

APPENDIX A

TRANSACTION COST ECONOMICS

This book concerns transactions between employers and employees. Employment is quite different from other economic exchanges, and we focus on the differences. However, labor exchanges do share some characteristics with other types of exchanges. *Transaction cost economics* is an important tool for studying economic exchanges in general, a tool we used in Chapter 4 and will be using intensively, if somewhat informally, to study labor-service exchanges. This appendix provides a somewhat more analytical and complete treatment of the general subject, in case you want one.¹

BASIC IDEAS

Textbook economics treats transactions as simple this-for-that exchanges; someone gives you twenty-five cents, you give her a donut; you give someone \$25,000, he gives you a new car. Both parties to these textbook transactions **know what they are giving and getting**, and the **transaction is completed virtually instantaneously**. Because classical microeconomics is built with this sort of transaction in mind, **clear** and virtually instantaneous transactions are called *classical*.

However, many transactions, perhaps the most important ones, are much more complex than this. The purchase of an education is one such transaction, in which every so often students write out tuition checks, and only much later do they discover what it is that they have bought. Stanford has purchased our labor services as professors and all parties to these transactions are continually learning what the transactions entail as time passes. If you hire a lawyer to represent you in a particular matter, exactly what you have bought is remarkably unclear. If you enter into an agreement with a manufacturer to be a licensed dealer of the manufacturer's products, much of what is material to your eventual profits (e.g., what the wholesale prices will be in the future, what quality the products will be, how often product innovations will be provided, how much national advertising the manufacturer will undertake) is unspecified. Even the simple example above of purchasing an automobile is not so clear and straightforward; you will eventually wish to buy spare parts and servicing for the car, and you don't have an explicit agreement about the price or availability of either when you enter into the contract, nor do you know how well the company will honor its warranties.

tinuous improvement, globalization, and HR concerns among his senior management team. At a conference where academics and practitioners (including some from his company) met to discuss research on these issues, he was aghast to learn from one of the presentations that the scorecard was being used at low levels of his organization not only as a basis for evaluating unit performance, but for evaluating particular managers.

16. Why the parenthetical *carefully selected*? Some of our reasons concerning experimental controls follow in the next paragraph. But you should also mind the aphorism that success breeds success. When an HR experiment succeeds in one facility, it can be held up as something of an example to others, making change easier at other locations. So the sites of such experiments should be selected, at least in part, with the probability of success in mind. Of course—and one of the points of this paragraph—the extent to which a successful innovation at Plant A serves as an example for employees and management at Plant B depends on how closely A resembles B. Trading off these two desiderata (chance of success versus similarity to other facilities) is the crux of the matter.

Textbook economics implicitly asserts that such transactions as these are simply more complex versions of classical transactions. But the complexities bear closer study. What complicates all these transactions (and others) is the confluence of two distinct sets of factors. One set of factors pertains to the parties involved. People are *boundedly rational*; they can't anticipate everything that might ever happen, and they can't perform complex tasks of optimization except at very high cost (if at all). People also have the potential for acting *opportunistically*; they may try to take advantage of others if the opportunity presents itself and economic or social forces don't prevent this.

These human factors are relatively unimportant when the transaction is clean and straightforward. But they become quite important when a second set of factors come into play—namely, when the transaction is complicated because it involves *time uncertainty*, and *privately held information*. When a transaction is complicated in this fashion, and when the parties to it are opportunistic, neither party can blithely trust the other's goodwill in meeting contingencies as they arise. To some extent, the parties will try to anticipate what contingencies may arise later and to specify, at the outset, how those contingencies will be met. However, because the parties are boundedly rational, there are *ex ante costs of negotiation*. Moreover, simply making an agreement *ex ante* is not the same as fulfilling it *ex post*, and opportunism by either side may lead to breach of the agreement. Each party will incur *ex post enforcement costs* in trying to ensure that the other side meets its obligations under the agreement. And because of their bounded rationality, the parties will probably be unable or unwilling to anticipate up front all the contingencies that will arise later. As time passes and contingencies unfold, the parties to the agreement will *renegotiate* their arrangement to meet circumstances; in other words, there will be *ex post (re)negotiation costs*.

