



**DEPARTMENT OF ECONOMICS
DISCUSSION PAPER SERIES**

RELATIONAL INCENTIVE CONTRACTS

James M. Malcomson

Number 508
October 2010

Manor Road Building, Oxford OX1 3UQ

Relational Incentive Contracts

James M. Malcomson*

University of Oxford

**Chapter for Handbook of Organizational Economics (edited by
Robert Gibbons and John Roberts), to be published by
Princeton University Press**

October 13, 2010

*I would like to thank the many people, too many to list individually, who have assisted me in writing this chapter. I am, however, particularly indebted to Robert Gibbons, Bentley MacLeod, John Roberts, Giancarlo Spagnolo and Michael Whinston, who read and commented on complete drafts. I would also like to thank Leverhulme Trust Major Research Fellowship F/08519/B for financial support of my research on relational contracts.

Abstract

This chapter reviews the literature on the theory of relational incentive contracts. It motivates the discussion by the classic applications of relational contracts to the GM–Fisher Body relationship and the relationships between Japanese automobile manufacturers and their subcontractors. It presents basic models with symmetric information to illustrate the fundamental issues and then goes on to consider specific investments, the role of legally enforceable contracts alongside relational contracts, private information, multiple suppliers, and issues of organization design.

Keywords: Relational contracts, informal enforcement, legal enforcement, incentives, private information, partnerships, vertical integration, organization design.

JEL classifications: D21, D23, D82, D86, L14, L22, L23, L24

Contents

1	Introduction	1
2	Classic applications of relational contracts	3
3	Two basic models	5
3.1	“Employment” model	5
3.2	Outsourcing model	11
4	Specific investments	14
5	Verifiable performance measures	17
6	Private information: non-persistent types	25
7	Private information: persistent types	33
8	Multiple suppliers	40
9	Partnerships	44
10	Organization design	51
11	Rent and relational contracts	59
12	Courts and empirical evidence	61
13	Concluding remarks	62

1 Introduction

This chapter is about relational contracts, agreements for which the on-going relationship between the parties plays an essential role in determining what happens. Its concern is with *relational incentive contracts*, those that involve incentives for taking actions, and not, for example, those involving only risk sharing, such as the relational contracts studied by Thomas and Worrall (1988) and surveyed in Malcomson (1999, Section 3). It focuses on supply relationships between firms, not purely financial contracts, sovereign debt (discussed in Obstfeld and Rogoff (1996, Chapter 6)) or employment (surveyed in Malcomson (1999, Section 5)). The same underlying principles apply to all these, though obviously details may differ, but because of the present focus the terminology of supply relationships between firms is used here even for results originally derived in other contexts. This chapter also focuses primarily on theory. Empirical evidence on interfirm contracts is surveyed in the chapter by Lafontaine and Slade (2011).

Some decades ago, Macaulay (1963) observed that many of the contractual relationships between firms in the US were not tightly specified in legal terms. A striking example is the so-called “battle of the forms”. Orders for supply of goods are made on standard forms with, on the front, details of what is to be supplied and, on the reverse, the purchaser’s standard terms of business. The supplier’s acknowledgement of the order also has standard terms of business specified on the reverse. But the purchaser’s and the supplier’s terms differ and no attempt is made to reconcile them. Such informality concerned lawyers interviewed for Macaulay (1963), some of them in-house lawyers for large firms employed to give legal advice on contracts who might thus be expected to influence the contracts used. Yet purchasing and supplying agents persisted in using loosely specified contracts with, seemingly, the view that there was so much business between the parties that they would work something out if either was dissatisfied. One purchasing agent went so far as to say: “One doesn’t run to lawyers if he wants to stay in business because one must behave decently” (Macaulay (1963, p. 61)).

The literature on relational contracts is concerned with the impact of the on-going nature of the relationship on trade between the parties, on their payoffs, on the nature of any legally enforceable contract that is used to supplement the relational contract, and on the design of organizations. Legal aspects of relational contracts have been explored by Macneil (1978). The underlying economic idea was explored early on by Telser (1980)

and Klein and Leffler (1981). If a supplier does not deliver something acceptable, the purchaser can look elsewhere in future. Provided the supplier anticipates a profit from continuing to trade with the purchaser, that acts as an incentive to provide what the purchaser wants. So there is a tension between providing something less costly, but less acceptable, now and keeping business in the future. It is a tension widely recognized in the business world. But the future incentive applies only if the supplier anticipates a strictly positive profit in the future, so the relationship has to be one in which free entry does not compete away all profits, perhaps because there are relationship specific investments, perhaps because the parties learn information about each other that is not available about other potential purchasers and suppliers, perhaps just because other potential purchasers and suppliers are wary of dealing with those who have had trading relationships come to an end with no apparent good economic reason. Moreover, the supplier will not incur more in costs to keep the relationship going than the value of any future profits. That limits what can be sustained. Those agreements that can be sustained in this way are said to be *self-enforcing* because they do not need courts to enforce them.

There are a number of strands to the literature on the economics of relational incentive contracts. One strand builds on the traditional principal-agent literature. In the most basic version, there is a principal (purchaser) who uses a single agent (supplier) to provide something with characteristics for which there are no verifiable measures that can be used to write a formal contract. That is the starting point for Section 3 of this chapter. This is extended to cases in which there is some verifiable measure but it does not capture all the parties care about and in which there are multiple agents (suppliers), issues taken up in Sections 4 to 8. In this strand, the principal does not contribute to the relationship except to pay for, and make use of, what is provided. In a second strand, all parties take actions that contribute to what is produced, as in a partnership, team or joint venture. This raises the issue, discussed in Section 9, of how each party is rewarded as a function of whatever measure of joint output is available. Yet another strand is concerned with organization design — whether the purchaser and supplier remain separate firms or are vertically integrated to produce components in-house, and who has the right to make what decisions, issues taken up in Section 10. The chapter concludes with some other issues arising from relational contracts. But first, in Section 2, it discusses some classic applications of relational contracts and the issues they raise for theory.

2 Classic applications of relational contracts

Although this chapter is primarily about theory, it is helpful to illustrate the issues that arise in relational contracts with two classic cases, the relationship between Fisher Body and General Motors (GM) in the period 1919-26 studied in detail in Klein et al. (1978) and Klein (2000), and the relationships between large Japanese automobile and electrical goods manufacturers and their suppliers studied by Asanuma (1989).

Klein (2000) describes the GM–Fisher Body relationship as follows. In 1919, GM agreed a 10-year exclusive dealing contract with Fisher Body for the supply of bodies for its automobiles. This was a relational contract, supplemented by a cost plus contract for each unit delivered specifying that GM would pay an amount equal to Fisher’s variable cost plus 17.6%. The relationship worked well up to 1924 but then became problematic for two reasons given by Alfred Sloan, the chief executive of GM. First, the cost-plus contract became unduly burdensome with the increase in the volume of bodies required by GM. Second, Fisher Body became unwilling to put up large amounts of capital to establish body plants near GM assembly plants, especially “Fisher’s refusal to make the investment in plant and equipment to move Fisher’s production of Buick bodies from Detroit to a site near the GM Buick assembly plant in Flint” (Klein (2000, p. 116)). The upshot was that GM took over Fisher Body. The interpretation Klein (2000) places on these reasons is that the *plus* element in the probably legally enforceable cost-plus contract became excessive in relation to capital and overhead costs given higher than anticipated volume and that the specific investments required to build body plants near GM assembly plants were too large to keep the arrangement within the zone of a self-enforcing relational contract.

Klein’s interpretation of these events has been disputed — see for example, the other papers in the issue of the *Journal of Law and Economics* in which Klein (2000) appears. But, whether or not correct, it raises interesting questions for theory. First, in what way might a larger specific investment make it more difficult to sustain a relational contract? Second, what is the role of an additional explicit contract, such as GM’s problematic cost-plus contract with Fisher Body, within a relational contract? And third, what does theory have to tell us about when vertical integration, which brings the incentive issues within a single firm, is to be preferred to a relational contract between two firms? All these are discussed in this chapter.

Asanuma (1989) carried out a detailed investigation of the contractual relationships between large Japanese companies in automobile and electrical machinery manufacturing and their suppliers. As he describes it, these large purchasers typically each use a small number of associated companies. They normally have a basic contract lasting initially for one year but renewed unless there are objections by either party. The basic contract specifies general obligations. It also specifies that orders (or monthly schedules in the case of automobiles) are to be regarded as contracts once accepted by the supplier and the frequency at which prices are to be renegotiated (typically 6 months). There is a general expectation that suppliers will reduce prices over time as a production run lengthens and that less highly regarded suppliers will be used as capacity buffers.

Asanuma (1989) attaches considerable importance to the rating system purchasers use for suppliers. To explain this, he considers suppliers in three categories. The first are those he terms *suppliers in general*, those who supply “off the shelf” goods (*purchased goods*), for example, from catalogues. For these, a long-term relationship may not be important. The other two categories are for suppliers of *ordered goods*. Here suppliers are differentiated between *common subcontractors* who produce components to designs provided by the purchaser (*drawings supplied*) and *excellent subcontractors* who produce components to their own design that has been approved by the purchaser (*drawings approved*). The latter are subcontractors who have shown themselves to be responsible suppliers in the past and are given a more responsible role in production. Typically, any change in status is associated with a change in design of the part supplied (with, for example, a new model of automobile), when the supplier would have to make new specific investments in any case. It is rare for a purchaser to switch supplier during the lifetime of a product model, so it is times at which a model changes that are most competitive.

The description provided by Asanuma (1989) raises some further questions for theory. When purchasers face, and want to find out more about, suppliers of potentially different types, what difference does it make to the relational contracts between the parties? And what does theory have to say about the process of dividing suppliers into different rating categories on the basis of past experience, giving more highly-rated ones additional responsibilities?

3 Two basic models

The role of relational contracts is isolated most clearly when a single purchaser buys from a single supplier and there is no measure of the supplier's performance that is verifiable in court. The purchaser's profit in period t from being supplied with quality $q_t \in [\underline{q}, \bar{q}]$ is $y_t(q_t) - P_t$, where P_t is the payment to the supplier. Quality can be multi-dimensional, with q_t a vector. It is essentially a short-hand term for anything under the control of the supplier that is of concern to the purchaser. Suppose that y_t is strictly increasing and concave (so, whenever $y_t(q_t)$ is twice differentiable, $y'_t(q_t) > 0$ and $y''_t(q_t) \leq 0$ for all q_t) and that $y_t(\underline{q}) = 0$ so that the purchaser would not be prepared to pay for quality at the lowest level. The supplier incurs cost $c_t(q_t)$ from supplying quality q_t and so makes profit $P_t - c_t(q_t)$ in period t if paid P_t for doing so. Conventionally, c_t is continuous, strictly increasing and strictly convex (so, whenever $c_t(q_t)$ is twice differentiable, $c'_t(q_t) > 0$ and $c''_t(q_t) > 0$ for all q_t), $c_t(\underline{q}) = 0$, and $\lim_{q_t \rightarrow \bar{q}} c_t(q_t) = \bar{c}$ with \bar{c} finite. If the parties choose not to trade with each other in period t , the purchaser receives payoff $\pi_t > 0$ (perhaps from producing the component in-house) and the supplier payoff $u_t > 0$ (from using the capacity to supply elsewhere). The parties have a common discount factor $\delta_t \in (0, 1)$. The purchaser checks the quality supplied with probability $\rho \in (0, 1]$.

When quality is unverifiable, it is not possible to write a legally enforceable contract that makes payment conditional on quality. However, as Bull (1987) and MacLeod and Malcomson (1989) observed, the parties can still use conditional payments provided the relationship is structured in such a way that it is worthwhile for the purchaser to make them. For that reason, it may be useful to have the payment P_t consist of two components, a legally enforceable component p_t that cannot be conditioned on quality and a voluntary bonus $b_t \geq 0$ that can. How these components relate to the types of payments actually made in practice is discussed later. Following Levin (2003), we can think of the purchaser and supplier agreeing informally a plan for q_t, p_t and b_t for the duration of their relationship. This is their *relational incentive contract*.

3.1 “Employment” model

In what has been termed the *employment model* in the literature, once the cost of supply has been incurred, the supplier cannot hold back supply until payment is made. The




Figure 1: Timing of events for period t in employment model

timing of events for period t in this model is shown in Figure 1. At the beginning of each period, the purchaser commits to the legally enforceable payment p_t .¹ The supplier then decides what quality to deliver. The purchaser monitors the quality delivered with probability ρ , pays p_t and decides whether to pay b_t . Finally in this period, the relationship may be terminated by either party, either because it is no longer even potentially mutually beneficial or because one or other party decides to end it for some other reason.

The term *employment* should not be taken too literally. Although this model is clearly appropriate for employment, it also applies to supply of many types of services. With this structure, no supply occurs if the relationship lasts for at most one period. With no future to be concerned about, the purchaser has no incentive to pay a positive bonus. Anticipating this, the supplier has no incentive to deliver quality above \underline{q} . With $q_t = \underline{q}$, the joint gains from the relationship, $y_t(\underline{q}) - c_t(\underline{q}) = 0$, are less than the joint payoffs $\pi_t + u_t$ the parties could get from making alternative arrangements, so there is no price p_t at which they would trade with each other. Thus the possibility of trade depends on the relationship being at least potentially on-going.

