Comments Welcome

Relational Contracts in Strategic Alliances

George Baker
Harvard Business School and NBER

Robert Gibbons
MIT and NBER

Kevin J. Murphy
USC's Marshall School

February 26, 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. The economics literature lacks a framework for analyzing this plethora of governance structures. 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 specific investment or hold-up), 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.

* We are very grateful for detailed discussions with several practitioners, especially Judy Lewent and Richard Kender of Merck, Inc. and Mark Edwards of Recombinant Capital, and also for research support from Harvard Business School (Baker and Gibbons), MIT's Sloan School of Management (Gibbons), and USC's Marshall School (Murphy).

Relational Contracts in Strategic Alliances

by

George Baker, Robert Gibbons, and Kevin J. Murphy

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 specific investment and hold-up 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. With this model, we 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), joint ventures (where a new entity is created and jointly owned by the parents), mutual divestitures (where an autonomous entity is created to pursue the joint project without parental ownership or direct control), acquisitions (where one parent acquires and controls the joint project), and mergers (where all the assets are merged into a single firm).[1] We examine the efficiency of these alternative governance structures in two cases: first, a static environment, and second, a relational environment (where the parties are engaged in ongoing relationships). Our analysis of the relational case identifies some of the key managerial challenges associated with each governance structure.

Our theory makes several predictions about which governance structure is efficient in various static environments. For example, in a static environment, if spillovers are small then alliances are relatively efficient, whereas if spillovers are large then joint ventures are relatively efficient. Similarly, our theory makes several predictions about which governance structure is efficient in various relational environments. For example, in a relational environment, if spillovers are small then joint ventures are relatively efficient, whereas if spillovers are large then mutual divestitures are relatively efficient. But comparing these two kinds of predictions generates a third kind: comparative predictions, which we derive by holding the basic parameters (such as the size of the spillovers) fixed and then comparing static to relational environments. For example, when spillovers are small then alliances are efficient under static governance but joint ventures are efficient under relational governance. We discuss how routine approaches might allow our static and relational predictions to be tested, and we describe a novel step towards testing our comparative predictions, based on the idea that the efficient governance structure may depend on a firm's network position.

In one way, this paper is a significant departure from the recent literature on vertical integration and the theory of the firm. While much of this literature has emphasized specific

[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.

investments, moral hazard, and hold-up, we emphasize "relational adaptation" (*i.e.*, the role of relationships in facilitating adaptation to states of the world, as they are realized). More specifically, we apply 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 idea in the relational-adaptation theory of the firm is simply that integration can either enhance or inhibit the parties' abilities to adapt as the state of the world is revealed. We use this idea in this paper to explore how different governance structures create different adaptation incentives, and thus partially remedy the ex post contracting problem emphasized in our discussions with practitioners.

In another way, this paper is a natural extension of the existing literature. Several recent papers (*e.g.*, Halonen (1994), Garvey (1995), and Baker, Gibbons, and Murphy (1999, 2002b)) have stressed the importance of ongoing relationships in enriching static property-rights models such as Grossman and Hart (1986), Hart and Moore (1990), and Hart (1995). In particular, the model in Baker, Gibbons, and Murphy (2002b) had only one asset and two parties, but in the Conclusion of that paper we conjectured that a model with more assets and more parties would provide a richer understanding of more complex governance structures such as the variety of strategic alliances described above. The present paper is our first attempt to explore that conjecture.

The paper is organized as follows. In the next section, we describe two economic environments: a simple environment that we use to develop intuition, and a richer setting that we use to show that the results from the simple case are robust. In Section 3 we define the five governance structures mentioned above and derive the efficient governance structure in various static environments. We also derive several testable predictions from this static model. Section 4 adds the possibility of ongoing relationships to each of these governance structures. We again derive first the efficient governance structure in various relational environments and then several testable predictions from this relational model. Finally, Section 5 contains the first step towards a novel empirical analysis of the pharmaceutical and biotechnology industries, in which we motivate the hypothesis that the efficient governance structure might vary with a firm's network position. Section 6 concludes.

2. Economic Environment

In this section, we first define a simple environment that allows us to build intuition about several common governance structures that firms use to coordinate their activities. We then define a richer setting, which we later use to show that the results from the simple case are robust.

2.1 A Simple Environment

There are four assets, {A, a, B, b}. Initially, there are two firms: 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, as described in more detail below.

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, we assume that the payoffs π_A, π_B, π_a, and π_b are observable but not verifiable.[2]

[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 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. We hope to pursue this approach in future work.

The coordinated use of {a, b} could either complement or compete with the core activities of one or both firms. To capture these possibilities, we allow the spillover payoffs to 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 case, 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 tradition of Grossman and Hart (1986), Hart and Moore (1990), and Hart (1995), 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 this contract guarantees only minimal performance, which may fall short of desired performance.

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 could imagine bargaining over asset ownership ex post. That is, because it is not possible to enforce a contract that directly influences the utilization decision by an asset's current owner, one could imagine selling the asset to a new owner whose self-interested utilization decision would be the one that could not be achieved by contract. But we rule out such renegotiation of asset ownership ex post, 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 in Section 3.

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 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 1 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}.

2.2 A Richer Environment

We again consider four assets, {A, a, B, b}, but now the asset-utilization decisions, states, and payoffs are much more general. Let d_a denote the decision taken regarding the utilization of asset a, where d_a is chosen from the finite set D_a, and likewise for d_b and D_b. We will write d for the vector of decisions (d_a, d_b) and D for $D_a \times D_b$. Let s be the state, drawn from the finite set S according to the probability density f(s). Let $\pi_a(d, s)$ denote the payoff to the owner of asset a if decisions d are taken in state s, and likewise for $\pi_b(d, s)$, $\pi_A(d, s)$, and $\pi_B(d, s)$. As in the simple case, these payoffs and the state are observable but not verifiable. Finally, the timing in this richer environment is the same as in the simple case, except that the decisions, states, and payoffs are much more general.

In the simple environment we determined the states in which it is efficient to use {a, b} jointly (producing $\pi_a = \pi_b = \pi > 0$) rather than separately (producing $\pi_a = \pi_b = 0$). We can now conduct the analogous exercise for this richer environment by defining the first-best asset-utilization decisions $d^{FB}(s)$ as the solution to

$$\max_{d \in D} \pi_A(d,s) + \pi_B(d,s) + \pi_a(d,s) + \pi_b(d,s).$$

In the next section we analyze whether any static governance structure can achieve these first-best decisions.

3. Static Governance

To begin our analysis of alternative governance structures, we first consider static models of five possibilities: mergers, alliances, acquisitions, joint ventures, and divestitures. Relative to the economics literature, this is a rich set of feasible governance structures, but it only begins to express the variety that one sees in the world. To document a bit of the latter, we analyzed data collected by industry specialists at Recombinant Capital on nearly 12,500 publicly disclosed multi-firm governance structures (*i.e.*, alliances, joint ventures, and the like) in the pharmaceutical and biotechnology industries from 1973 to 2001. The variety of governance structures (and purposes) found in these data is summarized in Table 1.

We know of no model that encompasses this range of governance structures (Rey and Tirole (1999) may be the closest), so we pause at the end of this section to ask which governance structure is efficient in various static environments. But our main interest in this paper is in the interaction between these formal governance structures and the ongoing relationships that we introduce in Section 4.

3.1 Alternative Static Governance Structures

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}. The owner of this merged firm will jointly utilize {a, b} whenever the total surplus is positive, therefore assets will be jointly used only in states s_{HH}, s_{HM}, and s_{MH} (see Panel A of Figure 1). In the other six states, it is efficient to forego the profit of 2π from coordinated 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}),

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

Since such mergers would completely internalize the externality problems that animate our model, we assume 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.

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 to coordinate activities involving the joint use of {a, b}. Since both payoffs and asset utilization are non-contractible, in this static analysis an arm's-length transaction amounts to unilateral decision-making: party 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 parties will be willing to participate is s_{HH} (see Panel B of Figure 1). Thus, the expected social surplus from an arm's-length alliance is

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

ACQUISITIONS. In an acquisition, the Firm A owns {A, a, b} and Firm B owns {B}. Given our assumption that $M + 2\pi < 0$, Firm A will choose to coordinate the use of {a, b} if and only if

$\pi_A(s) = H$ (see Panel C of Figure 1). The expected social surplus from an acquisition is therefore

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

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 $\pi_B(s) = H$, as shown in Panel D. Given our symmetry assumptions in this simple case, the expected surplus V^{AC} in (3) is the same whether Firm A acquires {b} or Firm B acquires {a}, but this need not be true in the richer environment.

