

MIT Open Access Articles

Relational Contracts and Organizational Capabilities

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Gibbons, R., and R. Henderson. "Relational Contracts and Organizational Capabilities." *Organization Science* 23.5 (2011): 1350–1364.

As Published: <http://dx.doi.org/10.1287/orsc.1110.0715>

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: <http://hdl.handle.net/1721.1/75010>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Relational Contracts and Organizational Capabilities

R. Gibbons and R. Henderson

MIT and HBS

For *Organization Science*

Version: July 18, 2011

Abstract

A large literature identifies unique organizational capabilities as a potent source of competitive advantage, yet our knowledge of why capabilities fail to diffuse more rapidly—particularly in situations in which competitors apparently have strong incentives to adopt them and a well developed understanding of how they work—remains incomplete. In this paper we suggest that competitively significant capabilities often rest on managerial practices that in turn rely on relational contracts (i.e., informal agreements sustained by the shadow of the future). We argue that one of the reasons these practices may be difficult to copy is that effective relational contracts must solve the twin problems of *credibility* and *clarity*, and that while credibility might in principle be instantly acquired, clarity may take time to develop and may interact with credibility in complex ways, so that relational contracts may often be difficult to build.

Acknowledgements: We would particularly like to thank our research assistants, Nancy Beaulieu and Tommy Wang, for their years of invaluable support in our research program. We are also grateful to our colleagues Nelson Repenning and John Sterman and MIT Sloan's Program on Innovation in Organizations and Markets for their sustained intellectual and financial support. Thanks for comments on this paper are due to the editors of this special issue and to David Garvin, Ranjay Gulati, Mauro Guillen, Sim Sitkin, Birger Wernerfelt and Sid Winter, and to seminar participants at Columbia GSB, Harvard Business School, Kellogg GSM, London Business School, LSE, Sydney and Tilburg.

1) Introduction

It is now widely accepted that there can be persistent performance differences among seemingly similar enterprises—be they work groups, plants, or firms (see Syverson (2011) for a recent survey). The strategy literature has long explored the possibility that such performance differences arise from organizational capabilities (e.g., Wernerfelt, 1984; Teece, Pisano and Shuen, 1997; Eisenhardt and Martin, 2000; Anand and Khanna, 2000; Aral and Weill, 2007). One of the central questions raised by this research is why such capabilities—if they are indeed a source of sustained competitive advantage—do not diffuse more rapidly.

The existing literature has proposed three answers to this question. First, incumbent managers may have problems of *perception*—they do not know they are behind because their cognitive frames blind them to new opportunities (e.g., Henderson and Clark, 1990; Christensen, 1997; Tripsas and Gavetti, 2000). Second, managers may have problems of *inspiration*—they know they’re behind, but they don’t know what to do about it because the organizational practices underlying key organizational capabilities involve either tacit knowledge (e.g., Winter, 1987, 1988) and/or complementarities (e.g., Levinthal, 1997; Rivkin, 2000) and so are hard to learn and communicate. Third, managers may have problems of *motivation*—they know they’re behind and they know what to do, but they don’t care because lack of competition in the market (or lack of incentives inside the firm) gives them insufficient incentive to adopt new practices (e.g., Reinganum, 1989; Bloom and Van Reenen, 2007; Bresnahan, Greenstein and Henderson, 2011).

While persuasive in many settings, these three answers to the problem of slow diffusion are less helpful in at least one important case: settings where managers acknowledge that they are behind and are spending heavily to catch up, and where there appears to be industry-wide agreement about best practice. For example, there have been more than 300 books and thousands of articles written about Toyota, yet until quite recently many automobile companies appeared to have great difficulty imitating its practices (Pil and MacDuffie; 1996). Similarly the practice of science-driven drug discovery was surprisingly slow to diffuse across the pharmaceutical industry despite widespread agreement about its effectiveness (Cockburn, Henderson & Stern, 2000). In this essay, we argue that such failures of diffusion may arise from difficulties in *administration*—managers know they’re behind, know what to do, and are motivated to do it, but they can’t get the organization to get it done.¹

Our argument is in two parts. First, we suggest that many organizational capabilities rest on managerial practices that in turn rely on *relational contracts*—an economist’s term for collaboration sustained by the shadow of the future, as

¹ We thank Jan Rivkin for bringing his “four ‘tion” labels to our attention (and allowing us to adapt them for our own purposes here).

opposed to formal contracts enforced by courts. Second, we suggest that many relational contracts are hard to build and to refine, and that this is often why managers “can’t get the organization to get it done.”

To unpack the first part of our argument—that key managerial practices frequently require relational contracts—we emphasize that such practices involve actions that cannot be fully specified in advance. Consider the following three examples we develop below. First, Lincoln Electric’s enduring success rests in part on the payment of bonuses that both managers and employees consider to be “fair,” but no manual can define exactly what constitutes a fair bonus for a particular worker in a particular year. Second, beyond compensation, the Toyota production system asks line workers to become “active problem solvers,” but cannot define in advance exactly which problems they might find or how they should be solved. Third, beyond manufacturing, Merck asks its researchers to behave “almost as if” they were academics, but what this means in terms of actual behaviors had to be worked out over many years and is still difficult to communicate today.

Because these managerial practices—each of which is fundamental to the success of these three firms—involve actions that cannot be specified in advance, it is typically impossible to motivate their performance via formal contracts (i.e., contracts that attach objective weights to objective measures).² Instead, if it is necessary to provide motivation for parties to take these actions, it will have to be done through informal agreements that involve subjective weights or subjective measures. In this essay we interpret such informal agreements as relational contracts.³

Of course, the idea that relational contracts are widespread and important has a long pedigree, across many disciplines and settings: in sociology, Macaulay (1963) studied informal relationships between firms; in anthropology, Geertz (1962, 1978) studied rotating credit associations and bazaar economies; and in political science, Ostrom (1990) studied communities of resource users. More recently, in the strategy literature, many authors have drawn on related ideas to analyze long, productive relationships between firms (e.g., Dyer, 1997; Dyer and Singh, 1998; Poppo and Zenger, 2002; Gulati and Nickerson, 2008).

Closer to our focus on the development of managerial practices within firms, there is also a long tradition in organization theory emphasizing the importance of informal understandings within organizations. For example, Blau and Scott (1962: 6) argued that “It is impossible to understand the nature of a formal organization

² Under special circumstances, it might be possible to write formal contracts based on outcomes rather than on actions and thereby motivate the desired actions. Much more often, however, available outcome measures are themselves incomplete and so formal contracts based on such measures induce gaming instead of consummate performance. See Gibbons (2005) for a review.

³ We do not mean to imply that the presence of relational contracts implies that the firm will not also rely on formal contracts. Indeed the two are often complementary.

without investigating the networks of informal relations and the unofficial norms as well as the formal hierarchy of authority and the official body of rules, since the formally instituted and the informal emerging patterns are inextricably intertwined.” Indeed, some have argued that high-performing organizations rely especially heavily on informal understandings, variously described as norms or cultures or contracts. For example, Barney (1986) suggested that an organization’s culture could be a source of competitive advantage, Rousseau (1989, 1995) studied psychological contracts in organizations, explicitly envisioning managers as parties to some of them, and a long tradition in the human-resources literature has documented the performance advantages of high-commitment work practices and the social contracts that sustain them (e.g., Hoffer-Gittell, 2002; Lincoln and Kalleberg, 1990; Kochan et. al., 1995).

To our knowledge, however, this literature has not explicitly linked these ideas to the idea that these informal understandings may be one of the reasons that competitively important practices are sometimes surprisingly slow to diffuse. Much of the capabilities literature, for example, has either made the implicit assumption that incentives within the organization are aligned, so that the adoption of new practices is primarily a problem of information transmission, or has labeled “appropriate incentives” as a distinct, separable capability. Similarly, while many writers have suggested that “trust” might be an important source of organizational performance (see for example, Bachmann and Zaheer, 2006 and Zaheer and Venkatraman, 1995), this literature is largely silent about how trust is built and so, as far as we know, has not studied how difficulties in building trust might be a key source of competitive advantage.

Here we explore this connection explicitly. We argue that many relational contracts are hard to build and refine and this is often why managers “can’t get the organization to get it done.” In particular, we argue that building and refining relational contracts requires solving two distinct problems: the problem of *credibility* and the problem of *clarity*. We believe that each of these problems can be quite difficult in isolation and that in combination they may create a substantial barrier to imitation.