In a potentially on-going relationship, the supplier's future payoff from t on for a relationship started at τ if both parties stick to their agreement consists of the payment $p_t + b_t$ less the cost of supply $c_t(q_t)$ in period t plus the expected payoff from the future. It can be written

$$U_t = p_t + b_t - c_t(q_t) + \delta_t U_{t+1}, \quad \text{for all } t \geq \tau. \quad (1)$$

If, on the other hand, the supplier delivers quality in breach of the agreement, this is undetected with probability $1 - \rho$ and everything then continues as if no breach had occurred, with the supplier receiving payoff $p_t + b_t + \delta_t U_{t+1}$. The breach is detected with

¹In the employment model of Section 3.1, p_t could be paid at this stage, thus eliminating the need for it to be legally enforceable. But in models discussed later, p_t is conditioned on verifiable information not available until after the supplier has set q_t and then the commitment is to a legally enforceable contract specifying these conditional payments.

probability ρ but the worst that can happen to the supplier is that the relationship comes to an end with payment of p_t , but without payment of the bonus b_t , because the supplier can always choose to quit at the end of period t . Let \tilde{U}_{t+1} denote the supplier's payoff from $t + 1$ on in those circumstances. Thus, by setting $q_t = \underline{q}$, the supplier can obtain a payoff from t on of no less than

$$p_t + (1 - \rho)(b_t + \delta_t U_{t+1}) + \rho \delta_t \tilde{U}_{t+1}, \quad \text{for all } t \geq \tau. \quad (2)$$

From (1) and (2), it is certainly better for the supplier to set $q_t = \underline{q}$ than to stick to any agreement with $q_t > \underline{q}$ unless the following incentive compatibility condition holds

$$\delta_t \{\text{future gain to supplier}\} \equiv \delta_t (U_{t+1} - \tilde{U}_{t+1}) \geq \frac{c_t(q_t)}{\rho} - b_t, \quad \text{for all } t \geq \tau. \quad (3)$$

The interpretation of this condition is most straightforward when $\rho = 1$: the future gains from having the relationship continue at the end of period t must then exceed the cost $c_t(q_t)$ of complying less the current period incentive to supply provided by the bonus b_t . Moreover, for the supplier to continue in the relationship when there has been no breach, the following individual rationality condition must also hold

$$U_t \geq \bar{U}_t, \quad \text{for all } t \geq \tau, \quad (4)$$

where \bar{U}_t denotes the supplier's payoff from t on if the relationship ends for reasons unrelated to the performance of either party. In principle, \bar{U}_t need not be the same as \tilde{U}_t .

The purchaser's future payoff from t on for a relationship started at τ if both parties stick to their agreement consists of the value $y_t(q_t)$ derived from supply less the payment $p_t + b_t$ in period t plus the expected payoff from the future. It can be written

$$\Pi_t = y_t(q_t) - p_t - b_t + \delta_t \Pi_{t+1}, \quad \text{for all } t \geq \tau. \quad (5)$$

If the purchaser does not pay the agreed bonus, the worst that can happen is that the relationship comes to an end because the purchaser can always choose to quit at the end of period t . Let $\tilde{\Pi}_{t+1}$ denote the purchaser's payoff from $t + 1$ on in those circumstances. Thus, by setting $b_t = 0$, the purchaser can obtain a payoff from t on of no less than

$y_t(q_t) - p_t + \delta_t \tilde{\Pi}_{t+1}$. Thus it is certainly better for the purchaser to pay no bonus than to stick to any agreement with $b_t > 0$ unless the following incentive compatibility condition holds

$$\delta_t \{\text{future gain to purchaser}\} \equiv \delta_t (\Pi_{t+1} - \tilde{\Pi}_{t+1}) \geq b_t, \quad \text{for all } t \geq \tau. \quad (6)$$

Moreover, for the purchaser to continue in the relationship when there has been no breach, the following individual rationality condition must clearly hold

$$\Pi_t \geq \bar{\Pi}_t, \quad \text{for all } t \geq \tau, \quad (7)$$

where $\bar{\Pi}_t$ denotes the purchaser's payoff from t on if the relationship ends for reasons unrelated to the performance of either party. In principle, $\bar{\Pi}_t$ need not be the same as $\tilde{\Pi}_t$.

Let $S_{t+1} = U_{t+1} + \Pi_{t+1}$ denote the joint payoff (or *surplus*) the parties receive from $t + 1$ on if the relationship continues and $\tilde{S}_{t+1} = \tilde{U}_{t+1} + \tilde{\Pi}_{t+1}$. Then (3) and (6) can be added to give the pooled incentive compatibility condition

$$\delta_t (S_{t+1} - \tilde{S}_{t+1}) \equiv \delta_t (U_{t+1} - \tilde{U}_{t+1} + \Pi_{t+1} - \tilde{\Pi}_{t+1}) \geq \frac{c_t(q_t)}{\rho}, \quad \text{for all } t \geq \tau. \quad (8)$$

It is clear from the way they have been derived that (4), (7) and (8) are necessary conditions for the parties to stick to any agreement with $q_t > 0$, that is, for the agreement to be self-enforcing. Following the argument in MacLeod and Malcomson (1989), they are also sufficient. Formally, provided an agreement satisfies (4), (7) and (8), there exist strategies implementing the agreement that form a subgame perfect equilibrium. Hence, (4), (7) and (8) provide a complete characterization of subgame perfect equilibria with qualities strictly greater than \underline{q} .²

The right-hand side of (8) is the cost of delivering quality q_t adjusted for the probability that cheating on the agreement is detected. That is sometimes referred to as the *reneging temptation*. The left-hand side of (8) is the joint future gain to the parties from continuing the relationship at the end of period t . Thus (8) requires that the joint future gain exceeds the reneging temptation. The joint future gain is independent of the payments made by

²This result does not depend on payment of p_t being legally enforceable for the reason given in footnote 1.

the purchaser to the supplier — these cancel when U_{t+1} and Π_{t+1} are added — so the only element of the relational incentive contract that affects whether (8) is satisfied is the sequence of qualities (q_t, q_{t+1}, \dots) from t on. For any sequence of qualities that satisfies (8), it is always possible to find a sequence of bonuses (b_t, b_{t+1}, \dots) such that (3) and (6) are both satisfied for all t . Only the joint future gain, not its division between purchaser and supplier, enters (8) for the following reason. Future gain to the supplier affects the quality that can be delivered directly via the supplier's incentive compatibility condition (3). But future gain to the purchaser can do so too because it can be used to induce the purchaser to pay a bonus via the purchaser's incentive compatibility condition (6) and that bonus affects the quality that can be delivered via (3).

As already explained, in a one-shot game the purchaser would never pay a positive bonus, so the supplier would always set quality at the lowest level. Thus the one-shot game has a prisoners dilemma type structure and one can apply the intuition for the standard Folk Theorem result in a repeated prisoners dilemma game. Suppose δ_t is constant at δ for all t . The Folk Theorem result tells us that, for a repeated prisoners dilemma game, any individually rational co-operation can be sustained for a discount factor δ sufficiently close to 1 by strategies that revert to those of the one-shot equilibrium if either party deviates. Consider any sequence of quality (q_t, q_{t+1}, \dots) from t on for which the left-hand side of (8) is positive for all t . As $\delta \rightarrow 1$, that left-hand side goes to infinity so any such sequence satisfies (8). Equality in (8) specifies the critical value of the discount factor for sustaining any particular sequence.³ Provided supply of this sequence of qualities is mutually beneficial, there also exists a sequence (p_t, p_{t+1}, \dots) such that (4) and (7) are satisfied for all t . Then the on-going nature of the relationship enables the purchaser to elicit quality above the minimum level from the supplier when that would not be possible in a one-period relationship.

Any sequences that satisfy (4), (7) and (8) can be sustained as equilibria. But some equilibria may be better than others. Let q_t^* denote efficient quality defined by

$$q_t^* = \arg \max_q y_t(q) - c_t(q), \quad \text{for all } t \geq \tau. \quad (9)$$

The purchaser and supplier can increase their joint payoff $U_\tau + \Pi_\tau$ at the start of the

³But see Blonski et al. (2009) on whether this value is actually critical for determining whether the parties sustain such a sequence.

relationship at τ unless, for each $t \geq \tau$, either (a) $q_t < q_t^*$ and (8) holds with equality, or (b) $q_t = q_t^*$. To see why, note that the component of the joint payoff at date t' , $U_{t'} + \Pi_{t'}$, arising at date $t \geq t'$ is just $y_t(q_t) - c_t(q_t)$, which q_t^* maximizes. Thus the joint payoff $U_{t'} + \Pi_{t'}$ is always increased by having q_t move towards q_t^* for $t \geq t'$. So moving q_t towards q_t^* increases the left-hand side of (8) for all dates prior to t . Suppose $q_t > q_t^*$. Then the parties can always increase their joint payoff at τ because reducing q_t increases the joint payoff for all dates up to t , so relaxing the constraint (8) for all dates prior to t by increasing the left-hand side, and also relaxes the constraint (8) for t by reducing the right-hand side. Appropriate choice of p_t and b_t can ensure that neither U_t nor Π_t is reduced. Suppose now $q_t < q_t^*$. If (8) held with strict inequality at t , q_t could be increased without violating that constraint and this would also increase the left-hand side of (8) for all dates prior to t , so relaxing the constraint for all dates prior to t . Thus increasing q_t would increase the joint payoffs at all dates up to t without resulting in any of the constraints (4), (7) and (8) being violated.

The argument just given establishes that any constrained efficient self-enforcing equilibrium (*ESE*) has, for each $t \geq \tau$, either (a) $q_t < q_t^*$ and (8) holding with equality, or (b) $q_t = q_t^*$. In the model in Ray (2002), any *ESE* also maximizes the supplier's payoff from some finite $t' > \tau$ on as long as $q_t < q_t^*$ for all $t \geq \tau$. In that model, the purchaser is committed to the agreement at each t , which allows no scope for a discretionary bonus b_t . The result in Ray (2002) then follows from (3) with $b_t = 0$ for all $t \geq \tau$. For any $t \geq \tau$ with $q_t < q_t^*$, quality can be closer to the efficient level if $U_{t+1} - \tilde{U}_{t+1}$ is increased, which must therefore be maximized in any *ESE*. But that is not necessary if discretionary bonuses are included in the relational contract agreement.

Pooling of incentive compatibility conditions as done to derive (8) from (3) and (6) is a standard feature in the literature that will recur in what follows. It is, though, important to recognize that it depends on the parties being risk-neutral. Also a standard feature in the literature is the limit on bonuses implied by (6) because the purchaser cannot receive a payoff from the future higher than is consistent with the supplier being prepared to continue the relationship. The result that conditions (4), (7) and (8) are necessary and sufficient does not depend on an environment, agreement or payoffs that are stationary over time. But, as Levin (2003) shows, in a stationary environment the parties cannot do better than use a stationary agreement, one with q_t, p_t and b_t the same for all t . Krantz




Figure 2: Timing of events for period t in outsourcing model

and Ohlendorf (2009) consider the renegotiation proofness of relational contracts.

3.2 Outsourcing model

In the employment model, once the cost of supply is incurred, the supplier cannot hold back supply until payment is made. But supply is sometimes of intermediate products produced on the supplier's own premises. That complicates the picture somewhat because the supplier could then withhold the product and use it as a bargaining chip for payment. This has been termed the *outsourcing model* in the literature. The timing of events for this model is shown in Figure 2. This differs from Figure 1 for the employment model because at the third stage in period t the supplier, as well as the purchaser, can decide whether to stick to their agreement.