MUTUAL DIVESTITURES. In a mutual divestiture, a third party (call it Firm C) purchases asset {a} from Firm A and asset {b} from Firm B. Since Firm C owns {a, b} and their joint use delivers profits of $2\pi > 0$ independent of the state, Firm C will choose to utilize the assets jointly in all states (see Panel E of Figure 1). The total expected surplus (now received by three parties) from divestitures is then

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

JOINT VENTURES. In a joint venture, Firm A owns {A} and Firm B owns {B}, but a new entity is created that owns assets {a, b}. Unlike a mutual divestiture, however, this new entity is jointly owned by Firms A and B. This joint ownership raises two governance issues: deciding and dividing; that is, how the owners reach decisions concerning asset utilization, and how the non-contractible payoffs from joint asset use (2π) are divided.

There are many possible mechanisms through which joint owners can decide asset utilization. Two simple examples are (1) unanimity is required to achieve joint use and (2) unanimity is required to prevent joint use. If unanimity is required to achieve joint use, the unique Nash equilibrium is to utilize {a, b} jointly only in state s_{HH}, yielding a utilization pattern identical to an alliance's. Alternatively, if unanimity is required to prevent joint use, the unique Nash equilibrium is to utilize {a, b} jointly in all states such that either $\pi_A(s) = H$ or $\pi_B(s) = H$.

In addition to these simple mechanisms, there may also exist more complex "message games" to determine utilization decisions under joint ownership, but we do not consider these in this simple case. Instead, we assume that a joint venture uses the second mechanism above: unanimity is required to prevent joint use.

In addition to a rule for deciding asset utilization, joint owners must also have a process for dividing non-contractible payoffs. Similar issues have been explored in the literature on politics in organizations, where the presence of non-contractible rents inside the organization induces organization members to engage in costly rent-seeking activities; see Skaperdas (1992) and Rajan and Zingales (2000) for models and further references. We do not explicitly model such rent-seeking activities, but instead take a reduced-form approach. Given the symmetry assumptions in this simple case, we assume that equilibrium rent-seeking results in an equal division of the rents (*i.e.*, π to each owner), but that the rent-seeking activities cost each owner the amount k.

Panel F in Figure 1 shows the utilization pattern and payoffs that result from the joint venture we have described. Assets {a, b} are used jointly if $\pi_i(s) = H$ for either owner, but joint use produces payoffs of $\pi_i(s) + \pi - k$ for owner i. Expected surplus is therefore

(6) $$V^{JV} = p_{HH}(2H+2\pi-2k) + 2p_{HM}(H+M+2\pi-2k) + 2p_{HL}(H+L+2\pi-2k).$$

<u>ADDITIONAL GOVERNANCE STRUCTURES.</u> There are many other possible governance structures in our model. 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. One of these hybrids can be seen as a joint venture with outside ownership, as follows.

Imagine a joint venture with three owners rather than two: Firm C as well as Firms A and B. This governance structure can be seen as adding Firm C as an outside owner to the (wholly owned) JV described above; it can also be seen as adding Firms A and B to the ownership group of the new firm created in the mutual divestiture. In either case, the same two issues arise as in a wholly owned JV: how does the ownership group decide on asset utilization, and how are non-contractible payoffs distributed? If asset utilization is decided by majority rule and non-contractible payoffs are distributed equally to all three owners then (in a static environment) asset utilization is identical to the utilization pattern in the wholly owned JV above. But even with majority rule and equal division, JVs with and without outside ownership produce different outcomes in a relational environment.

In addition to these limiting and hybrid cases of the asset-ownership configurations described above, one could also follow Maskin and Tirole (1999) by considering a parallel class of governance structures in which what we have so far called "asset ownership" would instead be called a "control right." That is, rather than envision the control of assets through ownership, one could reinterpret the model as concerning the control of decisions through contracts. We do not pursue this reinterpretation here, except to note that it might provide a new perspective on the empirical analysis of control rights in Lerner and Merges (1998).

3.2 Static Governance in the Richer Environment

It is simple to translate the definitions of alliances, acquisitions, mutual divestitures, and joint ventures from the simple case above to the richer environment described in Section 2.2. For example, in an alliance, $d_a(s)$ and $d_b(s)$ are chosen as a Nash equilibrium: Firm A chooses $d_a(s)$ to maximize $\pi_A(d_a, d_b^*(s), s) + \pi_a(d_a, d_b^*(s), s)$ given Firm B's choice $d_b^*(s)$, and likewise for Firm B. We do not pause here to define each of the governance structures in this richer environment, since our main interest is in the role of ongoing relationships in improving static governance structures. Thus, in Section 4 we take this richer environment more seriously, but here we simply assume what was true in the simple case above: that no static governance structure achieves the first-best decisions $d^{FB}(s)$ in every state.

Before considering efficient governance in static environments, there is one issue that we treated briefly above but can discuss more fully now: renegotiation of asset ownership ex post.

If such renegotiation were feasible in our simple environment then first-best asset utilization would be easy to achieve: the governance structure should begin as an alliance, but would be renegotiated to the appropriate acquisition (by Firm A or Firm B) if the HM or MH state arises. In our richer environment, however, there could easily be states in which no static governance structure produces first-best decisions. In such a version of the richer environment, even ex post renegotiation of asset ownership would not yield the first-best, so there would remain a role for ongoing relationships.[3]

3.3 Efficiency Comparisons under Static Governance

We now return to our simple case, in which coordinated use of {a, b} has both a benefit (π to each of the owners) and a potential cost (negative spillover effects on the core activities {A, B}). Merging all assets {A, a, B, b} into a single firm internalizes these externalities and results in first-best surplus, but none of the other governance structures achieves first-best efficiency (assuming that $p_{ij}>0$ and maintaining our assumptions about the relative magnitudes of H, M, L, and π). Some governance structures jointly utilize {a, b} when it is inefficient to do so; others do not jointly utilize {a, b} even when doing so would be efficient. The efficient static governance structure (excluding mergers) is the one that generates the highest surplus, which we define as

$$V^{ST} \equiv \text{MAX}\,[V^{AL}, V^{AC}, V^{DV}, V^{JV}]\,. \tag{7}$$

This determination of the efficient static governance structure will depend on both the magnitude of the inefficiencies (that is, the cost of joint-use in inefficient states, the lost benefits from joint-use in efficient states, and the deadweight loss from rent-seeking in joint ventures) and the state probabilities. Equations (3’) through (6’) restate the surplus calculations for

---

3 To reiterate, our feeling is that one simply does not see asset ownership flipping back and forth at high frequency, so either states are realized slowly or renegotiation is not occurring. On the other hand, the specific pattern predicted by the possibility of renegotiation in our simple environment does receive some support in the data: see Bleeke and Ernst (1995) on the fairly common occurrence that an alliance becomes a sale.

alliances, acquisitions, mutual divestitures, and (wholly owned) joint ventures relative to the first-best surplus (V^{FB} from (1)):

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

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

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

(6') $V^{JV} = V^{FB} - 2p_{HL}|H + L + 2\pi| - 2k(p_{HH}+2p_{HM}+2p_{HL})$.

Comparing (3') through (6') shows that each static governance structure is optimal for some parameter values. For example, if p_{HL} is small and p_{HM} is large, then either divestitures or (wholly owned) joint ventures will be optimal, depending on the size of k. This makes sense: when the probability of s_{HL} is small, those governance structures that inefficiently implement the project in this state (i.e., divestitures and JVs) will not suffer great efficiency losses, compared to structures that do not implement the project in the high-probability state s_{HM} where it would be efficient to do so (e.g, arm's-length alliances). Similarly, when L is sufficiently large, the alliance will be optimal. Again this is sensible: when the negative spillovers can be very large and negative, only a governance structure that guarantees that either party can stop implementation of the project will be efficient. Finally, an acquisition will be optimal when p_{HL} is moderately small, and k, the deadweight loss associated with the JV, is large.

Several related comparisons are summarized in Figure 2, which plots the efficient static governance structure as a function of H and π, for the case in which all the states are equally likely. As H increases (holding π fixed), the efficient governance structure moves from an alliance to an acquisition to a joint venture. The intuition behind this figure is that at low H there is not much surplus from joint use in the HM state, so an alliance is almost efficient, whereas at high H there is not much loss from joint use in the HL state, so a joint venture is almost efficient. When all the states are equally likely, mutual divestiture is not efficient. In Section 5 we briefly discuss how one might test predictions of this kind.