Each party to a potential transaction anticipates (at least, in some rough sense) that the transaction will involve these different costs. When deciding whether to enter into the transaction, the parties consider the benefits of the transaction, net of the transaction costs. When the transaction costs are very high, the parties may forgo the transaction altogether. When the costs are low enough so that the transaction can be undertaken, the parties will look for ways to arrange the transaction so that the transaction costs are relatively low.

Transaction cost economics studies these transaction costs and, in particular, how they vary with the form of the transaction, by which we mean the formal and informal institutional arrangements for adapting, adjudicating, and enforcing the transaction as time passes. (We will flesh out this vague definition over the next few pages.) The basic positive premise of transaction cost economics is that transactions will *tend* to take place in a form that minimizes the combined costs of the transaction. To understand all this, a concrete example may help.

Example of Design and Manufacture

Consider a transaction between two parties, A and B. Party A is a *designer* who is going to produce the plans for making some item, and party B is a *craft worker*

who will then manufacture the item. Party A's task requires certain pieces of capital equipment—a computer workstation, say—whereas party B's task requires various tools, jigs, and fixtures.

It is easy to think of a number of ways in which this transaction could be arranged. (1) Parties A and B could be distinct economic entities, each owning its own capital equipment, and party A, having completed the design, could hire B as an independent contractor to produce the pieces. (2) A could license the design to B. (3) B could purchase the design from A. (4) A could act as employer, owning all the capital and purchasing B's labor services. (5) B could own all the capital and purchase A's labor services. (6) Some third party C could own all the capital equipment and employ both A and B.

No doubt you can think of specific cases in which each of these transactional forms is employed. The question is: Why one and not another? Does it make a difference? To answer, we first have to be clear on how these different forms really vary.



The main implication of the form of a transaction is in dictating *the rights of the parties to make decisions and employ capital (including intellectual capital) as conditions change and contingencies arise*. In our example, who owns the right to use in the future the design created by party A, or the right to improve on that design, or to otherwise modify it? Who is able to decide how the jigs and fixtures that are used to make the items will be employed? When it comes to physical capital, the presumptive answer is that the owner of the capital has that right.³ If A owns all the capital and employs B as a skilled craft worker, then A can presumably fire B for cause and employ some substitute. If B owns all the capital and employs A as a designer, then B presumably has the right to continue production and adapt A's design. When A licenses a design to B, B retains the right to produce whatever quantity he wishes (subject to the original contractual arrangement, and typically with some royalty that is paid), although B will usually *not have the right to modify A's design and enter into competition with A*. At the same time, A will sometimes be restrained somewhat from making small changes in the design and entering into a licensing arrangement with some competitor of B. This is even more the case when B buys the design from A for a particular good.³

These rights are important because the immediate transaction is not the end of the story. The intellectual and physical capital created by this transaction will have value in the future, and the ability to command that value when the present agreement runs out or is voided by circumstances is a crucial part of the overall transaction. Moreover, each party is somewhat at the mercy of the other as the transaction evolves. Imagine, for example, that the two parties retain economic autonomy and each owns its own capital—specifically, B acts as an independent contractor for A. If B must make large investments in jigs and fixtures specific to this particular product, then B will make those investments in anticipation of recovering their cost in the price that A will pay for the work B does. But once B has made those investments, A could engage in a “holdup,” demanding that B do the work for a lower amount than was originally anticipated. Insofar as A and B can specify *ex ante* just what B will be paid for the work B does, no problem arises.

But when B's costs or A's quantity demands are uncertain (and especially when they are matters of proprietary information—for instance, when B can misrepresent costs), it becomes considerably more difficult to specify *ex ante* how the transaction will evolve. And what if A comes up with a new and improved model, for which the current jigs and fixtures can be used if modified? The production of this new and improved model is likely to be outside the bounds of the contract A and B originally made, so A is in a good position to proceed with the holdup of B.