If both parties decide to stick to their agreement at this stage, the supplier delivers the already-produced quality to the purchaser and the purchaser simultaneously pays the agreed bonus. If either decides to breach the agreement, the supplier retains the product and the purchaser pays no bonus. To complete the model, one must specify the alternative actions available to the parties in the event of breach. A common assumption in the literature is that the outcome of those actions corresponds to a Nash bargain but the precise formulation is unimportant here. For the present purpose, it suffices to denote by $z_t(q_t)$ the payment the supplier can get for quality q_t by withholding supply and bargaining. This payment will reflect not only what the product is worth to the purchaser but also what it can be sold for to some alternative purchaser. In this model, if the supplier were to decide not to provide the agreed quality and so risk breakdown of the relationship, it would not necessarily be optimal to set quality to \underline{q} because the product could still be sold for some payment. It simplifies the exposition here without compromising later use to set $\rho = 1$, so that substandard quality is always detected by

the purchaser. Then the optimal quality to which to deviate is

$$\tilde{q}_t = \arg \max_{q \geq \underline{q}} z_t(q) - c_t(q). \quad (10)$$

This is precisely the quality the supplier would produce if there were no relational contract and all transactions were via a spot market. The payoff to deviating from the agreed quality corresponding to (2) with $\rho = 1$ is then

$$p_t - c_t(\tilde{q}_t) + z_t(\tilde{q}_t) + \delta_t \tilde{U}_{t+1}, \quad \text{for all } t \geq \tau, \quad (11)$$

and the incentive compatibility condition corresponding to (3) with $\rho = 1$ becomes

$$\begin{aligned} \delta_t \{\text{future gain to supplier}\} &\equiv \delta_t (U_{t+1} - \tilde{U}_{t+1}) \\ &\geq c_t(q_t) - c_t(\tilde{q}_t) + z_t(\tilde{q}_t) - b_t, \quad \text{for all } t \geq \tau. \end{aligned} \quad (12)$$

In the outsourcing model, the supplier also has another decision to make, whether to supply the intermediate good once it has been produced. That will be worthwhile only if the bonus b_t combined with the benefit of having the relationship continue is more valuable than deviating by selling the good for $z_t(q_t)$ and having the relationship end. Formally, that requires

$$\delta_t (U_{t+1} - \tilde{U}_{t+1}) \geq z_t(q_t) - b_t, \quad \text{for all } t \geq \tau. \quad (13)$$

By the definition of \tilde{q}_t in (10), the right-hand side of (12) is necessarily at least as large as the right-hand side of (13), so the former is always the critical incentive compatibility condition for the supplier.

The purchaser will pay the bonus for agreed quality only if doing that and having the relationship continue is worth more than defaulting, paying $z_t(q_t)$ to get the supply despite defaulting and having the relationship end. Thus the incentive compatibility condition for the purchaser corresponding to (6) becomes

$$\delta_t \{\text{future gain to purchaser}\} \equiv \delta_t (\Pi_{t+1} - \tilde{\Pi}_{t+1}) \geq b_t - z_t(q_t), \quad \text{for all } t \geq \tau. \quad (14)$$

The overall incentive compatibility condition corresponding to (8) with $\rho = 1$ is obtained by adding (12) and (14) to get

$$\begin{aligned}\delta_t \left(S_{t+1} - \tilde{S}_{t+1} \right) &\equiv \delta_t \left(U_{t+1} - \tilde{U}_{t+1} + \Pi_{t+1} - \tilde{\Pi}_{t+1} \right) \\ &\geq c_t(q_t) - c_t(\tilde{q}_t) + z_t(\tilde{q}_t) - z_t(q_t), \quad \text{for all } t \geq \tau. \quad (15)\end{aligned}$$

To see the incentive differences between the employment and the outsourcing structures, compare (15) with (8) when $\rho = 1$ so that the environments are otherwise identical. The incentive compatibility condition (8) for the employment structure depends only on the quality q_t specified by the relational contract. That for the outsourcing structure (15) also depends on the quality \tilde{q}_t that the supplier would produce in the absence of a relational contract. Thus the structure of the relationship makes a difference to the quality that can be sustained. It is tempting to conclude that (15) is weaker than (8) whenever $z_t(q_t) \geq z_t(\tilde{q}_t)$. One must, however, be careful about this because it may be that \tilde{U}_t and $\tilde{\Pi}_t$ take different values in the outsourcing model from in the employment model. One obvious reason is that, in the outsourcing case, the supplier is willing to supply \tilde{q}_t even in a spot market without a relational contract because it can always be sold for $z_t(\tilde{q}_t)$. In contrast, in the employment model, no supply would occur without a relational contract.

The argument used above for the employment model to show that the purchaser and supplier can always increase their joint payoff if $q_t > q_t^*$ for any t does not work with the outsourcing model because reducing q_t may increase the right-hand side of (15) if it reduces $z_t(q_t)$ more than it reduces $c_t(q_t)$. But it remains the case that the purchaser and supplier can increase their joint payoffs $U_\tau + \Pi_\tau$ at the start of the relationship at τ unless, for each $t \geq \tau$, either (15) holds with equality or $q_t = q_t^*$. Otherwise moving q_t marginally in the direction of q_t^* would not violate (15) and would increase the joint payoffs for all dates up to t .

Because of its simpler structure, much of the discussion that follows is based on the employment model underlying (8) rather than the outsourcing model underlying (15). But many of the arguments can be extended to the outsourcing model and that model has been used in the literature to illustrate points that cannot be made using only the employment model.

4 Specific investments

Basic as it is, the employment model of Section 3.1 can be used to address the first question arising from the interpretation Klein (2000) places on the relationship between GM and Fisher Body: in what way might a larger specific investment make it more difficult to sustain a relational contract? But first the model must be adapted to include a specific investment. A specific investment by the supplier of $I \geq 0$ at τ in general affects both the output from, and the cost of, delivering quality, so write these as $y_t(q_t, I)$ and $c_t(q_t, I)$ respectively for $t \geq \tau$. It thus also affects the surplus from the relationship S_t . Furthermore, it may affect the ease of monitoring quality, which can be captured by writing the probability of monitoring $\rho(I)$. In addition, it may affect the supplier's payoff from breaching the relational contract. If, as in the GM–Fisher Body example, the specific investment involved moving the body plant, it may have reduced the payoff to trading with other potential purchasers. Denote that payoff by $\tilde{U}_{t+1}(I)$. Finally, if the supplier is to make the investment, the payoff in the period of investment τ must be large enough to make the investment worthwhile, that is $U_\tau \geq \bar{U}_\tau + I$.

Consider a stationary environment and a stationary relational contract, for which t subscripts can be dropped. Then $S = [y(q, I) - c(q, I)] / (1 - \delta)$. Define $\tilde{u}(I)$ and $\tilde{\pi}$ by $\tilde{U}(I) = \tilde{u}(I) / (1 - \delta)$ and $\tilde{\Pi} = \tilde{\pi} / (1 - \delta)$, so $\tilde{S} = [\tilde{u}(I) + \tilde{\pi}] / (1 - \delta)$. Then (8) becomes

$$\frac{\delta}{1 - \delta} [y(q, I) - c(q, I) - \tilde{u}(I) - \tilde{\pi}] \geq \frac{c(q, I)}{\rho(I)}. \quad (16)$$

It follows that a worthwhile specific investment that increases $y(q, I) - c(q, I)$ for given q can make it harder to sustain the relationship only if it increases

$$\frac{c(q, I)}{\rho(I)} + \frac{\delta}{1 - \delta} [\tilde{u}(I) + \tilde{\pi}]. \quad (17)$$

The specific investment would thus have to do at least one of the following: (a) increase the variable cost of producing the given quality, with the benefits of the specific investment arising from an increase in $y(q, I)$ of more than the increase in $c(q, I)$; (b) make monitoring quality more difficult, corresponding to a decrease in $\rho(I)$; or (c) increase the payoff $\tilde{u}(I)$ the supplier receives from breaching the relational contract. This chapter focuses on theory and so does not consider in detail whether the higher specific investments

that GM wanted Fisher Body to make could have operated in this way, but there are a number of straightforward considerations.

An obvious efficiency gain from siting a body plant close to an assembly plant would be a saving in the cost of transporting bodies to the assembly plant. If that were all, it would be worthwhile only if it involved a net saving after taking into account the cost of transporting materials, that is, only if it *reduced* the cost $c(q, I)$ for given q of delivering bodies to the assembly plant. Also plausible is that it would reduce the payoff $\tilde{u}(I)$ the supplier could get from breaching the relational contract because the plant would be further from other potential purchasers. Moreover, it is not obvious why moving the body plant should increase the difficulty of monitoring quality. Thus the reduction in transportation cost would seem to make it easier to sustain the relational contract. It may, of course, be that the specific investment increases the marginal benefit of quality $\partial y(q, I) / \partial q$ or decreases the marginal cost of quality $\partial c(q, I) / \partial q$, in which case it will change the efficient quality level defined by (9). Thus it may be that, at a lower level of specific investment, efficient quality is attainable with a relational contract but not at a higher level of specific investment. But that does not, in itself, make it harder to sustain the relational contract.

The outsourcing model of Section 3.2 may, however, be more appropriate to describe the relationship between GM and Fisher Body, though the value of a Buick body to another automobile assembler is questionable. Then the relevant incentive compatibility condition is (15) instead of (8). In that case, the specific investment might affect the quality \tilde{q}_t defined in (10) that the supplier would produce if planning to breach the relational contract. Then, in an obvious notation, (16) is replaced by

$$\begin{aligned} \frac{\delta}{1-\delta} [y(q, I) - c(q, I) - \tilde{u}(I) - \tilde{\pi}] \\ \geq c(q, I) - c(\tilde{q}(I), I) + [z(\tilde{q}(I), I) - z(q, I)] \quad (18) \end{aligned}$$

and so (17) is replaced by

$$c(q, I) - c(\tilde{q}(I), I) + [z(\tilde{q}(I), I) - z(q, I)] + \frac{\delta}{1-\delta} [\tilde{u}(I) + \tilde{\pi}].$$

There is then the additional consideration of how the specific investment affects the

difference in cost $c(q, I) - c(\tilde{q}(I), I)$ between producing q and producing $\tilde{q}(I)$, and how it affects the difference in the payment $z(\tilde{q}(I), I) - z(q, I)$ the supplier could get by withholding supply and bargaining.

These considerations apply to the continuation of the relational contract once the specific investment has been made. One of the issues arising between GM and Fisher Body, however, was the latter's unwillingness to make the specific investment in the first place. For a specific investment at τ to be worthwhile for the supplier, the supplier's payoff from τ on, U_τ , would have to increase by more than the cost of the investment. The previous section showed that, for any given sequence of qualities (q_t, q_{t+1}, \dots) it is feasible to support in a continuation equilibrium, there exist continuation equilibria with any feasible U_t and Π_t that satisfy $U_t \geq \bar{U}_t$ and $\Pi_t \geq \bar{\Pi}_t$. So, as long as the specific investment is jointly worthwhile in the sense of increasing $U_\tau + \Pi_\tau$ by more than it costs, there exists a continuation equilibrium with the increase in the supplier's payoff sufficient to cover the cost of the investment without reducing the purchaser's payoff. But it is possible that Fisher Body's concern was with whether GM would in fact play such a continuation equilibrium. Predicting which equilibrium will be played in a multiple equilibrium context is problematic. But the usual assumption in the relational contract literature is that the parties will play a Pareto superior equilibrium provided it is incentive compatible. Indeed, this is implicit in Klein's view that, up to 1924, the relational contract arrangement between GM and Fisher Body was efficient.

Whether, of course, they actually do is an empirical question that is not the subject of this chapter. But it is worth noting that there are many experimental studies indicating that participants co-operate even in the absence of third-party enforcement of agreements. For a survey, see Fehr and Falk (2002). Especially relevant here is Brown et al. (2004) in which the party acting in the role of purchaser could specify a payment corresponding to p_t and a performance level corresponding to q_t . In treatments in which these were a contract enforced by the experimenters, participants behaved very much in the way predicted by competitive market theory. But, when they were not an enforceable contract, repeated interaction resulted in huge increases in the performance level over that in one-off interactions, though not to the level in the enforceable contract treatment. The possibility of choosing to terminate the relationship was important for this. The experimental finding less consistent with the model used here is that, although each experiment lasted only a

known finite number of periods, performance did not drop to the minimum level possible even in the final period. Indeed, performance in the final period increased with the generosity that the partner had shown in previous periods. The authors term this a *reciprocity effect*. MacLeod (2007a) models this effect as arising from a small disutility of breaching an agreement. But, in an important sense, such an effect is complementary to the model here. It implies that co-operation does not necessarily unravel when the time horizon is actually finite and the experimental findings emphasise that repetition increases performance beyond what the reciprocity effect alone delivers. What an experiment lasting just a few hours obviously cannot pick up is how long lived such reciprocity effects are likely to be in practice.

Ramey and Watson (1997) analyse an employment type model in which the specific investment is made by the purchaser, not the supplier, at a stage when the future environment is uncertain. There are just two possible levels of quality, \underline{q} and \bar{q} . The purchaser can choose the level of specific investment I . Uncertainty about the future can be represented in the incentive compatibility condition (16) by re-interpreting ρ (independent of I in Ramey and Watson (1997)) as a random variable taking one of two values, $\bar{\rho}$ and $\underline{\rho}$ with $1 > \bar{\rho} > \underline{\rho} > 0$, whose realization is unknown at the time the investment decision is made. For I sufficiently high, (16) is satisfied for quality \bar{q} whether ρ takes the value $\bar{\rho}$ or $\underline{\rho}$ but for lower I it is satisfied only for $\bar{\rho}$. It is never satisfied for \underline{q} . When $\underline{\rho}$ is sufficiently unlikely, the purchaser's optimal level of I is such that (16) is not satisfied, and the relationship is not continued, if $\underline{\rho}$ occurs (termed a *fragile* contract). Ramey and Watson (1997) use this model to study the pattern of separations in a market equilibrium in which new trading partners cannot be immediately located and there are aggregate shocks. Improvements in the probability of rematching may result in greater use of fragile contracts. Ramey and Watson (2001) extend the model to the case in which both parties make specific investments and identify cases in which there is an optimal positive level of market friction that implements efficient investment.