With these characterizations of the static model as a starting point, we can now turn to our analysis of each of these governance structures when the parties can use ongoing relationships to remedy the inefficiencies of static governance.

4. Relational Governance

In this section we introduce the third element from our discussions with practitioners: ongoing relationships. Relationships in strategic alliances occur in a variety of different ways. First, alliances are often long-lived and involve continual interactions between the parties over extended periods of time. 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). Finally, firms often have multiple alliance partners who, in turn, have multiple alliance partners, creating a network of relationships that can facilitate information flows between firms that have not done a deal together (Gulati, 1995b).

The Recombinant Capital pharmaceuticals-biotechnology database introduced above does not offer complete information on the longevity of individual alliances. However, for the 372 alliances that were formally terminated between 1980 and 2001, the median time between the initial contract and the termination was 33 months.[4] In addition, for the 1,548 alliance contracts that were formally revised (but not terminated) during the sample period, the median time between the initial contract and the revision was 21 months, thus constituting a lower-bound on alliance longevity for these contracts. These data suggest that alliances are not one-shot transactions, but rather involve continual interactions spanning multiple years.

Table 2 presents evidence on repeat alliances between the same partners for alliances in the Recombinant Capital database. Most pairs of firms (9,462) do only one deal with each other,

---

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

but over a thousand pairs of firms do more than one deal together; 57 pairs do five or more deals together.

Table 3 offers some preliminary evidence on networks of relationships among 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. Figure 3 illustrates the network relationships among these top 12 pharmaceuticals and biotechs. On average, each firm among these 24 has alliance-relations with 15 of the other 23 firms. We interpret these data as supporting the importance of networks among large pharmaceutical and biotechnology firms.

To the extent that networks facilitate information flows, relationships can involve not only direct partners but partners of partners and so on. The 24 firms in Table 1 had alliance arrangements with a total of 1,308 partners outside of the 24; these 1,332 firms taken together were involved in 11,303 alliances (91% of the alliances identified by Recombinant). Including the partners of these partners yields a total of 3,421 firms (81% of the firms) who were involved in 98% of the reported alliances. These results 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. Overall, we interpret these data as providing strong circumstantial evidence for network-driven relational ties between firms that do not directly deal with each other.

In our model, relationships are used to remedy the ex post inefficiencies in the static governance structures analyzed in Section 3. In other models, relationships are used in a different way. As noted above, Halonen (1994), Garvey (1995), and Baker, Gibbons, and Murphy (1999, 2002b) add relationships to static property-rights models in the spirit of Grossman and Hart (1985), Hart and Moore (1990), and Hart (1995). In these static property-rights models, ex post surplus shares create ex ante incentives for non-contractible specific investments. Adding relationships to these models enriches the feasible set of ex post surplus shares, and so may improve ex ante incentives (and change the efficient structure of asset ownership, compared to the static model). In this paper, however, there are no ex ante actions in the economic environments described in Section 2, so relationships cannot be used to improve ex ante incentives.

As is now standard, we model an ongoing relationship as a repeated game. Following a large literature, we interpret an equilibrium in the repeated game as a "relational contract" (*i.e.*, an agreement between the parties that is so rooted in their shared experience that it 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 this paper, we restrict attention to first-best relational contracts, in which the asset-utilization decisions are $d^{FB}(s)$. The task is then to ascertain whether there exist payment schemes that induce the parties to make first-best decisions. These payments can exist in a variety of forms, regardless of the governance structure. First, the payments might be "efficiency wages," denoted by t and paid before the state or any decisions are observed. Second, the payments might be "bribes," denoted by $\tau(s)$ 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 $T(d, s)$ and paid depending on whether asset-utilization decisions are appropriately tailored to the state. More generally, the payments could include any combination of these three possibilities.

In a moral-hazard model (without asset ownership), MacLeod and Malcomson (1989) showed that any repeated-game equilibrium can be implemented either with efficiency wages or with subjective bonuses or with a combination of the two. Our initial analyses suggest that the same is true here: any equilibrium can be implemented with efficiency wages, bribes, subjective bonuses, or a combination of these. More specifically, for some of the governance structures defined in Section 3, we show below that if an equilibrium can be implemented at all then it can be implemented with a very intuitive choice of bonuses and efficiency wages. In this draft we assume (and expect that we will soon prove) that this result holds for all governance structures.

We analyze trigger-strategy equilibria: if any party reneges, the parties engage in static transactions thereafter. But we impose a form of renegotiation-proofness by assuming that if reneging occurs then the parties engage in efficient static governance structure thereafter. (Achieving efficient static governance will typically require a change in asset ownership, with an accompanying side-payment.) Thus, after reneging, the total expected surplus will be V^{ST} per period, as defined in (7).

This section proceeds in five steps. First, to lay the groundwork, we derive a familiar but useful feasibility condition for a repeated-game equilibrium: the maximum reneging temptation must be smaller than the present value of the net surplus from the relational contract. Second, for the simple environment, we derive the maximum reneging temptation for each of the governance structures defined above. Third, we show that the results from the simple environment are robust by analyzing relational contracts in the richer environment. Fourth, for the simple environment, we consider efficiency differences across relational governance structures. And finally, we compare static to relational governance, again for the simple environment.

4.1 Reneging Temptations in Relational Contracts

Under any of the ownership structures defined above, there are many reneging constraints that must be satisfied if a relational contract is to be a repeated-game equilibrium. For example, either party could refuse to pay or to accept an efficiency-wage payment, t. Similarly, either party could refuse to pay or to accept a bribe, $\tau(s)$. In addition, whichever party or parties own assets a and b could decide to deviate from first-best decisions, $d^{FB}(s)$. And finally, either party could refuse to pay or to accept a bonus, T(d, s).

In this section we derive a familiar but useful inequality that is necessary and sufficient for all the constraints just mentioned. In reducing the many reneging constraints to a single inequality, we are applying the arguments in MacLeod and Malcomson (1989) and Levin (2001) to our environment. Since our goal in this sub-section is only to provide the intuition for this important feasibility condition, we consider only the simplest governance structure: an Acquisition, in which Firm I (= A or B) owns assets {I, a, b}.

To express the various reneging constraints, we need the following notation, concerning payoffs, actions, efficient static governance, and potential asset sales after reneging. Let $U_I(d, s)$ denote Firm I's total payoff,

(8) $$U_I(d,s) \equiv \pi_I(d,s) + \pi_a(d,s) + \pi_b(d,s) \quad ;$$

similarly, let $U_J(d,s)$ denote Firm J's total payoff, $U_J(d,s) \equiv \pi_J(d, s)$. Let $d_I(s)$ denote Firm I's optimal action in state s, which solves

(9) $$\max_{d \in D} U_I(d,s) \quad ;$$

of course, $U_I(d_I(s), s) \geq U_I(d^{FB}(s), s)$. Finally, let V_I^{ST} denote Firm I's expected payoff per period under the efficient static governance structure, and likewise for V_J^{ST}, so that $V_I^{ST} + V_J^{ST} = V^{ST}$ in (7), and let p_{AC} denote the price (paid by Firm J to Firm I) for any asset sales necessary to achieve the efficient static governance structure. For example, suppose that an alliance is the efficient static governance structure. If reneging occurs in a relational contract in the acquisition above then Firm I will sell (say) asset b to Firm J, to achieve efficient static governance thereafter.[5]

[5] Two points should be noted here. First, the framework we present in this sub-section cannot handle the third party who would enter the analysis if the efficient static governance structure were mutual divestiture. This is an example of the kind of work that remains to be done on the theory in this paper. Second, it may seem inconsistent 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 {a, b} jointly is fleeting, and transferring ownership would take some time, so this is why there is no

Given this notation, the reneging constraints described above can be written as follows (where we adopt the convention that payments are from Firm J to Firm I). The constraints that the parties be willing to pay and to accept the efficiency wage payment t are

$$(10) \quad \left(1+\frac{1}{r}\right)\left[-t + E_s\{U_J(d^{FB}(s),s) - \tau(s) - T(d^{FB}(s),s)\}\right] \geq 0 + E_s\{U_J(d_I(s),s)\} + \frac{1}{r}V_J^{ST} - \frac{1}{1+r}p_{AC}$$

and

$$(11) \quad \left(1+\frac{1}{r}\right)\left[t + E_s\{U_I(d^{FB}(s),s) + \tau(s) + T(d^{FB}(s),s)\}\right] \geq 0 + E_s\{U_I(d_I(s),s)\} + \frac{1}{r}V_I^{ST} + \frac{1}{1+r}p_{AC} \quad .$$

The lefthand side of (10) is the expected present value of Firm J's payoffs on the equilibrium path. The righthand side of (10) consists of two payoffs this period if Firm J reneges on the efficiency wage payment t (namely, zero from not making the payment and $E_s\{U_J(d_I(s), s)\}$ from Firm I's optimal decision after J reneges) and two payoffs starting next period (namely, V_J^{ST} each period forever after from efficient static governance, but p_{AC} paid at the start of the next period for any asset sale necessary to achieve efficient static governance). For Firm J to be willing to pay t, the lefthand side of (10) must exceed the right, and analogously for Firm I in (11).

