 biotechnology firms, where "top" is defined by the number of alliances reported in the Recombinant Capital database. These 24 firms (defined as the surviving parent as of year-end 2001 in the case of mergers and acquisitions) comprised less than 1% of the 4,231 surviving parents (after mergers and acquisitions) in the sample, but were involved in 32% of the 12,451 reported alliances. In short, a few firms are doing lots of the alliances, raising the question: with whom?

Figure 4 shows the dense network of ties between these top pharmaceuticals and biotechs. On average, each firm among these 24 has at least one alliance with 15 of the other 23 firms. But far from all the alliances involving these 24 firms are with another of the 24 firms. To the contrary, the 24 firms in Table 2 had alliance arrangements with 1,308 partners outside of the 24, and these 1,332 firms entered 11,303 alliances (91% of the universe identified by Recombinant). Including the partners of these 1,332 firms yields a total of 3,421 firms (81% of the firms) who were involved in 98% of the reported alliances.

These results on indirect ties suggest that even the most peripheral firm was rarely more than "two phone calls" away from a "top 24" firm, who in turn was never more than two phone calls away from another peripheral firm. Again, we interpret these data as suggesting that relationships between alliance partners are potentially quite important, at least in the pharmaceutical and biotechnology industries.

And it turns out that these industries are far from unique. For example, the indirect ties between partners in alliances, joint ventures, and the like are again extremely dense in the internet sector (see http://www.orgnet.com/netindustry0104.gif) and the automotive industry. Based on suggestive evidence of this kind, we turn next to theoretical analyses of how governance structures aid or impede ongoing relationships.

3.2 Relational Governance in the Simple Model

In this sub-section we incorporate the third element evident from our discussions with practitioners: having introduced spillovers and ex post contracting problems in Section 2.1, we now add ongoing relationships. In our model, relationships remedy the ex post inefficiencies in the static governance structures analyzed in Section 2.2. As explained above, this emphasis on relational adaptation differs from the two earlier research streams, in which relationships were used in property-rights and agency models to enrich the feasible set of sharing rules, and hence influence ex ante incentives.

As is now standard, we model an ongoing relationship as a repeated game. As usual, the discount rate in the repeated game can be interpreted as reflecting the exogenous probability that the relationship will end. Thus, even the 372 alliances that were formally terminated may have begun life with the prospect of an ongoing relationship.

Following a large literature, we interpret an equilibrium in the repeated game as a "relational contract" (i.e., an agreement between the parties that cannot be enforced by a court, and so must be enforced by the parties' concerns for their reputations). Macaulay (1963) and Macneil (1978) introduced the idea of a relational contract to the sociological and legal literatures, respectively; early economic models of relational contracts include Klein and Leffler (1981), Telser (1981), and Bull (1987).

In our simple model, a relational contract is a set of payments and asset-utilization decisions, all of which may depend on non-contractible variables. We restrict attention to first-best relational contracts: the assets {a, b} are used jointly only in the efficient states s_{HH}, s_{HM}, and s_{MH}. The question then becomes whether there exist self-enforcing payment schemes that induce the parties to take these first-best asset-utilization decisions. These payment schemes can include positive or negative payments occurring before the state is revealed, payments after the state is revealed but before the utilization decisions are made, and payments after the utilization decisions are made.

We analyze trigger-strategy equilibria: if any firm reneges (on a payment or a utilization decision), the firms engage in static transactions thereafter. We view the punishment phase of trigger strategies as reflecting a reasonable tradeoff between the theoretical appeal of renegotiation and the intuitive appeal of spite. We impose another form of renegotiation by assuming that if reneging occurs then the parties engage in efficient static governance structure in all future periods. (Achieving efficient static governance will typically require a change in asset ownership at the end of the present period, with an accompanying side-payment.) Thus, after reneging, the total expected surplus will be the maximum surplus from the alternative static governance structures, V^{ST} in (5).

In Section 3.3, we analyze a much more general model than the simple one defined in Section 2.1. For this more general model, we derive a simple but powerful necessary and sufficient condition for whether a first-best relational contract exists for a given governance structure: the maximum total reneging temptation must be smaller than the present value of the difference in total surplus between first-best and efficient static governance. In this section, we informally apply this condition, as follows:

- First, for a given governance structure, identify the states in which the static outcome of this governance structure is not first-best.
- Second, for each such state, compute each firm's payoff from its static best response to first-best asset-utilization decisions by the other asset owner(s).
- Third, compute the difference between the firm's best-response payoff and its first-best payoff. Call this difference the firm's reneging temptation (in this state).
- Fourth, sum this reneging temptation across firms. Call the result the total reneging temptation (in this state).
- Fifth, search across all such states for the maximum total reneging temptation (under this governance structure). Call the result the reneging temptation for this governance structure.

Using this methodology, we now derive the reneging temptation for the three unique-control governance structures analyzed in Section 2.2: alliances, acquisitions, and mutual divestitures.

RELATIONAL ALLIANCES. A static alliance achieves the first-best in every state except s_{MH} and s_{HM}. In state s_{MH}, Firm A is tempted to deviate: first-best asset utilization requires joint use of {a, b} in this state, but yields the negative payoff $M+\pi$ for Firm A, compared to the payoff of zero that Firm A could earn from its static best response. Firm B, on the other hand, is not tempted to deviate in state s_{MH}, so the total reneging temptation in state s_{MH} is $|M + \pi|$. Similarly, the total reneging temptation in state s_{HM} (where Firm B is tempted but Firm A is not) is also $|M + \pi|$. Therefore, the reneging temptation for a first-best relational alliance is $R^{AL} = |M + \pi|$.

RELATIONAL ACQUISITIONS. A static acquisition (by Firm A of asset b) achieves the first-best in every state except s_{HL} and s_{MH}. In state s_{HL}, Firm A is tempted to deviate: first-best asset utilization forbids joint use of {a, b} in this state, but joint use would yield the positive payoff $H+2\pi$ for Firm A, compared to the payoff of zero to Firm A in the first-best outcome. Because Firm B owns neither of the relevant assets (a or b), Firm B has no reneging temptation, in any state. The total reneging temptation in state s_{HL} is thus $H+2\pi$. In state s_{MH}, in contrast, Firm A is again tempted to deviate, but for the opposite reason: first-best asset utilization requires joint use of {a, b} in this state, but yields the negative payoff $M+2\pi$ for Firm A, compared to the payoff of zero from Firm 's static best response. Hence, the total reneging temptation in state s_{HL} is $|M + 2\pi|$. By our assumptions that $M + 2\pi < 0$ and $H + M + 2\pi > 0$, we have that $|M + 2\pi| < H < H + 2\pi$. Therefore, the reneging temptation for a first-best relational acquisition is $R^{AC} = H + 2\pi$.

RELATIONAL DIVESTITURES. A static divestiture fails to achieve the first-best in the inefficient states: s_{HL}, s_{MM}, s_{ML}, s_{LH}, s_{LM}, and s_{LL}. Because Firms A and B own neither of the relevant assets, they have no reneging temptation, in any state. In all the inefficient states, Firm C is tempted to deviate: first-best asset utilization forbids joint use of {a, b} in each of these states, but joint use would yield the positive payoff 2π for Firm C, compared to the payoff of zero to

Firm C in the first-best. Therefore, the reneging temptation for a first-best relational divestiture is $R^{DV} = 2\pi$.

Even this simple analysis delivers three insights, which we generalize below. First, we can now identify the managerial challenges created by a given governance structure in a relational environment. In an alliance, for example, the challenge is to inspire joint use in states where it is efficient but one party is harmed by joint use. In an acquisition, in contrast, the challenge is to prevent joint use in states where it is beneficial to the acquiring firm but even more harmful to the divesting firm. Finally, in a divestiture, the challenge is to keep Firm C from implementing joint use without regard to the spillovers on Firms A and B.

Second, we can determine which of these governance structures is efficient in a relational environment. Simply put, the efficient governance structure minimizes the managerial challenges. More formally, governance structures with smaller reneging temptations can achieve the first-best at higher discount rates, and so are efficient for a broader range of parameters. Recall that the reneging temptations are $R^{AL} = |M + \pi|$ for first-best relational alliances, $R^{AC} = H + 2\pi$ for first-best relational acquisitions, and $R^{DV} = 2\pi$ for first-best relational divestitures. Thus, in this simple model, relational divestitures always dominate relational acquisitions, so relational divestitures are efficient if $|M + \pi| < 2\pi$, and relational alliances are efficient if $|M + \pi| > 2\pi$.

Third, we now see that relationships matter in two ways: first, relationships improve the performance of a given governance structure, relative to the static outcome; second, relationships often make it efficient to change the governance structure, away from the efficient structure in a static environment. As a trivial example of the latter, recall from Section 2.2 that, in a static environment, acquisitions are optimal if the probability p_{MH} is moderate. But acquisitions are dominated by divestitures in a relational environment.

3.3 Relational Governance in a Richer Model

Our simple model included two alienable assets, two (or three) firms, and nine states. We now extend the model to allow for arbitrary finite numbers of assets, firms, and states, as well as much more general payoffs and action spaces. Our goal is to generalize the three insights above: in a relational environment, (1) different governance structures create different managerial challenges, (2) the efficient governance structure minimizes these managerial challenges (i.e., reneging temptations), and consequently (3) the efficient governance structure typically differs from that in a static environment.

