Strategic Alliances: Bridges Between "Islands of Conscious Power"☆

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 90089-0804

Received 16 August 2004

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). We use this model to generate a large number of strategic alliance possibilities, including simple cooperative arrangements (coopetition), strategic divestitures, total divestitures, licensing agreements, and royalty agreements. We show that any of these possible strategic alliances could be optimal in a one-shot interaction, and that the possibility of future interaction between the parties changes the nature of these alliances in ways that alter the optimal organizational form and the management challenges faced.

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).

Strategic Alliances: Bridges Between "Islands of Conscious Power"

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 often cannot; and the importance of relationships in the successful implementation of these alliances. Standard ideas – such as inefficient hold-ups motivated by specific investments and inadequate investments motivated by bargaining over returns – 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 coopetition (where non-integrated parties compete and cooperate at the same time[1]), acquisitions (where one parent acquires and controls the joint project), total divestitures (where an autonomous entity is created to pursue the joint project without parental ownership or direct control), strategic divestitures (where an autonomous entity is create to own part of the project and to pursue it with the other parent), licenses (where some decision rights are extricated from their native assets and reallocated to new parties), and royalty contracts (where some payoff rights are extricated from their native assets and reallocated to new parties).

We see this paper as a contribution to the literature that seeks to describe and explain what Coase (1992) called the "institutional structure of production." For at least three-quarters of a century, the dominant view of this institutional structure has seen firms as "islands of conscious power … like lumps of butter coagulating in a pail of buttermilk" (Robertson, 1930: 85) —that is, firm boundaries are sharp, and within these boundaries the exchange transactions of markets are replaced by the authority transactions of firms. For example, Coase (1937: 388) quotes Robertson approvingly and then elaborates that, "Within a firm … market transactions are [replaced by] the entrepreneur–coordinator, who directs production." Simon's (1951) model of the employment relationship continues this tradition, as do Williamson's (1975, 1985) work on fiat within firms, Masten's (1988) "Legal Basis for the Firm," the property-rights model of Grossman-Hart-Moore, and the incentive-system model of Holmstrom-Milgrom-Tirole.[2]

While this "islands" view has been productive both theoretically and empirically, various dissenting and complementary views have occasionally surfaced. Even in 1937, Coase cautioned that "it is not possible to draw a hard and fast line which determines whether there is a firm or not "(p.392), and Alchian and Demsetz (1972) famously asserted that employers have

[1] See Brandenburger and Nalebuff (1996). The term "coopetition" was originally coined by Novell founder Ray Norda.

[2] See Grossman and Hart (1986), Hart and Moore (1990), and Hart (1995) on the property-rights model and Holmstrom and Milgrom (1991, 1994), Holmstrom and Tirole (1991), and Holmstrom (1999) on the incentive-system model.

no more authority over their employees than customers have over their grocers. In addition to theoretical discussions that have dissented from the focus on authority as the key to defining *what a firm is*, a complementary empirical strand of the literature has provided intriguing evidence about *what exists besides firms*. For example, Richardson (1972) gave a rich description of "industrial activity that our simple story, based as it is on a dichotomy between firm and market, leaves out of account … the dense network of co-operation and affiliation by which firms are inter-related." Cheung's (1983) description of contractual structures between firm and market, Eccles's (1985) work on quasi-firms, and Powell's (1990) discussion of networks ("Neither Market Nor Hierarchy") all continue this tradition of empirically grounded criticism of the simple dichotomy between markets and firms.

The many governance structures between firms and markets are sometimes summarized as "hybrid" governance structures; see Williamson (1985; 1996) and Menard (2004). One particular strand of this hybrids literature is especially relevant to our purposes: the (largely empirical) work in which firms have fixed boundaries but pass decision rights across these boundaries by contract. Early work along roughly these lines includes Cheung (1983), Goldberg and Erickson (1987), Joskow (1985), Masten and Crocker (1985), and Palay (1984). More recently, Arruñada, Garicano, and Vazquez (2001), Bajari and Tadelis (2001), Bidwell (2004), Elfenbein and Lerner (2003), Kaplan and Stromberg (2003), Klein (2000), Lafontaine and Masten (2002), Lerner and Merges (1998), Robinson and Stuart (2002), and Ryall and Sampson (2002) can be seen as analyzing contractual movements of decision rights across fixed firm boundaries. These contracts are the "bridges" in our title: firms may be islands, and the boundaries of these islands sometimes shift (via changes in asset ownership), but a useful map of the industrial terrain must include the "dense network of [bridges] by which firms are inter-related."

In this paper, we offer additional evidence on the importance of contractual bridges linking firms. More importantly, we develop a theoretical framework that captures a wide range of governance structures beyond the simple dichotomy between markets and firms. Our theoretical framework includes assets, contracts, and other hybrids. To do so, we expand on the Grossman-Hart-Moore (GHM) approach, as follows.

In the GHM approach, asset ownership conveys rights to make decisions ex post that were not contracted on ex ante. That is, ownership conveys the rights to residual decisions, or residual-decision rights. In the contractual literature, in contrast, we have seen decision rights moved across fixed firm boundaries (*i.e.*, without changing asset ownership). These decision rights can be extracted (or "alienated") from their native assets. Other decision rights cannot be alienated from their native assets. These inalienable decision rights are the residual – those that are left attached to the asset after all alienable decision rights have been removed (and perhaps reallocated). Our theoretical framework thus distinguishes between (and incorporates both) the residual-decision rights from the GHM approach and the residual decision-rights from the contracting approach.

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 (2004), 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 (and other governance structures) can either enhance or inhibit the parties' effort to use their relationship towards this end. We explain below how this 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.

The paper is organized as follows. We begin in Section 2 with an analysis of nearly 12,500 biotechnology alliances in the pharmaceutical and biotechnology industries. We document a dense network of contractual bridges among firms in the industry, and show that a plethora of governance structures is used for a wide range of purposes.

Section 3 introduces our theoretical framework. We define "governance structures" as an allocation of decision rights and payoff rights to parties through either contracts or asset ownership. We allow for a variety of observed governance structures for coordinating activities between firms, including acquisitions, coopetition, divestitures, licensing agreements, and royalty contracts.

Section 4 analyzes the efficient governing structure in one-shot transactions. We show that each of the governances structures enumerated in Section 3 can be optimal under some circumstances, but that (in general) none of them is first-best. Section 5 considers "relational governance," in which ongoing relationships help parties achieve efficient adaptation as states of the world are realized. 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 not only that relationships 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 4 and 5 restrict attention to governance structures with "unique control" (i.e., any given asset, decision right, or payoff right is owned by exactly one firm). An important type of strategic alliance—the joint venture—is an example of a governance structure that does not exhibit this type of unique control. Because our model emphasizes contracting problems ex post, joint control raises complex theoretical issues that have gone unnoticed in the property-rights approach to joint ownership. While a thorough analysis of joint control is beyond the scope of this paper, we offer a brief discussion of these issues before concluding in Section 6.

2. Bridges in Biotechnology

In this section, we offer suggestive evidence on the importance of bridges in the pharmaceutical and biotechnology industries. Our evidence comes from data collected by Recombinant Capital (specialists on biotechnology alliances) on nearly 12,500 publicly disclosed contracts between pharmaceutical and biotechnology firms from 1973 to 2001.

Table 1 shows the number of contracts (and the number of partners) for the top 12 pharmaceuticals and top 12 biotechnology firms, where "top" is defined by the number of contracts 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 publicly disclosed contracts. In short, a few firms are doing lots of the alliances, raising the question: with whom?

Figure 1 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 the remaining 23 firms. To the contrary, the 24 firms in Table 1 had contractual 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.

This evidence 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. We sketch such an industrial structure in Figure 2; the figure is stylized but gives some of the actual summary statistics in its legend.

The existence of this "dense network of [bridges] by which firms are inter-related" raises a question: what do these bridges do? The Recombinant data shed some light on this question as well, which we summarize in Table 2. The columns depict the prevalence of various observed governance structures used for coordinating activities across firms, including licensing agreements, investments, mergers, and acquisitions. The rows depict the activities that are being coordinated, including development, research, manufacturing, marketing, collaboration, or supply. Most of the observed alliances involve research and/or development activities (55%) and are structured as licensing agreements (66.5%).