The constraints that the parties be willing to pay and to accept the bribe τ(s) are

$$(12) \quad \left[-\tau(s) + U_J(d^{FB}(s),s) - T(d^{FB}(s),s)\right] + \frac{1}{r}\left[-t + E_s\{U_J(d^{FB}(s),s) - \tau(s) - T(d^{FB}(s),s)\}\right] \geq 0 + U_J(d_I(s),s) + \frac{1}{r}V_J^{ST} - \frac{1}{1+r}p_{AC}$$

and

$$(13) \quad \left[\tau(s) + U_I(d^{FB}(s),s) + T(d^{FB}(s),s)\right] + \frac{1}{r}\left[t + E_s\{U_I(d^{FB}(s),s) + \tau(s) + T(d^{FB}(s),s)\}\right] \geq 0 + U_I(d_I(s),s) + \frac{1}{r}V_I^{ST} + \frac{1}{1+r}p_{AC} \quad .$$

renegotiation of ownership ex post. On the other hand, we envision that the timing of events within a period has all the action up front, followed by the span of time that creates the need for one period of discounting, so this is why there can be asset sales after reneging.

There are two differences between (12) and (10): in (12), t has already been paid, so it does not appear in the first bracket on the lefthand side, and the state s has already been realized, so this period's payoffs on both sides are contingent on s rather than expectations. Otherwise, (12) and (13) mimic (10) and (11), respectively.

The constraint that Firm I be willing to take the first-best decisions $d^{FB}(s)$ is

$$(14)\qquad \left[U_I(d^{FB}(s),s)+T(d^{FB}(s),s)\right]+\frac{1}{r}\left[t+E_s\{U_I(d^{FB}(s),s)+\tau(s)+T(d^{FB}(s),s)\}\right] \geq U_I(d_I(s),s)+\frac{1}{r}V_I^{ST}+\frac{1}{1+r}p_{AC}\ .$$

This constraint is different than the others, because Firm I's decisions affect total surplus, rather than merely divide a fixed surplus as the payments do. Rearranging (14) yields

$$(15)\qquad \frac{1}{r}\left[t+E_s\{U_I(d^{FB}(s),s)+\tau(s)+T(d^{FB}(s),s)\}\right]-\left(\frac{1}{r}V_I^{ST}+\frac{1}{1+r}p_{AC}\right) \geq U_I(d_I(s),s)-\left[U_I(d^{FB}(s),s)+T(d^{FB}(s),s)\right]\ ,$$

which says that the present value of Firm I's future payoffs from staying on the equilibrium path (net of the future payoffs after reneging) must exceed the reneging temptation this period, $U_I(d_I(s), s) – U_I(d^{FB}(s), s) – T(d^{FB}(s), s)$. This expression for the reneging temptation this period, on the righthand side of (15), will play a key role below.

Finally, the constraints that the parties be willing to pay and to accept the bonus T(d, s) are

$$(16)\qquad -T(d^{FB}(s),s)+\frac{1}{r}\left[-t+E_s\{U_J(d^{FB}(s),s)-\tau(s)-T(d^{FB}(s),s)\}\right] \geq 0+\frac{1}{r}V_J^{ST}-\frac{1}{1+r}p_{AC}$$

and

$$(17)\qquad T(d^{FB}(s),s)+\frac{1}{r}\left[t+E_s\{U_I(d^{FB}(s),s)+\tau(s)+T(d^{FB}(s),s)\}\right] \geq 0+\frac{1}{r}V_I^{ST}+\frac{1}{1+r}p_{AC}\ .$$

There are two differences between (16) and (12): τ(s) has already been paid, so it does not appear in the first bracket on the lefthand side, and the decisions $d^{FB}(s)$ have already been taken, so the utility terms also do not appear on either side of (16). Otherwise, (16) and (17) mimic (12) and (13), respectively.

Miraculously, these seven reneging constraints can be reduced to a single inequality! (This result depends crucially on two elements of the model: risk-neutrality and the existence of the side-payment, t.) To sketch the proof of this result, consider the following specifications of the bonus, efficiency wage, and bribe:

(18) $$T(d^{FB}(s),s) = U_I(d_I(s),s) - U_I(d^{FB}(s),s)$$

and T(d, s) = 0 otherwise;

(19) $$t = V_I^{ST} + \frac{r}{1+r} p_{AC} - E_s\{U_I(d_I(s),s)\};$$

and τ(s) = 0 for all s. These payments cause an enormous simplification of the seven reneging constraints. In particular, (15) becomes 0 ≥ 0 and (16) becomes

(20) $$U_I(d_I(s),s) - U_I(d^{FB}(s),s) \geq \frac{1}{r}\left(V^{FB} - V^{ST}\right),$$

where $V^{FB} = E_s\{U_I(d^{FB}(s), s) + U_J(d^{FB}(s), s)\}$, and all the remaining constraints are tediously satisfied.

This derivation establishes sufficiency: if (20) holds then the payments defined in (18) and (19) imply that all seven reneging constraints hold. But (20) is also necessary, because adding (14) and (16) yields (20). Thus, the feasibility condition (20) becomes the central constraint in our repeated-game analysis.

4.2 Relational Governance in the Simple Environment

Cousins of the feasibility condition (20) can be used to calculate the reneging temptation for any governance structure, as follows. First, pick a governance structure. Second, identify states where any firm is tempted to renege on its part of first-best asset utilization. Third, compute the reneging temptation in each of these states for the tempted party. Finally, search across states for the maximum reneging temptation, which is then defined to be the reneging temptation for that governance structure. Using this methodology, we now derive the reneging temptations for relational contracts in alliances, acquisitions, joint ventures, and divestitures for our simple environment.

RELATIONAL ALLIANCES. Firm A is tempted to deviate in state s_{MH} (because first-best asset utilization requires joint use of {a, b}, but joint use has a negative payoff for Firm A) with reneging temptation $|M + \pi|$, while Firm B is tempted to deviate in state s_{HM}, also with reneging temptation $|M + \pi|$. Therefore the reneging temptation for relational alliances is $R^{AL} = |M + \pi|$.

RELATIONAL ACQUISITIONS. Firm A (as the acquirer of asset {b}) is tempted to deviate in state s_{HL} (because first-best utilization forbids joint use of {a, b}, but Firm A would receive a positive payoff from joint use), with reneging temptation $H + 2\pi$, and also in state s_{MH} (where the contract specifies joint utilization but A is tempted to refuse), with reneging temptation $|M + 2\pi|$. By our assumptions that $M + 2\pi < 0$ and $H + M + 2\pi > 0$, we know that $|M + 2\pi| < H < H + 2\pi$. Therefore, the reneging temptation for relational acquisitions (which, by symmetry, is the same if B acquires {a}) is $R^{AC} = H + 2\pi$.

RELATIONAL DIVESTITURES. Firm C (as the owner of {a, b}) is tempted to deviate (by jointly using the assets) in s_{HM}, s_{HL}, s_{MM}, s_{ML}, s_{LH}, s_{LM}, and s_{LL}. The short-run cost to Firm C of honoring the contract (by not utilizing) is 2π in each state. Therefore, the reneging temptation for relational divestitures is $R^{DV} = 2\pi$.

RELATIONAL JOINT VENTURES. In a joint venture governed by "unanimity to stop" (either party can force the joint use of {a, b}), Firm A is tempted to deviate in state s_{HL}, with reneging temptation $H + \pi - k$, while Firm B is tempted to deviate in state s_{LH}, also with reneging temptation $H + \pi - k$. Therefore the reneging temptation for relational joint ventures is $R^{JV} = H + \pi - k$.