In the spirit of the simple model, suppose that there are two kinds of assets. First, there are J alienable assets, indexed by by $j \quad J = \{1, \ldots, J\}$, akin to the assets {a, b} above. Let d_j denote the decision taken by the owner of asset j regarding the utilization of asset j, where d_j is chosen from the finite set D_j. Second, there are I inalienable assets, indexed by $i \quad I = \{1, \ldots, I\}$, akin to the assets {A, B} above. These inalienable assets belong to the I firms that may experience spillover payoffs from the utilization decisions concerning the J assets, akin to Firms A and B above. Let d_i denote the decision taken by firm i regarding the utilization of asset i, where d_i is chosen from the finite set D_i. (In the simple model, we ignored such decisions concerning inalienable assets, but it is easy to include them here.) We write **d** for the vector of decisions $((d_i, i \quad I); (d_j, j \quad J))$ and D for $\prod_{i \quad I} D_i \times \prod_{j \quad J} D_j$.

Let s denote the state, drawn from the finite set S according to the probability density f(s). There are then two kinds of payoffs, corresponding to the two kinds of assets. First, there are direct payoffs, which accrue to the owners of the alienable assets, akin to the payoffs π_a and π_b above. Let $\pi_j^D(\mathbf{d}, s)$ denote the direct payoff that accrues to the owner of asset j if decisions **d** are taken in state s. Second, there are spillover payoffs, which accrue to the owners of the inalienable assets, akin to the payoffs π_A and π_B above. Let $\pi_i^S(\mathbf{d}, s)$ denote the spillover payoff to firm i if decisions **d** are taken in state s. (It may be that one or more of the firms $i \quad I$ have spillover payments $\pi_i^S(\mathbf{d}, s) \equiv 0$, as Firm C did above.) As in the simple model, the state, the decisions, and the payoffs are observable but not verifiable.

The total payoff from decisions **d** in state s is the sum of the spillover payoffs and the direct payoffs,

$$(6)\qquad \sum_{i \in I} \pi_i^S(\mathbf{d}, s) + \sum_{j \in J} \pi_j^D(\mathbf{d}, s).$$

Let $\mathbf{d}^{FB}(s)$ denote the first-best decisions in state s, which maximize the total payoff in (6) over all $\mathbf{d} \in D$. Define

$$(7)\qquad V^{FB}(s) \equiv \sum_{i \in I} \pi_i^S(\mathbf{d}^{FB}(s), s) + \sum_{j \in J} \pi_j^D(\mathbf{d}^{FB}(s), s)$$

as the first-best total payoff in state s, and let $V^{FB} \equiv E_s[V^{FB}(s)]$.

A governance structure specifies which firms own which assets. Under unique control, each asset is owned by exactly one firm: there is neither joint ownership nor incomplete ownership. Formally, the unique-control governance structure g is the function G_g: J → I, where $G_g(j) = i$ means that firm i owns asset j under governance structure g. The payoff to firm i from decisions **d** in state s under governance structure g is then the firm's spillover payoff plus any direct payoffs that accrue to this firm in this governance structure,

$$(8)\qquad \pi_{ig}(\mathbf{d}, s) \equiv \pi_i^S(\mathbf{d}, s) + \sum_{\{j \in J: G_g(j) = i\}} \pi_j^D(\mathbf{d}, s) \quad .$$

In addition to defining payoffs, we must also define action spaces for governance structure g. Firm i controls inalienable asset i and also any alienable assets it owns in this governance structure, so firm i's action space is

$$(9)\qquad D_{ig} \equiv D_i \times \prod_{\{j \in J: G_g(j) = i\}} D_j \ .$$

We write $\mathbf{d}_{ig}$ as a typical element of D_{ig}.

We assume that, for each governance structure g, and for each state s, there is a unique Nash equilibrium, $\mathbf{d}_g^{NE}(s)$. That is, for each firm i, $\mathbf{d}_{ig}^{NE}(s)$ solves

(10) $$\max_{\mathbf{d}_{ig}\ \ D_{ig}} \pi_{ig}((\mathbf{d}_{ig},\mathbf{d}_{-ig}^{NE}(s)),s).$$

The expected payoff to firm i under the static governance structure g is then

(11) $$V_{ig}^{ST} \equiv E_s\left[\pi_{ig}(\mathbf{d}_g^{NE}(s),s)\right],$$

and we write V_g^{ST} for $\sum_{i\ \ I} V_{ig}^{ST}$.

Under a merger, a single firm would own all the assets, make all the decisions. and receive all the payoffs. In this case, $\mathbf{d}_M^{NE}(s) \equiv \mathbf{d}^{FB}(s)$ and $V^M = V^{FB}$. As before, we exclude mergers. (In fact, we have already done so by calling the assets i   I inalienable.) We now also assume that no other static governance structure achieves the first-best decisions $\mathbf{d}^{FB}(s)$ in every state. Therefore, the second-best static governance structure generates the expected total surplus

(12) $$V^{ST} \equiv \max_{g\ \ G} V_g^{ST} < V^{FB},$$

where G is the set of all governance structures (i.e., all functions mapping J into I). Let V_i^{ST} denote V_{ig}^{ST} under the second-best static governance structure that solves (12).

Because static governance is inefficient ($V^{ST} < V^{FB}$), it is natural to ask whether relational governance can improve on static governance. As in Section 3.2, we restrict attention to first-best relational contracts, in which the asset-utilization decisions are $\mathbf{d}^{FB}(s)$. Our task is then to determine whether there exist payment schemes that induce the parties to take these first-best decisions. As in our simple model, these payments can be positive or negative (i.e., they can be paid to or paid by a given firm), and they can occur at three different times. First, the payments might be "efficiency wages," denoted by t_{ig} and paid before the state or any decisions are observed. Second, the payments might be "bribes," denoted by $\tau_{ig}(s)$ and paid after the state is observed but before the parties make their asset-utilization decisions. Third, the payments might be "subjective bonuses," denoted by $T_{ig}(\mathbf{d}, s)$ and paid depending on whether asset-utilization decisions are appropriately tailored to the state.

Figure 5

Timing of Payments in a Relational Contract

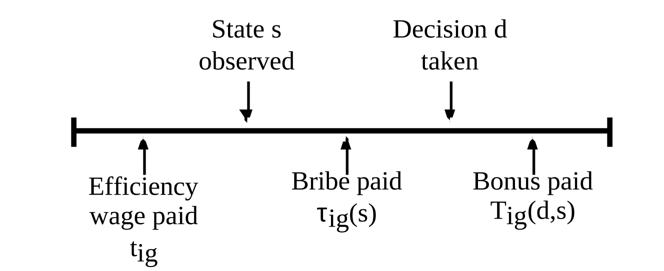

Figure 5 illustrates the timing of these potential payments within each period, relative to when the state is observed and the decisions are taken. We adopt the convention that a positive value of t_{ig}, $\tau_{ig}(s)$, or $T_{ig}(\mathbf{d}, s)$ is a payment to firm i, and a negative value a payment from firm i. Furthermore, we require that these payments balance: $\Sigma_{i \; I} t_{ig} = 0$, $\Sigma_{i \; I} \tau_{ig}(s) = 0$ for all s, and $\Sigma_{i \; I} T_{ig}(\mathbf{d}, s) = 0$ for all **d** and s.

As in Section 3.2, we analyze trigger-strategy equilibria, and we again assume that if reneging occurs then the firms engage in second-best static governance thereafter. Achieving second-best static governance typically requires reallocating ownership of the alienable assets, which typically requires a side-payment P_{ig} to firm i (or from firm i if P_{ig} is negative). To be concrete, one could imagine that these side-payments arise from the Nash bargaining solution. Instead, we impose only two weak constraints on these side-payments: balance (i.e., $\Sigma_{i \; I} P_{ig} = 0$) and individual rationality (i.e., $P_{ig} + \frac{1}{r}V^{ST}_{ig} \geq \frac{1}{r}V^{NE}_{ig}$, using some notation defined below).[6] Thus, after reneging, each firm receives P_{ig}, ownership of the alienable assets is reallocated, and the expected total surplus is V^{ST} per period thereafter, as defined in (12).[7]

Given a governance structure, there are many reneging constraints that must be satisfied if a relational contract is to be a repeated-game equilibrium. For a first-best relational contract to be an equilibrium under governance structure g, each firm i must: (a) be willing to pay (or receive) its efficiency-wage payment, t_{ig}; (b) be willing to pay (or receive) its bribe, $\tau_{ig}(s)$; (c)

---

6 Following Levin (2001), we could also allow these side-payments to depend on who reneged, which greatly simplifies the analysis of optimal equilibria.

7 It may seem strange that we rule out renegotiation of asset ownership ex post and yet allow for asset sales after reneging, but there is no inconsistency here. Again, our implicit assumption is that the opportunity to use assets jointly is fleeting, and transferring ownership would take some time, so this is why there is no renegotiation of ownership ex post.

be willing to take its first-best decisions, $d_{ig}^{FB}(s)$; and (d) be willing to pay (or receive) its bonus, $T_{ig}(\mathbf{d}, s)$.