The key to this potential holdup is that B's capital is most profitably employed at A's task; B's assets are specific to A. In the jargon of transaction cost economics, we say that this is a problem of *asset specificity*. If B's capital equipment could be employed just as profitably in transactions with various other parties, and if A tried to hold up B, B could just walk away from A with equanimity and employ his capital in some other way. In other words, if there are many As who could equally well use B's capital, then B can rely on competition among them as protection against a holdup.⁴ The problem of asset specificity seems to point toward an arrangement in which A pays for and thus owns the capital equipment. Such an arrangement seemingly economizes on the transaction costs incurred if A and B each owned their own capital.

But the story isn't quite so clean as this. If A owns the capital, she may still need to employ someone to use it. One easy-to-see problem with this sort of arrangement is that if B can damage A's capital equipment through negligence or abuse, then in the arrangement where A owns all the capital, A faces a moral hazard problem concerning whether B will take care of her (A's) equipment. In comparing the different ways the transaction could be arranged—(1) where A owns all the capital and employs B; (2) where each party owns some of the capital; and (3) where (say) B owns all the capital and employs A—we have to compare the *relative* transaction costs. Transaction cost economics says that the arrangement that is *relatively* the cheapest is the one that will tend to emerge and persist.⁵ If the costs incurred in avoiding a holdup are less than those incurred in dealing with the moral hazard problem, then we would expect to see separate ownership (or B's ownership of the capital) instead of A owning the capital and employing B. If the costs of preventing a holdup exceed those of dealing with the moral hazard problem, then we will tend to see A owning the capital (or, depending on the relative costs, B owning everything).

To see this in an example, consider that when a craft worker's tools are relatively general purpose, such as a carpenter's saw or hammer or drill, then there is a tendency for the tools to be owned by the craft worker. As the tools become more specific to a particular job, however, there is a tendency for them to be owned by the firm that employs the craft worker. Here we see the trade-off between two transaction costs: holdup possibilities if job-specific capital is owned by the craft worker, versus moral hazard costs of maintenance when tools are owned by the firm that employs craft workers to employ those tools. At the extremes, the force of minimizing transaction costs moves us clearly in a particular direction: Journeymen carpenters own their own hammers and saws, whereas Boeing owns

the jigs and fixtures for making 777 wings. It is in the middle ground where transaction cost analysis becomes most interesting and muddy.

Asset Specificity and the Push Toward Unified Governance

The story just told picks up one of the dominant themes in transaction cost economics, namely how the holdup problem or, more politely, the problem of asset specificity, pushes transactions toward a form in which some single entity owns all the assets.⁶ When physical capital is most economically employed in one and only one exchange, the owner of the capital is at risk relative to partners he or she may have in the exchange. If, for example, A manufactures bottles and B is a soft-drink bottler, it makes some sense for A and B to locate their two plants in close proximity to one another; this way they save on the expense of transporting empty bottles. But if they do this, both are potentially at risk. At any point, either party can try to renegotiate the agreement, seeking to extract a greater portion of the benefits of their physical proximity. If a manufacturer of automobile brakes locates its factory close to the Toyota assembly plant, and if this brake manufacturer engineers its brakes to fit in Toyotas and no other car, then the manufacturer is somewhat at the mercy of Toyota (the more so if Toyota enjoys a similar relationship with one or two other brake manufacturers). The costs of constant renegotiation and adaptation of the transaction to contingencies as they arise can be formidable in such instances, which may make it worthwhile to organize the transaction so that all the assets in question are commanded by some single economic entity.