This analysis exposes some of the key managerial challenges associated with these governance structures. For example, in an alliance, 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. In a joint venture, the challenge is similar to that in an acquisition, but now either party can be tempted, but only by half the project's payoffs (π rather than 2π) and at the haggle cost k. 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.

4.3 Relational Governance in the Richer Environment

All of the reneging temptations derived above carry over in the natural way to our richer environment, with analogous implications for the managerial challenges in these governance structures.

To be completed.

4.4 Efficient Relational Governance

The righthand side of the feasibility condition (20) is independent of governance structure and reflects the present value of honoring a relational contract that achieves first-best utilization and surplus. For a given governance structure, at sufficiently low discount rates, (20) is satisfied. However, at sufficiently high discount rates, (20) is not satisfied and the first-best outcome is not achievable, by any governance structure. For intermediate discount rates, (20) may be satisfied for some governance structures but not for others. For a given set of parameters, we define the efficient governance structure as the one that can satisfy (20) at the highest discount rate. That is, the efficient governance structure has the lowest reneging temptation, $\text{MIN}[R^{AL}, R^{AC}, R^{DV}, R^{JV}]$ as derived above.

In our simple environment, both relational joint ventures ($R^{JV} = H + \pi - k$) and relational divestitures ($R^{DV} = 2\pi$) dominate relational acquisitions ($R^{AC} = H + 2\pi$), since both $H + \pi - k$ and 2π are strictly less than $H + 2\pi$. But each of the other three relational governance structures can be efficient: relational alliances dominate relational divestitures if $|M + \pi| < 2\pi$; relational alliances dominate relational joint ventures if $|M + \pi| < H + \pi - k$; and relational divestitures dominate relational joint ventures if $2\pi < H + \pi - k$. These three comparisons are shown in Figure 4, which plots the efficient relational governance structure as a function of H and π, for the case in which all the states are equally likely.

The chief result in Figure 4 is that increasing π (while holding H fixed) causes the efficient relational governance structure to switch from mutual divestiture to an alliance (or perhaps to a joint venture). It is clear why mutual divestiture becomes inefficient as π grows: Firm C's payoff is limited to the 2π from the project and is independent of the spillover payoffs $\pi_A(s)$ and $\pi_B(s)$. It is also clear why an alliance becomes more efficient as π grows: the inefficiency in a static alliance is that firms do not achieve joint asset utilization in the HM and MH states, but the temptation to renege in these states in a relational alliance is $|M + \pi|$, which falls as π grows.

4.5 *Static versus Relational Governance*

Efficient governance in a static environment is determined by $V^{ST} \equiv \text{MAX}\,[V^{AL}, V^{AC}, V^{DV}, V^{JV}]$ in (7), which in turn depends on the probabilities of realizing states s_{ij} and on the magnitude of the resulting inefficiencies (that is, the cost of joint-use in inefficient states plus the lost benefits from lack of joint-use in efficient states). In contrast, efficient governance in a relational environment is determined not by surplus but rather by reneging temptations, $R^* \equiv \text{MIN}\,[R^{AL}, R^{AC}, R^{DV}, R^{JV}]$. Thus, holding fixed what one wants to achieve (such as the first-best, in our analysis), efficient governance in a relational environment is independent of the state probabilities and the static surpluses.

These differences in the determination of efficient governance structures in static versus relational environments produce the comparative predictions shown in Figure 5, which overlays the efficient static and relational governance structures as a function of H and π, for the case in which all the states are equally likely. The figure shows that, holding the model's basic parameters fixed, the efficient governance structure depends on whether the environment is static or relational. For example, when H and π are small, the efficient governance structure is an alliance in a static environment but a mutual divestiture in a relational environment.

5. Towards Testing

In Section 3 we derived several predictions from the static model, and likewise in Section 4 from the relational model. For example, in a static environment, if spillovers are small then alliances are relatively efficient, whereas if spillovers are large then joint ventures are relatively efficient. Similarly, in a relational environment, if spillovers are small then joint ventures are relatively efficient, whereas if spillovers are large then mutual divestitures are relatively efficient. Testing these static and relational predictions will require operationalizing the central variables in the model, such as the spillover payoffs H and L and the project size π. We believe that proxies for these variables may be available. For example, H and L seem more likely to be large (in absolute value) when a project is closely related to a firm's core business. Similarly, a project's size seems straightforward to approximate. Thus, we hold out some hope that tests of the predictions of our static and relational models may not be too difficult to produce.

In addition to these static and relational predictions, we also derived several comparative predictions, in which we held the model's basic parameters fixed and varied whether the environment was static or relational. For example, when spillovers are small then alliances are efficient under static governance but joint ventures are efficient under relational governance. These predictions seem to us to be probably the most interesting to test, in part because they will shed light on whether relationships matter, and in part because they may help build a link (or at least a dialogue) between sociological and economic ideas, as follows.

The first step in testing such comparative predictions is to operationalize the distinction between static and relational environments. In this section we propose a method for doing so, based on a reinterpretation of the central parameter of repeated-game models: r, the discount rate. Taken literally, the discount rate 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. But one can go further in reinterpreting r, by analyzing information flows in social structures. As a stark example, in a large population where players are randomly matched each period, the effective value of r is much higher if players observe only their own interactions, as opposed to observing everyone else's interactions as well. To investigate such a reinterpretation of r, we study the social structure of the pharmaceutical and biotechnology industries.

We return to the data from Recombinant Capital described above (which includes 12,451 alliances among 4,231 surviving-parent firms), but this time as network theorists. We begin by defining a firm to be *peripheral* if it has only one partner. Any firm that is not peripheral has more than one partner, so we call it a *hub*. About half (51%) of the firms in our data are peripheral, but the hubs come in several types. First, hubs can be either *isolated* or *connected*: an isolated hub is connected to only peripheral firms, whereas a connected hub is connected to at least one other hub. Connected hubs can then be broken into classes, depending on (1) how many partners they have, (2) how many of these partners are hubs rather than peripheral firms, and (3) how many partners these partner-hubs have.

We define three classes of connected hubs: an *entry* hub has only one partner that is not peripheral, an *intermediate* hub has between two and nineteen partners that are not peripheral, and a *core* hub has at least twenty partners that are not peripheral. These definitions of connected hubs allow for *isolated constellations*, such as two entry firms in a "dumbbell" (meaning that each entry firm has many peripheral partners but only the other entry firm as a non-peripheral partner), or three intermediate firms in a "triangle" (meaning that each entry firm has many peripheral partners but only the other two intermediate firms as non-peripheral partners), and so on. But these definitions of connected hubs also allow for a *central constellation* in which: core firms are densely interconnected and are involved in a disproportionate share of the alliances (with other core firms, with intermediate and entry firms, and with peripheral firms); intermediate firms are linked to a few core firms, but also conduct alliances with intermediate and entry firms, and perhaps many alliances with peripheral firms; and entry firms may conduct many alliances, but only one is with an intermediate or core firm.

Our analysis of the Recombinant Capital database shows that the pharmaceutical and biotechnology industries have a striking central constellation, with a handful of isolated constellations. This network structure is shown in Figure 6.

Having uncovered this network structure from the raw data, we hypothesize that core firms are more visible than peripheral firms: not only do core firms have more partners than peripheral firms have, but core firms are tied to firms that themselves have more partners than the partner of a peripheral firm typically has. For both these reasons, we hypothesize that interactions between core firms effectively have a lower discount rate than do interactions involving peripheral firms. That is, we hypothesize that the social structure that links core firms makes interactions between core firms more likely to be relational, whereas interactions involving peripheral firms are more likely to be static. In this sense, we are hypothesizing that efficient governance may depend on a firm's network position.

We reiterate that, in this paper, we cannot test the predictions that follow from this hypothesis: our comparative predictions hold constant the model's basic parameters (*e.g.*, H, L, π) and then compare static to relational environments; our network methodology offers hope for the latter, but our data do not permit the former. Nonetheless, we take a little satisfaction simply from having measured the social structure of an industry, and more satisfaction from having suggested how that social structure might be related to a key variable in a familiar economic model. Our approach thus blends a foundational idea from sociology with a widespread methodology from economics. Sociologists such as Gulati (1995a, 1995b) have been working in a similar spirit, but with different conceptions of networks and different conceptions of the causes and effects of relationships. It will be interesting to see whether our approach and the sociologists' can inform each other.