To simplify the formal statements of the reneging constraints (a) through (d), we introduce the following notation:

$\pi_{ig}^{FB}(s) \equiv \pi_{ig}(\mathbf{d}^{FB}(s),s)$	Payoff to firm i (excluding side payments) from first-best decisions in state s under governance structure g
$U_{ig}^{FB}(s) \equiv t_{ig} + \tau_{ig}(s) + T_{ig}(\mathbf{d}^{FB}(s),s) + \pi_{ig}^{FB}(s)$	Payoff to firm i (including side payments) from first-best decisions in state s under governance structure g
$V_{ig}^{FB} \equiv E_s\left[U_{ig}^{FB}(s)\right]$	Expected payoff to firm i (including side payments) from first-best decisions under governance structure g
$\pi_{ig}^{NE}(s) \equiv \pi_{ig}(\mathbf{d}_g^{NE}(s),s)$	Payoff to firm i (excluding side payments) from Nash equilibrium decisions in state s under governance structure g
$V_{ig}^{NE} \equiv E_s\left[\pi_{ig}(\mathbf{d}_g^{NE}(s),s)\right]$	Expected payoff to firm i (excluding side payments) from Nash equilibrium decisions under governance structure g
$\mathbf{d}_{ig}^{BR}(s) = \underset{\mathbf{d}_{ig} \in D_{ig}}{\text{argmax}}\ \pi_{ig}((\mathbf{d}_{ig},\mathbf{d}_{-ig}^{FB}(s)),s)$	Firm i's best response in state s under governance structure g to first-best decisions by all other firms
$\pi_{ig}^{BR}(s) \equiv \pi_{ig}((\mathbf{d}_{ig}^{BR}(s),\mathbf{d}_{-ig}^{FB}(s)),s)$	Payoff to firm i (excluding side payments) from best response in state s under governance structure g, when all other firms take first-best decisions

Let r denote the discount rate per period. Then, given the notation above, the reneging constraints (a) through (d) can be written as follows. Constraint (a), that firm i be willing to pay (or accept) its efficiency-wage payment, t_{ig}, becomes

(13) $$\left(1+\tfrac{1}{r}\right)V_{ig}^{FB} \geq V_{ig}^{NE} + P_{ig} + \tfrac{1}{r}V_i^{ST}$$

for all i. The left-hand side of (13) is the expected present value of firm i's payoffs on the equilibrium path, where all firms honor the relational contract. The right-hand side is the expected present value from reneging on the efficiency-wage payment.[8] If firm i reneges on the efficiency-wage payment, then the relational contract is broken, so no further payments (bribes or bonuses) will be made this period by any firm, all firms will take Nash equilibrium decisions in this period (generating expected payoff V_{ig}^{NE}), ownership of alienable assets will be reallocated (with payment P_{ig}), and second-best static governance will ensue forever after (generating expected present value $\frac{1}{r}V_i^{ST}$).

Constraint (b), that firm i be willing to pay (or accept) its bribe, $\tau_{ig}(s)$, becomes

(14) $$[\tau_{ig}(s)+\pi_{ig}^{FB}(s)+T_{ig}(\mathbf{d}^{FB}(s),s)]+\tfrac{1}{r}V_{ig}^{FB}\geq\pi_{ig}^{NE}(s)+P_{ig}+\tfrac{1}{r}V_i^{ST}$$

for all i and s. There are two differences between (13) and (14): in (14), t_{ig} has already been paid, so it does not appear in this period's payoffs (the terms in square brackets) on the left-hand side, and the state s has already been realized, so this period's payoffs are contingent on s, not expectations.

Constraint (c), that firm i be willing to take its first-best decisions, $d_{ig}^{FB}(s)$, becomes

(15) $$[\pi_{ig}^{FB}(s)+T_{ig}(\mathbf{d}^{FB}(s),s)]+\tfrac{1}{r}V_{ig}^{FB}\geq\pi_{ig}^{BR}(s)+P_{ig}+\tfrac{1}{r}V_i^{ST}$$

for all i and s. The left-hand side of (15) is the same as (14) except that $\tau_i(s)$ is omitted, because it has already been paid. The right-hand side of (15) is the same as (14) except that $\pi_{ig}^{BR}(s)$

---

8 To administer the efficiency-wage payments discussed here, and all the other payments discussed below, we envision an honest bank running a two-stage process involving an escrow account, as described below. The up-shot of this process is that the optimal way to renege on a payment is to do nothing – to pay nothing if $t_{ig} < 0$, and to accept nothing if $t_{ig} > 0$. Paying or accepting nothing is the implicit first term on the right-hand side of (13).

The two-stage escrow process is as follows: In the first stage, all parties have an opportunity to pay into the account. If $t_{ig} < 0$ then firm i is supposed to pay ½t_{ig}½into the escrow account; if $t_{ig} > 0$ then firm i is supposed to pay nothing into the account. If the anticipated payments are all made, then the bank moves to the second stage; otherwise, the bank makes a public announcement that it is returning all the payments and does so. In the second stage, if $t_{ig} > 0$ then the bank takes t_{ig} to firm i and asks whether the firm would like to accept this payment or not. If all the firms with $t_{ig} > 0$ are willing to accept these payments, then the bank disperses these amounts; otherwise, the bank makes a public announcement that it is returning all the payments to the firms in the first stage and does so.

replaces $\pi_{ig}^{NE}(s)$, because now firm i is deviating from $\mathbf{d}_{ig}^{FB}(s)$ to $\mathbf{d}_{ig}^{BR}(s)$, while the other firms choose $\mathbf{d}_{-ig}^{FB}(s)$.

Finally, constraint (d), that firm i be willing to pay (or accept) its bonus, $T_{ig}(\mathbf{d}, s)$, becomes

(16) $$T_{ig}(\mathbf{d}^{FB}(s),s) + \tfrac{1}{r} V_{ig}^{FB} \geq P_{ig} + \tfrac{1}{r} V_{i}^{ST}$$

for all i and s.

Summing each of (13) through (16) over i I (and recalling that all payments must balance across the parties) yields the necessary conditions

(13') $$\left(1+\tfrac{1}{r}\right)V_g^{FB} \geq V_g^{NE} + \tfrac{1}{r} V^{ST},$$

(14') $$\tfrac{1}{r} V^{FB} + \sum_i \pi_{ig}^{FB}(s) \geq \sum_i \pi_{ig}^{NE}(s) + \tfrac{1}{r} V^{ST} \text{ for all s,}$$

(15') $$\sum_i \pi_{ig}^{FB}(s) + \tfrac{1}{r} V_g^{FB} \geq \sum_i \pi_{ig}^{BR}(s) + \tfrac{1}{r} V^{ST} \text{ for all s, and}$$

(16') $$\tfrac{1}{r} V_g^{FB} \geq \tfrac{1}{r} V^{ST}.$$

Brief inspection shows that (13'), (14'), and (16') are trivially satisfied. For example, (13') holds because $V^{FB} > V^{ST} \geq V_g^{NE}$. But (15') is the key to our analysis: in the Appendix, we show that if (15') holds then there exist payments t_{ig} $\tau_{ig}(s)$, and $T_{ig}(\mathbf{d}, s)$ that satisfy (13) through (16). That is, (15') is necessary and sufficient for a first-best relational contract to exist under governance structure g.[9]

To interpret this necessary and sufficient condition, we rewrite (15') as

(17) $$R^g \equiv \max_s \left[\Sigma_i\left(\pi_{ig}^{BR}(s) - \pi_{ig}^{FB}(s)\right)\right] \leq \tfrac{1}{r}\left(V^{FB} - V^{ST}\right).$$

[9] Since (13) must hold for every i and (14) through (16) must hold for every i and every s, there are I + 3IS reneging constraints (where S is the number of possible states). The fact that all these reneging constraints can be reduced to a single inequality was discovered by MacLeod and Malcomson (1989) and generalized by Levin (2001).

In (17), $\pi_{ig}^{BR}(s) - \pi_{ig}^{FB}(s)$ is firm i's reneging temptation in state s under governance structure g – the generalization of step three of the procedure we used in Section 3.2. Summing this reneging temptation across firms was step four in that procedure, and finding the state in which this total reneging temptation is maximized was step five. Thus, the left-hand side of (17) is the generalization of the procedure from Section 3.2: R^g is the maximal total reneging temptation under governance structure g. Since we assume that no static governance structure achieves the first-best decisions in every state, R^g is strictly positive.

Fortunately, the right-hand side of (17) also is strictly positive, again because no static governance structure achieves the first-best (i.e., $V^{FB} > V^{ST}$). Thus, whether (17) holds comes down to the discount rate. Let r_g be the discount rate at which (17) holds with equality. For $r < r_g$, the present value of the net surplus from first-best decisions (the right-hand side) exceeds the maximal total reneging temptation (the left-hand side), so a first-best relational contract is feasible under governance structure g. That is, if the parties are patient enough, then asset ownership is irrelevant: the first-best can be achieved under any governance structure, because ongoing relationships have eliminated the incompleteness of contracts. For $r > r_g$, however, there is no relational contract that can induce first-best decisions under governance structure g.

Given (17), we can now revisit and generalize the three insights from the simple analysis in Section 3.2. That is, in a relational environment, (1) different governance structures create different managerial challenges, (2) the efficient governance structure minimizes these managerial challenges (i.e., reneging temptations), and consequently (3) the efficient governance structure typically differs from that in a static environment.