In the jargon of transaction cost economics, this push toward having all the assets controlled by a single entity is known as the push toward *unified governance*, where unified governance connotes the case of two distinct entities, each initially owning its own assets, which are subsequently unified into a single structure that owns all the capital. The first and most obvious reason for unified governance of assets, then, is that it economizes on the transaction costs of *ex post* renegotiation that would otherwise occur if specific assets were controlled by a number of distinct legal and organizational entities. Moreover, because there is presumably less opportunity for holdups with unified governance, fewer resources must be expended on bargaining and contract writing *ex ante* and on the enforcement of the transaction *ex post*. For all of these reasons, unified governance can save powerfully on some transaction costs.

Pushing in the other direction, however, are a great number of transaction costs that can increase as the ownership of physical assets is unified. Most important are the *agency costs* of labor. Physical assets can all be owned by one firm, but the know-how it takes to work them cannot be; human capital cannot legally be bought and sold. Moreover, the firm, while it owns the assets, employs people to use and care for those assets, and the firm must provide its employees with incentives to use and care for those assets in the firm's own best interests.⁷ At the same time, employment creates other specific assets, namely *firm-specific human assets* or *human capital*. The basic idea (developed at length in Chapter 4) is that a worker

develops skills and personal connections that are specific to his or her current employer. This leaves both the worker and the firm vulnerable to being held up. The employer may attempt to take advantage of the worker's inability to redeploy his or her skills elsewhere as advantageously as can be done with the current employer. At the same time, if the worker acquires a monopoly of knowledge or skills that are specific to the employer, the worker can hold up the employer by threatening to depart.

Thus, by bringing a set of activities in-house, an organization may alleviate holdup problems with external suppliers, only to exacerbate potential holdup problems vis-à-vis the firm's own employees. To repeat, there are trade-offs to consider in moving to an arrangement in which one party controls all the physical assets, and chief among these trade-offs are the transaction costs associated with employment relationships.

BETWEEN UNIFIED GOVERNANCE AND CLASSICAL TRANSACTIONS: RELATIONAL CONTRACTING

Between the extremes—having all the assets held by a single entity (unified governance) and having them controlled by many individuals who deal with each other in simple, discrete, and isolated classical transactions—are intermediate transactional forms known to economists as *relational contracting*. The idea is to have the assets in the hands of different entities, but to economize on transaction costs of negotiation and renegotiation by enlisting the beneficial effects of long-term relationships and reputation.

A good example arises in the case of automobile parts manufacture. In Japan, many of the parts of a car come from manufacturers that are economically distinct from the assembler. On a Toyota, the brakes, engine cylinders, and even the bumpers are manufactured by firms other than Toyota, firms that own their own manufacturing equipment and sell parts to Toyota.

On the face of it, this would seem to expose the Japanese to large transaction costs of the sort that normally push toward unified governance. At least some of the equipment for making bumpers and brakes for Toyotas is very specific to Toyotas; the bumper manufacturer would have a hard time turning to Nissan or Honda if Toyota tried a holdup, especially because physical proximity is very important in Japanese car assembly. (Toyota, on the other hand, is less liable to a holdup by its bumper supplier, because it typically has two or three suppliers for every major subcomponent that it buys.)

Because of precisely these considerations, auto assembly is organized much differently in the United States. The major U.S. auto assemblers are very highly integrated; they make internally many of the pieces that go into the cars they assemble (though this has been changing in recent years). Thus, they own many diverse assets for auto manufacture; in terms of the last section, they are exemplars of unified governance. Indeed, in the 1920s and 1930s, the degree of vertical integration in automobile assembly was extreme; the major manufacturers (General Motors and Ford) manufactured their own steel, glass, and tires!