6. Conclusion

A strategic alliance is a governance structure for coordinating activities among firms. Such governance structures are observed in many forms, including alliances, joint ventures, acquisitions, and mutual divestitures. 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 have developed a simple model that allows us to characterize and contrast alliances, joint ventures, acquisitions, and mutual divestitures. 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.

In addition to developing our model, we have also introduced suggestive empirical evidence drawn from Recombinant Capital's database of strategic alliances among pharmaceutical and biotechnology firms. We have documented a rich variety of governance structures and organizational objectives within these alliances. We have also established evidence on the importance of relationships and networks in pharma-biotech alliances.

We believe that our approach suggests a rich and tractable framework that could be used to analyze a wide variety of additional governance structures. In addition, we have analyzed only first-best relational contracts, but one could 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, and (based on this distinction) we have provided several predictions potentially refutable by existing or future data. In particular, we suggest that the network position of a firm yields information on the effective "discount rate" in relational contracts which, in turn, yields predictions regarding optimal governance structures in strategic alliances.

References

Baker, George, Robert Gibbons, and Kevin J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics* 109:1125-56.

__________, __________, and __________. 1999. "Informal Authority in Organizations." *Journal of Law, Economics, and Organization* 15: 56-73.

__________, __________, and __________. 2001. "Bringing the Market Inside the Firm?" *American Economic Review Papers and Proceedings* 91: 212-18.

__________, __________, and __________. 2002a. "s Happens: The (Long-Dormant) Relational-Adaptation Theory of the Firm." Working paper.

__________, __________, and __________. 2002b. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics* 117: 39-83.

Bleeke, Joel and David Ernst. 1995. "Is Your Strategic Alliance Really a Sale?" *Harvard Business Review* xx: yy-zz.

Bull, Clive. 1987. "The Existence of Self-Enforcing Implicit Contracts," *Quarterly Journal of Economics* 102:147-59.

Garvey, Gerald. 1995. "Why Reputation Favors Joint Ventures over Vertical and Horizontal Integration: A Simple Model." *Journal of Economic Behavior and Organization* 28: 387-97.

Grossman, Sanford and Oliver Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Ownership." *Journal of Political Economy* XCIV: 691-719.

Gulati, Ranjay. 1995a. "Does Familiarity Breed Trust? The Implications of Repeated Ties for Contractual Choice in Alliances." *Academy of Management Journal* 38: 85-112.

__________. 1995b. "Social Structure and Alliance Formation Patterns: A Longitudinal Analysis." *Administrative Science Quarterly* 40: 619-52.

Halonen, Maija. 1994. "Reputation and Allocation of Ownership." Working Paper, Helsinki School of Economics.

Hart, Oliver. 1995. *Firms, Contracts, and Financial Structure*. Oxford: Clarendon Press.

__________ and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy* XCVIII: 1119-58.

Klein, Benjamin. 1991. "Vertical Integration as Organizational Ownership: The Fisher Body-General Motors Relationship Revisited." In O. Williamson and S. Winter (eds.), *The Nature of the Firm: Origins, Evolution, and Development*. Oxford University Press.

__________, Robert Crawford, and Armen Alchian. 1978. "Vertical Integration, Appropriable Rents and the Competitive Contracting Process." *Journal of Law and Economics* XXI: 297-326.

__________ and Keith Leffler, "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, LXXXIX (1981), 615-641.

Lerner, Josh and Robert Merges. 1998. "The Control of Technology Alliances: An Empirical Analysis of the Biotechnology Industry." *Journal of Industrial Economics* 46: 125-56.

Levin, Jonathan. 2001. "Relational Incentive Contracts." Unpublished manuscript, Stanford University.

Macaulay, Stewart, "Non Contractual Relations in Business: A Preliminary Study," *American Sociological Review*, XXVIII (1963), 55-67.

MacLeod, Bentley, and James Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57: 447-80.

Macneil, Ian, "Contracts: Adjustments of long-term economic relations under classical, neoclassical, and relational contract law," *Northwestern University Law Review*, LCCII (1978), 854-906.

Maskin, Eric and Jean Tirole, "Two Remarks on Property Rights," *Review of Economic Studies* LXVI (1999), 139-49.

McQuade, Krista and Benjamin Gomes-Casseres. 1992. "Xerox and Fuji Xerox." Harvard Business School Case #9-391-156.

Rajan, Raghuram and Luigi Zingales. 2000. "The Tyranny of Inequality." *Journal of Public Economics* 76: 521-58.

Rey, Patrick and Jean Tirole, "Divergence of Objectives and the Governance of Joint Ventures," Working Paper, IDEI (1999).

Simon, Herbert. 1951. "A formal theory of the employment relationship." *Econometrica* 19: 293-305.

Skaperdas, Stergios. 1992. "Cooperation, Conflict, and Power in the Absence of Property Rights." *American Economic Review* 82: 720-39.

Telser, Lester, "A Theory of Self-Enforcing Agreements," *Journal of Business*, LIII (1981), 27-44.

Whinston, Michael. 2000. "On the Transaction Cost Determinants of Vertical Integration." Unpublished manuscript, Northwestern University.

Williamson, Oliver. 1971. "The Vertical Integration of Production: Market Failure Considerations." *American Economic Review* 63: 316-25.

__________. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York, NY: Free Press.

__________. 1985. *The Economic Institutions of Capitalism*, New York, NY: Free Press.

Figure 1

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

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. Acquisition (B buys)
B controls {a, b}

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

Panel E. 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

Panel F. Joint Venture
A,B jointly control {a, b}

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

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 1

Operational Objectives and Governance Structures for Pharmaceutical-Biotech Strategic Alliances, 1973-2001

	Governance Structure for Alliance					
Operational Objective of Alliance:	License	Investment	Merger or Acquisition	Joint Venture	Structure not Specifed	Total
Development	16.2%	4.6%	0.1%	0.7%	7.7%	29.4%
Research	13.3%	3.5%	0.1%	0.4%	7.3%	24.6%
Manufacturing or Marketing	4.7%	1.8%	0.4%	0.3%	10.6%	17.9%
Collaboration	7.3%	2.2%	0.0%	0.2%	6.9%	16.7%
Supply	4.3%	1.3%	0.3%	0.1%	3.1%	9.2%
Objective not specified	20.6%	4.9%	12.8%	2.1%		40.3%
Total	66.5%	18.4%	13.8%	3.8%	35.7%	

Note: 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. Totals sum to more than 100% because contracts frequently mention multiple objectives (e.g., research *and* development) and often note multiple governance structures (e.g., investment *and* license agreement).