By the problem of credibility we mean the problem of persuading others that one is likely to keep one’s promises. While credibility may derive from various sources, including what the literature has called “personal” or “institutional” trust (e.g., Zaheer and Venkatraman, 1995; Nooteboom, 1996), in this essay we focus instead on a strictly consequentialist logic, where one keeps one’s promises because things will go badly otherwise. Some have called this logic “calculative trust” (e.g., Williamson, 1993; Barney and Hansen, 1994; Nooteboom, 1996), and Rousseau et. al. (1998) call it “deterrence-based trust.” Like many of those authors, however, we see the consequentialist logic as distinct and different from trust; like Yamagishi

and Yamagishi (1994), we would prefer this logic to be labeled “assurance” rather than trust.⁴

By the problem of clarity we mean the problem of communicating the terms of the relational contract. The problem of clarity is less discussed and perhaps less evident than the problem of credibility, but we believe that clarity is at least as important a constraint on building and refining relational contracts. In brief, instead of asking whether others will *believe* one’s promises, we now ask whether others will *understand* one’s promises.

A rich literature has suggested that one reason why managerial practices may be slow to diffuse is that knowledge of such practices may be tacit or “embedded” and hence difficult to communicate (see, for example, Winter’s work cited with the inspiration problem above). Building on this literature, we argue that developing a shared understanding of a relational contract will be even more difficult because there is more to communicate. We suggest that building a relational contract requires developing a shared understanding of not only the necessary *task knowledge* (what each party is supposed to do) but also the necessary *relational knowledge* (what each party *could* do, either to break a promise or to punish someone who did, and what the payoffs from all these possible actions are).

In addition to the difficulties that these problems of credibility and clarity taken separately pose for building and refining relational contracts, we hypothesize that the biggest difficulties may arise from the interaction of these problems. For example, if one party acts in a way that is unexpected by the other, is miscommunication to blame, or gaming? More generally, the imperfect alignment of interests underlying the credibility problem creates significant new impediments to the communication necessitated by the clarity problem.⁵

The outline of the essay is as follows. To lay a foundation, we begin in Section 2 with a brief overview of relational contracts. In particular, we provide some basic models of when collaboration can be sustained by the shadow of the future, and we suggest that these models are broadly consistent with evidence drawn from both lab experiments and field data.

In Section 3, we develop our hypothesis that competitively significant management practices often rely on relational contracts by examining three familiar examples: subjective bonuses at Lincoln Electric, the production system at Toyota, and science-driven drug discovery at Merck. Our goal here is twofold. First, we hope to make credible the hypothesis that some competitively important managerial

⁴ In focusing here on the question of calculative trust we do not mean to imply that other forms of trust are unimportant. Indeed one of our hopes for this paper is that it might contribute to the ongoing discussion of the relationship among different forms of trust and their evolution over time.

⁵ See Spence’s (1973) signaling model and Crawford and Sobel’s (1982) cheap-talk model for an introduction.

practices are crucially dependent on relational contracts. While this idea may be familiar regarding subjective bonuses, it seems less appreciated regarding other managerial practices. Second, we begin to use our descriptions of these practices to illustrate the problem of clarity—i.e., the extensive amount of information that both employees and managers must hold in common if the necessary relational contracts are to be self-enforcing.

In Section 4 we expand on the question of clarity. We begin by distinguishing between “task knowledge” and “relational knowledge”, suggesting that many of the same mechanisms that make it difficult to learn how to do new tasks also make it difficult to learn about the relational contracts that might support them. We illustrate this idea through accounts of moments at Lincoln Electric and Credit Suisse where relational contracts threatened to break down, and discuss a number of experiments and models that may provide building blocks towards a theory of why relational contracts may be hard to build. Finally, Section 5 concludes.

2) A Primer on Relational Contracts: Theory and Evidence

In this section we provide a brief introduction to economic theory and evidence on relational contracts. The theory, in Section 2.1, begins with the simplest case and then sketches several enrichments. The evidence, in Section 2.2, includes both lab and field data.

2.1 Theory: Credibility from Incentive Compatibility

There is now a large theoretical literature on how relational contracts can facilitate efficient behaviors, both on their own (e.g., Bull (1987), MacLeod and Malcomson (1989), Levin (2003), Fuchs (2007), and Halac (2011)) and in combination with formal aspects of organizations and contracts (e.g., Baker, Gibbons, and Murphy (1994, 1999, 2002, 2011)); see Malcomson (2012) for a survey. As described in the Introduction, the essential intuition is straightforward and familiar: in some settings, one keeps one’s promises because things will go badly otherwise.

Kreps (1990) captures this intuition using the simple game shown in Figure 1. Although Kreps calls this the “Trust Game,” we emphasize that the relational-contracts argument is entirely consequentialist. Therefore, we would prefer “assurance” to “trust.” In short, when speaking of trust, we adopt March and Olsen’s (1989) view that “[t]he core idea of trust is that it is *not* based on an expectation of its justification. When trust is justified by expectations of positive reciprocal consequences, it is simply another version of economic exchange” (p. 27, emphasis added). In this essay, we are indeed considering “another version of economic exchange.”




Figure 1: The Trust Game

In a one-shot version of this game, player 1's initial choice amounts to either ending the relationship (and so receiving a payoff of zero) or trusting player 2 (and so receiving a payoff of -1, assuming that player 2 would maximize 2's payoff and so betray 1's trust). Since zero exceeds -1, player 1 should end the relationship. If the players' relationship is ongoing, however, very different outcomes are possible.

A classic formulation is to model a relationship between player 1 and player 2 as a repeated game that is equally likely to end after any period. The probability that the game will end influences the interest rate r per period that the players use in discounting their payoffs across periods. If the players are patient enough (i.e., r is sufficiently small, in part because the probability that the game will end is sufficiently small) then the following "trigger" strategies are an equilibrium of the repeated game.

Player 1: In the first period, play Trust. Thereafter, if all moves in all previous periods have been Trust and Honor, play Trust; otherwise, play Not Trust.

Player 2: If given the move this period, play Honor if all moves in all previous periods have been Trust and Honor; otherwise, play Betray.

The broader message from this example is that cooperation may be prone to defection, but in some circumstances defection can be met with enough punishment to motivate cooperation. To begin to generalize the example, imagine that player 2's payoffs (per period) are C from cooperation (generalizing the payoff 1 above), D from defection (generalizing the payoff 2 above), and P from punishment (generalizing the payoff 0 above), where $D > C > P$. The decision of whether to cooperate or defect then amounts to comparing two time-paths of payoffs: (C, C, C, \dots) versus (D, P, P, P, \dots) , as shown in Figure 2.




Figure 2: Payoffs from Cooperation Versus Defection and Punishment

The time-path of cooperation yields a higher present value than the time-path of defection if

$$(1) \quad \left(1 + \frac{1}{r}\right)C > D + \frac{1}{r}P ,$$

where $1/r$ is the present value of a dollar to be received every period (until the game ends) starting next period. Rearranging (1) yields

$$(2) \quad r < \frac{C - P}{D - C} ,$$

which is often restated as: if the players are sufficiently patient (*i.e.*, if r is sufficiently close to zero) then it is optimal to cooperate, foregoing the short-run temptation ($D - C$ now) for the long-term gain ($C - P$ thereafter).

This stylized model not only illustrates the main idea behind relational contracts but also suggests some limitations of the approach (at least as it has been developed and applied to date). One seeming limitation might be that we have cast the players – and, more broadly, the parties to any contract of this type as being motivated by “payoffs,” which may seem too narrow. But we take a broad view of such payoffs. In particular, we do not mean to suggest that money is the most important—or even necessarily an important—source of motivation inside firms. In Section 3 (where we discuss relational contracts within organizations), we therefore conceive of “payoffs” as including everything that might affect an individual’s experience of his or her job, including factors such as job assignment, degree of autonomy, status with the firm or work group, and other intangibles such as feelings of belonging or that one is making a difference.

A more serious limitation is that the analysis leading to (1) considers only the *payoffs* from cooperation, defection, and punishment, taking for granted not only that the parties know these payoffs but also that they know what *actions* constitute cooperation, defection, and punishment (such as Trust, Honor, and Betray in Figure 1). As we will see in the detailed descriptions of managerial practices in Section 3, taking it for granted that the parties know what these actions are may be a heroic assumption, especially when the parties are just beginning their relationship. For the rest of this section, however, we continue to make this assumption (thus implicitly focusing on steady-state relationships, where the parties have substantial shared experience about these actions and their payoffs).