How do the Japanese avoid the transaction costs that seemingly pushed the American auto manufacturers into being highly integrated? They do so by relying on long-term relationships between the major auto assemblers and their suppliers and, especially, on the reputation of the major auto assemblers. Toyota has dealt with its suppliers for many years, and each side anticipates that the relationship will continue.⁸ As contingencies arise, Toyota and its supplier "work out" how to modify their relationship, trusting that, to safeguard the relationship, neither party will attempt to take advantage of the other. In fact, the working-out in this case takes an extreme form: Toyota makes all the important decisions. As contingencies arise, Toyota pretty much tells its suppliers what they must do. And the suppliers feel protected in this relationship because they trust Toyota not to abuse its authority; Toyota won't abuse its authority because it has an overall reputation for not doing so with all its suppliers. If Toyota treated a brake supplier badly, this would be noticed by the suppliers of its bumpers, lights, cylinders, and so on, and they would begin to demand contractual safeguards. Toyota, to avoid the costs of negotiating and enforcing such contracts, will treat all its suppliers fairly.⁹ Thus, by relying on long-term relationships and on reputation, Toyota and its suppliers can avoid the transaction costs normally associated with diffuse governance of specific assets. And Toyota and its suppliers gain by avoiding unified governance and some of the other transaction costs that unified governance would occasion.¹⁰

Three Pure Forms of Relational Contracting

The foregoing is just an example, but it gives the basic idea of relational contracting. Parties deal repeatedly with one another, relying on the value of their ongoing relationship or the value of their reputation as insurance against opportunistic behavior. Consider the following three general types of relational contracting.

In *balanced bilateral relationships*, two parties to a transaction both retain a fair degree of autonomy, and contingencies are met as they arise through joint consensus. One might imagine that this would greatly increase the transaction costs of *ex post* renegotiation, as the two parties haggle over how to proceed in every instance (given that consensus is required to proceed). However, each party has a powerful incentive to be reasonable in its dealings with the other because it values the ongoing relationship, and any unreasonable behavior might lead to a rupture of the relationship. A good example of this is any ongoing joint venture between two corporations, where each is expected to add resources in the form of knowledge or capital to the joint venture.

In *hierarchical (bilateral) relationships*, there are two parties to a transaction (the bilateral part), and one of the two, called the hierarchical superior, is chiefly responsible for deciding how contingencies will be met as they arise. The other party, called the hierarchical subordinate, retains the right to rupture the relationship but is otherwise largely under the power of the superior. The hierarchical superior can be trusted because the superior has a reputation to protect, either with this single hierarchical subordinate or more generally with a large number of hierarchical subordinates, all of whom monitor how the superior acts in its relations

with its many subordinates. An example here is the previously mentioned relationship between Toyota and its suppliers.

In *trilateral relationships*, there are two primary parties to the transaction and a third outside authority is invoked when the two cannot decide how to meet contingencies as they arise. Any time two parties sign an agreement that provides for binding arbitration in the event of a dispute, they are engaging in a trilateral relationship. The arbitrator is normally trusted by each side because she has her reputation as a fair arbitrator to protect; an outside authority that ruled capriciously or in a manner that was generally perceived as unjust would soon lose her customer base. In a sense, in any bilateral exchange taking place within a society with a well-developed system of contract or administrative law, the courts (or quasi-judicial authorities, such as the NLRB in the United States) act as trilateral authority.

Let us return briefly to a couple of examples mentioned earlier. The first involves you (if you are a student), contracting with your university. As noted above, you pay tuition to the university without getting in return a very detailed contract of what you will get as a function of the contingencies that arise. The reason a detailed contract isn't written in this instance is obvious: The contingencies are wide and varied and hard to anticipate. (Will you like Finance? Will you find the need for another course in Accounting? What if you get a failing grade in a required course? What if a professor whose courses you were eager to take becomes ill and cannot teach?) Writing a contract that covered all contingencies, or even just the ones that could be thought of, would be exorbitant. What you buy instead is a loosely defined procedure by which you and your university will collectively decide how those contingencies will be met. This procedure largely consists of: You ask, and the Dean (and administration) decides. Of course, you reserve certain rights of appeal, both to higher university authorities and to the courts. But the presumption is that the school's administration calls the tune, as long as you are treated according to standards of due process as they are employed at your university. You can petition the administration concerning your desires, but the administration decides whether to grant your petitions, and there is very little you can do or say if the answer comes back "No." In other words, what we have here is a very hierarchical relationship, with you as the subordinate.