Overall, the evidence confirms that a plethora of governance structures are used for a wide range of purposes. Furthermore, the biotechnology industry is far from unique in this regard. For example, the indirect ties between partners in alliances, joint ventures, and the like are again extremely dense in the internet sector and the automotive industry.[3] Based on suggestive evidence of this kind, we turn next to theoretical analyses of how governance structures aid or impede ongoing relationships.

3. Defining Governance Structures

3a: Economic Environment

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

[3] For an example of networks in the internet sector, see http://www.orgnet.com/netindustry0104.gif.

together, in coordinated fashion. The decisions about how to use a and b in the joint project are denoted $d_a \in D_a$ and $d_b \in D_b$ respectively. Coordinated use of the assets {a, b} produces profits π_a and π_b, both positive; any other uses of {a, b} produce profits of zero.[4] We assume, initially, that the profits π_a and π_b accrue to the owners of assets a and b.

In addition to producing profits π_a and π_b, 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.) We assume that the payoffs π_A and π_B are private benefits, and are observable but not verifiable. All of this is summarized in Figure 3.

Figure 3
Economic Environment

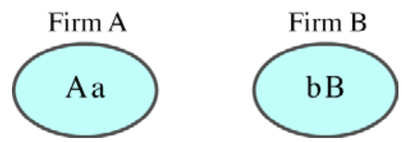

If the holders of both d_a and d_b agree, coordinated use of {a,b} yields:

Direct profits: π_a, π_b alienable
Spillovers: π_A, π_B private benefits

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 $\pi_B(s)$ have finite support of $\overline{\pi_A}$, $\underline{\pi_A}$, and $\overline{\pi_B}$, $\underline{\pi_B}$ respectively, and are drawn from the joint distribution $F(\pi_A,\pi_B)$.

We define assets a and b to be a combination of payoff rights π_a and π_b and decision rights D_a and D_b. In Section 3c below, we assume that the decision rights are inextricably linked to the payoff rights, so that decision rights and payoff rights can be transferred only as a bundle, by transferring ownership of the assets. In Section 3d, in contrast, we introduce the possibility of transferring either the decision rights or the payoff rights separately, by contract.

---

[4] More formally, one could imagine a state variable σ such that coordinated use of {a, b} in state σ means that the decisions $d_a*(\sigma)$ and $d_b*(\sigma)$ were chosen. We suppress the state σ for notational simplicity.

The timing of the model is as follows. Initially, the parties choose a "governance structure," which is an allocation of decision rights and payoff rights, possibly allocated through asset ownership. This allocation may be accompanied by state-independent side payments. Next, the state of the world (s) is revealed, after which the parties make decisions. Finally, after decisions are made, payoffs are realized by the parties holding the payoff rights.

We assume that the opportunities presented by the revelation of s are fleeting, in the sense that decisions must be made immediately after the state is revealed: there is no time to re-contract on either decision rights or payoff rights between the time that s is revealed and decisions must be made.

We also make the important assumption that decisions (as distinct from decision rights) are not contractible either before or after the state of the world is known. As mentioned in the Introduction, this issue of contracting problems ex post was emphasized to us by practitioners. This assumption rules out the possibility of *ex post* renegotiation to achieve efficient decision making. Whoever holds the decision right *ex ante* will make the decision that is in his best interest *ex post*, and no Coasian bargaining will occur to achieve *ex post* efficiency.

3b: First Best Implementation and Mergers

We begin our discussion of feasible governance structures by defining first-best decision making in this model. Simply put, the project should be implemented whenever the total payoffs are positive, that is whenever $\pi_A(s)+\pi_B(s)+\pi_a+\pi_b>0$. This could be achieved by merging firms A and B, thereby internalizing the spillover effects.

In this paper, we are interested in governance structures that do not involve the transfer of ownership of either A or B. While we acknowledge that mergers are important empirically, and are often used to internalize spillovers between firms, we believe that they are not efficient solutions in many of the situations that we attempt to model in this paper, where the assets a and b (and the profits and spillovers associated with their use) are small compared to assets A and

B. We therefore assume that it is not worth combining A and B to solve the spillover problems associated with a and b. Our argument rests on the assumption of some (unmodeled) costs associated with combining assets. We assume that these costs – which might include the costs of integrating control systems and cultures, overcoming communications barriers, and the costs of moving decision-makers farther from the consequences of their actions – are small compared to the value of assets being combined. Therefore, while we assume that these costs can be ignored in governance structures that involve combining assets a and b, they would loom large if we considered integrating the parent firms A and B solely to achieve coordinated use of assets a and b.

3c: Governance Structures Involving Asset Ownership of a and b

We now analyze feasible governance structures that involve transferring (or not transferring) ownership of assets a and b. There are four possible arrangements, which we label coopetition, an acquisition, a total divestiture, and a strategic divestiture.

Coopetition involves separate ownership of assets a and b by parties A and B: either A owns a and B owns b, or they swap a and b, possibly with a side payment. In one-shot interactions between the parties, this means that the project will be implemented only when it is in each of their interests to proceed. Of course, in a repeated interaction, coopetition can induce different implementation patterns.

In an *acquisition*, party A (for instance) owns both a and b. In this case, the externalities imposed on party B will be irrelevant to the implementation decision, at least in one-shot interactions. In a *total divestiture*, both A and B divest a and b to a third party, C. Since the project payoffs π_a and π_b are always positive, party C will always proceed with the project, ignoring (in one-shot interactions) the spillovers on A and B. In a *strategic divestiture*, one party (say, B) divests b to a third party. A strategic divestiture could be accompanied by an asset swap.

The four possible governance structures that involve the transfer of assets a and b are shown in the table below. Recall that ownership of an asset confers both the decision right over that asset and the payoff right flowing from that asset.

Table 3: Governance Structures Involving Asset Ownership

Governance Structure	Party A holds:	Party B holds:	Party C holds:
Coopetition	(d_a, π_a)	(d_b, π_b)	-
	(d_b, π_b)	(d_a, π_a)	-
Acquisition	(d_a, π_a), (d_b, π_b)	-	-
	-	(d_a, π_a),(d_b, π_b)	-
Total Divestiture	-	-	(d_a, π_a), (d_b, π_b)
Strategic Divestiture	(d_a, π_a)	-	(d_b, π_b)
	-	(d_a, π_a)	(d_b, π_b)
	(d_b, π_b)	-	(d_a, π_a)
	-	(d_b, π_b)	(d_a, π_a)

3d: Governance Structures Involving Contracting

We now explore the possibility that certain decision rights and payoff rights can be allocated by contract rather than by asset ownership.[5] We do this not only to examine the theoretical possibilities that such contracts introduce, but also because our discussion with practitioners and our reading of the empirical literature suggest that contracts that allocate decision rights and payoffs are a common feature of strategic alliances.

We maintain the assumption throughout that *decisions* are not contractible. We assumed above that decision rights over an asset could only be transferred by transferring ownership of

5 For our purposes, allocating decision rights through contracting allows us to separate decision rights from payoff rights (recall that we defined assets as an inseparable coupling of decision and payoff rights). More broadly, Maskin and Tirole (1999) show that theories of integration through asset ownership are formally equivalent to theories of contracting.

the asset. However the *decision right* over whether or not to use an asset could be allocated to another party, without transferring ownership. Consider, for instance, the right to market a product that is the result of a development effort by two firms. These marketing rights could be licensed to either party, giving them sole right to decide whether and how to market the product.

We define the transfer of a decision rights (without the transfer of payoffs) to be a license agreement. Such contractual arrangements offer many new governance structures. In Table 4 below, we only consider structures that give parties different incentives from those induced by governance structures shown in Table 3. (We also ignore license agreements that give only decision rights to party C, who would then have decision rights with no payoffs.) The first row shows an arrangement that gives A all of the decision rights over the project, but none of the profits (other than the inalienable spillovers).