To conclude this sub-section, we sketch some simple enrichments of the basic theory, to prepare for the discussion of evidence in the next sub-section. First, and informally, Figure 1 could be enriched in various ways, such as by allowing two levels of cooperation: full and partial. Full cooperation yields payoff C but has defection temptation D, as above, whereas the analogous values for partial cooperation are c and d, where $C > c > 0$ and $D > d > 0$. Given plausible assumptions about these payoffs, parties that are patient enough can sustain full cooperation, as in (2), whereas parties that are somewhat more impatient can sustain only partial cooperation (and, as always, parties that are too impatient cannot sustain any cooperation).

Second, and shifting attention from cross-sectional comparisons across relationships to longitudinal comparisons within an ongoing relationship, imagine that there is not only a probability that the game will end after any period, but also an independent probability after each period that the payoffs in Figure 2 will change permanently from (C, D, P) to (C, D', P), where $D' > D$. (A parallel argument holds if the payoffs change to (C, D, P'), where $P' > P$.) If the parties are impatient enough that

$$(3) \quad \frac{C - P}{D' - C} < r ,$$

then cooperation will end once the payoffs change. Nonetheless, if the parties are not too impatient (i.e., r is below an upper bound not shown in (3)), then it is an equilibrium for the parties to begin the game by cooperating, planning to cooperate until either the payoffs change or the game ends.

Third, the payoffs might neither be constant nor change permanently (as in the first enrichment), but instead fluctuate randomly across periods. That is, suppose that the payoffs are independently and identically distributed over time and that the parties observe the current payoffs before taking actions each period. In period t, the current payoffs from cooperation and defection are C_t and D_t , and the expected future payoffs from cooperation and punishment are $E(C_t)$ and $E(P_t)$, so (1) becomes

$$(4) \quad C_t + \frac{1}{r} E(C_t) > D_t + \frac{1}{r} E(P_t) .$$

In this setting, if the long-term gain of $E(C_t) - E(P_t)$ in each future period outweighs all possible realizations of the short-run temptation $D_t - C_t$, then (4) implies the critical value of r below which the parties can sustain permanent cooperation despite fluctuating payoffs. Alternatively, if there are extreme realizations of the short-run temptation that violate (4) then cooperation will end once the short-run temptation hits a high enough value, but if the parties are not too impatient then it is an equilibrium to begin by cooperating, planning to cooperate until either an extreme temptation arrives or the game ends.

Finally, in the most sophisticated of these enrichments, where the parties must repeatedly adjust their collaboration over the course of their relationship, suppose that there are two levels of cooperation (full and partial, with payoffs $C > c$ and $D > d$, as above) and that the payoff from defecting on full cooperation fluctuates randomly across periods (e.g., D_t can be either high or low, $D_H > D_L$). If the parties are impatient enough that

$$(5) \quad \frac{C - P}{D_H - C} < r ,$$

then the parties cannot sustain full cooperation in periods when the defection payoff is high. On the other hand, if the parties are not too impatient, then it is an equilibrium for the parties to achieve full cooperation in periods when the defection payoff is low and partial cooperation when it is high. Strikingly, there is evidence for this sophisticated behavior (as well as the others described here), as we discuss next.

2.2 Evidence from the Lab and from Relationships Between Firms

We now briefly introduce evidence that models like these capture important aspects of behavior. We begin with the basic model, summarized by equations (1) and (2), after which we turn to the enrichments, summarized by equations (3) through (5). In this section, we restrict our attention to lab experiments and to field data on relationships between firms, saving discussion of relationships within firms for Section 3. All the evidence we describe is only illustrative; many further examples could also be given.

Condition (1) predicts that cooperation is less likely as the returns to defection rise (i.e., as D or P increases).⁶ As evidence in this spirit, consider the field study by McMillan and Woodruff (1999) and the lab study by Brown and Serra-Garcia (2010). McMillan and Woodruff use a sample of firms in Vietnam during 1995-97 to study trade credit (i.e., the extent to which a buyer's payment is made significantly after the supplier delivers the goods). They find that suppliers grant buyers more trade credit if there are fewer similar suppliers within 1 kilometer and if the

⁶ For fixed parameters, either (1) holds or it does not, so cooperation is either possible or not. If we imagine data with, say, heterogeneous values of C , however, then an increase in D or P can be said to reduce the likelihood of cooperation, such as in the sense of a probit.

supplier's main competitor is located more than 1 kilometer away (i.e., when punishment payoffs are lower for buyers). Turning to lab evidence, Brown and Serra-Garcia conduct an experiment that varies a borrower's ability to expropriate loaned funds by allowing (or not allowing) the borrower to reinvest funds following default. They find that the volume of trade decreases and that borrowers are more likely to default (in initial periods) when expropriation is possible (i.e., when defection payoffs are higher).

As direct evidence for the idea that cooperation is more likely in environments with lower discount rates, as predicted by (2), consider the experiment by Dal Bo (2005), who conducts a repeated Prisoners' Dilemma and varies the probability that the game will be continued. Dal Bo finds that higher probabilities of continuation (which imply lower values of r) lead to higher rates of cooperation. Engle-Warnick and Slonim (2004) provide a related experimental finding.

There are also indirect examples of (2), based on the idea that there may be more than two parties available and different pairings of parties may occur over time. As a bleak baseline case, suppose that there are many parties, that pairings are random each period, and that each party knows only its own experience (i.e., there is no information about what other parties did while in other pairings). In this case, it is likely to take many periods before a given pair meets each other again, so r is high and cooperation is unlikely. But there are various alternatives to this bleak baseline case, each of which can be interpreted as reducing r and so making (2) more likely to hold, as follows.

In lab studies, Brown and Zehnder (2007) allow parties to choose their partners each period, leading to both long-standing relationships and improved cooperation, and Duffy and Ochs (2009) find greater cooperation in fixed pairs than in random pairs. Using field data, Robinson and Stuart (2007) study biotechnology alliances, finding that alliance partners with closer proximity to each other in the industry network are less likely to use equity (and use less equity when it is involved in the deal) and use more funding pledged upfront rather than funding based on milestones. Similar results hold when either of the firms is more central in the overall network (as distinct from more proximate to its partner). These findings are consistent with the idea that firms that are more proximate (and hence more likely to encounter each other) or more central (and hence more likely to have their actions visible to others) are less likely to defect.

Turning from cross-sectional evidence to longitudinal, recall the idea that an increase in the defection temptation can induce defection, as discussed in connection with (3) and (4). In this spirit, Guriev, Kolotilin, and Sonin (forthcoming), study nationalizations of foreign-owned oil assets in 161 countries from 1960-2006. During this period there were 98 nationalizations in 42 countries. Controlling for country fixed effects, they find that nationalization is more likely when oil prices are high (i.e., defection payoffs are high) and when political institutions are weak (i.e., punishment costs are low, so the payoff during punishment, P' discussed with (3), is

high). These independent variables—the oil price and political institutions—may be somewhere between the permanent change envisioned in (3) and the temporary fluctuation envisioned in (4).

Finally, as evidence on the sophisticated behavior in (5), where the parties repeatedly adjust their collaboration over the course of their relationship, consider Rotemberg and Saloner (1986), who study variations in the extent of collusion over the business cycle. In their model, full collusion (where all firms charge the monopoly price) produces moderate reneging temptations in low-demand periods but higher reneging temptations in high-demand periods, so the colluding firms can implement only partial collusion in the high-demand periods, as discussed in connection with (5). Empirically, Rotemberg and Saloner study both price-cost margins over the business cycle for a variety of industries, as well as actual price wars in specific industries, and in both cases find evidence broadly consistent with their theory: margins are smaller in booms (i.e., cooperation is reduced at moments when defection temptations are higher).

To conclude this section, we recap the argument we have made thus far. First, repeated-game models of collaboration predict that cooperation can occur if defection can be met with sufficient punishment (where punishment occurs over time, so the prospect of punishment has a greater impact on parties' current decisions when they care more about the future). Second, in both lab studies and field data on relationships between firms, there are empirical findings that are consistent with such models. As noted above, however, both the models and the evidence we have described focus on the problem of *credibility* (i.e., whether one actor should believe another's promises) and ignore the problem of *clarity* (i.e., whether one actor can understand another's promises). They thus implicitly focus on steady-state relationships, where the parties have substantial shared experience about their actions and payoffs, and avoid the problem of how these steady-state relationships come about.

In the next section, we add two aspects to our argument. First, we consider relational contracts within firms rather than between. In particular, we describe three important management practices that we see as relying crucially on relational contracts. Second, we provide sufficiently detailed descriptions of these relational contracts within firms to suggest that it cannot have been easy to develop a shared understanding of exactly what these contracts entailed, thus laying the groundwork for a more extended discussion of the problem of clarity in section 4.