So why did you enter into such a powerless position? Why did you trust your university not to abuse its power over you? (Consider how much you put at risk by going to school, and you will see just how much you have at jeopardy.) In part, we assume, you chose to come to your university because of its reputation (and its fit with your interests and desires for an educational program). The university's interest in maintaining its reputation both with you as a future alumnus and with other students is the best guarantee you have that you will not be exploited as a student.

This asymmetrical relational form is no accident, according to transaction cost economics; it precisely parallels the balance of reputational stakes in the relationship. A university's reputation is far more at stake than is yours, because you are apt to be a student only once, whereas a university looks forward to many generations of students. So it is efficient to give the university administration the greater

share of authority in adapting a specific transaction to circumstances.¹¹ Note well: We aren't saying that your education isn't as important to you as it is to your university. In fact, your educational experience *per se* probably matters a lot more to you than it does to your dean. But the argument isn't that authority should be given to the party with the most at stake. Rather, authority should be given to the party (or parties) whose *reputation* is most at stake, because it is a reputation stake that can serve to protect the interests of the other party.

To take another example, a textbook author is obliged to guarantee to her publisher that she will not write another book that competes with the one being published. But publishers do not guarantee that they will not publish some other, competitive textbook. Why the asymmetry? Because authors can rely on publishers guarding their reputation. If a publisher were to publish X's book and then that of a competitor (and if the publisher slighted X's book in advertisements, and so on), word would quickly spread among the close-knit community of potential academic authors, and the publisher would have a hard time signing contracts for other books. For a publisher, acquiring a bad reputation would be catastrophic, as publishers rely on an ongoing flow of new books for their profits. Authors can rely on their publisher's desire to maintain a good reputation to safeguard their interests. In contrast, a publisher can't rely as well on an author's desire to maintain a reputation to safeguard the publisher's interests; an author is unlikely to have very many books in her, and so she is much more likely to want to milk any one for everything possible. Thus, authors must give publishers binding and written guarantees of their behavior, whereas authors can rely on implicit guarantees from publishers.

Carrying this example a step further, contracts to publish a book are roughly 95% about the author's responsibilities and only 5% about the publisher's responsibilities. The contract reserves to the publisher unfettered rights to decide on price, advertising, sampling policy, and so on.¹² Indeed, the boilerplate contract that publishers send out gives the publishing company the power to call for revised editions whenever it suits the company (or not at all), regardless of the author's desires. But this clause is rarely if ever invoked (except via jaw-boning), because the culture of academic authors holds that this is unfair, and a publisher invoking this clause would suffer in reputation, even though the publisher has a full legal right to insist on a revised edition.

When authors negotiate with publishers, they typically negotiate on things like the retail price of the book and sampling policy, and they exchange letters recording their informal agreements. But those agreements don't turn up in the contract, and the contract is quite clear that such letters are not made part of the contract. Nonetheless, those letters are of value, given what the publisher has at stake in terms of reputation; although they have no legal force, those letters can be shown to other potential authors. Again, we have a case of a hierarchical relationship that is hierarchical because of the relative reputation stakes that are involved.

Most relational contracting situations do not fit neatly into one of the three pure categories above. Indeed, in the hierarchical form, the hierarchical subordinate generally retains the right to sunder the relationship, which can be a very sig-

nificant right indeed. Most real-life relational contracts are a mixture in various degrees of the three forms. It is the degree to which the three are mixed, and even more the *ways* in which the mix is concocted (who retains what rights), that determine the transaction costs associated with any specific form.

EMPLOYMENT RELATIONS AND TRANSACTION COST ECONOMICS

Now that we have the general ideas and jargon of transaction cost economics at our disposal, we can quickly reconstruct and review the essential argument of Chapter 4.