Figure 2

Efficient Organizational Forms under Static Governance

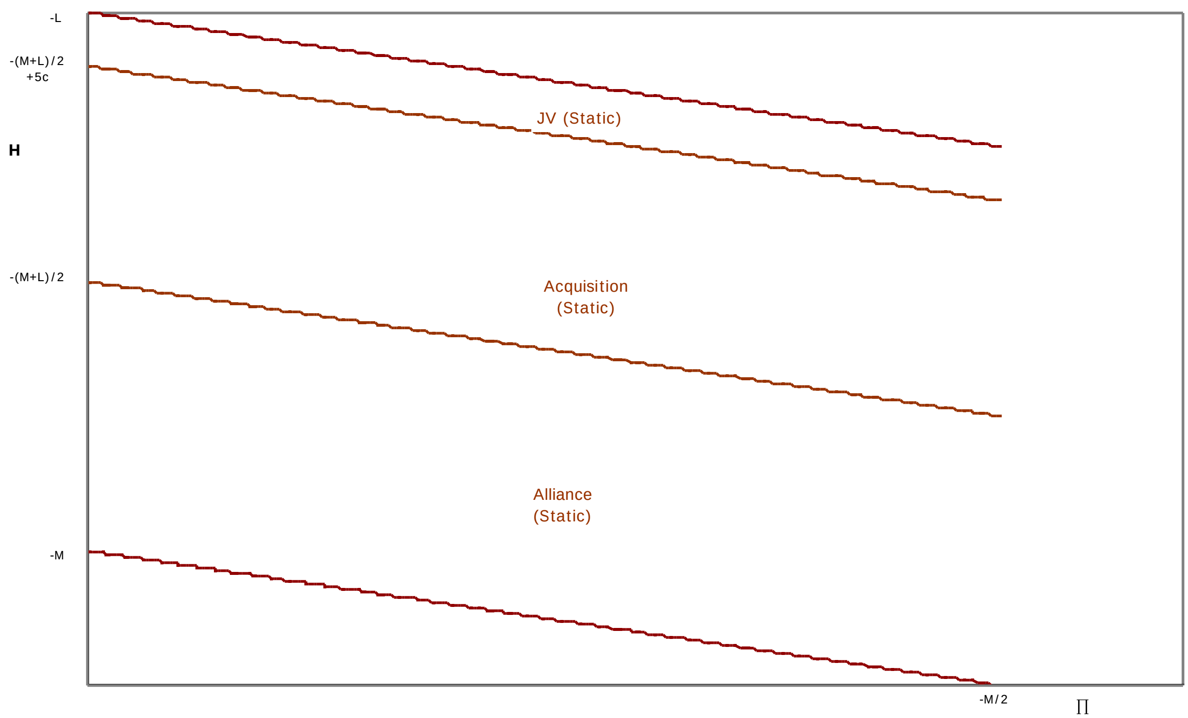

Table 2

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 3

Strategic Alliances Among the Top 12 Pharmaceuticals and Top 12 Biotechs

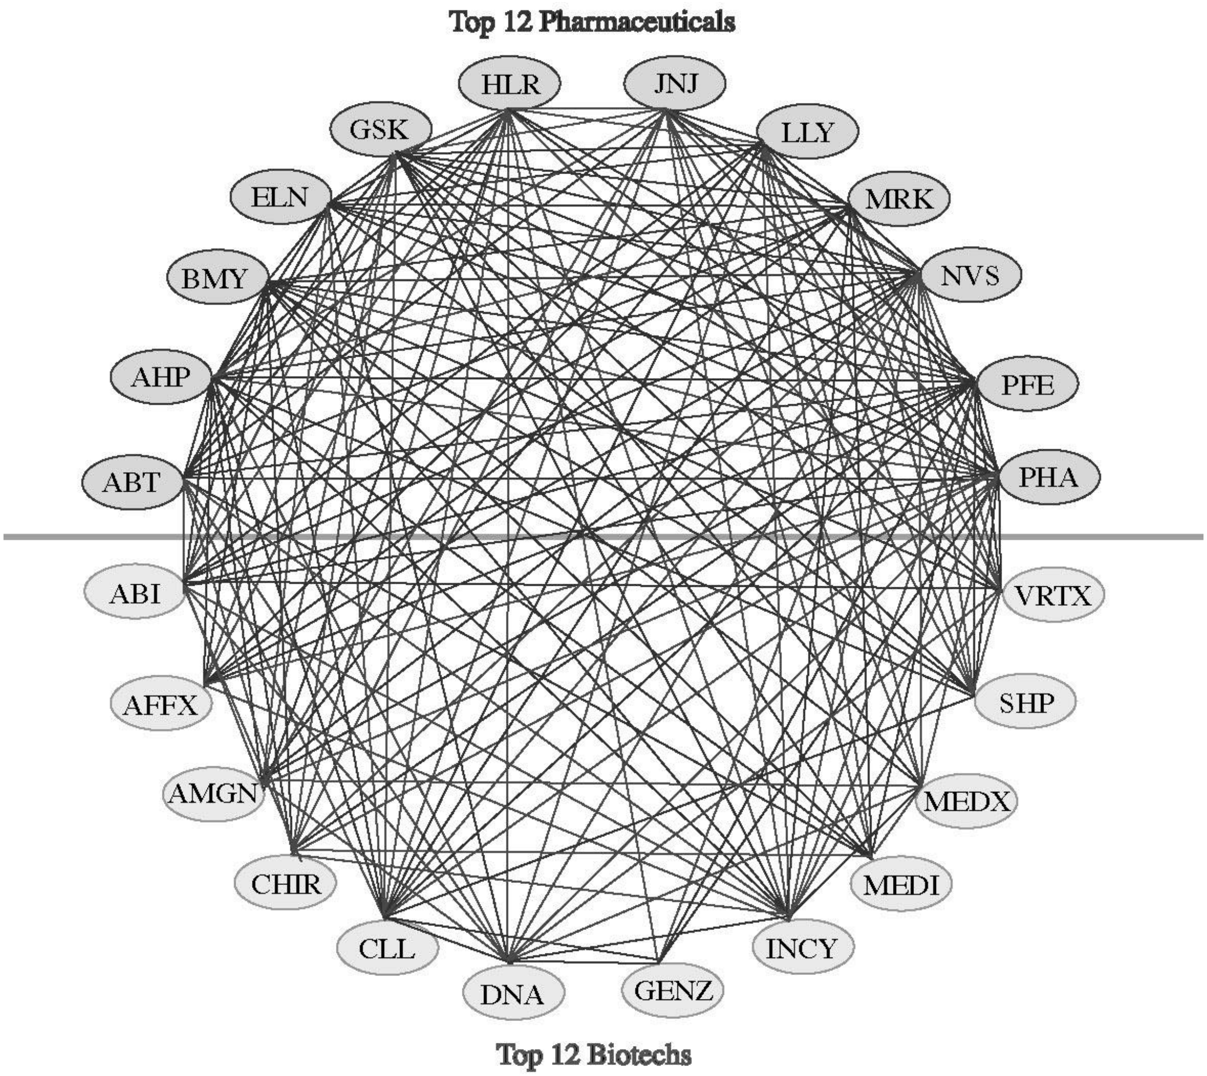

Note: Ticker symbols correspond to companies included in Table 3. Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, based on approximately 12,500 publicly disclosed contracts and arrangements from 1973-2001. Contracts are assigned to the surviving parent as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

Figure 4

Efficient Organizational Forms under Relational Governance

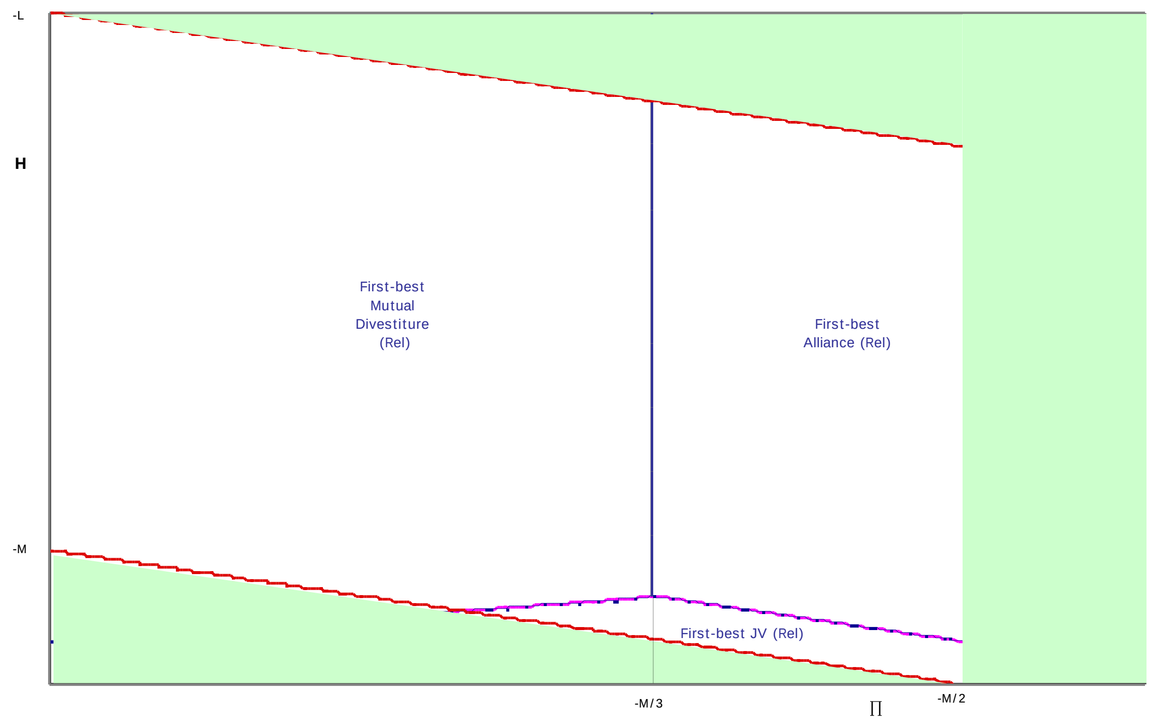

Figure 5

Comparison of Efficient Organizational Forms under Static and Relational Governance