3) Managerial Practices and Relational Contracts

We now describe important managerial practices within three leading firms: subjective bonuses at Lincoln Electric, the production system at Toyota, and science-driven drug discovery at Merck. In each of these settings, we focus on key actions that we believe could not be motivated by formal contracts and that we thus hypothesize are the subject of relational contracts. We treat Lincoln briefly, because

it corresponds reasonably well to the model in Section 2.1, where one party relies on the other to keep its promises, but we discuss Toyota and especially Merck in more detail, because they illustrate the more typical and more interesting case where each party is relying on the other.

In addition to presenting a plausible case that competitively significant managerial practices rely on relational contracts, our goal in these descriptions is to begin to illustrate the idea that building a relational contract may require addressing problems of *clarity* as well as of *credibility*. We show that these relational contracts require not only a high level of task knowledge—i.e., of the actions that constitute cooperation, but also a great deal of relational knowledge—i.e., of the payoff to cooperation for each party, of each party's ability and incentive to defect, and of the actions and payoffs that constitute punishment. Section 4 builds on these descriptions to explore the problems inherent in building relational contracts in more detail.

Lincoln Electric makes arc welders. Their productivity, innovation, and profitability have made them arguably the world's leader in their industry, and their management practices have brought them substantial scrutiny—from Fast and Berg (1975) through Hastings (1999) and beyond. For brevity, we focus on the firm's incentive scheme, which involves both a piece rate (specifying the rate of pay per task completed) and a subjective bonus (based on factors that can be difficult to measure, such as dependability, quality, and ideas/cooperation). See Milgrom and Roberts (1995) for descriptions of how Lincoln's other practices complement the subjective bonus we describe.

Lincoln's bonus is very important for both the workers and the firm: over several decades, the average bonus to an individual was approximately as large as the individual's total piece-rate pay (depending on the person and the year) and the average bonus pool was about half of the firm's pre-tax, pre-bonus earnings (again with substantial variation across years). Crucially, however, neither the individual bonus received by a worker nor the aggregate pool paid by the firm is determined by a formula (i.e., an objective weight attached to an objective performance measure). Instead, both are discretionary.

In the language of Section 2, Lincoln's workers may earn high payoffs from working hard, as bonuses can approach half of total compensation, but since payment of the bonus is discretionary, Lincoln's management can in principle defect by paying too small a bonus or none at all. Several formal models have explored the conditions under which such a discretionary bonus plan satisfies a credibility constraint such as equation (1) above.⁷ But these models ignore the degree to which such contracts must also solve a potentially severe clarity problem: Lincoln's

⁷ See Bull (1987), MacLeod and Malcomson (1989), and Levin (2003) for repeated-game models of such relational incentive contracts and Baker, Gibbons, and Murphy (1994) for a two-part pay plan like Lincoln's, combining an objective piece rate with a discretionary bonus.

relational contract rests on a number of shared understandings that may well be difficult to imitate. For example, the size of the bonus is contingent not only on the productivity of an employee, but also on the employee's "dependability," the "quality" of their work and the degree to which they contributed new ideas and "cooperated" in the improvement of the production process—none of which is easy to define or measure.

A similar reliance on relational contracts—and an associated reliance on the development of an extensive shared understanding of concepts that are intrinsically hard to define without shared experience—is evident in our next example, from Toyota's production system (TPS).

Many researchers have documented that the TPS relies on innovative contributions by the workforce through shop-floor problem-solving (see, for example, Ohno, 1995 and Womack, Jones and Roos, 1991). The kinds of behaviors asked of the workforce include "raise problems when you see them" and "be an effective member of problem-solving teams." As with the aspects of worker performance that are rewarded in Lincoln's subjective bonus, we conjecture that these desired behaviors in Toyota's production system cannot be specified precisely enough to be measured and rewarded in a formal contract.

As one example, a key discretionary behavior (by both workers and supervisors) involves the "andon cord" (a rope on the assembly line that, when pulled by a worker, sends an alert to supervisors that there may be a problem on the line). If the supervisor fails to resolve the potential problem, then pulling the andon cord may result in stopping the line—an enormously disruptive event in many continuous-flow production systems. Building an effective relational contract around the use of the andon cord—and around participation in problem solving more broadly—requires not only solving the credibility problem, but also developing a shared understanding of a host of subjective ideas. For example, both employees and managers must develop a shared understanding of questions such as which types of problems are worth pulling the cord for and how supervisors should respond in those cases in which the cord has been pulled inappropriately. They must also learn at what point, if any, will supervisors penalize workers (financially or socially) for mistakes in pulling the cord, and what kinds of rewards are appropriate when the cord has been used well.

Notice that even this simplified discussion of the use of the andon cord raises an important issue not present in the Trust Game in Section 2 or in our discussion of Lincoln's bonus. In the Trust Game, Player 1 chooses whether to trust Player 2, who then chooses whether to honor / cooperate or betray / defect. Likewise, in Lincoln's bonus, the worker is subject to the firm's discretion over what bonus to pay. With the andon cord, however, each party is being given discretion by the other: the worker to stop the line and the manager to implement worker suggestions and to support problem solving in "appropriate" ways. We suspect that such mutual dependence is more the rule than the exception. But in such settings—where both

parties can cooperate, defect, and punish—even more knowledge must be held in common if the relational contract is to be sustainable. To illustrate this mutual dependence, and more broadly the complexity of the knowledge on which relational contracts often rely, Table 1 sketches some of the potential actions by workers and managers that must be mutually understood in order to build a relational contract around the use of the andon cord.

Table 1: Cooperation, Defection, and Punishment in the use of the Andon cord at Toyota.

	Cooperate	Defect	Punish
Worker	Actions: 1. Pull the andon cord when you see a problem 2. Offer suggestions on improvements to the production process (that might make your job redundant)	Actions: 1a. Never pull the andon cord (out of fear of being punished) 1b. Pull the andon cord to stop the line and avoid work when there is no true problem 2. Keep improvements hidden from co-workers and managers	Actions (in response to perceived defection by manager): 1. Sabotage the manufacturing line 2. Pull andon cord frequently 3. Engage in absenteeism
Supervisor	Actions: 1. Recognize potential problem when andon cord pulled and aid in problem-solving 2. Implement improvements without necessarily cutting jobs 3. Accept authority of work teams to make some shop-floor decisions	Actions: 1. Punish workers for pulling andon cord (even appropriately) 2. Cut workforce once they discover potential innovations 3. Interfere in work teams and override their decisions	Actions (in response to perceived defection by worker): 1. Penalize workers (financially or socially) for pulling andon cord 2. Remove the andon cord

Just as our discussion of Lincoln’s bonus ignored a host of complementary practices that sustain performance at Lincoln, here our discussion of Toyota’s andon cord focuses on only one small aspect of the TPS. Nonetheless, even this small aspect illustrates the range of knowledge that may be required to build the necessary relational contract. Furthermore, while these issues are important in manufacturing settings such as Lincoln and Toyota, they are perhaps even more important in knowledge-work settings, which we illustrate next by discussing science-driven drug discovery at Merck.

For many decades, pharmaceutical research firms attempted to discover new drugs through a large-scale process of trial and error. For example, several successful psychotropic drugs were discovered by putting large numbers of potentially biologically active compounds into distressed rats and chemically tinkering with the few compounds that seemed to have some effect on the rats’

behavior (Henderson, 1994). This process of drug discovery relied primarily on the skills of highly trained medicinal chemists who could design and construct chemical compounds that were likely to have pharmacological effects. Even for very successful drugs, however, the specific biochemical mechanisms responsible for the drug's therapeutic effects were often not well understood.

Beginning in the 1980s, advances in the scientific literature offered new understandings of the biochemical mechanisms underlying many diseases, making it possible for drug candidates to be targeted at specific mechanisms. Pharmaceutical firms seeking to change from the old process of drug discovery (large-scale trial and error) to a new one (based on the mechanisms of disease) needed to undergo several significant changes (Henderson, 1994). First, completely new kinds of scientists had to be hired, such as molecular biologists. Second, the new process required rich communication across scientific disciplines and therapeutic areas, whereas the old process had performed well when conducted in disciplinary and functional silos. Finally (and most important, for our purposes), the new process required the firm's scientists not only to stay current with the promising mechanisms being discovered in the academic literature but also to conduct such leading-edge science themselves, in-house.