- Employment is a complex transaction involving time and uncertainty.
- The bounded rationality of the parties involved implies that the terms of the employment "contract" cannot be spelled out *a priori* in their entirety, but must be formed as time passes and contingencies arise.
- Also because employment is long-term, both employer and employee develop transaction-specific assets, leaving each liable to holdup by the other party as time passes and the relationship develops.
- Unified governance of the human assets involved is not possible, because of laws against indentured servitude.
- Hence, the parties must resort to relational contracting. It may be that one party is given the preponderance of authority to determine how the relationship will proceed; in fact, a caricature image of employment is that it is hierarchical, with the boss deciding and the worker acceding. But other forms, such as balanced bilateral or trilateral governance, can be found, with hybrids also appearing.
- As with any relational contract, either a balance of power between the two parties involved or the stake of one party (or both parties) in a general reputation provides each with assurances against abuse by the other.

ENDNOTES

1. For even more complete and analytical treatments, see Oliver E. Williamson, *The Economic Institutions of Capitalism* (New York: Free Press, 1985); or David M. Kreps, *A Course in Microeconomic Theory* (Princeton, NJ: Princeton University Press, 1990), Chapter 20.
2. This is subject to obvious caveats. It characterizes capitalist societies to greater or lesser extents, and it is subject to the state's rights of eminent domain.
3. An example of this arises in the case of publishing. A textbook author is obligated by her agreement with her publisher not to write a competitive product and sell it to a different publisher. In contrast, the publishing company does not have any formal requirement not to publish some book that competes directly with another book it is publishing. We will get back to this asymmetry in a bit.
4. Even if there are many As, B must still worry if there is the chance that they will form a cartel or otherwise collude against B's interests. He must also worry about develop-

ing a reputation as someone who will walk away from an agreement, thereby reducing the propensity of other potential exchange partners to enter into agreements with him (see Appendix B).

5. This is subject to an important caveat. In comparing the costs of different arrangements, you often must consider more than the individual transaction. Institutional arrangements are rarely fine-tuned for each individual transaction, but will tend to economize on an array of related transactions for the sorts of reasons discussed in Chapter 3. See Appendix B for more on this.
6. Note that this can't work completely in the context of human assets, as it is essentially precluded by laws against slavery and indenture.
7. Note also that to the extent that managers of specific organizational subunits have divergent interests, with decision-making rights delegated from the owner(s) to those subunit managers, there are ample holdup possibilities under unified governance.
8. These long-term economic ties are often supported by personal and social ties—for instance, the management of a Toyota supplier often includes individuals who once worked at Toyota and so have personal relationships with managers of the parent company.
9. You might object that we've slighted the social forces that keep Toyota from abusing its suppliers. This objection is completely correct—we have slighted the important social forces—but then this appendix concerns transaction cost *economics*.
10. It is too far off the point to catalog all those transaction cost savings here, but one is worth mentioning: By not integrating, Toyota can segment its "workforce" in ways that Ford and GM cannot. Because the employees of the bumper manufacturer are not Toyota workers, they needn't be extended all the benefits and safeguards that go to Toyota permanent employees, for example. And this sort of arrangement permits the assembler to impose sharper incentives on the parts manufacturer when necessary; the argument is a bit involved, but it comes down essentially to consistency and distinctions as outlined in Chapter 3. Indeed, American assemblers, noting the power and advantages of relational contracting in Japan, have in the recent past begun to mimic the Japanese, disintegrating vertically in some respects.
11. Compare this to instances in which a school engages in contract executive education for a large corporation. In this case, both the school and the corporation will endure, and each has something at stake in a continued relationship. Hence, in these cases, disputes can be (and often are) settled not by the dictates of one side but instead by negotiation and consensus.
12. Sampling policy means the number of books that will be distributed free of charge to colleagues of the author and potential adopters of the book.