Table 4: Governance Structures Involving Contracting Over Decision Rights

Governance Structure	Party A holds:	Party B holds:	Party C holds:
License Agreement	d_a, d_b	π_a, π_b	-
	π_a, π_b	d_a, d_b	-
	d_a	-	$(d_b,\pi_b),\pi_a$
	-	d_b	$(d_a,\pi_a),\pi_b$

Contractual transfers of payoff rights from a particular asset without the transfer of decision rights is more complex. We consider two possible assumptions. The first assumption, which is parallel to our assumption about decision rights, is that payoffs themselves are not contractible, but payoff rights are. That is, it might be possible to transfer the payoffs from an asset to another party in their entirety, but since the actual amount of the payoff is not contractible, payoffs are not divisible. We imagine that such a situation might arise when it is possible for one party (say A) to siphon-off cash flows from a project. Thus if a contract promised party B 10% of the profits, these profits would mysteriously fail to materialize.[6] Only by giving all of the payoffs to the B, along with the rights to siphon, can payoffs be transferred between parties. If one makes the reasonable assumption that these siphoning rights are residual control rights then transferring payoff rights is equivalent to transferring ownership in the asset, while licensing back the implementation decision rights. Since we have already considered this case above, we will not analyze governance structures that allocate payoff rights.

Another possible assumption is that payoffs themselves are contractible. This assumption allows for "royalties" to be paid to any party when a project is implemented. Our discussions with practitioners, our reading of the empirical literature, and our examination of the data presented in the introduction, suggest to us that such royalties are often a feature of strategic alliances. While a full analysis of how royalty rates could be added to the set of governance structures enumerated above is beyond the scope of this paper, we will suggest one possibility. The asset ownership specified in the Coopetition structure, combined with a royalty rate which allows all possible allocations of the total payoffs ($\pi_a+\pi_b$) between A and B is one possible type of Royalty contract. It is shown in Table 5 below.

Table 5: Governance Structures Involving Contracting Over Payoff Rights

Governance Structure	Party A holds:	Party B holds:	Party C holds:
Royalty Contract	$d_a, \alpha(\pi_a+\pi_b)$	$d_b, (1-\alpha)(\pi_a+\pi_b)$	-

4. Efficient Governance of One-Shot Transactions

Having defined a set of governance structures—that is, an allocation of decision rights and payoff rights to the parties—in Section 3, we now examine which of these structures will generate the highest surplus when the parties are interacting in a one-shot transaction. This

6 For example, Hollywood is replete with examples of people who have contracted for a share of the "net profits" from a movie which turn out to be small or zero even after the movie is highly successful (Weinstein, 1998).

requires calculating the expected value of each governance structure. We will show that each of the governances structures enumerated in Section 3 can be optimal under some circumstances, but that (in general) none of them is first-best.

Much of the analysis in this section will be graphical: we will show how each governance structure results in a different pattern of implementation across possible states of the world. This will allow us to prove both that none of the structures is first best, and that each can be second best.

Figure 4

First-best implementation

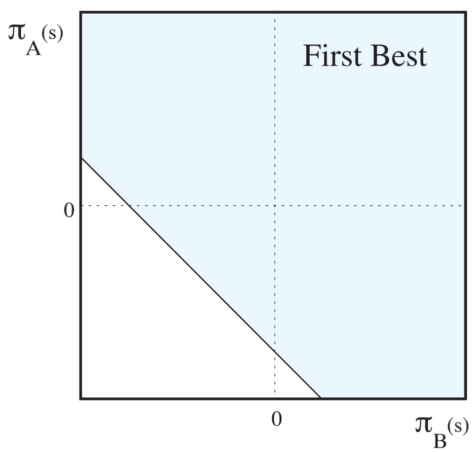

The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The first-best area is defined as outcomes for which $\pi_A(s) + \pi_B(s) + \pi_a + \pi_b > 0$.

Since the pair $\{\pi_A(s),\pi_B(s)\}$ completely characterizes the state in this model, we will show each governance structure as a shaded region of implementation in the (π_A,π_B) plane. Figure 4 shows first best implementation, defined as all realizations of the state s such that $\pi_A(s) + \pi_B(s) + \pi_a + \pi_b > 0$. Expected surplus from first best implementation equals:

$$V^{FB} = \int_{\underline{\pi_A}}^{-\underline{\pi_B}+\pi_a+\pi_b} \int_{-x+\pi_a+\pi_b}^{\overline{\pi_B}} (y + x + \pi_a + \pi_b) f(x,y)\,dx\,dy + \int_{-\underline{\pi_B}+\pi_a+\pi_b}^{\overline{\pi_A}} \int_{\underline{\pi_B}}^{\overline{\pi_B}} (y + x + \pi_a + \pi_b) f(x,y)\,dx\,dy$$

Figure 5

Implementation under Coopetition

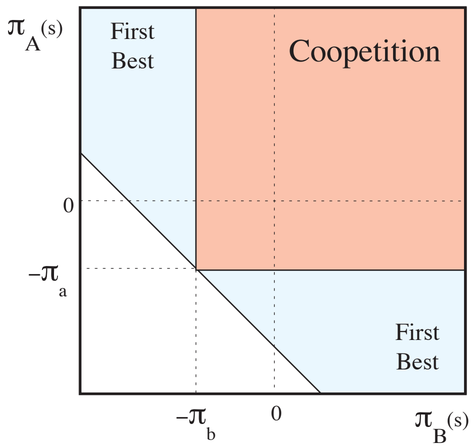

The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The first-best area is defined as outcomes for which $\pi_A(s) + \pi_B(s) + \pi_a + \pi_b > 0$. The Coopetition area is defined as outcomes where $\pi_A(s) + \pi_a > 0$ *and* $\pi_B(s) + \pi_b > 0$.

Under one-shot Coopetition, each party will choose to implement only when the sum of his spillovers and his payoff from the project is positive. Thus party A will choose to implement whenever $\pi_A(s)+\pi_a>0$. B will only choose to implement when $\pi_B(s)+\pi_b>0$. This leads to the implementation pattern shown in Figure 5. Surplus under one-shot Coopetition is:

$$V^{CO} = \int_{-\pi_a}^{\overline{\pi_A}} \int_{-\pi_b}^{\overline{\pi_B}} (x + y + \pi_a + \pi_b) f(x, y) dx\, dy$$

As is clear from the shading, Coopetition is inefficient since there are states in which the project should be implemented but is not. This inefficiency results from the inability of the parties to bargain *ex post*, and devise a set of side payments (in the one-shot game) that would lead to efficient implementation.

Figure 6

Implementation under Acquisitions

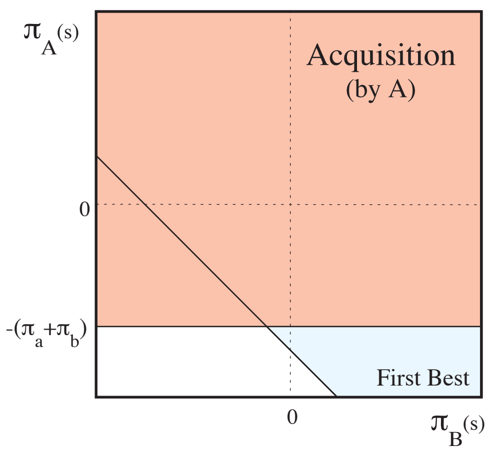

The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The first-best area is defined as outcomes for which $\pi_A(s) + \pi_B(s) + \pi_a + \pi_b > 0$. The Acquisition area is defined as outcomes where $\pi_A(s) + \pi_a + \pi_b > 0$.

In an Acquisition (by party A), A has all of the decision rights, and all of the project payoffs. But he ignores party B's spillovers. Thus he will choose to implement the project whenever $\pi_A(s) + \pi_a + \pi_b > 0$. This pattern of implementation is shown in Figure 6. Surplus in an Acquisition by A is:

$$V^{ACQ(A)} = \int_{-(\pi_a+\pi_b)}^{\overline{\pi_A}} \int_{\underline{\pi_B}}^{\overline{\pi_B}} (y + x + \pi_a + \pi_b) f(x,y)\,dx\,dy$$

An Acquisition is also not first best, as is clear from the figure: it leaves projects unimplemented in states where they should be, and implements projects when they should not be. In addition, either an Acquisition or Coopetition could be the second-best one-shot governance structure. Note than an Acquisition implements in states when Coopetition does not, but it also fails to implement in some when Coopetition does. Whether an Acquisition or Coopetition is more efficient depends on the relative likelihood of these outcomes.