5 Verifiable performance measures

The GM–Fisher Body relational contract contained a 10-year exclusive dealing arrangement supplemented by a cost-plus contract for each unit delivered. Klein (2000, p. 125) is

of the view that a court would presumably have enforced the cost-plus contract. For it to be legally enforceable, the number of units delivered and the cost of provision would have to be verifiable. Of course, if everything the parties care about is verifiable, there would be no need for a relational contract. But there is an obvious question of when, if only some of the relevant information is verifiable, it is advantageous to use a supplementary legally enforceable contract using that information within a broader relational contract.

Outside one-shot, simultaneous action games, it may, as Bernheim and Whinston (1998) show, be actively harmful to constrain the actions of the parties when there is verifiable information to do so. Although their formal results apply only to finite horizon games, the essential point, as they show, carries over to infinite horizon games. Indeed, it can be seen from the basic model set out above. In that model, provided payments are verifiable, it would be perfectly possible to fix by formal contract the bonus b_t that the purchaser pays to the supplier. But to do so would prevent the purchaser tailoring the payment to the unverifiable quality delivered by the seller.

Use of a supplementary formal contract can, nevertheless, enable a relational contract to work under a wider set of conditions, as Baker et al. (1994) show. This can be seen in the employment model of Section 3.1 with the addition of a verifiable outcome measure $x_t = \mu_t q_t$, where $\mu_t > 0$ is the realization of a random variable that is observed by the supplier before choosing q_t but not observed by the purchaser. For simplicity, let $\rho = 1$ since that plays no role in the discussion. In a one-period relationship, basing payment on x_t alone distorts incentives but enables quality above the minimum level, see Baker (1992). Specifically, with legally enforceable payment $p_t(x_t)$ assumed by Baker to have constant slope $p'_t(x_t) = \beta_t > 0$, actual quality depends on μ_t even though efficient quality does not. With an on-going relational contract, suppose the purchaser commits to the enforceable payment just described. (Baker et al. (1994) impose $b_t = 0$ but that is not essential to the argument.) Incentive compatibility for the purchaser is unchanged because $p_t(x_t)$ has to be paid in any case. However, if the supplier is going to deliver substandard quality, it is optimal to do so in a way that maximizes the current period variable payoff $p_t(\mu_t q_t) - c(q_t)$. That is, the supplier will deliver quality $\tilde{q}_t(\beta_t \mu_t)$ defined by $c'_t[\tilde{q}_t(\beta_t \mu_t)] = \beta_t \mu_t$. Thus, the supplier's incentive compatibility condition (3) becomes

$$\delta \left(U_{t+1} - \tilde{U}_{t+1} \right) \geq c_t(q_t) - c_t[\tilde{q}_t(\beta_t \mu_t)] - \beta_t \mu_t [q_t - \tilde{q}_t(\beta_t \mu_t)] - b_t, \text{ for all } t \geq \tau.$$

Added to the purchaser's incentive compatibility condition (6), this implies that the pooled incentive compatibility condition (8) is replaced by

$$\delta \left(S_{t+1} - \tilde{S}_{t+1} \right) \geq c_t(q_t) - c_t[\tilde{q}_t(\beta_t \mu_t)] - \beta_t \mu_t [q_t - \tilde{q}_t(\beta_t \mu_t)], \text{ for all } t \geq \tau. \quad (19)$$

For any $q_t > \tilde{q}_t(\beta_t \mu_t)$, the right-hand side of (19) is smaller than the right-hand side of (8), so higher quality at t is sustainable for any given future gain $S_{t+1} - \tilde{S}_{t+1}$ and this makes it feasible to sustain quality closer to the efficient level whenever that level is not sustainable without the formal contract. But β_t high enough to achieve efficient quality q_t^* when μ_t is low may result in $\tilde{q}_t(\beta_t \mu_t) > q_t^*$ when μ_t is high. Quality can still be kept to the efficient level, provided that level satisfies (19), by the purchaser not paying the bonus if quality is too high as well as if it is too low, but it may not be possible to choose β_t to sustain efficient quality for both high and low realizations of μ_t . Even so, appropriate choice of β_t enables the purchaser and supplier, at least on average, to get closer to the efficient level. It is thus optimal for the parties to use a formal contract to supplement their relational contract even though it is based on imperfectly verifiable information.

The existence of such imperfectly verifiable measures is not, however, necessarily an unmixed blessing. It may increase the payoffs \tilde{U}_{t+1} and $\tilde{\Pi}_{t+1}$ the purchaser and the supplier receive after breaching because it may enable them to guarantee to other potential partners that an alternative match will be productive even if they are no longer trusted to stick to a relational contract. Baker et al. (1994) show that, for certain configurations, it may even destroy the incentive compatibility of the relational contract. Di Tella and MacCulloch (2002) provide a related application in which better outside alternatives can crowd out informal incentives. In their case, more generous state provision of unemployment benefit reduces informal intra-family insurance more than one-for-one, so total transfers to the unemployed fall.

This model brings out a general point: a supplementary enforceable contract affects a relational contract through its impact on the payoffs the parties receive in the event one of them defaults on what they agree. In it, however, the verifiable outcome measure x_t is of no direct concern to the purchaser. It does not, therefore, actually capture the contractual arrangement between the Japanese manufacturers and their suppliers described by Asanuma (1989), and between GM and Fisher Body, where the quantities ordered were treated as a formal contract and quantity was clearly of direct concern to the

purchasers. That situation is captured by Schmidt and Schnitzer (1995) using the multi-task agency approach of Holmstrom and Milgrom (1991). Suppose $q_t = (q_t^1, q_t^2)$ has two components, q_t^1 measuring quantity that is verifiable and q_t^2 measuring quality that is not, with $c_t(q_t^1, q_t^2)$ increasing in each argument. In the absence of a supplementary formal contract, the only difference to the analysis in Section 3.1 is that \tilde{U}_{t+1} and $\tilde{\Pi}_{t+1}$, and hence \tilde{S}_{t+1} , need to reflect the possibility that the parties can trade without a relational contract using a legally enforceable contract on q_t^1 alone if that is worthwhile when q_t^2 is set at its lowest level \underline{q}^2 . But this applies whether or not the parties supplement their relational contract with an enforceable contract on q_t^1 and so does not affect whether, given the possibility of a supplementary enforceable contract, it is worthwhile using one. If the parties specify q_t^1 in a supplementary enforceable contract, it is optimal for the supplier, if defaulting, to deliver quality $q_t^2 = \underline{q}^2$ in quantity q_t^1 . Thus, the incentive compatibility condition for the supplier (3) when $\rho = 1$ becomes

$$\delta \left(U_{t+1} - \tilde{U}_{t+1} \right) \geq c_t(q_t^1, q_t^2) - c_t(q_t^1, \underline{q}^2) - b_t, \text{ for all } t \geq \tau.$$

Added to the purchaser's incentive compatibility condition (6), this implies that the pooled incentive compatibility condition (8) is replaced by

$$\delta \left(S_{t+1} - \tilde{S}_{t+1} \right) \geq c_t(q_t^1, q_t^2) - c_t(q_t^1, \underline{q}^2), \text{ for all } t \geq \tau.$$

This is weaker than (8) with $\rho = 1$ and so enlarges the set of q_t that can be sustained as equilibria whenever the efficient level is otherwise unattainable. Thus, specifying quantity in an enforceable contract is certainly worthwhile. As in Baker et al. (1994), however, by improving the payoffs that can be attained without a relational contract, the possibility of specifying quantity in an enforceable contract may make it harder to sustain a relational one. Iossa and Spagnolo (2009) show that, even if q_t^1 is of no value to the purchaser (as a result of which delivery is not enforced in equilibrium), contracting on it can help sustain a relational contract by enabling an additional penalty to be imposed on a defaulting supplier. Daido (2006) analyses a related model in which neither component of q_t is verifiable but y_t has two components, each randomly related to the corresponding component of q_t and one of which is verifiable.

The model in Schmidt and Schnitzer (1995) does not fully capture the relationship

between GM and Fisher Body because the enforceability of the cost-plus contract would require that Fisher's cost was verifiable. In the model, that would remove the need for a relational contract altogether so, to retain the role for a relational contract, it would need to be enriched further. But it re-emphasizes the result in Baker et al. (1994) that it may be valuable to supplement a relational contract with a legally enforceable contract conditioned on verifiable information.

A different role for a legally enforceable contract to supplement a relational contract arises when verifiable measures of performance are imprecise and the supplier is risk averse. Some individual suppliers are much smaller than purchasers such as GM and Toyota and so might be expected to be more risk averse. Asanuma and Kikutani (1992) show that the relations between the big Japanese automobile manufacturers and their satellite (as opposed to intermittent) subcontractors contain mechanisms through which the manufacturer absorbs the risks involved in their transactions to a substantial degree. Using a technique from Kawasaki and McMillan (1987), they estimate that four large Japanese auto manufacturers each bear close to 90% of the risk and estimate a corresponding degree of absolute risk aversion for the satellites.

Pearce and Stacchetti (1998) adapt the employment model to a risk averse supplier in a relational contract. The purchaser remains risk neutral. The environment and the contractual arrangements are stationary. The outcome y in period t (from a finite set) is now a random variable with verifiable realization y_t that is *iid* conditional on the quality q_t (also from a finite set) delivered by the supplier in period t . The distribution of y can thus be written $F(y | q)$. Quality remains unverifiable but is observed by the purchaser. (In this respect, the model differs from the repeated principal-agent model in Radner (1985), in which the agent's effort is not observed by the principal.) With y_t verifiable, the parties can use a legally enforceable contract that specifies payment $p(y)$, required by Pearce and Stacchetti (1998) to be non-negative. But they can improve on that by adding a relational element with voluntary bonus $b(y)$ conditional on the agreed quality being delivered. The incentive compatibility condition for the purchaser to pay the bonus, corresponding to (6), then becomes

$$\delta \left[\Pi(y) - \tilde{\Pi} \right] \geq b(y), \text{ for all } y, \quad (20)$$

where $\Pi(y)$ is the future payoff to the purchaser in the continuation equilibrium when y

occurs. Pearce and Stacchetti (1998) show that, if total compensation $P(y) = p(y) + b(y)$ is not independent of y (as it would be with efficient risk sharing), it is optimal to have (20) hold with equality for every y for which $p(y) > 0$. The intuition is that, if incentive compatibility for the purchaser were not a concern, lowering $p(y)$ and increasing $b(y)$ by the same amount would strengthen the supplier's incentive to deliver the agreed quality while leaving unchanged the overall payoffs to both parties when y occurs. This would provide scope for adjustment of payments to decrease the variation in total compensation to the risk-averse supplier and shift risk to the risk-neutral purchaser, a Pareto improvement that remains feasible until either $p(y) = 0$ or (20) holds with equality. Pearce and Stacchetti (1998) also show that a higher current total payment to the supplier [$P(y') > P(y'')$] is associated with: (a) a higher current legally enforceable payment [$p(y') > p(y'')$]; (b) a higher future payoff to the supplier [$U(y') > U(y'')$]; and (c) a lower current bonus [$b(y') < b(y'')$]. To see why, suppose it is desirable for incentive reasons to reward the agent more for outcome y' than for outcome y'' . To smooth rewards intertemporally, it is then desirable to have both the present payoff to the supplier $P(y)$ and the future payoff $U(y)$ higher for y' than for y'' . Because shocks are *iid* and hence the future joint surplus is independent of y , $U(y') > U(y'')$ implies $\Pi(y') < \Pi(y'')$. With (20) a binding constraint, that in turn implies $b(y') < b(y'')$. But then $P(y') > P(y'')$ only if $p(y') > p(y'')$. This result yields the empirically testable prediction that a higher total payment at t' than at t'' consists of a higher fixed payment but a lower bonus payment.

The discussion so far in this section applies to the employment model. With the outsourcing model, a formal contract based on the verifiability of the quantity supplied can play a further role, as Itoh and Morita (2009) show. In the absence of a formal contract, the overall incentive compatibility condition is (15). Suppose the parties supplement their relational contract with a formal contract that the supplier will provide the intermediate product for the payment p_t and a court will enforce specific performance of the contract so that, if the intermediate product is produced, it must be delivered to the purchaser and the purchaser must pay p_t . (Quality q_t cannot, of course, be the subject of a legally enforceable contract because it is unverifiable.) Then the supplier cannot extract a payment other than p_t for delivery whatever quality is produced and the situation is equivalent to the employment model with overall incentive compatibility constraint (8) and $\rho = 1$. Clearly, this condition is weaker than (15) for given q_t if

$z_t(\tilde{q}_t) - c_t(\tilde{q}_t) > z_t(q_t)$, for which a necessary condition is obviously $z_t(\tilde{q}_t) > z_t(q_t)$. That might be the case when, as discussed by Itoh and Morita (2009), producing quality q_t instead of \tilde{q}_t makes the intermediate product sufficiently more specific to the purchaser to reduce the payment the supplier could get for it by bargaining.