Perhaps not surprisingly, staying current with the academic literature involved more than sitting in the firm's lab and reading the journals. For example, there was great value in also attending conferences, but one typically could not fully understand (or, in some cases, even be invited to) many conferences without having one's own active research in the area—a classic instance of “absorptive capacity” (Cohen and Levinthal, 1990). As a result, some pharmaceutical firms encouraged their scientists to publish—even though of course an ultimate goal was to secure patents on drug candidates. Henderson and Cockburn (1994) call a firm “PROPUB” if an individual scientist's career inside the firm depended on the scientist's standing in the outside scientific community, and they find that patent output is significantly greater in PROPUB firms (even with dummy variables for therapeutic area and firm).

This new process has been labeled “science-driven drug discovery” (Cockburn, Henderson, and Stern, 2000). Note that it is not the mere presence of more or different scientists within pharmaceutical firms that prompts this label; for example, there were many synthetic chemists involved in the old process. Rather, the label reflects the fact that the new process within the firm relies so heavily on interactions with and assessments from the outside scientific community. It is this role for the scientific community (and the possible tension of this role with the firm's goal of patents on drug candidates) that we explore as a final example of both the presence of relational contracts underlying key managerial practices and the difficulties in building a relational contract, as follows.

Firms pursuing science-driven drug discovery needed scientists to behave almost as if they were academic scientists—to explore questions of fundamental

scientific interest, to publish in the refereed literature, and to attend academic conferences. At the same time, however, these scientists could not act completely like academics. They had to also take actions that increased the likelihood of the discovery of new drugs. Making a major scientific breakthrough—even winning a Nobel prize—without simultaneously generating knowledge that could be quickly applied to the search for new drugs would not constitute full success. Those firms that successfully developed this shared understanding, of whom Merck is the most well known, introduced more “significant” drugs and grew faster than any of their competitors, but the practice was relatively slow to diffuse across the industry (Cockburn, Henderson and Stern, 2000).

We see science-driven drug discovery as a complex managerial practice that cannot be sustained without a relational contract. No mechanistic formula could tell a research supervisor whether a particular scientist was appropriately straddling the fine line between behaving like an academic scientist and behaving like a commercially oriented drug hunter, and thus no formal contract alone could ensure that researchers were behaving as the firm desired. Staying connected to the academic world required publishing in the journals and attending conferences—but no one could specify how many papers a scientist should publish and which conferences he or she should attend, particularly when every scientist’s research trajectory was likely to be different. When was going to conferences vital to the research, and when was it consumption? And on the manager’s side, when a scientist had published extensively in a vital new field but failed to receive an increased research budget, was her manager reneging on the relational contract or responding appropriately to the fact that the scientist’s lab still had not produced any plausible drug candidates?

Table 2 sketches our conception of cooperation, defection, and punishment for the scientist and the manager in the relational contract underlying science-driven drug discovery. As with Lincoln’s bonus and Toyota’s production system our description here is necessarily partial, but it nonetheless illustrates the extent and complexity of the information that must be understood by all parties if the necessary relational contract is to be self-enforcing.

Table 2: Cooperation, Defection, and Punishment in Science-Driven Drug Discovery

	Cooperation	Defection	Punishment
Scientist	Action: behave almost like an academic scientist, but be sure to develop useful knowledge for discovering new drugs.	Action: either shirk (represent lack of results as unlucky research) or behave like an academic scientist (pursue problems for their own sake, build external reputation)	Action (in response to perceived defection by manager, and perhaps depending on nature of that defection): behave like an academic scientist, or ignore research and become a drug hunter

Manager	Action: reward the scientist who displays high-science behavior even if no new drugs result.	Action: fail to increase resources for scientists who publish; reward only those who produce drugs	Action (in response to perceived defection by scientist, and perhaps depending on nature of that defection): fire the scientist, or cut funding
---------	--	--	---

4) Building and Refining Relational Contracts

Our descriptions of some of the relational contracts in place at Lincoln Electric, Toyota and Merck have highlighted both that many significant managerial practices may rely on relational contracts and that the “relational knowledge” that is required to sustain these contracts may be very extensive.

In this section we build on these two ideas to focus more squarely on the key question of why relational contracts may be hard to build and refine—and hence hard to imitate. The three cases in section 3 constitute sampling on success, in the sense that those relational contracts did get built. For a sharper view of the difficulties the clarity problem can create, we now sample from the other end of the distribution, when failure to communicate and the resulting lack of shared understanding inhibited the development or refinement of relational contracts.

4.1 Building Clarity: Task and Relational Knowledge

In our descriptions of Lincoln, Toyota, and Merck, we saw that developing a shared understanding of the desired behaviors was not easy because there was uncertainty about both appropriate actions and (although we did not emphasize it as much) associated payoffs. We can state these difficulties abstractly using Figures 1 and 2. To participate in the relational contract involving those Figures, player 1 needs to know (a) what behaviors constitute cooperation by her (Trust rather than Not Trust), (b) what behaviors are then available to player 2 as cooperation or defection by him (Honor and Betray, respectively), (c) what payoffs player 2 would receive from those available behaviors (1 versus 2 in Figure 1, or C versus D in Figure 2), and (d) what payoffs player 1 would receive if everyone cooperates versus not (1 versus 0 in Figure 1).⁸

Of all the actions and payoffs described in (a) through (d), only (a) seems to us akin to what is sometimes called *task knowledge* (i.e., how one should behave in the organization), so we hereafter call (b) through (d) the additional *relational knowledge* needed to participate in a relational contract. As noted in the Introduction, other research traditions (e.g., Winter (1987, 1988) on tacit knowledge or Levinthal (1997), Milgrom and Roberts (1995) and Rivkin (2000) on complementarities) have made important strides exploring the difficulties of

⁸ Recall from Section 2.1 that player 1 needs to know (b) and (c) to assess whether player 2 will cooperate; if so, player 1 then needs to know (a) and (d) to decide whether to cooperate herself.

learning and communicating task knowledge. As hard as it may be to learn or communicate one's own task knowledge in (a), however, it must be at least as hard and presumably harder to learn someone else's task knowledge in (b), but player 1 cannot participate in the relational contract without this understanding of player 2's available behaviors, as well as the associated payoffs in (c). In this sense, there is simply more (and different and probably harder) knowledge to learn and communicate in the case of a relational contract than in the case of one's own task knowledge. Notice, too, that the task may be further complicated by the fact that if all goes well and both parties cooperate with each other, events off the equilibrium path—i.e., the behaviors and payoffs associated with both defection and punishment—cannot be observed but must be inferred.

As one striking example of the difficulty of communicating relational knowledge, we reconsider Lincoln Electric. After decades of high productivity and apparently appropriate bonuses in their Cleveland operations, Lincoln expanded overseas, with initially very negative results (Hastings, 1999). In fact, earnings in Cleveland were at record levels, but overseas losses were so large that the company as a whole was in the red. A new question thus arose about what constituted a "fair" bonus for Cleveland workers: should the bonus be based on Cleveland profits or on those of the firm as a whole? Naturally, when the firm had operations only in Cleveland, this distinction never arose, but it had suddenly become crucial. Formally, there was now deep ambiguity about which of management's actions would be considered "cooperation".

We provide the resolution of Lincoln's story in the next sub-section. For now we simply note that even a longstanding relational contract may not entail shared understanding of all the desired behaviors in all the possible circumstances that could arise. In a similar spirit, Ellison and Holden (2009) analyze a model where a principal attempts to teach an agent how to respond to fluctuating circumstances. Formally, the timing of each period is: (1) the agent observes the state of the world, (2) the agent chooses an action, (3) the principal observes the state and sends a message to the agent, and (4) both parties receive a common payoff that is positive if the agent choose the appropriate action for the state and zero otherwise. The novel assumption in the Ellison-Holden model is that there are some states of the world that the parties cannot discuss until they have experienced them together. In this setting, there can be more or less useful early realizations of the state, allowing more or less useful instructions from the principal to the agent. To the degree that this phenomenon is real it suggests that the development of relational contracts may be significantly path-dependent, with some states of the world enabling the creation of much more robust contracts than others; see also our discussion of Chassang (2010) below.