Figure 7

Implementation under Total Divestiture

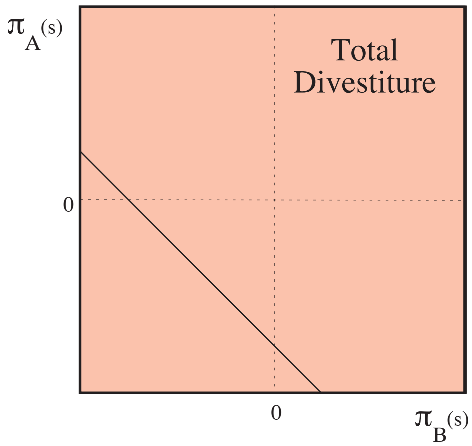

The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The Total Divestiture area is defined as outcomes where $(\pi_a + \pi_b) > 0$ (which, given our assumptions, is always true).

A Total Divestiture implements in all states, since the payoffs π_a and π_b are always positive, and so party C will always choose to go ahead with the project. This implementation pattern is shown in Figure 5, and results in surplus

$$V^{TD} = \int_{\underline{\pi_A}}^{\overline{\pi_A}} \int_{\underline{\pi_B}}^{\overline{\pi_B}} (y + x + \pi_a + \pi_b) f(x,y) dxdy$$

Here again, a Total Divestiture is not first best, since it implements in many states where implementation is not optimal. However, it also implements whenever implementation is efficient. Thus, it could be the second best governance structure.

A Strategic Divestiture can result in four possible implementation patterns in the one-shot game, as shown in Table 3. Figure 8 below shows only one: the divestiture by B of b to C. In this case, C will always choose to go ahead with the project (because $\pi_b > 0$) and A will go ahead with the project only if $\pi_A(s) + \pi_a > 0$. The surplus is:

$$V^{SD(B)} = \int_{-\pi_a}^{\overline{\pi_A}} \int_{\underline{\pi_B}}^{\overline{\pi_B}} (x + y + \pi_a + \pi_b) f(x,y) dx\, dy$$

Similarly, the surplus associated with a divestiture by A of a to C is given by:

Figure 8

Implementation under Strategic Divestiture

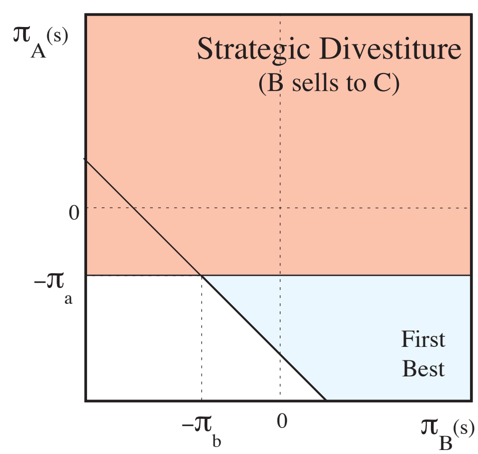

The figure depicts a strategic divestiture in which B sells b to C. The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The Strategic Divestiture area is defined as outcomes where $(\pi_A + \pi_a) > 0$.

$$V^{SD(A)} = \int_{\underline{\pi_A}}^{\overline{\pi_A}} \int_{-\pi_n}^{\overline{\pi_B}} (x + y + \pi_a + \pi_b) f(x, y) dx\, dy$$

Clearly, these governance structure do not achieve the first best, but could be second best.

A License Agreement, in which B receives the payoffs from both π_a and π_b but "licenses" the decision rights d_a and d_b to A is shown in Figure 9. In this case, A will wish to go ahead with the project whenever $\pi_A(s) > 0$, independent of B's payoffs. The surplus associated with this particular licensing agreement is given by

$$V^{L(A)} = \int_{-(\pi_a+\pi_b)}^{\overline{\pi_A}} \int_{0}^{\overline{\pi_B}} (x + y + \pi_a + \pi_b) f(x, y) dx\, dy$$

Figure 9

Implementation under License Agreement

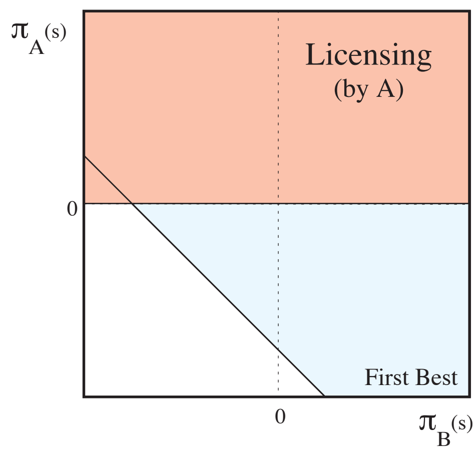

The figure depicts a licensing agreement in which B receives the payoffs from both π_a and π_b but "licenses" the decision rights d_a and d_b to A. The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The Licensing area is defined as outcomes where A's private benefit is positive, $\pi_A > 0$.

A Royalty Contract, in which A makes decisions over d_a and receives a share α of the gains from implementation ($\pi_a + \pi_b$), and B makes decisions over d_b and receives a share ($1 - \alpha$), is shown in Figure 10. Under this governance structure, A will wish to implement the project if $\pi_A(s) + \alpha(\pi_a + \pi_b) > 0$, while B will wish to implement if $\pi_A(s) + (1-\alpha)(\pi_a + \pi_b) > 0$. Since both parties must choose to implement, the resulting surplus is:

Figure 10

Implementation under Royalty Contract

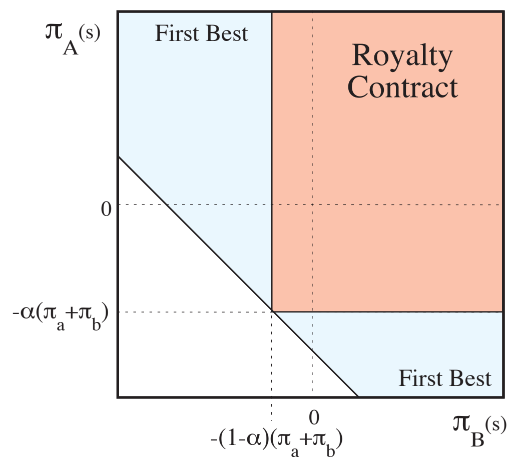

The figure depicts a royalty agreement in which A and B receives α and ($1 - \alpha$) respective shares in the payoffs from implementation, $\pi_a + \pi_b$. The vertical axis shows the private benefit to A, $\pi_A(s)$, and the horizontal axis shows the private benefit to B, $\pi_B(s)$. The Royalty area is defined as outcomes where $\pi_A + \alpha\,(\pi_a + \pi_b) > 0$ *and* $\pi_B + (1-\alpha)\,(\pi_a + \pi_b) > 0$.

$$V^{R(\alpha)} = \int_{-\alpha(\pi_a+\pi_b)}^{\overline{\pi_A}} \int_{-(1-\alpha)(\pi_a+\pi_b)}^{\overline{\pi_B}} (x + y + \pi_a + \pi_b) f(x,y)\,dx\,dy$$

Note that, by varying the royalty payment α in Figure 10, implementation will occur in different states. The implementation will not be first best (since there are first-best regions in Figure 10 that cannot be implemented with a simple royalty payment), but can be second best depending on the relative probabilities of different states.

As is evident from inspection of Figures 4-10, none of these governance structures dominates any other: each is a possible second best structure in a one-shot game. This result demonstrates that, even in a simple set-up like this one, many possible structures could be optimal. Thus the plethora of forms that we see in the data on strategic alliances is not surprising. Optimal governance in the one-shot game requires choosing, from this plethora of possible structures, the one that maximizes total surplus. Thus the optimal one-shot governance structure solves:

$$V^{ONE\text{-}SHOT} \equiv \max[V^{CO}, V^{ACQ}, V^{TD}, V^{SD}, V^{L}, V^{R}]$$

It is important to recall that we have not examined a merger (combining A and B) as a possible solution. The fact that none of these one-shot governance structures is first best suggests that mergers will sometimes be optimal. Specifically, when assets a and b are large relative to A and B, so that the inefficiency of the second best one-shot governance structure is large relative to the costs of integrating A and B, then it may be worth bearing these costs and merging the two firms.