First, under governance structure g, the managerial challenges in state s are to induce firms with $\pi_{ig}^{BR}(s) - \pi_{ig}^{FB}(s) > 0$ to take first-best decisions. Note that this notion of management applies between as well as within firms. This usage is completely consistent with what we heard from practitioners, who frequently described strategic alliances as "lots of work."

Second, for a given value of r, if governance structure g satisfies (17) but governance structure g' does not, then the first-best can be achieved under g but not under g', so g' is inefficient. We therefore define the efficient governance structure to be the one that can achieve the first-best at the highest possible discount rate. Since the right-hand side of (17) is independent of the governance structure, maximizing r_g amounts to minimizing the left-hand side. That is, the efficient governance structure solves

$$\min_{g \in G}\left\{ \max_{s \in S} \Sigma_i \left(\pi_{ig}^{BR}(s) - \pi_{ig}^{FB}(s) \right) \right\} . \tag{18}$$

Third, the efficient governance structure in a relational environment is determined by very different considerations than in a static environment. As in (12), the latter solves

$$\max_{g \in G}\left\{ E_s \left(\Sigma_i \pi_{ig}^{NE}(s) \right) \right\} \tag{19}$$

Clearly, (18) and (19) will typically have different solutions. For example, while (19) depends on Nash equilibrium decisions, (18) depends on first-best decisions and on best-responses to first-best decisions. Also, while (19) is an expectation that depends on the probability distribution across states, (18) depends on only one state (the one where the reneging temptation is the largest) and is independent of the probability distribution.[10]

In sum, the evidence in Section 3.1 suggested that strategic alliances may be rooted in a dense network of relationships, and our discussions with practitioners indicated that such relationships are important in determining the performance of the alliances. We therefore developed a model of two ways that ongoing relationships matter in strategic alliances. First, for any given governance structure, relational contracts can improve upon the static outcome, achieving the first-best if the discount rate is sufficiently small. Second, and more importantly,

[10] We have restricted attention to first-best relational contracts. If the discount rate is too high to achieve the first-best under any governance structure, it would be natural to consider second-best relational contracts, as an improvement over the efficient static governance structure. It is easy to construct examples where the second-best relational contract involves departing from the first-best in one high-temptation, low-probability state. In such an example, the second-best relational contract would not be independent of the probability distribution across states, but could easily be independent of most of these probabilities, whereas the (19) would continue to involve the entire distribution.

the efficient governance structure in a relational environment differs from that in a static environment, because asset ownership affects reneging temptations in the ongoing relationship.

4. Joint Control

The models in Sections 2 and 3 can be used to analyze a wide variety of governance structures. The simple model in Section 2 allowed us to analyze alliances, acquisitions, and divestitures, and the general model in Section 3 could be used to analyze many more asset-ownership patterns. But there is one familiar governance structure that even our general model cannot analyze: a joint venture (JV).

In this section, we explain why we treat JVs (and other governance structures involving joint ownership) as qualitatively different from the unique-control governance structures analyzed above. In Section 4.1, we summarize the GHM approach to modeling joint ownership, and we explain why we cannot graft this approach onto our theory in Sections 2 and 3. In Section 4.2, of necessity, we develop a new approach to joint ownership that is consistent with our theory. Finally, in Section 4.3, we relate our approach to the literature on property rights and rent seeking, exploring its implications not only for joint ownership between firms (such as JVs), but also for "joint ownership" within firms (such as "matrix" organizations).

4.1 The GHM approach to joint ownership

Suppose there are two parties and two assets. Then there are three kinds of governance structures: integration (where one party owns both assets), non-integration (where each party owns one asset), and joint ownership. Because GHM-style models assume that asset utilization is contractible ex post, the parties can bargain over asset utilization ex post. Such bargaining occurs under non-integration and under joint ownership, but not under integration.

Under joint ownership, neither party can use either asset without the other's consent, so the parties' threat points in ex post bargaining are independent of any asset-specific investments that the parties might have made. For concreteness, suppose the bargaining yields a 50-50 split of the joint surplus (over and above these threat points). Then both parties have half-strength incentives to make asset-specific investments.

Under non-integration, the parties again bargain over asset utilization ex post, but the threat points are different: each party can use its asset without the other's consent, so a party's threat point now depends on investments specific to its asset. Now a 50-50 split of the joint surplus (over and above these new threat points) creates stronger investment incentives than under joint ownership, because each party now considers how investments affect its threat point, as well as how investments affect its bargained share above the threat point.

Because incentives are stronger (but not too strong) under non-integration, the GHM approach predicts that joint ownership is sub-optimal (Hart and Moore, 1990). But joint ventures clearly do exist in the world, so there has been some interest in modifications of the GHM approach that accommodate joint ownership (e.g., Hart, 1995; de Meza and Lockwood, 1998; Rajan and Zingales, 1998, and Halonen, 2002). But all of these modifications preserve what is, for our purposes, the key assumption in the GHM approach to joint ownership: the only difference between joint ownership and non-integration is in the threat points for ex post bargaining.

This difference in threat points seems real enough: when party 1 owns asset 1, party 1 has more control over that asset than if both parties jointly own that asset. In our view, however, there are other important differences between joint ownership and non-integration. We turn next to our new approach to joint ownership, which surfaces some of these differences.

4.2 A new approach to joint ownership

As we have emphasized throughout this paper, our theory assumes that asset utilization is not contractible ex post. One implication of this assumption is that parties cannot bargain over asset utilization ex post. Instead, in unique-control governance structures, asset owners simply take whatever utilization decision they desire (be it their short-run optimum in the static case, or perhaps the first-best decision in the relational case). But what happens under joint control?

A full answer to this question will require a paper of its own. In this section, therefore, we develop only a reduced-form model, but we emphasize two key questions raised by a model of joint ownership without ex post contractibility: how owners decide on the utilization of jointly owned assets, and how they divide the non-contractible payoffs accruing to jointly owned assets). We resolve both the deciding and dividing issues by analogy to the theory of rent-seeking, as follows.

Figure 6

Joint Venture

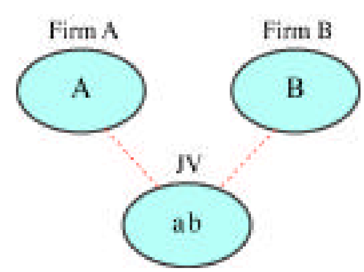

Joint ownership of {a.b} involves contested control, and perhaps contested consumption

Imagine a joint venture in the simple static model from Section 2.2: Firm A owns {A} and Firm B owns {B}, but a new entity is created that owns assets {a, b}. Unlike a mutual divestiture (Figure 4d), however, this new entity is jointly owned by Firms A and B, as shown in Figure 6. This implies that the two parent firms must decide, together, what actions the joint venture should take. We assume that control is contested ex post; this approach is guided in part by our discussions with practitioners, who told us that joint ventures, in comparison with other forms of strategic alliance, consumed significant amount of management time and energy.[11] More specifically, after observing the state s_{ij}, the parties simultaneously choose how much to expend in a battle for control. If Firm A expends k_A and Firm B k_B, then

[11] We emphasize that the contest is not over whether the assets are jointly utilized (since this would violate our maintained non-contractibility assumption) but rather over which party gets to make the utilization decision.

Firm A gains control over the jointly owned assets with probability $q(k_A, k_B)$ and Firm B gains control with probability $1 - q(k_A, k_B)$. Once a party has gained control, that party decides the utilization of the jointly owned assets {a, b}, but this control lasts only for the current period.

After decision comes division (of the non-contractible payoffs accruing to the jointly owned assets). Two polar approaches to the division issue are: (1) gaining control conveys not only decision rights but also division rights (i.e., the firm that gains control also receives all the payoffs, as in an acquisition, but again for one period only); (2) gaining control over decision rights has no impact on the subsequent battle for division rights (i.e., after the contest for control, there is an independent contest for consumption). Our guess is that reality typically lies between these two extremes.

Given some resolution of the division issue, one could then compute the equilibrium expenditures $k_A^*(s_{ij})$ and $k_B^*(s_{ij})$, and hence the state-dependent probability that Firm A wins the contest for control, $q(k_A^*(s_{ij}), k_B^*(s_{ij}))$. It would be straightforward to flesh out the model so that the parties' equilibrium expenditures in the contest for control are zero in states where the parties agree on asset utilization (such as s_{HH}, s_{MM}, s_{ML}, s_{LM}, and s_{LL}), but in other states are increasing in the extent to which parties disagree about asset utilization (such as s_{HM} versus s_{HL}, where the disagreement is stronger in the latter of these two states, so the equilibrium expenditures would be higher). For our purposes, however, the key point is that joint ownership creates a contest for control, and the parties' endogenous participation in this contest creates state-dependent control (not possible in the unique-control cases modeled above), at the cost of the equilibrium expenditures $k_A^*(s_{ij})$ and $k_B^*(s_{ij})$. If the benefits of state-dependent control outweigh the social costs of the expenditures, a joint venture could be the second-best static governance structure, producing higher expected total surplus than any of the static unique-control governance structures analyzed in Section 2.2.[12]

[12] Again, we assume that the opportunity to utilize the assets ex post is fleeting, so there is not time to reallocate asset ownership ex post. Much discussion with practitioners (as well as simple introspection) suggests that there are often fights between JV partners over near-term decisions, and that changing the ownership structure of the JV would not be a practical way to avoid these fights.