As a second case about lack of shared understanding, consider Stewart's (1993) account of how Credit Suisse, a large European bank, bought the US investment bank First Boston, taking the company private under the name CS First Boston (CSFB). Roughly speaking, in the first two years that Credit Suisse controlled

CSFB, all firms in the investment-banking industry performed poorly and employees received bonuses that were lower than the historical average but comparable to bonuses paid at other investment banks. In the third year, however, CSFB performed better than previously and yet worse than its competitors, and Credit Suisse paid bonuses that were above bonuses at CSFB in the first two years but lower than those at other firms in the third year. A crisis ensued. In brief, CSFB asserted that the bonus policy in their industry was match-the-market, meaning that bonuses should be competitive with bonuses at other top-bracket firms. In contrast, the Swiss asserted that in *their* industry the bonus policy was pay for performance, meaning that a banker's bonus depended on how he and his bank performed. Note that these two policies make identical pay prescriptions when all firms in the industry have the same performance, as was broadly true in the first two years. More generally, parties with different understandings may not appreciate that this difference exists until key events occur—with the further complication that behavior by one party during such an event may be intended as cooperation but nonetheless be viewed by the other as defection.

For inspiring evidence that different groups can indeed reach different shared understandings in similar environments, consider the lab studies by Weber and Camerer (2003) and Selten and Warglien (2007). Both explore common-interest games, where two players receive the same positive payoff if they successfully coordinate, but they receive no payoff if they fail to coordinate. More specifically, the players are attempting to build a shared language during repetitions of a game like the following: (1) player 1 observes the state of the world, $s \in S$; (2) player 1 sends a costless message to player 2, $m \in M$; (3) player 2 chooses an action $a \in S$; (4) payoffs to each player are $U(s, a) = 1$ if $a = s$ and $U(s, a) = 0$ otherwise; (5) player 2 observes s . Thus, player 1 would like to send a message $m(s)$ that tells player 2 that the state is s (so that player 2 will then choose the action $a = s$). The problem is that, at least in the early going, player 2 has little basis for understanding player 1's messages.

Both papers find that different pairs of players develop different languages, even though these pairs are playing in the same environment (except for the random realizations of the states of the world over time). Because these are common-interest games, we interpret the knowledge being communicated as task knowledge, not relational knowledge. In particular, there is no concern with defection or punishment in a common-interest setting. Nonetheless, consistent with the large literature on the difficulties of communicating tacit knowledge, different pairs of players take different lengths of time to develop a shared understanding, and different pairs hold different shared understandings once they reach them. From this lab evidence, we find it only a small stretch to imagine that similar forces could cause reasonable people to hold different understandings in situations like those at Lincoln or CSFB. Of course, the issue at these two companies involved bonuses, so these were not common-interest settings; rather, the problem of credibility also arose, as we discuss next.

4.2 *The Dynamics of Credibility*

Beyond the fact that implementing a relational contract requires communicating not only task knowledge but also relational knowledge, there is another difficulty that complicates learning and communicating a relational contract compared to learning and communicating task knowledge: relational contracts are relevant only when goal alignment is imperfect.⁹ Thus, whatever the information that needs to be communicated (task and/or relational), it may not be in someone's interest to reveal that information. Instead, information might be withheld or distorted. There is of course an enormous economics literature on these issues, and we note that discussion along these lines started in organization theory before information economics or organizational economics existed, such as in Cyert and March (1963).¹⁰ The question we face, however, is how imperfect alignment of interests affects parties' efforts to build a relational contract (as part of a managerial practice).

Returning to our accounts of Lincoln Electric and Credit Suisse First Boston, we find two different scenarios unfolding after a lack of shared understanding was discovered. At Lincoln Electric, where the question was whether Cleveland's bonuses should be based on Cleveland profits or on those of the firm as a whole, the firm ultimately decided in favor of the former principle (even though doing so required the firm to borrow in order to pay the bonus). At CSFB, in contrast, CS stuck to its policy of pay for performance (rather than the bankers' preferred policy of match the market), leading to the departure of many prominent bankers.

There are many ways to relate the dynamics of these cases to theories of the kind we have considered in this essay. Our point here is not to endorse particular theoretical interpretations of these two cases, but rather to surface the general issue that these cases present: how will parties respond when a possible misunderstanding surfaces? We now sketch three complementary approaches to this issue, in the hope that further theory and evidence on the dynamics of relationships will ensue.

One approach, in the spirit of Section 2.1, is to imagine that the critical moments in these cases correspond to sudden changes in payoffs. For example,

⁹ The part of organizational economics called "team theory," initiated by Marschak and Radner (1972), can be interpreted as exploring information acquisition, communication, and decision-making when all parties have the same interests (but gathering and communicating information are costly activities). See Garicano and Van Zandt (2012) for a recent discussion of approaches in this vein.

¹⁰ For example: "Where different parts of the organization have responsibility for different pieces of information relevant to a decision, we would expect some bias in information transmitted due to ... attempts to manipulate information as a device for manipulating the decision. ... [But] we cannot reasonably introduce the concept of communication bias without introducing its obvious corollary – 'interpretive adjustment'" (pp. 79 and 85, Cyert and March, 1963).

Credit Suisse may have reached a point where the defection temptation permanently increased from D to D' , along the lines of (3), causing the bank to defect. (To put this point less abstractly, it may be that what the Swiss wanted all along was a global financial supermarket, for which it needed *an* investment bank, but not necessarily a top-bracket investment bank, so they gave up on the latter when its price increased.) But this and any other analysis from Section 2.1 assumes that the parties have a shared understanding of what actions (in what states) constitute cooperation in their relationship, whereas the bankers at CSFB and their Swiss owners apparently did not have such a shared understanding. We therefore turn to models where the parties have something to learn.

Many authors have enriched the credibility models from Section 2.1 to include private information about players' payoffs, so that parties learn about their partners as an equilibrium progresses; see MacLeod and Malcomson (1988), Watson (1999, 2002), and Halac (2011), for example. As an illustration, in Figures 1 and 2, suppose that player 1 is uncertain about player 2's discount rate, which player 1 believes could be high (r_H) or low (r_L). In models like these, it can be optimal to terminate a relationship after learning enough bad news about one's partner's type (perhaps as at CS). Likewise, it can be optimal to enrich a relationship, such as moving from partial to full cooperation, after learning enough good news about one's partner's type. As a result, one could imagine player 2 searching for ways to signal that she is the patient (or "high trust") type, r_L , so as to induce greater cooperation from player 1.

Models of learning about one's partner often suggest empirical approaches based on unmeasurable managerial attributes. Such empirical work has a strong tradition, from at least Mundlak (1961) through at least Bertrand and Schoar (2003). But there is an alternative approach that also sheds light on performance differences, based on path-dependence rather than unmeasured heterogeneity. Chassang (2010) offers one model in this alternative spirit, in which a Principal and an Agent build a relational contract, as follows.

In each period, the Principal first chooses whether to invest or not, where investing imposes a cost k on the Principal but delivers a benefit b to the Agent (and not investing delivers zero to both parties and ends that period). If the Principal does invest then different actions from the feasible set A randomly become available that period, and both parties observe which actions are available that period.

In the feasible set A there are two kinds of actions, unproductive and productive. An unproductive action costs nothing for the Agent to take but produces no output for the Principal, whereas a productive action costs c to take and produces a given positive output with probability q and zero with probability $1-q$. It is common knowledge what the number of productive actions is and what a given productive action produces when it produces positive output, but initially only the Agent knows which actions are the productive ones.

In the first period, to induce the Agent to take a productive action (if one is available) instead of an unproductive action, the Principal threatens not to invest in several future periods if this period's output is zero. Note that this punishment will occur on the equilibrium path, because a productive action could produce zero output. In this sense, learning (i.e., identifying a new action as productive) is expensive. On the other hand, later in the game, if an action has already produced a positive output then the Principal knows that the action is productive, so if the Agent now takes this action and it produces zero output then the Principal does not need to punish the Agent.

Because learning is expensive (in the sense of punishments, and also in the sense of opportunity cost after at least one productive action has been identified), it can be optimal to stop learning before all productive actions are identified. Because opportunities to learn arrive randomly, otherwise identical dyads may stop learning after identifying different sets of productive actions. That is, each dyad converges to a steady-state relational contract, but different dyads can converge to different relational contracts that produce varying degrees of cooperation. Thus, Chassang's model can produce persistent performance differences among otherwise similar dyads because of path-dependence in building a relational contract—very much in the spirit of the administration problem that animated this essay.

5) Conclusion

An extensive literature has suggested that organizational capabilities are difficult to imitate both because they require the communication of task knowledge that is often deeply embedded in organizational routines and because problems of complementarity and local search mean that the processes of incremental learning that characterize many organizations make it difficult to communicate this knowledge. Here we have attempted to develop a complementary explanation for the often slow diffusion of competitively significant capabilities by focusing on the fact that many key managerial practices rely on relational contracts, and by suggesting that building these relational contracts requires moving beyond task knowledge to the development of “relational knowledge.” We suggested that relational knowledge may be substantially more difficult to develop than task knowledge both because there is much more of it and because its acquisition is complicated by incentive problems.