The results above that depend on verifiable delivery should not, however, be overplayed. With verifiable delivery, it may be possible to achieve efficient quality even without a relational contract. Suppose, in the spirit of the discussion of specific investments in Hart and Moore (1988), the parties were to contract on a fixed price $p_t = y_t(q_t^*)$ for delivery, where q_t^* is the efficient quality defined in (9), but the purchaser has the right to refuse to accept delivery. Then, even in a single-period model with no relational contract, it is optimal for the supplier to supply quality q_t^* because the purchaser will not accept delivery of quality less than q_t^* at price p_t but will accept delivery of quality q_t^* at that price. An upfront payment by the supplier can be used to divide the gains from trade between the parties in any proportions. In Itoh and Morita (2009), a random variable whose realization is not observed until after quality has been determined also affects the value y_t to the purchaser so it is not possible to determine the exact price in advance in the way just specified. But even then, provided delivery is verifiable, an alternative mechanism can still achieve efficient quality without a relational contract. Specifically, the parties can agree a legally enforceable contract giving the supplier the right to set any price at all for delivery after the uncertainty is resolved and the purchaser the right to choose whether to accept delivery at that price. Then it is optimal for the supplier to produce whatever quality is efficient given the information available at the time the quality decision is made and to set the price at the purchaser's valuation of that quality given the realization of the random variable. Evans (2008) shows how powerful contracts of this type can be at achieving efficiency under quite general conditions.

The discussion thus far has taken as given what is verifiable and so can be made legally enforceable. But it may be possible for the parties to influence, at a cost, the extent to which an agreement is legally enforceable. In Kvaløy and Olsen (2009), the probability that a breach of the relational contract will be verifiable to the courts can be increased by incurring an enforcement cost. They use the employment model with either reliance or expectation damages. For spot market trading to be worthwhile, the parties must incur the enforcement cost. A relational contract reduces the benefit of incurring the

enforcement cost though, under appropriate conditions, it is still worth incurring some enforcement cost. However, the resulting equilibrium quality may be lower, despite the joint gains being higher because of the saving on the enforcement cost. Moreover, under a relational contract, higher δ (which Kvaløy and Olsen (2009) interpret as a proxy for greater trust between the parties) may actually result in lower quality.

Battigalli and Maggi (2002) develop a framework for analysing when it is worth incurring the cost of describing contingencies in order to enable the parties to write a legally enforceable contract on them. Battigalli and Maggi (2008) extend that framework to repeated interaction with relational contracts. They consider an outsourcing model in which q_t is a vector of qualities, each of which is binary, with $c_t(q_t)$ additively separable in these qualities. All qualities are potentially verifiable at a cost but efficient quality decisions may change over time because of exogenous events. The formal contracts the parties can write can contain a different type of clause for each element of q_t : a *contingent clause* specifying that element conditional on exogenous events, a *spot clause* specifying that element for period t only after the exogenous events for period t are known (formally this is a non-contingent clause that is modified after the events are observed and corresponds to settling for $z_t(q_t)$ in the outsourcing model), and an *enrichment clause* that starts off as a non-contingent clause but is changed to a contingent clause when a relevant contingency first occurs. Different types of clauses have different writing costs. The choice between clause type depends on the cost of describing contingencies relative to the cost of describing the quality to be delivered. For a contingent clause, both costs must be incurred, but only once. For a spot clause, the cost of describing contingencies is avoided but the cost of describing the quality to be delivered must be incurred each period. An enrichment clause incurs the cost of describing the quality to be delivered at the start, postponing the cost of describing contingencies until an event occurs for which the non-contingent quality is inappropriate but then incurring the cost of describing quality a second time. Alternatively, the parties can allow an element of q_t to be enforced informally by a relational contract. Naturally, for δ high enough to satisfy a condition corresponding to (15) whatever exogenous events occur, it is optimal to rely on relational contracting because it saves on writing costs. But for lower δ , it becomes optimal to specify those elements of q_t with high cost of supply in a formal contract. Contingent clauses have a single upfront cost and so do not affect incentive compatibility

in the future, whereas spot and enrichment clauses have expected future costs that have to be deducted from the future gains to continuing the relationship, the left-hand side of (15). In this situation, an increase in writing costs or uncertainty leads to an increase in the number of elements of q_t enforced by the relational contract. For δ sufficiently low, a relational contract is ineffective and all elements of q_t are optimally governed by a formal contract.

6 Private information: non-persistent types

It is clear from the discussion in Asanuma (1989) that Japanese auto and electrical machinery manufacturers are concerned to find out about the differences in characteristics of their suppliers and treat them accordingly. From a theoretical perspective, we can think of suppliers being of different inherent types that the purchaser does not observe directly but that are important to the relationship. This is not captured by the models of Section 3 because there it is assumed that the purchaser knows everything about the supplier that is relevant, even though some things are not verifiable. In this section and the next, we consider a supplier with private information about type (the cost of supply) that is not known to the purchaser. We start in this section with the case in which type is not persistent so that learning about the supplier's type in one period reveals no useful information about the supplier's type in the next period. The supplier's cost of maintaining quality may be higher in one period because of unusually high demand from other customers, because critical employees are off sick, or because quality control equipment is not working properly. That is not the situation Asanuma (1989) had in mind but it is the case most fully discussed in the literature, notably by Levin (2003) and MacLeod (2003).

In Levin (2003), as in Pearce and Stacchetti (1998), the outcome y in period t in the employment model is a random variable that is *iid* conditional on q in period t and distributed $F(y | q)$, with corresponding density function f . Differences from Pearce and Stacchetti (1998) are that, as in the employment model of Section 3.1, the realization y_t is unverifiable and the supplier is risk neutral. Moreover, before choosing q_t , the supplier observes a parameter of the cost function $\theta_t \in [\underline{\theta}, \bar{\theta}]$ that is an *iid* draw for period t . The cost of delivering quality q_t , now denoted $c(q_t, \theta_t)$, is differentiable and increasing in θ_t ,

and satisfies the standard single crossing property, corresponding in the case it is twice differentiable to $\partial^2 c(q_t, \theta_t) / \partial q_t \partial \theta_t > 0$. Levin (2003) considers two cases. In one, the *hidden information* case, the purchaser observes q_t but not θ_t . In the other, the *moral hazard* case, the purchaser observes θ_t but not q_t . Both cases can be handled together by specifying the performance outcome as a subset $\phi_t \subseteq \{\theta_t, q_t, y_t\}$ that is observed by both purchaser and supplier and allowing the voluntary bonus to be conditioned on ϕ_t . In this stationary environment, there is no loss from restricting attention to stationary contracts in which total payment to the supplier can be written $P(\phi_t) = p + b(\phi_t)$ but it adds flexibility in this stochastic environment to no longer require that the bonus be non-negative.

The crucial difference from the models discussed above is that it is now optimal to agree quality as a function $q(\theta)$ of type, with efficient quality $q^*(\theta)$ defined by

$$q^*(\theta) = \arg \max_q E_y[y \mid q] - c(q, \theta), \text{ for all } \theta \in [\underline{\theta}, \bar{\theta}],$$

where E_y is the expectations operator over the random variable y . But the purchaser observes only one of q_t and θ_t and so does not know whether the supplier actually delivered what was agreed. Thus the supplier cannot be punished directly for breaching the agreement about quality, so the incentive compatibility condition for the supplier to deliver quality has to take the form familiar from one-period principal-agent models

$$q(\theta) \in \arg \max_q E_y[P(\phi) \mid q] - c(q, \theta), \text{ for all } \theta \in [\underline{\theta}, \bar{\theta}]. \quad (21)$$

The incentive compatibility conditions for the parties not to breach the agreement about the bonus depend on whether both purchaser and supplier observe the outcome y_t or only the purchaser does. Start with the former, which Levin (2003) calls the case of an *objective performance measure*. Since the agreed bonus can be either positive (in which case the purchaser pays the supplier) or negative (in which case the supplier pays the purchaser), we require incentive compatibility conditions for both. For the purchaser not to default on the largest positive bonus requires, analogously with (6),

$$\delta \left(\Pi - \tilde{\Pi} \right) \geq \sup_{\phi} b(\phi). \quad (22)$$

Similarly, for the supplier not to default on the largest negative bonus requires

$$\delta \left(U - \tilde{U} \right) \geq - \inf_{\phi} b(\phi). \quad (23)$$

These two conditions can be added to give a pooled condition that, like (8), is independent of the distribution of the gains from the relationship and that Levin (2003) calls the *dynamic enforcement (DE)* constraint

$$\delta \left(S - \tilde{S} \right) \equiv \delta \left(U + \Pi - \tilde{U} - \tilde{\Pi} \right) \geq \sup_{\phi} b(\phi) - \inf_{\phi} b(\phi). \quad (24)$$

This constraint imposes a restriction on the maximum difference in the bonuses that can be used and hence on the incentives that can be provided. This maximum difference is determined by the need to make payment of the bonus incentive compatible given the limited future surplus that is available to induce the parties to pay it. Levin (2003) shows that a quality schedule $q(\theta)$ generating surplus S can be implemented if and only if (21) and (24) are satisfied. These replace (8) for the employment model of Section 3.1 which, in effect, combines the incentive compatibility constraint (21) and the dynamic enforcement constraint (24) for θ a fixed value known to the purchaser as well as the supplier. To see this, note that (3) corresponds to the constraint (21). From that, the highest bonus required to induce the supplier to provide quality q_t is $c_t(q_t)/\rho$. The lowest bonus is zero. Thus the right-hand side of (8) corresponds to the difference between the highest and the lowest bonuses exactly like the right-hand side of (24) but with the supplier's incentive compatibility condition substituted for the highest bonus.

To understand the *hidden information* case (in which the purchaser observes q_t but not θ_t), suppose there were only the highest cost type $\bar{\theta}$. Then, with $\rho = 1$ as assumed in Levin (2003), the incentive compatibility condition (8) would be

$$\delta \left(S - \tilde{S} \right) \geq c \left(q(\bar{\theta}), \bar{\theta} \right).$$

With a continuum of types as in Levin (2003), type θ must not prefer the quality for some other type to $q(\theta)$. The conditions for that are familiar from standard, one-period adverse selection models, see Baron (1989) or Milgrom and Segal (2002): $q(\theta)$ must be non-increasing and type θ must receive the information rent over the highest cost type $\bar{\theta}$

of $\int_{\theta}^{\bar{\theta}} c_{\theta} \left(q \left(\tilde{\theta} \right), \tilde{\theta} \right) d\tilde{\theta}$. So the incentive compatibility constraint for type θ becomes

$$\delta \left(S - \tilde{S} \right) \geq c \left(q \left(\theta \right), \theta \right) + \int_{\theta}^{\bar{\theta}} c_{\theta} \left(q \left(\tilde{\theta} \right), \tilde{\theta} \right) d\tilde{\theta}. \quad (25)$$

With $q \left(\theta \right)$ non-increasing, the right-hand side of (25) is non-increasing in θ , so the constraint is tightest for $\theta = \underline{\theta}$. Levin (2003) shows that $q \left(\theta \right)$ non-increasing and (25) satisfied for $\theta = \underline{\theta}$ are necessary and sufficient for $q \left(\theta \right)$ to be implemented. He then shows that, for a concave cost distribution, if the efficient $q^* \left(\theta \right)$ does not satisfy (25) for $\theta = \underline{\theta}$, an optimal $q \left(\theta \right)$ satisfies one of the following:

1. *pooling*: $q \left(\theta \right) = \tilde{q}$ for some $\tilde{q} < q^* \left(\theta \right)$, for all θ ;
2. *partial pooling*: for some $\hat{\theta} \in \left(\underline{\theta}, \bar{\theta} \right)$, $q \left(\theta \right) = \tilde{q} < q^* \left(\theta \right)$ for all $\theta \in \left[\underline{\theta}, \hat{\theta} \right]$ and $q \left(\theta \right) < q^* \left(\theta \right)$ and strictly decreasing for all $\theta \in \left(\hat{\theta}, \bar{\theta} \right)$.

Thus either quality is the same for all types, or it is the same for an interval of the lowest cost types and decreasing in type for higher cost types. As in the employment model of Section 3.1, quality is never above the efficient level $q^* \left(\theta \right)$ for *any* type.