In a relational JV, the parties agree not to contest control or consumption, thus avoiding the social costs of these contests. Therefore, in addition to the reneging temptations over side payments and decisions in Section 3.3, new reneging temptations arise—the temptations to fight for control (and subsequently make a "best response" decision as defined in Section 3.3) and to fight for consumption. While a complete analysis of the relative efficiency of relational JV's is beyond the scope of this paper, we expect that a version of the necessary and sufficient condition (17) will continue to hold, where $\pi_{ig}^{BR}(s)$ is modified to include the social costs and private benefits of contest expenditures.

4.3 Joint ownership, property rights, and rent seeking

At a theoretical level, the GHM approach assumes that asset utilization is contractible ex post, whereas our approach assumes the opposite. Our assumption surfaced two issues that the GHM approach resolves through bargaining ex post: deciding (on asset utilization) and dividing (the non-contractible payoffs from asset utilization). We resolved these issue by modeling contests for control and for consumption. We now interpret these contests in terms of the large literature on rent-seeking, and explore their implications for joint ventures, matrix organizations, and other forms of "joint ownership."

Our interpretation of the contests for control and for consumption begins with Coase (1960). For most purposes, the key lesson from this classic paper is an irrelevance result: in a world with no transaction costs, efficient outcomes will arise as long as property rights are fully specified, regardless of which parties own which rights. For our purposes, however, the key idea is the (usually unstated) alternative to Coase's Theorem: if property rights are not fully specified then inefficiency is likely to result. In our approach, joint ownership is an incomplete specification of property rights, and inefficiency results.

The inefficiency in our model arises from the parties' expenditures in the contest for control (and the contest for consumption, if it occurs). These expenditures are similar to the rent-seeking behaviors analyzed in the large literature beginning with Buchanan and Tullock

(1962). More recently (and more formally), Skaperdas (1992) has explored a contest model of rent-seeking between individuals (or individual firms) in a world with no property rights, and Rajan and Zingales (2000) have explored a similar contest model of rent-seeking between individuals (or divisions) within a firm. In both settings, incomplete property rights cause the Coase Theorem not to apply.

Our strong impression is that contests for control are an important issue for JVs, and would expect the same in any governance structure involving joint ownership. Furthermore, we would expect the same inside firms, where property rights are incompletely specified. For example, we believe that contests for control are an important issue for matrix and related organizations, where a worker has more than one boss, but it is not (and cannot be) clearly specified which boss controls which aspects of the worker's activities.

5. Contracting for Control

In the simple static model in Section 2.2, what we called an alliance is really just non-integration: Firm A owns assets {A, a} and Firm B {B, b}, and the firms' asset-utilization decisions are not influenced by any contract (formal or informal). There may be some alliances as simple as this, but we know of none.

The relational models in Section 3 breathe some life into an alliance, by adding relational contracting to the asset-ownership considerations of the static model. As we have emphasized, a wide range of practitioner accounts, case studies, and empirical projects have convinced us that ongoing relationships are an important element of many alliances. But most alliances we know do not rely solely on such informal contracting; instead, the parties also utilize a formal (i.e., court-enforceable) contract. In the words of one practitioner, "In our best alliances, we put the formal contract in a drawer and build a relationship on top."

In this section, we explore such combinations of formal and informal contracting. We begin by following Maskin and Tirole (1999), who observed that certain static models of asset ownership could be reinterpreted as static models of formal contracting. This observation can be applied directly to our unique-control models in Sections 2 and 3, but needs some care when applied to our joint-control model in Section 4. In both cases (unique- and joint-control), we discuss the empirical implications of this formal-contracting reinterpretation of our asset-ownership models.

In the simple static model in Section 2.2, {a, b} could represent decision rights allocated by contract, rather than assets allocated by ownership.[13] For example, in an alliance between a small biotech company and a large pharmaceutical firm, {a} might be the right to manage clinical trials and {b} the right to market the final product. Similarly, in the general model in Section 3.3, the J assets could be reinterpreted as J decision rights, again allocated by contract, with no assets involved whatsoever. These formal-contract reinterpretations of our unique-control static models seem close to the theoretical foundation implicitly assumed in some empirical work on alliance contracts, such as Lerner and Merges (1998), Elfenbein and Lerner (2002), and Ryall and Sampson (2002); see Section 6 for more discussion.

Contractible and non-contractible decision rights can, of course, co-exist. For example, the set of decision rights {1, …, J} could consist of rights {1, …, J'} that are contractible, and are not associated with any assets, and rights {J'+1, …, J} that are non-contractible, and so must be allocated by asset ownership. Such a bifurcation of the set of decision rights should be straightforward to handle in empirical work, as long as the same care is paid to recording and analyzing both asset ownership and contract terms. For example, in an alliance between Firms A and B (where the firms own their core assets {A} and {B}, respectively), suppose that decision right {a} is contractible but decision right {b} is non-contractible. Then an empirical analysis may give badly distorted results if it focuses solely on who controls {a} by contract, ignoring who controls {b} by ownership.

---

[13] We emphasize again the difference between contractible *decision rights* (who makes the decision) and contractible *decisions* (what decisions get made). We have the former; GHM models have the latter (ex post).

Adding relational contracts to formal-contracting models of unique control is straightforward, given the analysis in Section 3.3. The function $G_g: J \rightarrow I$ still describes which parties control which decision rights under governance structure g, regardless of whether these decision rights are contractible, non-contractible, or flexible, and regardless of whether these rights are allocated by asset ownership or by contract. Given G_g, the relational-contracting analysis proceeds as before, culminating in the necessary and sufficient condition (17).

Allowing for joint control in formal-contracting models is also straightforward, at least for the theory, but may have noteworthy implications for empirical work. Recall that joint ownership of an asset means that the right to determine asset utilization is incompletely specified. Similarly, joint control of a decision right means that none of the parties involved has a well-specified right to make the decision. Empirical analyses of decision rights in alliance contracts, notably Lerner and Merges (1998), reveal that, while many rights are uniquely allocated to one of the contracting parties, other rights are unspecified. Our assumption is that the alliance partners jointly control all project-related rights that are not explicitly specified in the contract, and we envision contests for control and consumption analogous to those analyzed in Section 4. As in Section 4, the potential advantage of leaving rights unspecified is that they can facilitate state-dependent control; these benefits must be weighed against the social cost of the contest expenditures.

6. Towards Testing

Neither the “simple” or “general” models developed in Sections 2 and 3 is ideally suited to empirical testing: our simple model is too specific, making predictions that depend crucially on the simple structure and assumed parameter values, while our general model is too general, offering few testable predictions in its current form. We believe that empirical tests of a model like ours will require placing the theory in a specific context, where the structure and size of the payoffs can be ascertained with more precision. While the detailed specification of such an empirical strategy is beyond the scope of this paper, in this section we examine the existing empirical work in this area to show how our modeling approach can help inform future empirical work. We also discuss some strategies for testing our model, and some of the difficulties and opportunities raised by these strategies.

6.1 Existing Literature

Three empirical papers on alliance form would seem to confirm the importance of relational contracting in strategic alliances. Gulati (1995a), Robinson and Stuart (2002), and Ryall and Sampson (2002) all test whether “relationships” affect the structure of strategic alliances. Gulati, who examines 2400 alliances in the biopharmaceuticals, new materials, and automotive industries, shows that the likelihood that a given alliance will involve the sharing of equity ownership between the parties is negatively related to how often these two parties have joined in an alliance in the past. He concludes from these findings that trust between the parties affects the form that the alliance will take.

Robinson and Stuart analyze the network of alliances between and among biotech and pharmaceutical firms in the period 1976-1998. Using measures from network theory, they assess the “proximity” and “centrality” of firms in the network, and ask whether these measures affect the form of contracts. Similar to Gulati, they find that firms who are better connected in the network are less likely to use equity as part of the mechanism for governing their strategic alliances. They argue that close links to other firms in the network provide an alternative to equity as a control mechanism, because “embedding individual alliances in a network of past alliances gives actors an opportunity to levy long-term reputational penalties against their counterparties” (p. 31).

Ryall and Sampson, using a case-study methodology, look in detail at alliance contracts to show that the form of the contracts varies depending on whether the parties engage in repeat interactions. They argue that informal interactions can serve as a substitute for contractual safeguards.

All of these papers confirm the main result from our general model: relationships matter. When firms are able to engage in relational contracting, they choose a different alliance form than when they engage in one-shot dealings. More specifically, our model predicts that as the discount rate r changes, the optimal form of alliance will change. We interpret the key right-hand-side variables in each of these studies as measuring changes in r. Taken literally, r reflects the time value of money. In many repeated-game models, however, r is also understood to reflect the probability that the relationship will be (exogenously) terminated before the next period begins; the present value of a payoff stream under the literal interpretation of r becomes the expected present value when termination probabilities are considered. Thus, the repeated interactions measured by Gulati and Ryall-Sampson can be seen as increasing the likelihood of a future together, and thus as decreasing r.