Furthermore, the fact that none of these governance structures is first best in the one-shot game suggests that relationships—which allow self-enforcing relational contracts to solve the *ex post* bargaining problem and achieve efficient adaptation to the state-of-the-world—could be efficient. In Section 5, we turn to the analysis of such relational contracts.

5. Relational Governance

In this section, we begin by presenting evidence that ongoing relationships between firms are potentially important. Turning to theory, we then describe our approach to modeling relational contracts, and show that relational contracting in the strategic alliances presented in Section 3 can be modeled as a special case of a more general model of relational governance.

5a. 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).

The Recombinant Capital database summarized in Section 2 provides evidence on 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.[7] 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 6 presents evidence on repeat alliances between the same partners. In the Recombinant Capital database, most pairs of firms (9,462) do only one 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, as discussed in Section 2 and depicted in Table 1 and Figures 1 and 2, the dense

---

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

network of inter-relations in the biotechnology industry clearly suggests indirect ties between firms in the industry.

5b. Modeling Relational Governance

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 3, we now add ongoing relationships. In our model, relationships reduce the ex post inefficiencies that arise under the static governance analyzed in Section 4.

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, 2002), and Halonen (2002) enriched static models in the Grossman-Hart-Moore tradition 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 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 enriching the feasible set of sharing rules and hence improving ex ante incentives. Instead, our focus is on the complementary problem of ex post adaptation.

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).

We analyze trigger-strategy equilibria: if any firm reneges, 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 negotiating a change in governance structure at the end of the present period.)

In Baker, Gibbons, and Murphy (2004), henceforth BGM, we derive a simple but powerful necessary and sufficient condition for whether a given decision rule can be supported as a relational contract (*i.e.*, exists as a trigger-strategy equilibrium in the repeated game) under a given governance structure. Our goal in this section is to apply this condition, in an informal and intuitive way, to the governance structures defined and analyzed in Sections 3 and 4, in order to motivate the following three insights: in a relational environment, (1) different governance structures create different managerial challenges, (2) the efficient governance structure minimizes these managerial challenges, and consequently (3) the efficient governance structure typically differs from that in a static environment. To conduct this informal discussion, however, we require a few formal preliminaries; see BGM for details.

An Enriched Economic Environment: In BGM, we analyze a more general model than Section 3's, which allows for arbitrary numbers of (a) parties, (b) assets (*i.e.*, inseparable bundles of decision and payoff rights), (c) alienable decision rights not linked to any payoff rights, (d) alienable payoff rights not linked to any decision rights, and (e) inalienable private decision rights. Formally, we assume there are I parties, J assets, K decision rights not linked to payoff rights, and M payoff rights not linked to decision rights. Party $i \in I$ receives inalienable private benefit π_i and controls inalienable decision rights $d_i \in D_i$. Asset $j \in J$ consists of the inalienable pair of decision rights $d_j \in D_j$ and payoff rights π_j. Decision right $k \in K$ is not linked to any payoff right and is denoted $d_k \in D_k$. Payoff right $m \in M$ is not linked to any decision rights and is denoted π_m. We continue to denote the state by s, drawn from the finite set S according to the probability density f(s). We write **d** for the vector of decisions, chosen from a set D with domain $D \equiv \prod_{i \in I} D_i \times \prod_{j \in J} D_j \times \prod_{k \in K} D_k$. These decisions affect both the inalienable private benefits and the payoffs associated with alienable payoff rights and alienable assets.

As above, we define a "governance structure," $g \in G$, as an assignment of assets, decision rights, and payoff rights across parties. Let G be the set of feasible governance structures. We define $J(i,g) \subset J$ as the assets held by party i under governance structure g, $K(i,g) \subset K$ as the decision rights (not attached to payoff rights) held by party i under governance structure g, and $M(i,g) \subset M$ as the payoff rights (not attached to decision rights) held by party i under governance structure g.

We define $\pi_{ig}(\mathbf{d},s)$ as the total payoff to party i under governance structure g in state s; this total payoff includes private benefits π_i, plus payoffs from assets $j \in J(i,g)$, plus payoffs from payoff rights not associated with decision rights $m \in M(i,g)$:

$$(1) \qquad \pi_{ig}(\mathbf{d},s) \equiv \pi_i(\mathbf{d},s) + \sum_{j \in J(i,g)} \pi_j(\mathbf{d},s) + \sum_{m \in M(i,g)} \pi_m(\mathbf{d},s).$$

Similarly, we define D_{ig} as the decision space for party i under governance structure g; this decision space includes inalienable decision rights D_i, and alienable decision rights $k \in K(i,g)$, and decision rights associated with assets $j \in J(i,g)$:

$$(2) \qquad D_{ig} \equiv D_i \times \prod_{j \in J(i,g)} D_j \times \prod_{k \in K(i,g)} D_k.$$

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 party i, $\mathbf{d}_{ig}^{NE}(s)$ solves:

$$(3) \qquad \max_{\mathbf{d}_{ig} \in D_{ig}} \pi_{ig}((\mathbf{d}_{ig}, \mathbf{d}_{-ig}^{NE}(s)), s).$$

The expected payoff to Party i under the static (i.e., spot) governance structure g is then

$$(4) \qquad V_{ig}^{ST} \equiv E_s\left[\pi_{ig}(\mathbf{d}_g^{NE}(s), s)\right].$$

and we write V_g^{ST} for the total expected surplus, $V_g^{ST} \equiv \sum_{i \in I} V_{ig}^{ST}$. The optimal (second-best) governance structure solves

$$(5) \qquad V^{ST} \equiv \max_{g \in G} V_g^{ST}.$$

Relational Governance in the Enriched Environment: We analyze whether an arbitrary decision rule, $\mathbf{d}^{RC}(s)$, can be supported as a relational contract under a given governance structure. Our task is then to determine whether there exist payment schemes that induce the parties to take these decisions.

Figure 11

Timing of Payments in a Relational Contract

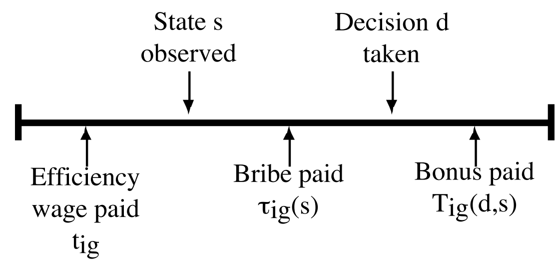

Payments can occur at three different times. First, the payments might be "efficiency wages," denoted by t_{ig} and paid before the state or any decisions are observed. Second, the payments might be "bribes," denoted by $\tau_{ig}(s)$ and paid after the state is observed but before the parties make their 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 11 illustrates the timing of these potential payments within each period, relative to when the state is observed and the decisions are taken. These payments can be positive or negative (i.e., they can be paid to or paid by a given party). We require that these payments balance: $\Sigma_{i\in I}\, t_{ig} = 0$, $\Sigma_{i\in I}\, \tau_{ig}(s) = 0$ for all s, and $\Sigma_{i\in I}\, T_{ig}(\mathbf{d}, s) = 0$ for all **d** and s.