While we hope that this argument is compelling as a hypothesis, there is clearly much that remains to be done to prove its validity and determine its boundary conditions. For example, while it is well established that organizations are replete with relational contracts, and while we believe that many competitively significant managerial practices rely on relational contracts, careful empirical work that put these ideas to the test would clearly be useful. Similarly, empirical work that explored the development of relational contracts over time as an integral part of the development of managerial practices would also be of significant value. One

approach might be to begin with careful qualitative studies and then progress to more systematic, potentially survey-based work that could enable comparisons across firms over time.

We also suspect that these ideas, if they do indeed prove robust to careful empirical investigation, may have significant implications for managerial action. A considerable literature has explored the processes that enable firms to become “learning organizations” (see, for example, Argote, 1999 and Senge, 1993). One could imagine a complementary focus on the processes that enable firms to build relational contracts. Extensive and credible communication is almost certainly important, and we suspect that in many circumstances the ability to communicate convincingly that one is a “high trust” type would be very useful. Another possibility is that managers who excel in the development of relational contracts take advantage—or perhaps even create—situations that test the limits of the current contract to then take actions that refine the contract, such as by communicating credibly about otherwise unrevealed payoffs and preferences. Similarly it might be the case that successful managers develop the ability to change relational contracts without triggering the perception that they are reneging on existing agreements; Kaplan and Henderson (2005) suggest that the latter perception may be a significant barrier to organizational change.

Another potential extension of these ideas is to explore their implications for the role of organizational culture and ritual. One point of connection, for example, may flow from the fact that an effective relational contract tends to achieve cooperation, rather than defection. As a result, it may be easy to forget (or for newcomers never to have observed) what someone else’s opportunities for defection might be, or how tempting these opportunities might be, or what someone else’s opportunities for punishment might be, or how effective these punishments might be. This situation is akin to one where “means become ends,” where a person or group forgets *why* it does something, and instead remembers only *what* it is currently supposed to do, thus leaving itself unprepared to respond to fluctuating circumstances. Parties to a relational contract therefore might tell stories: to remind each other of what could occur, and to sketch appropriate behavior if it does.

More broadly we believe that our framing opens up some intriguing issues as to the relationship between the dynamics of conventional relational contracts (i.e., purely “calculative” trust) and other forms of social capital. As the extensive literature on trust has demonstrated, organizations are shot through with beliefs and expectations derived from personal and social dynamics that give rise to many different kinds of trust. We hope that our attempt to clarify the role of relational contracts may contribute to the development of a broader understanding of the relationship among different forms of trust and their role within the firm. Within the firm, for example, are calculative and affective and/or social trust complements or substitutes? Might it be the case that firms characterized by high levels of personal trust find it significantly easier to build and refine new relational contracts?

Organizational researchers have long suggested that the informal structures of firms are critically important to their performance—and that in some circumstances high-commitment work practices or the ability to sustain high levels of trust may be very powerful. Our hope is that by linking these ideas to the analytical construct of a relational contract we will be able to catalyze further research in the area and, ultimately, to support practicing managers in building effective organizations. If building relational contracts is as important as we believe it to be, research that yields insight into the mechanisms behind their development and the strategic choices on which they rest could make an important contribution on a wide variety of fronts.

References

Anand, B. N., T. Khanna. 2000. Do Firms Learn to Create Value? the Case of Alliances. *Strategic Management Journal* **21**(3) 295 –315.

Aral, S., P. Weill. 2007. IT assets, Organizational Capabilities, and Firm Performance: How Resource Allocations and Organizational Differences Explain Performance Variation. *Organization Science* **18**(5) 763–780.

Argote, Linda, 1999. *Organizational Learning: Creating, retaining and transferring knowledge*. Springer.

Bachmann, R., A. Zaheer. 2006. *Handbook of Trust Research*. Edward Elgar Publishing, Northampton Ma.

Baker, G., R. Gibbons, K. J. Murphy. 1994. Subjective Performance Measures in Optimal Incentive Contracts. *Quarterly Journal of Economics* **109**(4) 1125 –1156.

Baker, G., R. Gibbons, K. J. Murphy. 1999. Informal Authority in Organizations. *Journal of Law, Economics, and Organization* **15**(1) 56 –73.

Baker, G., R. Gibbons, K. J. Murphy. 2002. Relational Contracts and the theory of the Firm. *Quarterly Journal of Economics* **117**(1) 39–84.

Baker, G., R. Gibbons, K. J. Murphy. 2011. Relational Adaptation. Working Paper.

Barney, J. B. 1986. Organizational Culture: Can It Be a Source of Sustained Competitive Advantage? *The Academy of Management Review* **11**(3) 656–665.

Barney, J. B., M. H. Hansen. 1994. Trustworthiness as a Source of Competitive Advantage. *Strategic Management Journal* **15** 175–190.

Bertrand, M., A. Schoar. 2003. Managing with Style: The Effect of Managers on Firm Policies." *Quarterly Journal of Economics* **118**: 1169-1208

Blau, P. M., W. R. Scott. 1962. *Formal Organizations: A Comparative Approach*. Chandler Publishing Company, San Francisco, Ca.

Bloom, N., J. V. Reenen. 2007. Measuring and Explaining Management Practices Across Firms and Countries. *the Quarterly Journal of Economics* **122**(4) 1351 –1408.

Bresnahan, T., Greenstein, S. and R. Henderson, 2011. Schumpeterian Competition and Diseconomies of Scope: Illustrations from the Histories of Microsoft and IBM. Forthcoming in *The Rate and Direction of Inventive Activity, 50th Anniversary*

Volume, edited by Josh Lerner and Scott Stern, NBER and University of Chicago Press.

Brown, M., M. Serra-Garcia. 2010. Relational Contracting Under the Threat of Expropriation - Experimental Evidence. *Ebc Discussion Paper, 2010-18* 1–65.

Brown, M., C. Zehnder. 2007. Credit Reporting, Relationship Banking, and Loan Repayment. *Journal of Money, Credit and Banking* **39**(8) 1883–1918.

Bull, C. 1987. The Existence of Self-Enforcing Implicit Contracts. *The Quarterly Journal of Economics* **102**(1) 147 –159.

Christensen, C. M. 1997. *The Innovator's Dilemma*. Harvard Business Press.

Cockburn, I. M., R. M. Henderson, S. Stern. 2000. Untangling the Origins of Competitive Advantage. *Strategic Management Journal* **21**(Fall) 1123–1145.

Cohen, W. M., D. A. Levinthal. 1990. Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly* **35**(1) 128–152.

Crawford, Vincent P., Sobel, Joel. 1982. Strategic Information Transmission. *Econometrica*, Vol. 50(6):1431-51.

Cyert R., March J. 1963. *A Behavioral Theory of the Firm*. Oxford, UK: Blackwell.

Dal Bo, P. 2005. Cooperation Under the Shadow of the Future: Experimental Evidence from infinitely Repeated Games. *American Economic Review* **95**(5) 1591–1604.

Duffy, J., J. Ochs. 2009. Cooperative Behavior and the Frequency of Social Interaction. *Games and Economic Behavior* **66**(2) 785–812.

Dyer, J. H. 1997. Effective Interfirm Collaboration: How Firms Minimize Transaction Costs and Maximize Transaction Value. *Strategic Management Journal* **18**(7) 535–556.

Dyer, J.H., and H. Singh. 1998. The Relational View: Cooperative Strategy and Sources of Interorganizational Competitive Advantage. *Academy of Management Review* **23**(4) 660-679.

Eisenhardt, K. M., J. A. Martin. 2000. Dynamic Capabilities: What Are They? *Strategic Management Journal* **21**(Fall) 1105–1121.

Engle-Warnick, J., R. L. Slonim. 2004. The Evolution of Strategies in A Repeated Trust Game. *Journal of Economic Behavior & Organization* **55**(4) 553–573.

Fast, N., N. Berg. 1975. The Lincoln Electric Company. Harvard Business School Case 376--028.

Fuchs, W. 2007. Contracting With Repeated Moral Hazard and Private Evaluations. *The American Economic Review* **97**(4) 1432–1448.

Garicano, Luis and Timothy Van Zandt. 2012. "Hierarchies and the Division of Labor." Forthcoming in R. Gibbons and J. Roberts (eds.), *The Handbook of Organizational Economics*. Princeton, NJ: Princeton University Press.