One can go further in reinterpreting r, by analyzing information flows in social structures (Kandori, 1992; Gulati 1995b). Consider a model in which players from a large population are matched in pairs each period. If they are randomly matched to anyone in the network, and players do not communicate with each other, then the likelihood that people will have to live with the consequences of their past actions is small. But in an environment in which players are more likely to see each other again, or in which parties can share their experiences with other players, the effective discount rate is lower. We interpret Robinson-Stuart as showing that when parties are more densely connected in the network, their effective r is lower, and they choose a different form of strategic alliance.

While we take Gulati, Ryall-Sampson, and Robinson-Stuart as confirmatory of the broad predictions of our theory of relational contracting, we readily acknowledge that they are far from a stringent test of our particular model. Indeed, our model highlights many of the difficulties of testing theories about alliance form and relationships.

Furthermore, as Whinston (2002) points out, theories suggest to empiricists what they should look at, and new models often highlight problems with both dependent and independent variables in empirical work done prior to a new theory. Several other papers in the literature join Gulati and Robinson-Stuart by using equity participation as their dependent variable. Pisano (1989), and Oxley (1997) examine characteristics of the joint project's technology and market to predict whether the parents will make an equity investment in the alliance. But as several of these papers acknowledge, the real dependent variable of interest is control over decisions about the project. Whether and how equity participation confers control to the parties is unmodeled. As Section 4 above highlights, there are complex questions about how control is exercised under joint ownership, and several key parameters of the alliance (such as the distinction between voting control and economic interests) are unmeasured in these studies. Furthermore, as noted by Robinson-Stuart, equity ownership could be playing a number of other roles in the alliance: the exploitation of perceived under-valuation, a demonstration of trust, or mutual hostage-taking.

As an alternative to equity participation as a dependent variable, several other papers in the literature do careful analyses of the contracts themselves to determine control rights. Ryall and Sampson examine a small number of technology alliance contracts, and envision counting contract clauses in a wider sample. Lerner and Merges (1998) and Elfenbein and Lerner (2002) both examine which parties hold a variety of decision rights. Lerner-Merges look at how technological and market factors affect the allocation of control rights in a sample of R&D alliances, while Elfenbein-Lerner (using data from internet portal alliances) test predictions of several model in the property-rights literature, showing how differences in the importance of parties' actions, and their bargaining power, affect the allocation of twelve specific control rights. Our model, and especially the extensions presented in Section 5 above, suggest that this sort of careful accounting of decision rights will be necessary to understand fully the nature of relational contracting in strategic alliances.

6.2 Other Possible Tests of our Model

As discussed above, we interpret several of the papers in this literature as testing the proposition that relationships matter to the form of strategic alliances. In this subsection, we suggest that our modeling approach has some additional implications about what sorts of factors should affect the type of strategic alliances observed in practice. Specifically, our focus on the role of spillovers suggests that an important determinant of alliance form is whether the alliance is in a business similar to that of the parents. When the alliance is in a closely related business, the likelihood that there will be large spillovers onto the parents is greater. According to our model, this should affect the form of the strategic alliance. Again, the specific effect of spillover size on alliance form may vary across contexts, but we give one possible implementation of this idea below.

Similarly, our focus on contracting problems ex post, as opposed to issues of asset specificity and hold-up, suggests that factors like the difficulty of writing contracts that bind parties to their promises in all foreseeable states should affect how alliances are structured. Tests based on contract completeness are difficult, however, and require finding good proxies for exogenous changes in contracting costs that are not correlated with other characteristics of either the market or the technology. In addition, as emphasized by Whinston (2002) and Baker and Hubbard (2001), predictions made by models of this type depend crucially on details about the specific context: the incentives of the parties, the states of nature that are likely to occur, how their decisions interact, and the costs and benefits that they face.

We believe that tests of our theory of relational contracting in strategic alliances will require detailed modeling of a specific context. In this spirit, we conclude this section by sketching out how one might test the most novel prediction of our model: that the managerial challenges posed by one-shot static alliances differ from those of relational alliances. Recall that structuring one-shot alliances requires careful attention to the *expected* payoffs to all of the parties. In contrast, managing relational alliances requires that attention be paid to the possibility of extreme payoffs to a single party, since these extremes are what determine the parties' reneging temptations in particular states.

This distinction, between focusing on extremes for one party rather than on expectations across all parties, suggests possible empirical tests of our theory. Consider a specific context (such as an alliance between a small biotech firm and a large pharmaceuticals company) in which the two firms are trying to decide how to pursue a possibly profitable joint project. Suppose that the plausible governance structures are an acquisition of the assets required for the project by the pharmaceutical firm (an "acquisition" in our simple model above), or an agreement to cooperate, with each party retaining ownership of their relevant assets (an "alliance" in our simple model).

In a one-shot transaction, the trade-off between an acquisition and an alliance comes down to estimating which governance structure produces the highest expected value. The problem with the one-shot alliance is under-exploitation of the opportunity, since either party can effectively veto the project, and will do so without regard to positive spillovers to the other party. The problem with a one-shot acquisition is that the acquirer (the pharmaceutical firm) will not consider the negative spillovers on the biotech firm in states in which its own payoff is trivially positive, and will similarly ignore positive spillovers on the biotech firm in states in which its own payoff is trivially negative. The trade-off between an alliance and an acquisition in a one-shot transaction will thus depend on which states are more likely, and the relative costs and benefits across these states for each firm.

According to our model, the trade-off between an acquisition and an alliance in a repeated interaction is quite different. Here, the challenge is to get the relevant party to implement the project in states where his own payoff tells him not to, or to veto the project when he would like to proceed but the negative spillovers on the other party are large. Note that this trade-off involves payoffs to single parties in specific states, rather than the expected payoff to both parties across all states.

In the simple context just described, our model can make some surprisingly direct predictions. Imagine that the likelihood of a state with a small negative payoff to the pharmaceutical firm and a large positive payoff to the biotech firm is negligible. Now the problem with a static acquisition is that it always over-implements the project, while the static alliance under-implements. Thus, all else equal, more profitable projects will tend to favor acquisitions in the one-shot transactions. But project profitability has exactly the opposite effect on the trade-off in the repeated interaction. A more profitable project makes the temptation for the pharmaceutical firm to over-implement greater, and reduces the temptation for the alliance to under-implement; this implies that relational alliances are more efficient for more profitable projects. This situation thus provides a clear prediction about the differences between static and relational interactions: more profitable projects are more likely to be structured as acquisitions when the interaction between the parties is one-shot, but more likely to be structured as alliances in ongoing relationships.

We expect that successful empirical tests of our model (and models like ours) will require substantial knowledge of the context in which the data are gathered, and detailed information on the parties, their motivations, costs and benefits. As noted above, such empirical work is already beginning to be done, and we look forward to its continuation.

7. Conclusion

A strategic alliance is a governance structure for coordinating activities among firms. Such governance structures are observed in many forms, including alliances, acquisitions, mutual divestitures, and joint ventures. In an effort to understand this plethora of governance structures, we conducted detailed conversations with practitioners who design, implement, consult to, and negotiate terms for such alliances. These practitioners emphasized three issues as crucial to the design and performance of strategic alliances: spillovers, ex post contracting problems, and relationships.

In this paper we developed a simple model that allows us to characterize and contrast a variety of governance structures—including alliances, acquisitions, and mutual divestitures—

where each asset is owned (or each decision right is controlled) by exactly one firm, and one governance structure—a joint venture—where assets are jointly owned (or decision rights jointly controlled) by multiple firms. We have identified the inefficiencies associated with each of these governance structures in static environments. More importantly, we have identified the managerial challenges and the efficiency consequences of these governance structures in relational environments. We have established two important ways in which relationships are important: relationships improve the performance of any given governance structure, and the efficient governance structure under relational governance differs from the efficient structure under static governance.

We believe that our approach suggests a rich and tractable framework that could be used to analyze a wide variety of governance structures. In addition, we have analyzed only first-best relational contracts, but one could also analyze the second-best governance structures that are efficient at intermediate interest rates (i.e., where r is too high to allow any governance structure to achieve the first-best, but not so high as to render infeasible all relational contracts beyond the trivial spot contract). Perhaps more importantly, we have suggested ways to operationalize our distinction between static and relational governance (i.e., empirical proxies for the discount rate, the importance of spillovers, and ex post contracting problems), and we have provided several predictions potentially refutable by existing or future data.

References

Aghion, P., Dewatripont, M., Rey, P., 2002. Partial contracting, control allocation, and cooperation. Unpublished working paper. Harvard University.

Baker, G. P., 1992. Incentive contracts and performance measurement. Journal of Political Economy 100, 598-614.

Baker, G. P., Gibbons, R., Murphy, K. J., 1994. Subjective performance measures in optimal incentive contracts. Quarterly Journal of Economics 109, 1125-1156.

Baker, G. P., Gibbons, R., Murphy, K. J., 1999. Informal authority in organizations. Journal of Law, Economics, and Organization 15: 56-73.

Baker, G. P., Gibbons, R., Murphy, K. J., 2002a. The relational-adaptation theory of the firm. Unpublished working paper. MIT.

Baker, G. P., Gibbons, R., Murphy, K. J., 2002b. Relational contracts and the theory of the firm. Quarterly Journal of Economics 117, 39-83.