Given a governance structure, there are many reneging constraints that must be satisfied if a given decision rule is to be a relational contract (*i.e.*, a repeated-game equilibrium). Specifically, each Party i must be willing to: (a) pay (or receive) its efficiency-wage payment, t_{ig}; (b) pay (or receive) its bribe, $\tau_{ig}(s)$; (c) take its decisions, $d_{ig}^{RC}(s)$; and (d) pay (or receive) its bonus, $T_{ig}(\mathbf{d}, s)$. To simplify the statements of these reneging constraints, we introduce the following notation:

$\pi_{ig}^{RC}(s) \equiv \pi_{ig}(\mathbf{d}^{RC}(s), s)$	Payoff to Party i (excluding side payments) from relational-contract decisions in state s under governance structure g

$\mathbf{d}_{ig}^{BR}(s) = \arg\max_{\mathbf{d}_{ig} \in D_{ig}} \pi_{ig}((\mathbf{d}_{ig}, \mathbf{d}_{-ig}^{RC}(s)), s)$ — Party i's best response in state s under governance structure g to relational-contract decisions by all other parties

$\pi_{ig}^{BR}(s) \equiv \pi_{ig}((\mathbf{d}_{ig}^{BR}(s), \mathbf{d}_{-ig}^{RC}(s)), s)$ — Payoff to Party i (excluding side payments) from best response in state s under governance structure g, when all other parties take relational-contract decisions

Let r denote the discount rate per period. Then our necessary and sufficient condition for the decision rule $\mathbf{d}^{RC}(s)$ to be supported as a relational contract under governance structure g is the following:

$$(6) \qquad R_g^{RC} \equiv \max_s \left[\Sigma_i \left(\pi_{ig}^{BR}(s) - \pi_{ig}^{RC}(s) \right) \right] \leq \tfrac{1}{r} \left(V^{RC} - V^{ST} \right).$$

In (6), $\pi_{ig}^{BR}(s) - \pi_{ig}^{RC}(s)$ is firm i's reneging temptation in state s under governance structure g. Summing this reneging temptation across firms and finding the state in which this total reneging temptation is maximized yields R_g^{RC}, the maximal total reneging temptation produced by decision rule $\mathbf{d}^{RC}(s)$ under governance structure g. For this decision rule to be a relational contract, there must be enough surplus created from abiding by the decision rule (and its associated payments), relative to the payoffs generated by efficient static governance. The present value of this surplus is the righthand side of (6); it must be large enough to deter all the parties from reneging in every state, hence the maximal total reneging temptation on the righthand side of (6).

Recall that our interest in relational contracts arises because we seek to reduce the ex post inefficiencies that arise under the static governance analyzed in Section 4. Thus, both the left- and the right-hand sides of (6) are positive, for any relational contract of interest. Therefore, whether (6) holds for a given relational contract comes down to the discount rate. Let r_g^{RC} be the discount rate at which (6) holds with equality. For r < r_g^{RC}, the present value of the net surplus from relational-contract decisions (the right-hand side) exceeds the maximal total reneging temptation (the left-hand side), so the decision rule $\mathbf{d}^{RC}(s)$ can be supported as a relational contract under governance structure g.

Insights from Relational Governance: Given (6), we can now revisit the three insights listed 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.

First, under governance structure g, the managerial challenges in state s are to induce firms with $\pi_{ig}^{BR}(s) - \pi_{ig}^{RC}(s) > 0$ to take relational-contract 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 (6) but governance structure g' does not, then the decision rule $\mathbf{d}^{RC}(s)$ can be supported as a relational contract under g but not under g'. For purposes of implementing the decision rule $\mathbf{d}^{RC}(s)$, therefore, g' could be said to be inefficient. We therefore define the efficient governance structure – for purposes of implementing the decision rule $\mathbf{d}^{RC}(s)$ – as the one that can implement $\mathbf{d}^{RC}(s)$ at the highest possible discount rate. Since the right-hand side of (6) is independent of the governance structure, maximizing r_g^{RC} amounts to minimizing the left-hand side:

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

Third, the efficient governance structure in a relational environment is determined by very different considerations than in a static environment. The latter solves (5), or

(8) $$\max_{g \in G} \left\{ E_s \left[\Sigma_i \pi_{ig}(\mathbf{d}_g^{NE}(s), s) \right] \right\}$$

Clearly, (7) and (8) will typically have different solutions. For example, while (8) depends on

Nash equilibrium decisions, (7) depends on relational-contract decisions *and* on best-responses to relational-contract decisions. Also, while (8) is an expectation that depends on the probability distribution across states, (7) depends on only one state – the one where the total reneging temptation is the largest.

Of course, all of this discussion has taken the decision rule $\mathbf{d}^{RC}(s)$ as given and then asked which governance structure can best support this decision rule. In some settings, this is exactly the right approach: if there is any governance structure that can support the first-best decision rule, then one can identify it via the analysis above. But in other settings (namely, those in which the first-best is not feasible for any governance structure), one should optimize not only over the governance structure g but also over the decision rule $\mathbf{d}^{RC}(s)$. For purposes of this paper, however, we restrict attention to implementing the first-best decision rule, as shown in Figure 4. The question then becomes: which of the governance structures defined in Section 3 can support the first-best decision rule as a relational contract at the highest discount rate?

To answer this question, we must first compute r_g^{FB} – the value of r that solves

$$(9) \qquad R_g^{FB} \equiv \max_s \left[\Sigma_i \left(\pi_{ig}^{BR}(s) - \pi_{ig}^{FB}(s) \right) \right] \leq \tfrac{1}{r} \left(V^{FB} - V^{ST} \right)$$

with equality (where the variable definitions are analogous to those given above), and we must then compare r_g^{FB} across the governance structures defined in Section 3 – coopetition, acquisition, total divestiture, strategic divestiture, license agreements, and royalty contracts. This analysis can be done graphically, by comparing the shaded implementation patterns in Figures 5 through 10 to the first-best implementation pattern in Figure 4, as follows.

Under coopetition, for example, inefficiencies arise when one party fails to utilize its asset in coordinated fashion, because the resulting payoff to that party is negative, even though the total payoff would be positive if the assets were used in coordinated fashion (*i.e.*, the resulting positive payoff to the second party more than outweighs the negative payoff to the first). In an extreme case, suppose that $\pi_A(s) = \overline{\pi_A}$ but $\pi_B(s) = \underline{\pi_B}$, where $\overline{\pi_A} + \underline{\pi_B} + \pi_a + \pi_b > 0$ (as in

Figure 5). Then coordinated use would be first-best, but this will not occur under the coopetition governance structure, at least in a static environment. In a relational environment, however, we can use (9) to solve for the highest value of r such that party B is willing to undertake coordinated use in this state. Indeed, this state has the largest reneging temptation for party B, and the reneging temptation in this state for party A is zero, so $\underline{\pi_B}$ is the total reneging temptation in this state. Analogously, $\underline{\pi_A}$ is the total reneging temptation when $\pi_A(s) = \underline{\pi_A}$ but $\pi_B(s) = \overline{\pi_B}$. Thus, the maximal total reneging temptation for the coopetition governance structure is max $\{|\underline{\pi_A}|, |\underline{\pi_B}|\}$, and from this we can compute the highest value of r at which coopetition can support the first-best.

One can compute analogous maximal total reneging temptations for the other governance structures defined in Section 3 and analyzed in Section 4. Consistent with the three insights above, one finds that the maximal temptation arises in different states for different governance structures. For example, under an acquisition, the temptation is for party A to ignore the externality on party B and implement when doing so is not first-best. More importantly, the maximal reneging temptation thus has different magnitudes for different governance structures. The problem of governance-structure choice thus reduces to the analog of (7): find the governance structure that minimizes the maximal total reneging temptation from the first-best decision rule.

To summarize: The evidence in Section 5a 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 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 (as is familiar from repeated-game analyses of oligopoly, for example). Second, and more importantly, the efficient governance structure in a relational environment differs from that in a static environment, because the governance structure affects the parties' reneging temptations in the ongoing relationship.

6. Discussion and Conclusion

6a. Joint Control

Before summarizing the contribution of this paper, we first provide a brief discussion of an important type of strategic alliance that we have not analyzed: the joint venture (JV). As discussed in the introduction, our focus on contracting problems ex post raise novel issues with joint control. For this reason, we view JVs (and other governance structures involving joint ownership) as qualitatively different from the unique-control governance structures analyzed above.

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. But what happens under joint control?

A full answer to this question will require a paper of its own. However, we offer a sketch of an answer by analogy to the theory of rent-seeking. Imagine a joint venture in the simple static model from Section 3c: Firm A owns {A} and Firm B owns {B}, but a new entity is created that owns assets {a, b}. Unlike a total divestiture (see Table 3) 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 amounts of management time and energy.[8] 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 $k_A(s)$ and Firm B $k_B(s)$, then Firm A gains control over the jointly owned assets with probability $q(k_A(s), k_B(s))$ and Firm B gains control with probability $1 - q(k_A(s), k_B(s))$. 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.