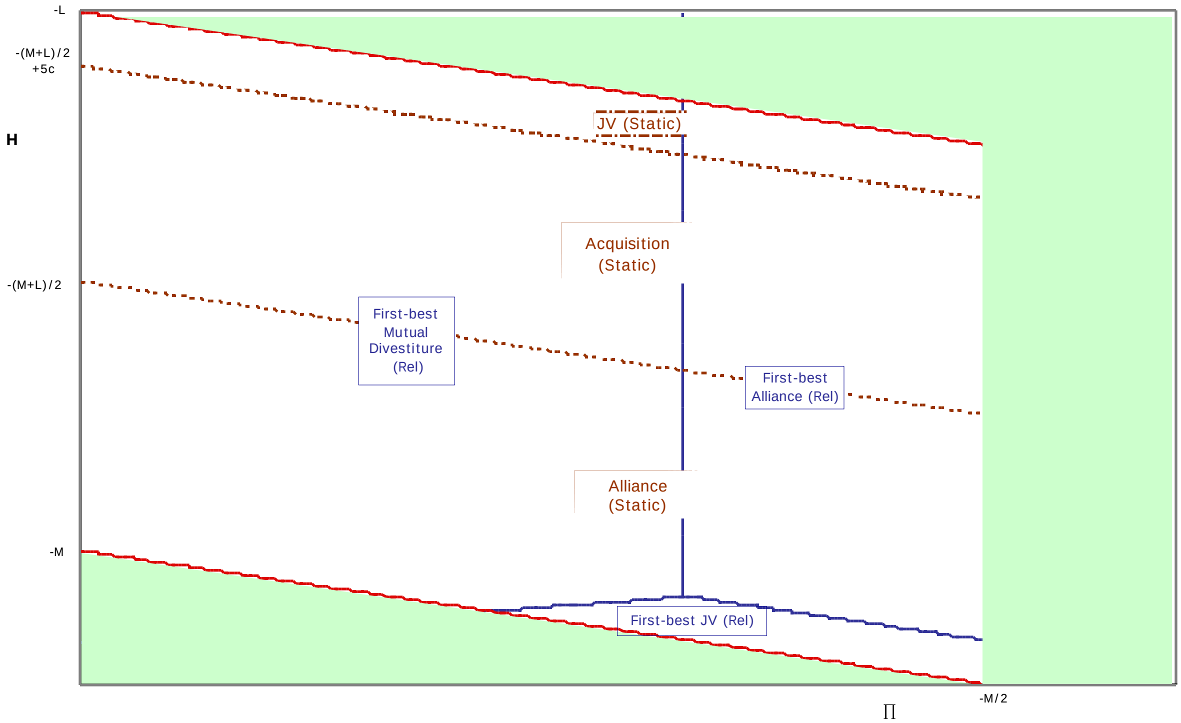

Figure 6

Networks in Recombinant Capital Database of Pharmaceutical-Biotech Alliances

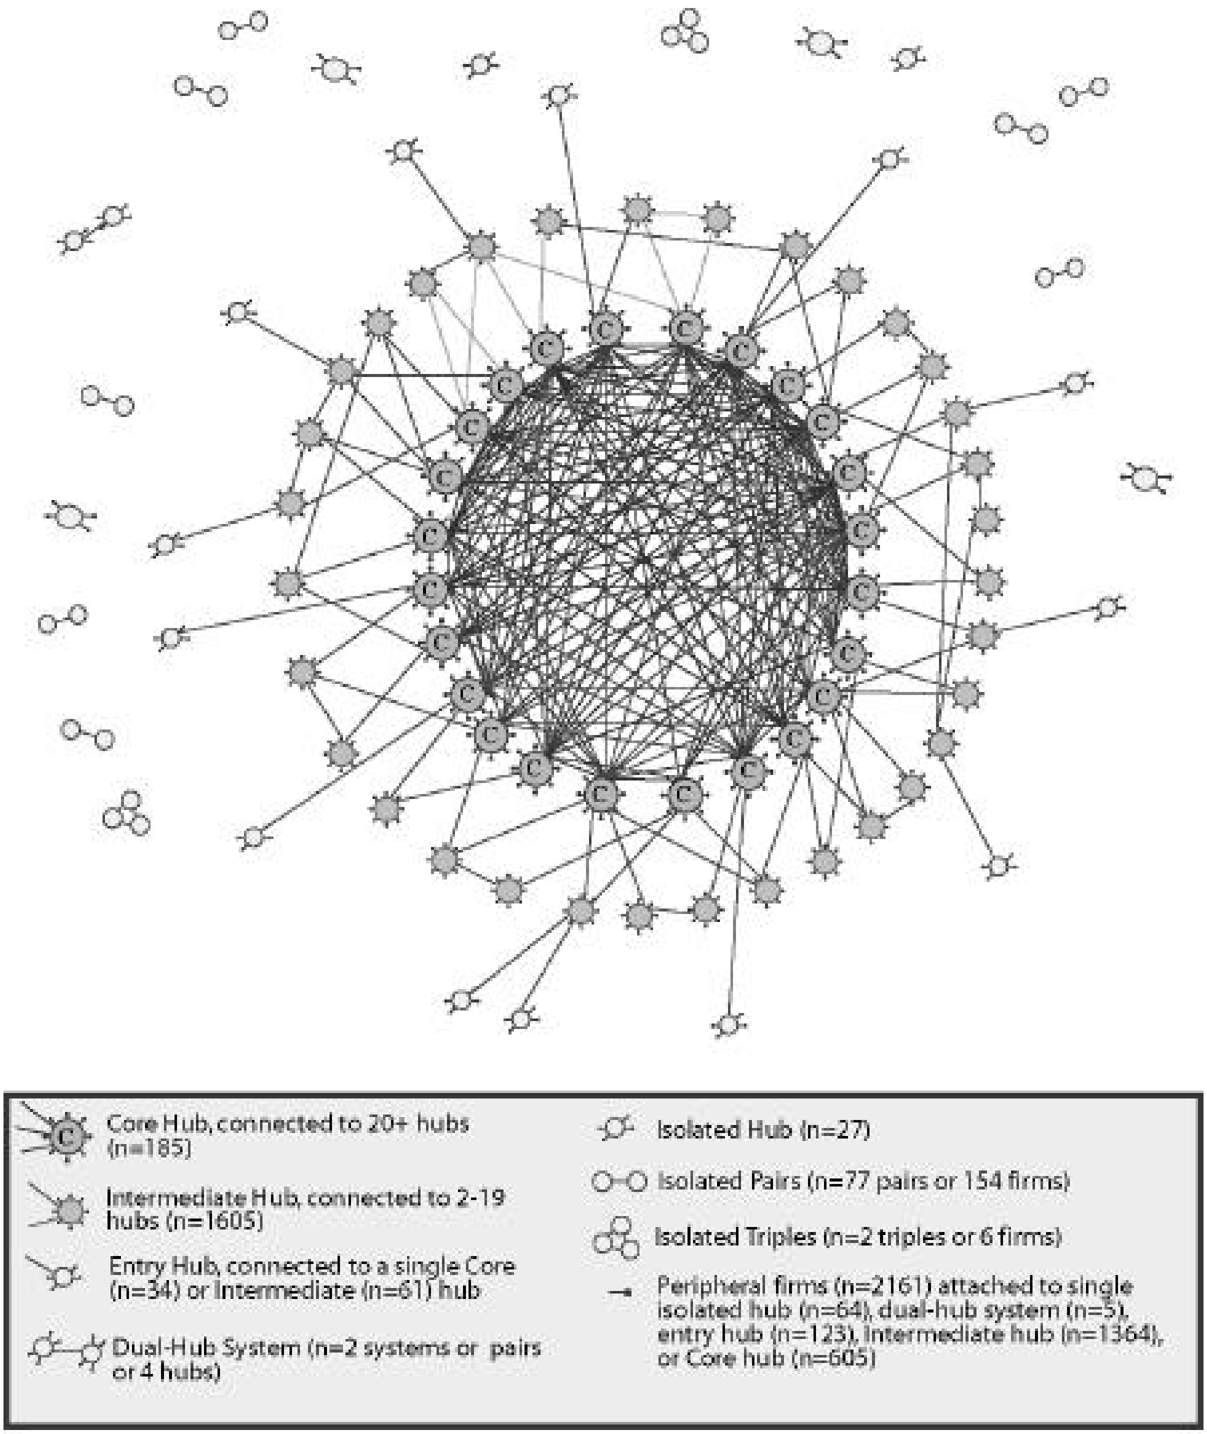

Note: Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, which includes 4,231 unique entities (surviving parents as of year-end 2001).