Baker, G. P., Hubbard, T. N., 2001. Empirical strategies in contract economics: information and the boundary of the firm. American Economics Review Papers and Proceedings 91, 189-194

Buchanan, J., Tullock, G., 1962. The Calculus of Consent. University of Michigan Press, Ann Arbor.

Bull, C., 1987. The existence of self-enforcing implicit contracts. Quarterly Journal of Economics 102, 147-59.

Che, Y., Yoo, S., 2001. Optimal incentives for teams. American Economic Review 91, 525-541.

Coase, R. 1960. The problem of social cost. Journal of Law and Economics 3, 1-44.

DeMeza, D., Lookwood, B., 1998. Does asset ownership always motivate managers? Outside options and the property rights theory of the firm. Quarterly Journal of Economics 113, 361-386.

Elfenbein, D., Lerner, J., 2002. Ownership and control rights in internet portal alliances, 1995-1999. Unpublished working paper. Harvard Business School.

Garvey, G. 1995. Why reputation favors joint ventures over vertical and horizontal integration: a simple model. Journal of Economic Behavior and Organization 28, 387-397.

Grossman, S. J., Hart, O. D., 1986. The costs and benefits of ownership: a theory of vertical and lateral ownership. Journal of Political Economy 94, 691-719.

Gulati, R., 1995a. Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances. Academy of Management Journal 38, 85-112.

Gulati, R., 1995b. Social structure and alliance formation patterns: a longitudinal analysis. Administrative Science Quarterly 40, 619-652.

Halonen, M., 2002. Reputation and allocation of ownership. Unpublished working paper. Helsinki School of Economics.

Hart, O., 1995. Firms, Contracts, and Financial Structure. Clarendon Press, Oxford.

Hart, O., Holmstrom, B., 2002. Vision and firm scope. Unpublished working paper. Harvard University.

Hart, O., Moore, J., 1990. Property rights and the nature of the firm. Journal of Political Economy 98, 1119-1158.

Holmstrom, B., 1982. Moral hazard in teams. Bell Journal of Economics 13, 324-340.

Holmstrom, B., Milgrom, P., 1991. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design. Journal of Law, Economics, and Organization 7, 24-52.

Kandori, M., 1992. Social norms and community enforcement. Review of Economic Studies 59, 63-80.

Klein, B., Leffler, K., 1981. The role of market forces in assuring contractual performance. Journal of Political Economy 89, 615-641.

Lerner, J., Merges, R., 1998. The control of technology alliances: an empirical analysis of the biotechnology industry. Journal of Industrial Economics 46, 125-56.

Levin, J., 2001. Relational incentive contracts. Unpublished working paper. Stanford University.

Macaulay, S., 1963 Noncontractual relations in business: a preliminary study, American Sociological Review 28, 55-67.

MacLeod, W. B., Malcomson, J., 1989. Implicit contracts, incentive compatibility, and involuntary unemployment. Econometrica 57, 447-480.

Macneil, I., 1978. Contracts: adjustments of long-term economic relations under classical, neoclassical, and relational contract law. Northwestern University Law Review 72, 854-906.

Maskin, E., Tirole, J., 1999. Two remarks on property rights. Review of Economic Studies 66, 139-149.

McQuade, K., Gomes-Casseres, B., 1992. Xerox and Fuji Xerox. Harvard Business School Case #9-391-156.

Oxley, J., 1997. Appropriability hazards and governance in strategic alliances: a transactions cost approach. Journal of Law, Economics and Organization 5.

Pisano, G., 1989. Using equity participation to support exchange: evidence form the biotecholonogy industry. Journal of Law, Economics and Organization 5, 109-126.

Rajan, R., Zingales, L., 1998. Power in a theory of the firm. Quarterly Journal of Economics 113, 387-432.

Rajan, R., Zingales, L., 2000. The tyranny of inequality. Journal of Public Economics 76, 521-558.

Rayo, L., 2002. Relational team incentives and ownership. Unpublished working paper. Stanford University

Rey, P., Tirole, J., 2001. Divergence of objectives and the governance of joint ventures, Unpublished working paper. IDEI.

Robinson, D. T., Stuart, T. E., 2002. Just how incomplete are incomplete contracts? Evidence rom biotech strategic alliances. Unpublished working paper. Columbia University.

Ryall, M., Sampson, R., 2002. The effects of repeated interaction on the organization and performance of R&D alliances. Unpublished working paper. New York University.

Simon, H., 1951. A formal theory of the employment relationship. Econometrica 19, 293-305.

Skaperdas, S., 1992. Cooperation, conflict, and power in the absence of property rights. American Economic Review 82, 720-739.

Telser, L., 1981. A theory of self-enforcing agreements. Journal of Business 53, 27-44.

Whinston, M. D., 2002. On the transaction cost determinants of vertical integration. Unpublished working paper. Northwestern University.

Williamson, O., 1975. Markets and Hierarchies: Analysis and Antitrust Implications. Free Press, New York.

Appendix

THEOREM. Suppose that (14) holds for all states s for a given governance structure g.[14] Then, there exists payments t_i $\tau_i(s)$, $T_i(\mathbf{d}, s)$, and P_i that satisfy constraints (10) through (13).

PROOF.

1. Let P_i be any market-clearing price for asset reallocation following reneging, satisfying $P_i + \frac{1}{r}V_i^{ST} \geq \frac{1}{r}V_i^{NE}$ for all i and $\Sigma_i P_i = 0$. For example, let $P_i = \frac{1}{r}\left(V_i^{NE} - V_i^{ST}\right) + \frac{1}{rI}\left(V^{ST} - V^{NE}\right)$.

2. Set $T_i(\mathbf{d}^{FB}, s) = \pi_i^{BR}(s) - \pi_i^{FB}(s) + P_i + \frac{1}{r}V_i^{ST} - \frac{1}{r}V_i^{FB} + \varepsilon(s)$, where
$$\varepsilon(s) = -\frac{1}{I}\left(\sum_i \left(\pi_i^{BR}(s) - \pi_i^{FB}(s)\right) - \frac{1}{r}\left(V^{FB} - V^{ST}\right)\right)$$

3. Note that $\Sigma_i T_i(\mathbf{d}^{FB}, s) = 0$ (since $\Sigma_i P_i = 0$). Also, given (14), note that $\varepsilon(s) \geq 0$.

4. Substituting $T_i(\mathbf{d}^{FB}, s)$ into (13) yields $\pi_i^{BR}(s) - \pi_i^{FB}(s) + \varepsilon(s) \geq 0$. Since $\pi_i^{BR}(s) - \pi_i^{FB}(s) \geq 0$ and $\varepsilon(s) \geq 0$, (13) is satisfied.

5. Substituting into (12) yields $\varepsilon(s) \geq 0$, so (12) is satisfied.

6. Next, set $\tau_i(s) = \pi_i^{NE}(s) - \pi_i^{BR}(s) + \lambda(s)$, where $\lambda(s) = -\frac{1}{I}\left(\sum_i \left(\pi_i^{NE}(s) - \pi_i^{BR}(s)\right)\right)$

7. Note that $\Sigma_i \tau_i(s) = 0$. Also, since $\Sigma_i \pi_i^{NE}(s) \leq \Sigma_i \pi_i^{FB}(s) \leq \Sigma_i \pi_i^{BR}(s)$, note that $\lambda(s) \geq 0$.

8. Substituting $\tau_i(s)$ and $T_i(\mathbf{d}^{FB}, s)$ into (11) yields $\varepsilon(s) + \lambda(s) \geq 0$, so (11) is satisfied.

9. Substituting $\tau_i(s)$ and $T_i(\mathbf{d}^{FB}, s)$ into $V_i^{FB} \equiv E_s\left[t_i + \tau_i(s) + T_i(\mathbf{d}, s) + \pi_i^{FB}(s)\right]$ yields $V_i^{FB} \equiv E_s\left[t_i + \pi_i^{NE}(s) + P_i + \frac{1}{r}V_i^{ST} - \frac{1}{r}V_i^{FB} + \varepsilon(s) + \lambda(s)\right]$, or $\left(1 + \frac{1}{r}\right)V_i^{FB} \equiv t_i + V_i^{NE}(s) + P_i + \frac{1}{r}V_i^{ST} + E_s[\varepsilon(s) + \lambda(s)]$.

10. Since $E_s[\varepsilon(s) + \lambda(s)] \geq 0$, (10) is satisfied at $t_i = 0$ for all i.

[14] The "g" subscript is suppressed throughout this Appendix.