The intuition for these results is as follows. If the relationship is sufficiently productive and the discount factor sufficiently close to 1, then (25) with $\theta = \underline{\theta}$ is not a binding constraint at efficient quality $q^* \left(\theta \right)$ for each θ . Since $q^* \left(\theta \right)$ is non-increasing, it can then be implemented and it is certainly optimal to do that. But, when (25) with $\theta = \underline{\theta}$ is binding for $q \left(\theta \right) = q^* \left(\theta \right)$, it makes sense to reduce quality $q \left(\theta \right)$ below the efficient level $q^* \left(\theta \right)$ for all θ . It has to be below the efficient level for some θ and reducing $q \left(\theta \right)$ marginally below $q^* \left(\theta \right)$ has only a second-order effect on joint surplus, whereas increasing $q \left(\theta \right)$ that is substantially below $q^* \left(\theta \right)$ has a first-order effect. Thus, rather than have $q \left(\theta' \right) = q^* \left(\theta' \right)$ and $q \left(\theta'' \right) < q^* \left(\theta'' \right)$ for some θ' and θ'' , it is better to reduce $q \left(\theta' \right)$ marginally and increase $q \left(\theta'' \right)$ marginally. Moreover, when (25) with $\theta = \underline{\theta}$ is binding, it must be the case that the requirement for $q \left(\theta \right)$ to be non-increasing is also binding at $\theta = \underline{\theta}$. If that were not so, (25) with $\theta = \underline{\theta}$ would be relaxed by having $q \left(\theta \right)$ jump upwards immediately above $\underline{\theta}$ to enable $c \left(q \left(\underline{\theta} \right), \underline{\theta} \right)$ to be reduced and, since that would affect output for only one of a continuum of values of θ that has no probability mass, it would certainly be optimal. Thus, if efficient quality is not achieved for all θ , there must be pooling for θ sufficiently close to $\underline{\theta}$. If the constraint (25) with $\theta = \underline{\theta}$ is sufficiently tight, this pooling will apply

to all θ . But when it is less tight, the intuition for not pooling higher cost types before lower cost types arises from the limit on incentives. Starting from the minimum quality level, the bonus does not need to be increased much to raise quality uniformly. But at higher quality levels, it requires increasingly higher bonuses to raise the quality supplied by lower θ types and so it is optimal to focus incentives on the higher θ types.

Payment in this model depends only on the quality delivered. So an implication is that, if either the discount factor δ or the total gain from cooperation $S - \tilde{S}$ is sufficiently low for full pooling of types to be optimal, quality and payment will be independent of day-to-day changes in the supplier's cost. For higher values of $\delta (S - \tilde{S})$ such that partial pooling is optimal, there will be some variation in quality and payment, but only when day-to-day costs are particularly high. In practice, this might take the form of the purchaser paying less when delivered a batch of components with a particularly high proportion of defective items. Thus one would expect to see quality more closely tailored to the supplier's cost conditions, and greater payment variability, in relationships in which the total gain from co-operation $S - \tilde{S}$ is greater and/or trade orders are more frequent so that the discount factor δ relevant to the time between orders is higher.

Malcomson (2011b) analyses the hidden information model in Levin (2003) with the addition that the value of quality to the purchaser changes from period to period without persistence and that the purchaser (but not the supplier) observes its value at the beginning of the period. There is thus private information about the purchaser's type, as well as about the supplier's, and efficient quality depends on both. The purchaser can signal her type to the supplier by making the payment p a function of that type and offering a bonus conditional on both it and the quality actually delivered. For those payments to be incentive compatible, the purchaser must expect an information rent when the value of quality is above its lowest level, which imposes restrictions on how p can vary with the purchaser's type. (It is an *expected* rent because the purchaser does not know the supplier's type at the stage p is offered.) Pooling across both purchaser and supplier types is always an equilibrium (as is pooling across supplier types in Levin (2003)). Efficient quality is increasing in the value of quality to the purchaser and decreasing in the cost of quality to the supplier but, as in Levin (2003), efficient quality may not be sustainable by a relational contract. When it is not, it turns out that, if quality is increasing in its value to the purchaser for some purchaser types, it must actually be *decreasing* in its value to

the purchaser for some other purchaser types. This reinforces the advantages of pooling over purchaser types and thus of making payment and quality independent of day-to-day changes in the purchaser's circumstances.

For the *moral hazard* case (the purchaser observes θ_t but not q_t), Levin (2003) uses two assumptions that also play a role in models discussed later.

Assumption 1 (MLRP) *The probability density function f corresponding to F satisfies the monotone likelihood ratio property that*

$$\frac{\partial f(y | q) / \partial q}{f(y | q)} \text{ is increasing in } y. \quad (26)$$

Assumption 2 (CDFC) *The distribution function F satisfies the convexity of the distribution function condition that $F(y | q = c^{-1}(x; \theta))$ is convex in x for any θ .*

Levin (2003) shows that, for the moral hazard case under these assumptions,

1. an optimal contract implements $q(\theta) \leq q^*(\theta)$ for all θ ;
2. payments $P(\theta, y)$ are *one-step* in the sense that $P(\theta, y) = \underline{P}$ for all $y < \hat{y}(\theta)$ and $\bar{P} = \underline{P} + \frac{\delta}{1-\delta} (S - \tilde{S})$ for all $y \geq \hat{y}(\theta)$, where $\hat{y}(\theta)$ is the point at which the likelihood ratio $(f_q/f)(y | q(\theta))$ switches from negative to positive as a function of y .

The intuition is that, with a risk-neutral supplier, it makes sense to use the limited bonuses in the way that gives the strongest incentives, so the maximum incentive is given for outcomes that are more likely to arise with high quality. Thus, as in the employment model of Section 3.1, only one level of bonus is required.

Levin (2003) also considers the case in which the purchaser observes y_t but the supplier does not, the case of a *subjective performance measure*. This adds a further complication to contracting. Now not only does the purchaser, not observing one of q_t and θ_t , not know whether the supplier has breached an agreement $q(\theta)$ on quality but also the supplier, not observing y_t , does not know whether the purchaser has breached any agreement about how the bonus is to vary with y_t . So, if the bonus agreement specifies that the amount of the bonus is to depend on y_t , the purchaser can make a short-term gain by claiming that y_t is at the level corresponding to the lowest bonus payment. Formally, suppose

after privately observing y_t , the purchaser delivers a message m_t to the supplier with a bonus $b_t(m_t)$ based on this. Then the purchaser will make distinct reports m' and m'' in response to distinct outcomes y' and y'' only if the two reports yield the same payoffs. With a stationary contract, this requires $b(m') = b(m'')$. But in a stationary environment with a stationary contract and hence a stationary future surplus S , this conflicts with providing incentives for the supplier because the supplier's payoffs following y' and y'' will have to be the same. As an alternative to a stationary contract, Levin (2003) considers a *termination contract* with a cutoff threshold \hat{y} such that, if $y_t < \hat{y}$, $P_t = p$ and the relationship terminates whereas, if $y_t \geq \hat{y}$, $P_t = p + b$ and the relationship continues. A termination contract is a special case of what Levin (2003) calls a *full review contract* defined as follows: given any history up to t and compensation offer at t , any two distinct outputs $y'_t \neq y''_t$ generate distinct reports $m'_t \neq m''_t$. The implication is that, in equilibrium, the purchaser maintains no private information from period to period. (This restriction can also be stated in terms of the equilibrium concept, namely perfect public equilibria of the repeated game.) For the case in which q_t is not observed by the purchaser and θ is a fixed, known parameter, Levin (2003) shows that, if an optimal full review contract exists, a termination contract can achieve an optimum. Moreover, the termination contract preserves the property that there is just a single level of bonus, the one-step property. Since the stochastic term relating outcome to quality is *iid*, the gain from continuing the relationship at $t + 1$ is independent of the outcome at t , so termination is always inefficient. It can nevertheless occur in equilibrium because the stochastic shock will sometimes result in $y_t < \hat{y}$. When that happens, the parties are effectively “burning money”. This is a price that has to be paid to induce the purchaser to report the outcome truthfully while still providing incentives for the supplier to deliver a better expected outcome.

It is natural to ask whether there might be some way in which the inefficiency resulting from termination might be reduced. One possibility arises because the inefficiency is the result of the purchaser and supplier having different information about the outcome. So, if it were possible to use a mediator to narrow the information gap, that might reduce the probability of the relationship being terminated. Another possibility is to adopt the procedure suggested in Abreu et al. (1991) of lengthening the period between performance reviews and using the *average* outcome over the longer period to determine

whether termination is triggered. This enables the threat of later punishments to provide incentives in earlier periods and reduces the probability that termination occurs at a review date. Fuchs (2007), in a model with two potential output levels and two potential quality levels, shows that a contract with the properties that the supplier receives a fixed payment until the relationship is terminated, with the purchaser providing no feedback on observed output in the meantime, is optimal. To see why, consider a *T-period review contract* in which the purchaser waits T periods before considering terminating the relationship even if output is consistently low. Increasing T reduces the expected loss from termination along an equilibrium path. But increasing T beyond a certain point results in the limited penalty provided by termination providing insufficient incentives to the supplier in earlier periods. Fong and Li (2010*b*) consider the implications of the agent having limited liability in this setting. Fong and Li (2010*a*) show that garbling of output signals intertemporally, with past outputs having an enduring effect on future signals, can increase the efficiency of contracting. Their model has two output levels. Incentive compatibility requires a bonus for the higher. When higher output results in good outcome signals spread out over time, the same total reward for high output can be given with a lower bonus, thus relaxing the dynamic enforcement constraint (24).

MacLeod (2003) explores the case of a subjective performance measure when the supplier is risk averse, as in Pearce and Stacchetti (1998). While using a one-period employment model, and hence not formally a relational contract, he investigates limits on rewards similar to (24) that arise from a relational contract. In his model, as in the subjective information case in Levin (2003), the purchaser’s information about performance is private and satisfies MLRP (Assumption 1). However, unlike in Levin (2003), the supplier receives a private signal about performance that is correlated with the purchaser’s information. To allow for the possibility of “burning money” (which we know from Levin (2003) is essential with no objective performance measure) in the context of a single period model, MacLeod (2003) introduces a third party to act as a budget breaker, that is, to enable the payment made by the purchaser to be different from that received by the supplier. He shows that, if the supplier’s signal is perfectly correlated with the purchaser’s, the optimal contract is the same as if the purchaser’s signal were verifiable. Moreover, the joint loss from imperfect correlation goes to zero as the correlation gets close to perfect. In contrast, if the supplier’s signal is uncorrelated with the purchaser’s,

the purchaser's payments are independent of his own signal but depend on the supplier's. Moreover, the supplier's compensation is the same for all signals of the purchaser except the worst, for which it is lower, because with MLRP the lowest signal is the most informative about whether effort is low. Thus, the one-step property applies not just to the risk neutral supplier in the model in Levin (2003). Moreover, the supplier will normally (that is, for most outcomes) receive the higher level and only occasionally, for the very worst performance, receive the lower. To get useful results for the case in which the supplier's signal is imperfectly correlated with the purchaser's, MacLeod (2003) uses the additional assumption that the supplier's signal is, with known probability, either the same as the purchaser's or completely uninformative. In that case, if the joint punishment to the purchaser and the supplier (that is, the net payment to the third party) is bounded, the supplier's compensation increases with the signal up to a certain level but is the same for all signals above that level. There are also some results on the effect of the purchaser having biased perceptions of the supplier's performance. For further discussion of efficient contractual arrangements under these types of circumstances, see MacLeod (2007b).

Many of these models share a common characteristic. With relational contracts, the only thing that induces the parties not to breach their agreement is the potential gain from cooperation in the future. Because that potential gain is finite, the gap between the highest and the lowest future payoffs, and hence the maximum incentive that can be provided, is limited. When there is private information about non-persistent types and the potential gains from future cooperation are insufficient to sustain efficient quality for all types, that typically results in more pooling across types than when there is full information. This characteristic is encouraging for trying to understand supply relationships that, as documented by Carlton (1986), have prices that are relatively rigid over time despite apparent short-term changes in production conditions. But there remains research to be done to assess the extent to which those rigidities can genuinely be attributed to pooling across types that are private information.

7 Private information: persistent types

The Japanese manufacturers studied by Asanuma (1989) used information about past performance of suppliers to decide the terms for future trades. There is no gain in

productive efficiency from this when suppliers' types are, as in the models in Section 6, *iid* draws each period, though there may be strategic reasons for it in order to punish a supplier who has breached a relational contract. It is different when types are persistent over time and have to be discovered by the other party.

In Bull (1987), it is the purchaser's type that is persistent but unknown to the supplier, rather than the other way round. The issue there is how a relational contract can be sustained between an infinitely-lived purchaser and a supplier who operates for only a finite period of time that is known in advance. In the basic model of Section 3, the purchaser would not pay a bonus in the final period of the relationship, with no future of concern the supplier would then set quality at the lowest level in that period, and so the relationship would be terminated in the final period because there is no possibility of mutually beneficial trade. Given this, the same argument can be repeated successively for each previous period to show that the relational contract never gets off the ground in the first place. But, if there is a succession of short-lived suppliers who observe how their predecessors have been treated, the relationship can be sustained in the way described by Crémer (1986). In Bull (1987), there are two types of purchasers, those who never pay a bonus in the final period and those who do if the reputation so acquired is sufficiently valuable. The latter are induced to pay because they do not wish to be thought to be the former. In Mukherjee (2008), the supplier has an unchanging type unknown to *both* parties in advance and only a finite (two-period) lifetime. The purchaser can choose whether to make information about performance available to other potential purchasers, in which case career concerns of the type discussed by Holmström (1999) provide incentives for effort, or not to do so and rely on a voluntary bonus to provide incentives. Depending on the precise specification, career concerns can be either substitutes or complements for relational incentives. There is now a substantial game-theoretic literature on models in which one long-lived player faces a succession of short-lived players, see Mailath and Samuelson (2006).