Figure 6
Joint Venture

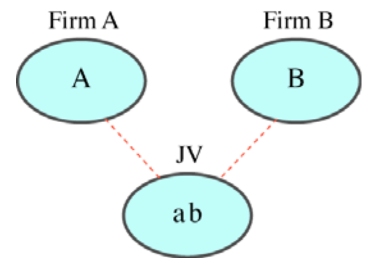

Joint ownership of {a.b} involves contested control, and perhaps contested consumption

The key point of this simple model 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 Sections 3 and 4.[9]

In a relational JV, the parties agree not to contest control, thus avoiding the social costs of these contests. Therefore, in addition to the reneging temptations over side payments and decisions in Section 5, new reneging temptations arise—the temptations to fight for control. 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 (6) will continue to

8 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.

9 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.

hold, where $\pi_{ig}^{BR}(s)$ is modified to include the social costs and private benefits of contest expenditures.

6b. Concluding Remarks

A strategic alliance is a governance structure for coordinating activities between non-integrated firms. Such governance structures are observed in many forms, and constitute the contractual and non-contractual "bridges" between firms in a market economy. While many of these structures have been studied in isolation, there has not been a unified theoretical framework through which to view them simultaneously.

In this paper, we have provided a framework that allows us to begin to understand the plethora of strategic alliances that exist in the world, and to explore when different types of asset ownership and contractual arrangements would be optimal. Our theory is strongly influenced by our conversations with practitioners, who stressed the importance of spillovers, ex post contracting problems, and relationships. We show how and why these factors are important.

Our simple model allows us to characterize and contrast a variety of governance structures, including coopetition, acquisitions, mutual divestitures, strategic divestitures, license contracts, royalty agreements, and (preliminarily) joint ventures. We have identified the inefficiencies associated with each of these governance structures in one-shot interactions. In addition, we have identified the managerial challenges and the efficiency consequences of these governance structures in long-term relationships. 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 often differs from the efficient structure under one-shot governance.

We believe that our approach provides a rich and tractable framework that could be used to analyze the wide variety of governance structures that exist. Of course, formal testing of the framework would demand a significantly more detailed understanding of the actual parameters of the model. How profitable are the joint opportunities? What are the spillovers on the parent companies? How likely are the different possible outcomes? What are the problematic extreme states that might hinder the establishment of a relational contract? To answer these questions probably requires that such empirical studies be conducted at the industry level, with a thorough understanding of the institutional details and the costs and constraints in the industry. As discussed above, such studies (including several in this volume) are enjoying a resurgence. Our goal in this paper is to help this effort along.

References

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

Arruñada, Benito, Luis Garicano, and Luis Vázquez. 2001. Contractual Allocation of Decision Rights and Incentives: The Case of Automobile Distribution. Journal of Law, Economics, and Organization 17: 257-84.

Bajari, Patrick and Steven Tadelis. 2001. Incentives versus Transaction Costs: A Theory of Procurement Contracts. Rand Journal of Economics 32: 387-407.

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., 2002. Relational contracts and the theory of the firm. Quarterly Journal of Economics 117, 39-83.

Baker, G. P., Gibbons, R., Murphy, K. J., 2004. Contracting for Control: Decision Rights, Payoff Rights, and Relationships in Firms, Contracts, and Other Governance Structures. Unpublished working paper.

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

Bidwell, Matthew. 2004. What Do Firms Do Differently? Comparing the Governance of Internal and Outsourced IT Projects. Unpublished manuscript, MIT's Sloan School.

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.

Cheung, Steven. 1983. The Contractual Nature of the Firm. Journal of Law and Economics XXVI: 1-21.

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

Coase, R. H. 1992. The Institutional Structure of Production. American Economic Review 82: 713-19.

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.

Elfenbein, Daniel and Josh Lerner. 2003. Ownership and control rights in Internet portal alliances, 1995-1999. RAND Journal of Economics 34: 356-69.

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.

Goldberg, V., Erickson, J., 1987. Quantity and Price Adjustment in Long-Term Contracts: A Case Study of Petroleum Coke. Journal of Law and Economics XXX, 100-100.

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.

Joskow, Paul. 1985. Vertical Integration and Long-Term Contracts: The Case of Coal-Burning Electric Generation Plants. Journal of Law, Economics, and Organization 1: 33-80.

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

Kaplan, Steven and Per Strömberg. 2003. Financial Contracting Theory Meets the Real World: An Empirical Analysis of Venture Capital Contracts. Review of Economic Studies 70: 281-315.

Klein, B., 2000. The Role of Incomplete Contracts in Self-Enforcing Relationships. Revue D'Economie Industrielle 92: 67-80.

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.

Lafontaine, F., and S. Masten. 2002. Contracting in the Absence of Specific Investments and Moral Hazard: Understanding Carrier-Driver Relations in US Trucking. Unpublished manuscript.

Levin, J. 2003. Relational incentive contracts. American Economic Review 93: 835-57.

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.

Masten, S., __________ and Keith Crocker. 1985. Efficient Adaptation in Long-Term Contracts: Take-or-Pay Provisions for Natural Gas. American Economic Review 75: 1083-93.

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.

Palay, T. 1984. Comparative Institutional Economics: The Governance of Rail Freight Contracting. Journal of Legal Studies 13:265-87.

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.

Robertson, D. H. (1940) Control of Industry. London: Nisbet & Co.

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.

Weinstein, M. I. Profit sharing contracts in Hollywood: evolution and analysis, Journal of Legal Studies 27.

Whinston, Michael. 2003. On the transaction cost determinants of vertical integration/ Journal of Law, Economics, and Organization 19: 1-23.

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

Table 1

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

Panel A *Top 12 Pharmaceutical Firms*	Number of Contracts	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 Contracts	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 1