APPENDIX B

RECIPROCITY AND REPUTATION IN REPEATED INTERACTIONS

In Chapter 4 and Appendix A, we used some notions from game theory concerning how self-interest can lead to trust and cooperation between employer and employee, based on either balance-of-power or reputation considerations. This appendix reviews the essential game-theoretic ideas and then gives some illustrative applications.

THEORY

Background: Games and Nash Equilibria

By way of background, we begin by discussing the formal techniques of analysis that we will use, the techniques of *noncooperative game theory*.¹

A *strategic form game* is a formal description of a competitive situation consisting of: a list of participants or *players*; for each player a list of *strategies* that the player can use; and for each *strategy profile* or combination of strategies, one strategy for each player, a *payoff* for each player. The payoffs are abstract (utility) measures of the desirability to each player of the outcomes, where a higher payoff is better, and we assume players choose strategies in order to maximize the expected value or mean of their payoff.

For two participants, we sometimes use pictures like Figure B-1 to describe one of these situations. The two participants are called *player 1* and *player 2*. The rows of the matrix in Figure B-1 correspond to player 1's strategies; there are three rows in the matrix, so player 1 has three strategies, labeled N, M, and S (for North, Middle, and South, respectively). There are four columns, so player 2 has four strategies, labeled W, CW, CE, and E (for West, Center West, Center East, and East, respectively). The participants must simultaneously and independently choose strategies. Their choice determines one of the twelve cells in the matrix, and the pair of numbers in the corresponding cell describes the payoffs to the two players. The first number represents the payoff to player 1, and the second number is the payoff to player 2. So, for example, if player 1 chooses M and player 2 chooses CW, player 1's payoff is 4 and player 2's payoff is 2.

		PLAYER 2			
		W	CW	CE	E
PLAYER 1	N	8,0	3,5	0,4	6,6
	M	5,0	4,2	0,1	5,0
	S	9,9	2,10	0,0	4,10

Figure B-1 An example of a two-player game.

We want to predict how the players might act in situations described in this fashion. Our chief formal tool for doing this is a *Nash equilibrium*. This is a list of strategies, one for each player, such that each player is doing as well for herself as possible, measured by her expected payoff, if she (correctly) anticipates that the other players will follow their parts of the Nash equilibrium. So, for example, M and CW is a Nash equilibrium of the game in Figure B-1: If player 1 anticipates that player 2 will choose CW, then M is better than N or S; while if player 2 anticipates that player 1 will choose M, then CW is better than W, CE, or E. Note that N and E is another Nash equilibrium of the game in Figure B-1.

In what way is a Nash equilibrium helpful in predicting how the game will be played? The idea is that *if* each of the players is able to anticipate how the other(s) will act, then it must be that their joint behavior constitutes a Nash equilibrium. The reason for this is that *if* all the players anticipate what will happen, and if each is trying to maximize his or her payoff, then each must be doing as well as possible against the strategies of the others.

Note well, this isn't saying that in every situation we study each participant will actually be able to anticipate how others will act. We must have some reason to believe that the participants can correctly anticipate the actions of their fellows. It may be that correct anticipations are the product of prior discussion among the participants, or experiences that they share, or sometimes the result of logical reasoning. As we go to *apply* the ideas of game theory to real-world situations, we will need to ask, more or less constantly, whether it is reasonable in the situation being modeled to suppose that each participant can anticipate fairly accurately how the others will act.

This is the sort of formal analysis we will use in the rest of this part of the appendix. We close this background section with two remarks. First, in what follows we will look at situations that are somewhat more complex than the simple sort of game given in Figure B-1. In particular, we will be describing situations in which the list of strategies available to a participant is never made explicit, but instead is implicit in a description of actions the participant can take and information the participant will have when the actions must be taken. Second, because of these complications, we will *not* be mathematically precise in our arguments. You should be prepared for a fair bit of sloppy hand-waving. If you want to see things done absolutely formally, with all the mathematical *i*'s dotted and *l*'s crossed, you will need to look at an advanced textbook.²