The focus in this chapter is on long-lived suppliers. Hörner (2002), in a model in which purchasers are consumers, allows for both noise and privately-known persistent types but simplifies by having just two potential outcomes and suppliers of just two types. Types are not specific to a particular purchaser, so any information revealed publicly affects what other potential purchasers are prepared to pay that supplier. *Bad* types always

choose low quality. *Good* types choose between high and low quality. Whichever quality is chosen, the outcome y_t to consumers may be good or bad, though there is a higher probability it is good when quality is high. As a result, type is not fully revealed in the first period. Consumers are indistinguishable to suppliers, so it is never incentive compatible for them to pay a bonus and all payment is thus via p_t . Hörner (2002) focuses on *non-revealing equilibria* in which consumers leave a supplier as soon as they experience a bad outcome, all suppliers who have been supplying for the same length of time have the same price, and newly-entering good suppliers earn zero expected profits. In such an equilibrium, the probability that a supplier is good increases with the time that supplier has been in the market and hence so also does price. For newly-entering good suppliers to have zero expected profits, price must start out negative. Yang (2009) analyses a similar setup with a single purchaser. Halac (2009) considers a model in which the purchaser's payoff π_t if not trading with the supplier is private information. When the purchaser receives all the gain from continuing the relationship, there is an incentive for the purchaser to behave as if π_t is low even when it is high in order to induce the supplier to believe that a higher bonus is incentive compatible and accordingly supply high quality, but then to renege by taking the high outside option π_t and not paying the bonus. Conversely, when the supplier receives all the gain, the purchaser has an incentive to behave as if π_t is high even when it is low in order to induce the supplier to believe that there is less gain to appropriate.

The alternative to two types analysed in the literature is a continuum of types. In MacLeod and Malcomson (1988), as in Hörner (2002), types are not specific to a particular purchaser. Formal differences from the employment model of Section 3.1 are that the environment is stationary, the cost of quality is $c(q)/\theta$ where θ is the supplier's type that is known to the supplier but not the purchaser, the probability ρ that breach is detected is 1, and there are no bonuses so that the only payment is the fixed payment. Suppose the purchaser were to group supplier types into a set of ranks r (1 being the lowest), assumed observable to potential purchasers, with payment $p^r > p^{r-1}$, minimum quality performance $q^r > q^{r-1}$ to stay in r and, as described by Asanuma (1989), those in rank r who deliver quality below q^r demoted to $r - 1$ (or dropped as a supplier and used by another purchaser in the equivalent of rank $r - 1$). The expected utility from staying in

rank r for ever is then

$$U^r(\theta) = p^r - \frac{c(q^r)}{\theta} + \delta U^r(\theta), \quad \text{for all } r \geq 1, \text{ all } \theta, \quad (27)$$

which corresponds to (1) in the employment model of Section 3.1 with the bonus set to zero. Because other potential purchasers observe a supplier's rank, the price for rank r is determined by the zero-profit condition $p^r = y(q^r)$, so (27) can be solved to give

$$U^r(\theta) = \left[y(q^r) - \frac{c(q^r)}{\theta} \right] / (1 - \delta), \quad \text{for all } r \geq 1, \text{ all } \theta. \quad (28)$$

It is convenient to define $U^0(\theta)$ as the future payoff to type θ if dropped from the lowest rank (rank 1) permanently with payoff $U^0(\theta) = u / (1 - \delta)$.

Consider which types θ will deliver quality high enough to stay in rank r if they reach it. The incentive compatibility condition for this, corresponding to (3) with no bonus, is

$$\delta [U^r(\theta) - U^{r-1}(\theta)] \geq \frac{c(q^r)}{\theta}, \quad \text{for all } r \geq 1, \text{ all } \theta. \quad (29)$$

Define θ^r as the lowest value of θ that satisfies this condition. Substitution from (28), use of the convention that $q^0 = 0$ and $c(q^0) = -u$, and re-arrangement gives

$$\theta^r = \frac{1}{y(q^r) - y(q^{r-1})} \left[\frac{c(q^r)}{\delta} - c(q^{r-1}) \right], \quad \text{for all } r \geq 1. \quad (30)$$

A *equilibrium hierarchy* is a triple (q^r, p^r, θ^r) that satisfies this condition and the zero-profit requirement that $p^r = y(q^r)$ for each r . Competition between purchasers ensures all supplier types who can be profitably employed will be. In particular, θ^1 must be at the lowest value that is incentive compatible, so q^1 is determined to minimise the right-hand side of (30) for $r = 1$, given the conventions $q^0 = 0$ and $c(q^0) = -u$. This determines q^1 and θ^1 . Given q^1 and θ^1 , q^2 is determined to minimise the right-hand side of (30) for $r = 2$. This determines q^2 and θ^2 and so on up to rank R such that $\theta^{R+1} > \bar{\theta}$, the highest type there is. This iterative process determines a unique equilibrium hierarchy.

In the description in Asanuma (1989, p. 14), purchasers “select a number of firms that have relatively good qualities from among those suppliers that have already been tried” for more responsible treatment. Suppose all supplier types start in rank 1 and are promoted to rank 2 if they deliver sufficiently high quality, then to rank 3 if they continue

to perform sufficiently well, and so on. Can the purchaser set promotion criterion \hat{q}^r from rank r to $r + 1$ such that supplier types eventually sort themselves into equilibrium ranks? To sort properly, it must be that types $\theta \geq \theta^{r+1}$ deliver sufficiently high quality for promotion. They will do that only if meeting the quality criterion \hat{q}^r to gain promotion to rank $r + 1$ and then staying there has a higher payoff than staying in rank r . That is, only if

$$p^r - \frac{c(\hat{q}^r)}{\theta} + \delta U^{r+1}(\theta) \geq U^r(\theta), \quad \text{for } \theta \geq \theta^{r+1}, r \geq 1.$$

But, for selection to occur, it must also be that those of type $\theta < \theta^{r+1}$ do not find it worthwhile to deliver sufficiently high quality for promotion so

$$p^r - \frac{c(\hat{q}^r)}{\theta} + \delta U^{r+1}(\theta) < U^r(\theta), \quad \text{for } \theta < \theta^{r+1}, r \geq 1.$$

With θ continuous, these two conditions imply the sorting condition

$$p^r - \frac{c(\hat{q}^r)}{\theta^{r+1}} + \delta U^{r+1}(\theta^{r+1}) = U^r(\theta^{r+1}), \quad \text{for all } r \geq 1. \quad (31)$$

Substitution for $U^r(\theta^{r+1})$ from (27) gives

$$\delta [U^{r+1}(\theta^{r+1}) - U^r(\theta^{r+1})] = \frac{c(\hat{q}^r)}{\theta^{r+1}} - \frac{c(q^r)}{\theta^{r+1}}, \quad \text{for all } r \geq 1. \quad (32)$$

Because θ^{r+1} is, by definition, the lowest θ for which the incentive compatibility condition for the supplier (29) is satisfied for rank $r + 1$, it follows from (29) that

$$\delta [U^{r+1}(\theta^{r+1}) - U^r(\theta^{r+1})] = \frac{c(q^{r+1})}{\theta^{r+1}}, \quad \text{for all } r \geq 1. \quad (33)$$

Equating the right-hand sides of (32) and (33) gives

$$c(\hat{q}^r) = c(q^{r+1}) + c(q^r), \quad \text{for all } r \geq 1,$$

which implies $\hat{q}^r > q^{r+1}$, for $r \geq 1$. So the quality required to be promoted from rank r to rank $r + 1$ is higher than that required to stay in rank $r + 1$ once there. This process, if implemented by the purchaser, eventually results in a supplier with θ between θ^r and θ^{r+1} being sorted into rank r . That is, all types reach the highest rank in which it is

worthwhile for them to deliver the quality required to stay there. Moreover, implementing this promotion process is worthwhile for the purchaser. Those providing just sufficient quality to stay in a rank generate no profit for the purchaser. But those striving for promotion set quality higher while being paid the same, so it is clearly more profitable for the purchaser to set up the promotion system.

Malcomson (2011*a*) analyses a development of the hidden information model in Levin (2003) in which the supplier's type is unknown to the purchaser but, instead of being *iid*, is persistent over time. As in MacLeod and Malcomson (1988), there is a continuum of supplier types, with the outcome y_t a deterministic function of quality. But, in contrast to MacLeod and Malcomson (1988), Malcomson (2011*a*) allows for bonuses. Moreover, the supplier's type is specific to the purchaser, which makes for a larger gain from continuing the relationship over alternative opportunities once information about type has been revealed. Provided the match is sufficiently productive, pooling of those supplier types who continue the relationship is always an equilibrium, as a result of which quality is not tailored to the supplier's type and the purchaser does not learn about that type. When some separation of types is feasible but efficient quality is not attainable, the parties can, however, do better in an equilibrium that initially partitions types into bands, followed by an optimal continuation equilibrium in which quality for those types in a band is at the highest level consistent with the commonly-known information. That level is the highest that satisfies a condition corresponding to (16), with I (now an argument of only the supplier's cost function $c(q, I)$) denoting the highest cost type in the band. But, with such a continuation equilibrium, separation may not be feasible and full separation of types never is. These are also characteristics of the two-period ratchet effect model of Laffont and Tirole (1988) and thus seem robust in models with a continuum of privately known persistent types and the inability to commit to future contractual arrangements. Here, incentive compatibility for the purchaser restricts the spread of future rewards to different supplier types which, as in Levin (2003), limits the potential for separating them. But because here, unlike in Levin (2003), types are persistent, a supplier can obtain an information rent for the indefinite future, not just the current period, by pretending to be a higher cost type. In contrast, by revealing type, the supplier loses that information rent in an optimal continuation equilibrium. The implication is that quality must jump by a discrete amount between adjacent bands of types despite type being continuous. Thus,

the characteristic of MacLeod and Malcomson (1988) that, as described in Asanuma (1989), suppliers of different types are grouped into categories is preserved, despite the use of bonuses and despite supplier types being specific to the purchaser. In this model, however, among supplier types for which efficient quality is not sustainable, quality in the first period is strictly lower than in the continuation equilibrium for all but the least productive for which the relationship can be sustained. Unlike in the model of Laffont and Tirole (1988), it never exceeds the efficient level. Thus the purchaser learns about the supplier by cautiously building up the relationship with levels of quality lower than can eventually be sustained.

Although types are banded in both MacLeod and Malcomson (1988) and Malcomson (2011a), the differences in what the bands provide are differences in quality, not fundamental differences in tasks that might correspond to the differences between the *drawings supplied* and the *drawings approved* subcontractors discussed by Asanuma (1989). If it is efficient for different types of supplier to do different tasks, there is an additional dimension on which incentive compatibility can operate because there is a gain to the purchaser in having a supplier do the task for which it is efficient given its type. Suppose for example, the purchaser commits to a fixed payment p_t for drawings approved work that is higher than that for drawings supplied work by more than the difference in the cost of supply. Then a high-type supplier has an additional incentive to reveal its type because, if promoted to drawings approved work, future profits will be greater. Moreover, the purchaser has an incentive to promote sufficiently high types because it is more efficient to have them do drawings approved work. Gietzmann and Larsen (1998) have a formal model of this based on the promotion model in Fairburn and Malcomson (1994) and developed further in Fairburn and Malcomson (2001). Though not an infinite horizon model, the underlying mechanism would seem to apply to the relational contract approach. Prendergast (1993) develops a model based on a similar underlying idea, though again not formally in a relational contract.⁴

⁴Chassang (2010) explores a model in which private information takes the form of the supplier knowing which of many actions are potentially productive to the purchaser and which not, whereas the purchaser does not know this. Because the model does not permit payments between the parties (which puts it into a different category from most of the literature discussed here), the possibility of non-cooperation for a while is the only way for the purchaser to influence the supplier's incentive to take costly action.