Figure 3

Payouts from Joint Utilization of Assets {a, b} for Mergers, Markets, Acquisitions, and Mutual Divestitures

Panel A. Merger of A and B
Utilization maximizes Total Surplus

$\pi_A(s)$ \ $\pi_B(s)$:	H	M	L
H	A+B: 2(H+π)	A+B: H+M+2π	A+B: H+L+2π
M	A+B: H+M+2π	A+B: 2(M+π)	A+B: M+L+2π
L	A+B: H+L+2π	A+B: M+L+2π	A+B: 2(L+π)

Panel B. Arm's Length Alliance
A controls {a}, B controls {b}

$\pi_A(s)$ \ $\pi_B(s)$:	H	M	L
H	A: H+π B: H+π	A: H+π B: M+π	A: H+π B: L+π
M	A: M+π B: H+π	A: M+π B: M+π	A: M+π B: L+π
L	A: L+π B: H+π	A: L+π B: M+π	A: L+π B: L+π

Panel C. Acquisition (A buys)
A controls {a, b}

$\pi_A(s)$ \ $\pi_B(s)$:	H	M	L
H	A: H+2π B: H	A: H+2π B: M	A: H+2π B: L
M	A: M+2π B: H	A: M+2π B: M	A: M+2π B: L
L	A: L+2π B: H	A: L+2π B: M	A: L+2π B: L

Panel D. Mutual Divestiture
C controls {a, b}

$\pi_A(s)$ \ $\pi_B(s)$:	H	M	L
H	A: H B: H C: 2π	A: H B: M C: 2π	A: H B: L C: 2π
M	A: M B: H C: 2π	A: M B: M C: 2π	A: M B: L C: 2π
L	A: L B: H C: 2π	A: L B: M C: 2π	A: L B: L C: 2π

Note: Shaded region denotes the control-specific utilization decision for {a, b}. The shaded decisions are illustrated assuming that $0 < 2\pi < -M < H + 2\pi < -L$.

Table 2

Repeated Strategic-Alliance Transactions Between Unique Pairs of Organizations, 1973-2001

Number of Transactions Between Unique Partner-Pairs	Number of Transactions	Total Number of Alliances	% of Total Alliances
1	9,462	9,462	76.0%
2	805	1,610	12.9%
3	182	546	4.4%
4	60	240	1.9%
5 or More	57	360	2.9%
Alliances between organizations ultimately merged or combined		912	7.3%

Note: Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, based on publicly disclosed contracts and arrangements from 1973-2001. Alliances are assigned to the surviving parent, regardless of whether the parent was involved in the original arrangement. Totals sum to more than 100% because some alliances have more than two partners.

Table 3

Pharmaceutical and Biotech Firms Most Active in Strategic Alliances, 1973-2001

Panel A *Top 12 Pharmaceutical Firms*	Number of Alliances	Number of Partners	Pharma Partners	Biotech Partners	Partners in Top 24
1. GlaxoSmithKline (GSK)	373	248	11.7%	58.5%	20
2. Pharmacia (PHA)	370	271	12.2%	44.1%	21
3. Pfizer (PFE)	287	194	14.4%	57.7%	19
4. Novartis (NVS)	230	167	16.2%	54.5%	18
5. Elan (ELN)	228	153	22.2%	38.6%	14
6. Hoffmann-La Roche (HLR)[a]	224	164	11.7%	62.0%	17
7. Johnson & Johnson (JNJ)	212	170	16.5%	37.6%	16
8. Abbott (ABT)	201	174	13.3%	49.7%	14
9. American Home Products (AHP)	175	124	21.0%	56.5%	19
10. Lilly (LLY)	164	132	13.6%	62.9%	16
11. Merck (MRK)	164	118	16.1%	58.5%	16
12. Bristol-Myers Squibb (BMY)	150	128	10.9%	57.8%	15

Panel B *Top 12 Biotech Firms*	Number of Alliances	Number of Partners	Pharma Partners	Biotech Partners	Partners in Top 24
1. Applera (ABI)	214	183	13.7%	38.3%	15
2. Chiron (CHIR)	172	136	20.0%	31.1%	12
3. Genentech (DNA)	124	92	14.1%	54.3%	14
4. Genzyme (GENZ)	122	102	14.7%	32.4%	6
5. Shire Pharmaceuticals (SHP)	119	85	24.7%	36.5%	12
6. Incyte Genomics (INCY)	107	90	25.8%	42.7%	17
7. Celltech (CLL)	106	89	25.8%	37.1%	15
8. Affymetrix (AFFX)	91	69	26.1%	30.4%	10
9. Medarex (MEDX)	88	73	16.4%	41.1%	10
10. Medimmune (MEDI)	86	67	22.4%	25.4%	10
11. Vertex (VRTX)	79	63	25.8%	32.3%	12
12. Amgen (AMGN)	78	66	21.2%	42.4%	12

Note: Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, based on approximately 12,500 publicly disclosed contracts and arrangements. Companies ranked (and "top companies" defined) by number of alliances. The number of alliances reported excludes alliances with entities that ultimately became wholly owned subsidiaries of the companies in the table. Contracts are assigned to the surviving parent, regardless of whether the parent was involved in the original arrangement.

[a]Hoffmann-La Roche is a wholly owned subsidiary of privately held Roche Holdings.

[b]Applera, formed by the combination of Applied Biosystems and Celera Genomics, trades under two tracking stocks, ABI (Applera-Applied Biosystems) and CRA (Applera-Celera Genomics).

Figure 4

Strategic Alliances Among the Top 12 Pharmaceuticals and Top 12 Biotechs

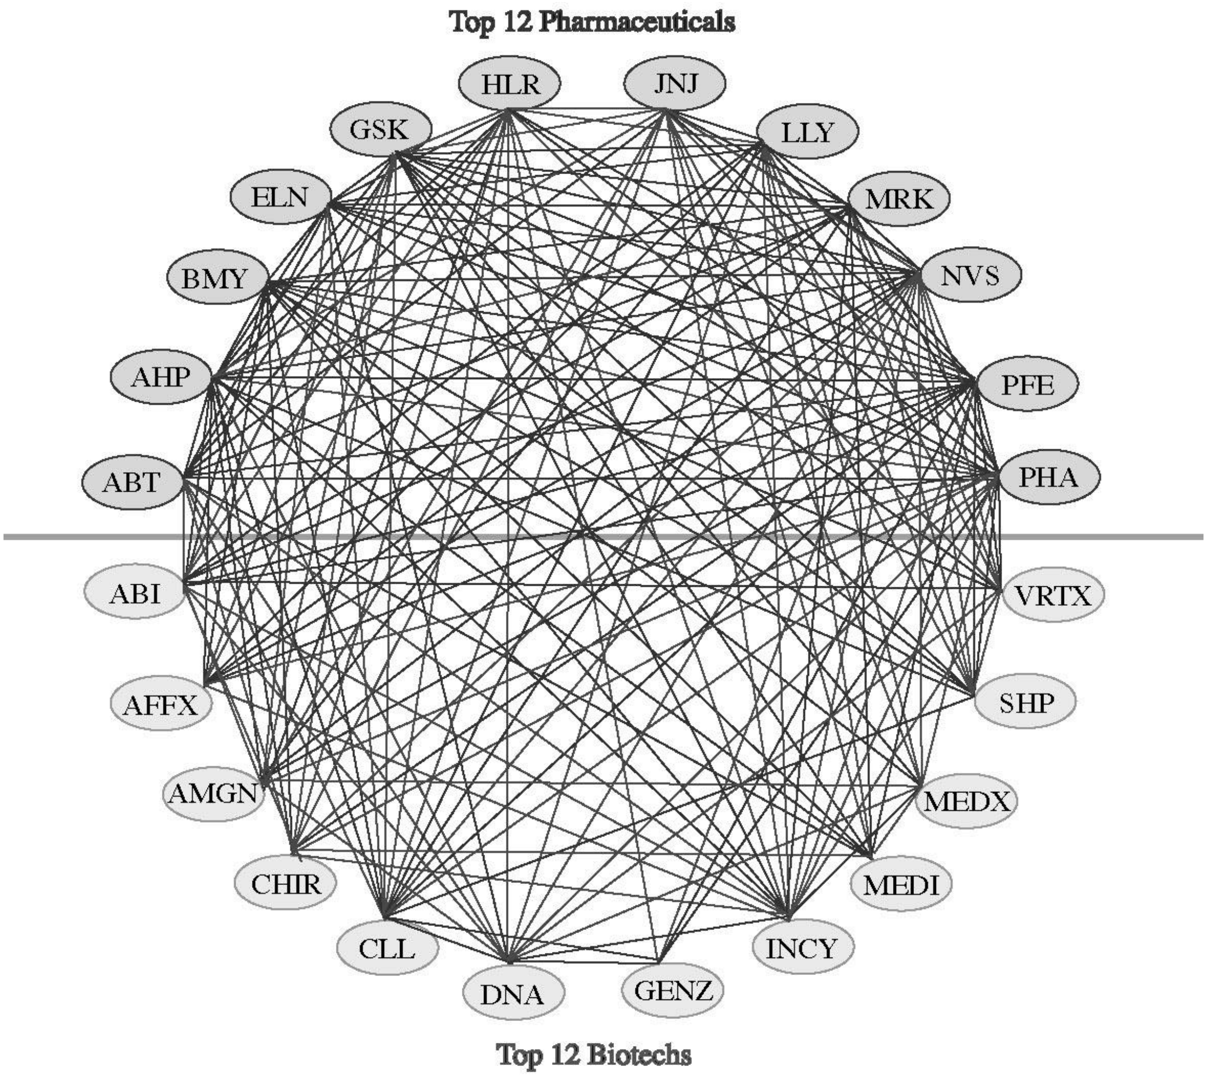

Note: Ticker symbols correspond to companies included in Table 3. Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, based on approximately 12,500 publicly disclosed contracts and arrangements from 1973-2001. Contracts are assigned to the surviving parent as of year-end 2001, regardless of whether the parent was involved in the original arrangement.