Strategic Alliances Among the Top 12 Pharmaceuticals and Top 12 Biotechs

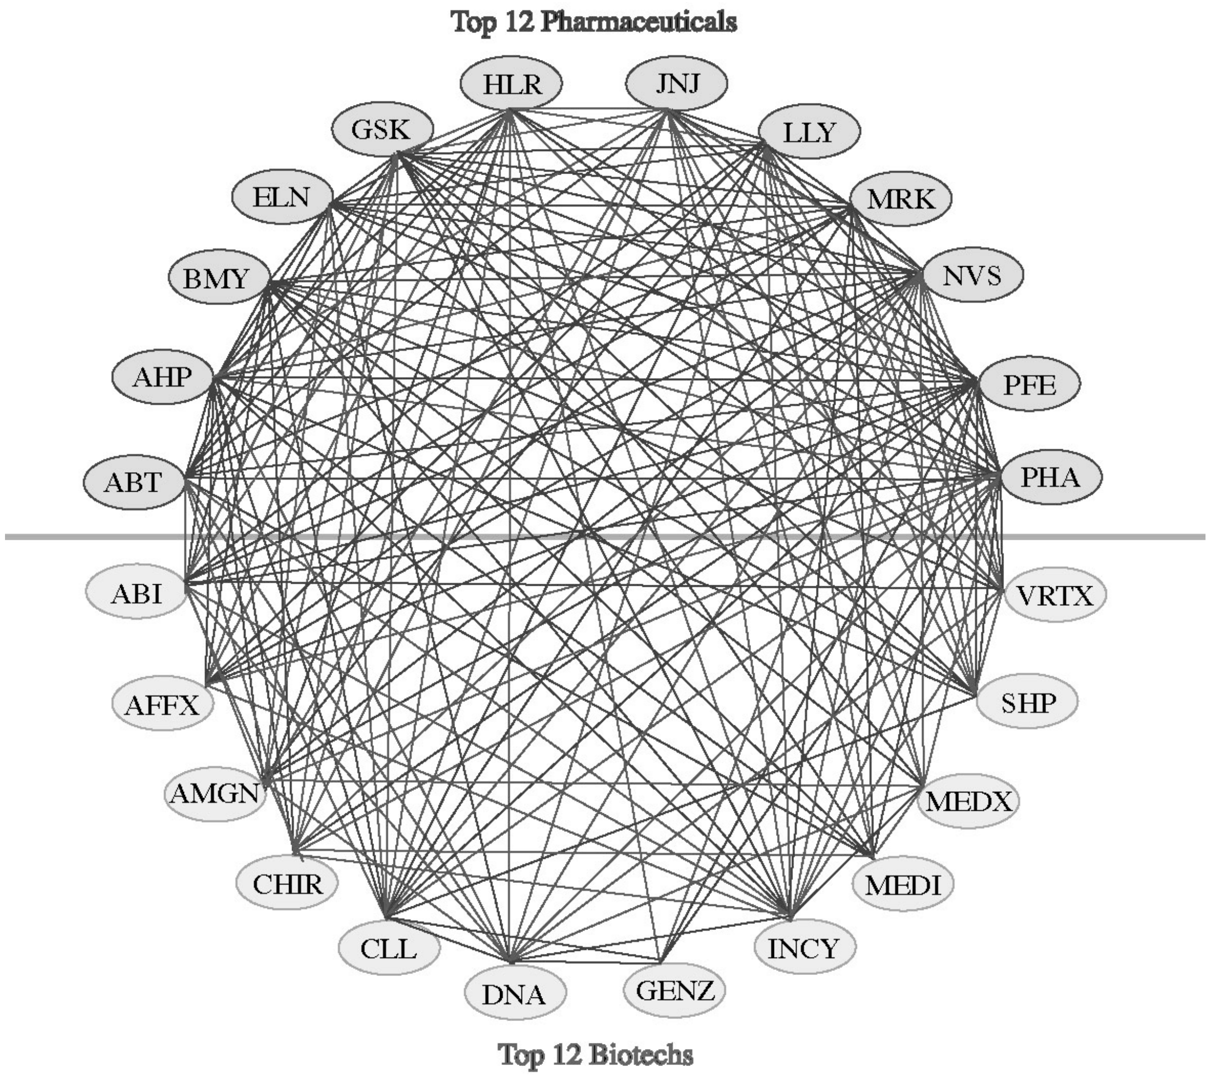

Note: Ticker symbols correspond to companies included in Table 1. Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, based on approximately 12,500 publicly disclosed contracts and arrangements from 1973-2001. Contracts are assigned to the surviving parent as of year-end 2001, regardless of whether the parent was involved in the original arrangement.

Figure 2

Networks in Recombinant Capital Database of Pharmaceutical-Biotech Alliances

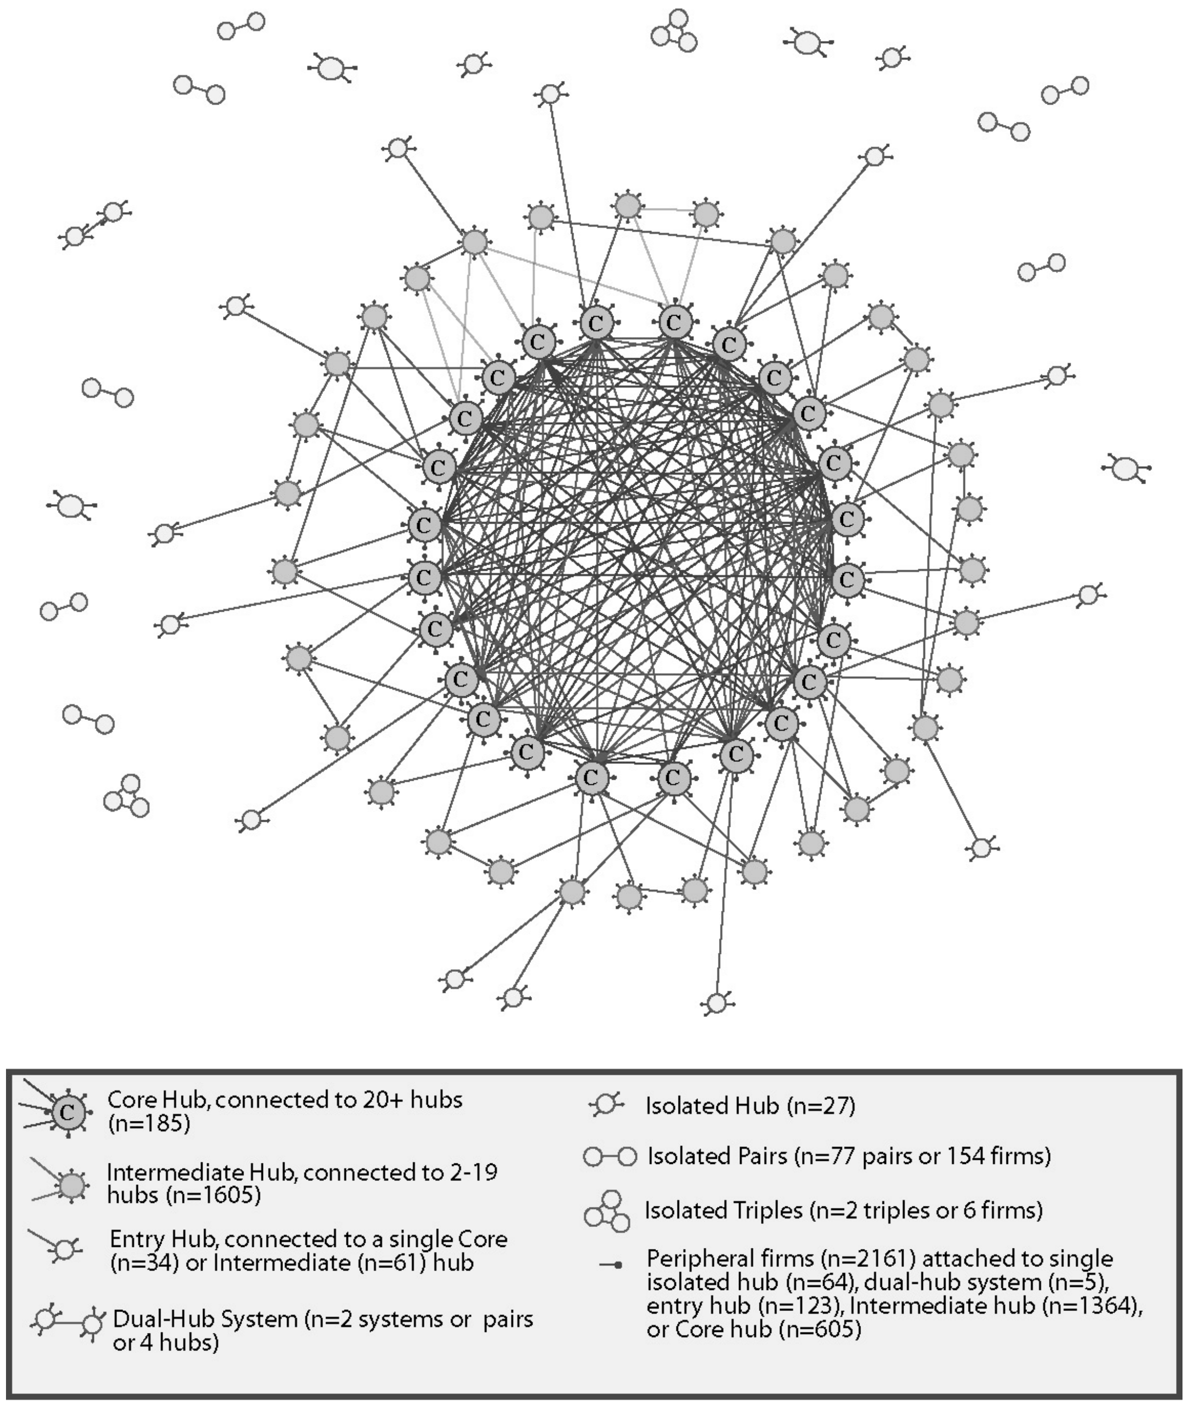

Note: Data extracted from Recombinant Capital database of alliances in the pharma-biotech industry, which includes 4,231 unique entities (surviving parents as of year-end 2001).

Table 2

Operational Objectives and Governance Structures for Biotech 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).

Table 6

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.