Geertz, C. 1962. The Rotating Credit association: A "Middle Rung" in Development. *Economic Development and Cultural Change* **10**(3) 241–263.

Geertz, C. 1978. The Bazaar Economy: information and Search in Peasant Marketing. *The American Economic Review* **68**(2) 28–32.

Gibbons, R. 2005. Incentives Between Firms (and Within). *Management Science* **51**(1), pp. 2-17.

Gulati, R., J. A. Nickerson. 2008. Interorganizational Trust, Governance Choice, and Exchange Performance. *Organization Science* **19**(5) 688–708.

Guriev, S., A. Kolotilin, K. Sonin. 2011. Determinants of Nationalization in the Oil Sector: A Theory and Evidence from Panel Data. *Journal of Law, Economics, and Organization* Forthcoming.

Halac, M. 2011. Relational Contracts and the Value of Relationships. *American Economic Review* Forthcoming.

Hastings, D. F. 1999. Lincoln Electric's Harsh Lessons from International Expansion. *Harvard Business Review* **77** 162–180.

Henderson, R. 1994. The Evolution of integrative Capability: innovation in Cardiovascular Drug Discovery. *Industrial and Corporate Change* **3**(3) 607 –630.

Henderson, R., I. Cockburn. 1994. Measuring Competence? Exploring Firm Effects in Pharmaceutical Research. *Strategic Management Journal* **15**(S1) 63–84.

Henderson, R. M., K. B. Clark. 1990. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly* **35**(1) 9–30.

Hoffer Gittell, Jody, 2002. *The Southwest Airlines Way*, McGraw-Hill

Kaplan, Sarah and Rebecca Henderson, 2005. "Bridging Organizational Economics and Organizational Theory", September-October 2005, Vol 16., No 5., pp 509-521.

- Kochan, T. A., H. C. Katz, R. B. Mckersie. 1995. *The Transformation of American Industrial Relations*. Cornell University Press.
- Kreps, D. M. 1990. Corporate Culture and Economic Theory. J. Alt, K. Shepsle, Eds., *Perspectives On Positive Political Economy*. Cambridge Univ Pr.
- Lazzarini, S., G. Miller, T. Zenger. 2004. Order with Some Law: Complementarity vs. Substitution of Formal and Informal Arrangements. *Journal of Law, Economics, and Organization*, Issue 20(2), 261-298.
- Levinthal, D. 1997. Adaptation on rugged landscapes. *Management Science*, vol. 43: 934-950.
- Levin, J. 2003. Relational Incentive Contracts. *American Economic Review* **93**(3) 835-857.
- Lincoln, James and Arne Kalleberg *Culture, Control and Commitment*, Cambridge University Press, 1990
- Macaulay, S. 1963. Non-contractual relations in business: a preliminary study. *American Sociological Review*, 28: 55-67.
- Macleod, W. B., J. M. Malcomson. 1989. Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment. *Econometrica* **57**(2) 447-480.
- Malcomson, J. M. 2012. Relational Incentive Contracts. R. Gibbons, J. Roberts, Eds., *the Handbook of Organizational Economics*. Princeton University Press, Princeton, Nj, Forthcoming.
- March, J. G., J. P. Olsen. 1989. *Rediscovering Institutions: The Organizational Basis of Politics*. Free Pr, New York.
- Marschak, Jacob, and Roy Radner. 1972. *Economic Theory of Teams*. New Haven, CT: Yale University Press.
- Mcmillan, J., C. Woodruff. 1999. Interfirm Relationships and Informal Credit in Vietnam. *Quarterly Journal of Economics* **114**(4) 1285-1320.
- Milgrom, Paul and John Roberts. 1995. Complementarities and Fit: Strategy, Structure, and Organizational Change in Manufacturing. *Journal of Accounting and Economics*, vol.19: 179-208.
- Mundlak, Y. 1961. Empirical Production Function Free of Management Bias. *Journal of Farm Economics* 43: 44-56.

- Nooteboom, B. 1996. Trust, Opportunism and Governance: A Process and Control Model. *Organization Studies* **17**(6) 985–1010.
- Ohno, T. 1988. *Toyota Production System: Beyond Large-Scale Production*. Productivity Press.
- Ostrom, E. 1990. *Governing the Commons: The Evolution of institutions for Collective Action*. Cambridge University Press.
- Pil, F.K. and J.P. MacDuffie. 1996. The Adoption of High-Involvement Work Practices. *Industrial Relations*. **35**(3), 423-455.
- Poppo, L., T. Zenger. 2002. Do Formal Contracts and Relational Governance Function as Substitutes Or Complements? *Strategic Management Journal* **23**(8) 707–725.
- Reinganum, J.F., 1989, "The Timing of Innovation: Research, Development and Diffusion", in Schmalensee, R. and R. Willig, eds., *Handbook of Industrial Organization, Volume 1*, Amsterdam: North Holland.
- Rivkin, J. W. 2000. Imitation of Complex Strategies. *Management Science* **46**(6) 824–844.
- Robinson, D. T., T. E. Stuart. 2007. Financial Contracting in Biotech Strategic Alliances. *Journal of Law and Economics* **50**(3) 559–596.
- Rotemberg, J. J., G. Saloner. 1986. A Supergame-theoretic Model of Price Wars During Booms. *The American Economic Review* **76**(3) 390–407.
- Rousseau, D. M. 1989. Psychological and Implied Contracts in Organizations. *Employee Responsibilities and Rights Journal* **2**(2) 121–139.
- Rousseau, D. M. 1995. *Psychological Contracts in Organizations: Understanding Written and Unwritten Agreements*. Sage Publications, inc, London Uk.
- Rousseau, D. M., S. B. Sitkin, R. S. Burt, C. Camerer. 1998. Not So Different After All: A Cross-Discipline View of Trust. *Academy of Management Review* **23**(3) 393–404.
- Selten, Reinhard and Massimo Warglien. 2007. "The emergence of simple languages in an experimental coordination game." *Proceedings of the National Academy of Sciences* 104: 7361-66.
- Senge, Peter. 1993. *The Fifth Discipline: The art and practice of the learning organization*. Crown Business Press.

- Spence, A. Michael. 1973. Job market signaling. *Quarterly Journal of Economics*, 87:355-374.
- Stewart, J. 1993, July 26. Taking the dare. *The New Yorker* 34-39.
- Syverson, C. 2011. What Determines Productivity? *Journal of Economic Literature* Forthcoming.
- Teece, D. J., G. Pisano, A. Shuen. 1997. Dynamic Capabilities and Strategic Management. *Strategic Management Journal* **18**(7) 509–533.
- Tripsas, M., G. Gavetti. 2000. Capabilities, Cognition, and Inertia: Evidence from Digital Imaging. *Strategic Management Journal* **21**(Fall) 1147–1161.
- Watson, J. (1999), 'Starting small and renegotiation,' *Journal of Economic Theory*, 85, 52–90.
- Watson, J. (2002), 'Starting small and commitment,' *Games and Economic Behavior*, 38, 176–199.
- Weber, Roberto and Colin Camerer. 2003. "Cultural Conflict and Merger Failure: An Experimental Approach." *Management Science* 49: 400-15.
- Wernerfelt, B. 1984. A Resource-Based View of the Firm. *Strategic Management Journal* **5**(2) 171–180.
- Williamson, O. E. 1993. Calculativeness, Trust, and Economic Organization. *Journal of Law & Economics* **36**(April) 453–486.
- Winter, S. G. 1987. Knowledge and Competence as Strategic assets. *Handbook On Knowledge Management*. Ballinger, 159–184.
- Winter, S. G. 1988. On Coase, Competence, and the Corporation. *Journal of Law, Economics, and Organization* **4**(1) 163–180.
- Womack, J. P., D. T. Jones, D. Roos. 1991. *The Machine That Changed the World: the Story of Lean Production*. Harperbusiness, New York.
- Yamagishi, T., M. Yamagishi. 1994. Trust and Commitment in the United States and Japan. *Motivation and Emotion* **18**(2) 129–166.
- Zaheer, A., N. Venkatraman. 1995. Relational Governance as An interorganizational Strategy: An Empirical Test of the Role of Trust in Economic Exchange. *Strategic Management Journal* **16**(5) 373–392.

Zaheer, A. and N. Venkatraman, "Relational governance as an interorganizational strategy: An empirical test of the role of trust in economic exchange", *Strategic Management Journal*, 16(5), (1995), pp. 373-392.