8 Multiple suppliers

In the models discussed so far, the purchaser treats each supplier independently. New issues arise when there are multiple suppliers. They can, for a start, be rewarded according to their performance relative to other suppliers. Che and Yoo (2001) investigate this possibility. As in Pearce and Stacchetti (1998), the outcome y_t for each supplier, which depends on both quality and a random shock, is verifiable so the principal (who makes a “take it or leave it” contract offer) can commit to a contract conditioned on y_t for each supplier. The outcome and the quality delivered are each assumed to take one of just two values, with the probability of each outcome strictly positive whatever quality is chosen. Unlike in Pearce and Stacchetti (1998), quality is not observed by the purchaser and, because either outcome can occur whatever quality is chosen, quality cannot be directly inferred from the outcome. Moreover, the suppliers are risk neutral and have limited liability that is assumed to rule out payments from them to the purchaser. The important addition is that the shock is common to the two suppliers. Thus, even when the outcome for one supplier is independent of the quality provided by the other, their outcomes are positively correlated. When the relationship can last at most one period, the optimal contract takes the form that a supplier is paid only when her outcome is better than that of the other supplier. That is, payment corresponds to an extreme form of *relative performance evaluation*. The intuition is that, because the shock is common to both suppliers, a high outcome is a stronger signal of high quality when the other supplier’s outcome is low than when it is high. Thus, this payment scheme gives the biggest incentive to deliver high quality. In a repeated relationship, the contract itself does not depend on whether a deviation has occurred in the past because the principal, who determines the contract, cannot observe deviations in quality and all payments are legally enforceable. When each supplier observes the quality delivered by the other in the past, the optimal payment scheme for a discount factor sufficiently close to one is *joint performance evaluation* (a supplier is paid more the better the other supplier’s outcome). The intuition is that, with a suitably chosen contract, it is an optimal response to joint performance evaluation for a supplier to deliver low quality in the future if the other supplier delivers low quality. That provides both suppliers with the strongest incentive not to deviate by delivering low quality. Che and Yoo (2001) also consider the case in which a single task can be allocated to either one supplier or two suppliers, but in the

latter case the purchaser observes only a joint outcome.

Kvaløy and Olsen (2006) extend the model in Che and Yoo (2001) to the case in which the outcomes are not verifiable but there are no common shocks. In that case, as in the employment model discussed in Section 3.1, payments by the purchaser that are conditional on outcomes have to be made incentive compatible. As in Che and Yoo (2001), an advantage of joint performance evaluation in a repeated relationship is that it can make use of suppliers' abilities to monitor each other. But with the outcome unverifiable, an advantage of relative performance evaluation is that it reduces the purchaser's incentive to breach the relational contract because higher payments to one supplier correspond to lower payments to the other. Kvaløy and Olsen (2006) show that, for given differences between the high and low outcomes and between costs of high and low quality, a higher discount factor favours joint performance evaluation because the higher value of the future when the discount factor is high relaxes the incentive compatibility constraint for the purchaser. Conversely, a lower discount factor favours relative performance evaluation. This result contrasts with that of Che and Yoo (2001) in which joint performance evaluation is always preferable in a relational contract with verifiable output.

Levin (2002) considers n risk-neutral suppliers and continuous quality, not just two suppliers and two possible levels of quality as in Che and Yoo (2001) and Kvaløy and Olsen (2006). As in those other papers, the quality $q_t^i \in [\underline{q}, \bar{q}]$ chosen by supplier i in period t at cost $c^i(q_t^i)$ is unverifiable but is observed by the purchaser and the other suppliers. Otherwise the assumptions are as for a stationary version of the employment model of Section 3.1 with $\rho = 1$. Suppose the outcome is additively separable across suppliers. In a bilateral relational contract in which each supplier ceases to perform only if individually cheated by the purchaser, the incentive compatibility constraint for each relationship i corresponding to (8) is then

$$\delta \left[S^i(q^i) - \tilde{S}^i \right] \geq c^i(q^i), \text{ for all } i, \quad (34)$$

where $\delta \left[S^i(q^i) - \tilde{S}^i \right]$ is the joint gain from continuing the relationship with supplier i . But with multiple agents, there can be a multilateral relational contract in which all suppliers cease to perform if the purchaser breaches on any one of them. Then the

incentive constraints (34) can be pooled over suppliers to give the single constraint

$$\delta \left[S(q) - \tilde{S} \right] \geq \sum_{i=1}^n c^i(q^i), \quad (35)$$

where $\delta \left[S(q) - \tilde{S} \right]$ is the joint gain from continuing all n supplier relationships. (34) allows the gain from one relationship to be used to sustain only that relationship. (35) allows the gain from a relationship to be used in a cross-subsidizing way to sustain other relationships and, in particular, those that are more productive at the margin. That can increase the total surplus sustainable. The mechanism is essentially the same as in the multi-market oligopoly model of Bernheim and Whinston (1990), where gains from collusion in one market can help sustain collusion in another. Levin (2002) also shows that, if q_t^i is known only to i but generates an observable signal with error, an optimal multilateral contract is payment by relative performance (a tournament) with a fixed prize to at most one supplier. The underlying reason is the same as in Kvaløy and Olsen (2006).

There is, however, a potential downside to multilateral commitments: they are harder to back away from when outside circumstances change. Suppose production methods change in such a way that it is no longer efficient to use one of the suppliers. With bilateral relational contracting, the relationship with that supplier can be terminated without affecting the relationship with the other suppliers. With multilateral relational contracting, other suppliers may interpret termination with one supplier as the purchaser breaching the multilateral agreement. Suppose they do so with probability α . For concreteness, suppose there are just two suppliers, that quality is a binary choice (\underline{q} or \bar{q}) and that supplier 2 becomes redundant with Poisson probability $1 - \gamma$. Then, with α chosen at the lowest level consistent with the purchaser's incentive compatibility, multilateral contracting improves on bilateral contracting if and only if

$$(1 - \delta) \left[S^2(\bar{q}) - \tilde{S}^2 \right] \geq \frac{1 - \gamma}{\gamma} c^2(\bar{q}).$$

The same idea extends to more than two suppliers in a way that fits in with the distinction in Asanuma (1989, p. 17-18) between suppliers who are to be kept operating more or less continuously and those used as a capacity buffer. Levin (2002) shows that, if there are

two groups of suppliers all of whom are to be retained after a shock, the purchaser can always do at least as well by using a multilateral contract with them all. If there are two groups both containing suppliers to be dropped if requirements decrease, the purchaser will have to re-contract with any suppliers in both groups who are to be retained in the event of an adverse shock. Thus again the purchaser can do at least as well by combining the two groups. It follows that there will be an optimal arrangement with at most one group that includes suppliers to be released and (potentially at least) another group that contains only suppliers to be retained. In this model, as in most of the models discussed in Section 7, there is, however, no reason for the two groups to be doing different tasks.

Other papers explore different avenues when there are multiple suppliers. Schöttner (2008) analyses a model of task assignment with three tasks to be divided between two suppliers that is based on the formulation in Baker et al. (1994) and derives results on the characteristics that determine which tasks should be purchased from the same supplier. Board (2008) analyses a model in which the purchaser must incur a specific investment each period to enable a supplier to produce, as a result of which the purchaser invests in at most one supplier in each period. The amount of specific investment required varies across time and across suppliers, so it is efficient to invest in different suppliers in different periods. There is, however, a benefit to having a set of preferred suppliers to whom the purchaser is *loyal*, in the sense of choosing only from this subset of possible suppliers. The underlying reason is that the higher the probability of future trades with a supplier, the weaker is the current period's incentive compatibility condition. Increasing that probability acts like increasing ρ in (8).

Finally, in Calzolari and Spagnolo (2010), as in Board (2008), the purchaser trades with only one supplier in each period. However, the fixed cost of each potential supplier is an *iid* random draw each period that is not observed by the purchaser. If quality were contractible, the purchaser could achieve efficient quality by inviting potential suppliers to bid in an auction. But, with non-contractible quality, the purchaser has to induce the chosen supplier to deliver quality through the relational contract. In Calzolari and Spagnolo (2010), the purchaser holds auctions but may choose not to do so every period to avoid paying the fixed cost of organizing one and, when it does so, to restrict competition to a set of preferred suppliers for the same reason as in Board (2008). Reducing the number of competitors increases the expected informational rent to the winning supplier,

so it reduces the purchaser's future gain from not cheating on the relational contract. Of course, it also increases the incumbent supplier's future gain by the same amount *conditional on* that supplier winning the next auction. So, conditional on the incumbent supplier winning the next auction, the joint future gains from not defaulting on the relational contract are unchanged. But reducing the number of competitors increases the probability that the incumbent supplier *will* win the next auction, so the net effect is to increase the joint gains from not defaulting, which acts like increasing the left-hand side of (8) and so enabling higher quality to be sustained. Since the auction ensures the purchaser gets the gains from increases in efficiency, there is a trade-off between increasing the number of competitors to get the lowest price and reducing it to increase quality. So, if the number of potential competitors is sufficiently large, it is optimal to restrict the number entering the auction. Reducing the frequency of auctions can also increase sustainable quality. Calzolari and Spagnolo (2010) give conditions under which the purchaser actually does better by having suppliers collude in the auction process.

Underdeveloped in the relational contract context are the implications of the common agency situation in which multiple suppliers, each with private information about his own type, produce simultaneously. Martimort and Stole (2009) analyse common agency in the absence of relational contracts when the principal can choose which of a set of agents to contract with.

9 Partnerships

In all the models discussed so far, there are clearly identified suppliers whose decisions about quality affect the value of the output to a purchaser. In some contracts between firms, for example joint ventures, the two parties are on a more symmetric footing in that the measurable output involves simultaneous inputs from both. In many models of such *partnerships*, the parties are, initially at least, entirely symmetric. To reflect the more neutral position, this section refers simply to “parties”, $i = 1, 2$. It is concerned only with partnerships involving relational contracts. Section 10 on organization design considers issues arising from the existence of assets and who should own them.

For a basic model of partnerships, let q_t^i denote the quality chosen by party i at time t , $y(q_t^1, q_t^2)$ the value of the joint output, $b_t^i(q_t^1, q_t^2)$ the amount received by party i

conditional on performance as long as the relational contract has not been breached, and $z_t^i(q_t^1, q_t^2)$ the amount received by party i if the relational contract has been breached. The functions $b^i(\cdot)$ are, like the voluntary bonuses in the models discussed above, in general not legally enforceable. The sum over i of $b_t^i(\cdot)$ equals the joint output, as does that of $z_t^i(\cdot)$. There can also be fixed payments p_t^i made under the relational contract; the sum of these across i must be zero. Qualities are chosen simultaneously by both parties, with the choice of each subsequently observed by the other. With a superscript i attached to the previous notation to denote party i , the payoff to party i , corresponding to that in (1) for the employment model, if both parties stick to their relational contract then becomes

$$U_t^i = p_t^i + b_t^i(q_t^1, q_t^2) - c^i(q_t^i) + \delta U_{t+1}^i, \quad i = 1, 2; \text{ for all } t \geq \tau. \quad (36)$$

Because the quality chosen by party i affects the payoff $z_t^i(q_t^1, q_t^2)$ party i receives from breaching the agreement, optimal breach behaviour is similar to that for the outsourcing model in Section 3.2. Specifically, if i were planning to breach the agreement about quality, it would be optimal to deliver quality

$$\tilde{q}_t^i = \arg \max_{q_t^i \geq \underline{q}} z_t^i(q_t^1, q_t^2) - c^i(q_t^i), \quad i = 1, 2; \text{ for all } t \geq \tau, \quad (37)$$

which corresponds to (10) for the outsourcing model. The payoff of party 1 from doing that, corresponding to (11) for the outsourcing model, would be no less than

$$p_t^i + z_t^i(\tilde{q}_t^1, q_t^2) - c^i(\tilde{q}_t^i) + \delta \tilde{U}_{t+1}^i, \quad i = 1, 2; \text{ for all } t \geq \tau. \quad (38)$$

The requirement that the payoff in (36) exceed that in (38) gives the incentive compatibility condition corresponding to (12) for the outsourcing model

$$\delta (U_{t+1}^1 - \tilde{U}_{t+1}^1) \geq [c^1(q_t^1) - c^1(\tilde{q}_t^1)] + z_t^1(\tilde{q}_t^1, q_t^2) - b_t^1(q_t^1, q_t^2), \quad \text{for all } t \geq \tau. \quad (39)$$

A corresponding condition holds for party 2

$$\delta (U_{t+1}^2 - \tilde{U}_{t+1}^2) \geq [c^2(q_t^2) - c^2(\tilde{q}_t^2)] + z_t^2(q_t^1, \tilde{q}_t^2) - b_t^2(q_t^1, q_t^2), \quad \text{for all } t \geq \tau. \quad (40)$$

Recall that the sum over i of $b_t^i(q_t^1, q_t^2)$ equals $y(q_t^1, q_t^2)$. Thus, (39) and (40) can be added

to give the pooled incentive compatibility condition

$$\delta \left(S_{t+1} - \tilde{S}_{t+1} \right) \geq \sum_{i=1}^2 [c^i(q_t^i) - c^i(\tilde{q}_t^i)] - y(q_t^1, q_t^2) + z_t^1(\tilde{q}_t^1, q_t^2) + z_t^2(q_t^1, \tilde{q}_t^2),$$

for all $t \geq \tau$, (41)

where $\delta \left(S_{t+1} - \tilde{S}_{t+1} \right)$ is, as before, the joint gain from continuing the relationship at the end of period t . To this must be added the individual rationality conditions corresponding to (4) and (7) and conditions to ensure that it is incentive compatible for each party i to make payment of any negative $b_t^i(\cdot)$.

The functions $b_t^i(\cdot)$ do not appear in (41). How the parties divide the joint output under the relational contract is important only for individual incentive compatibility via (39) and (40) and for satisfying the individual rationality conditions corresponding to (4) and (7), a property that resurfaces later in the discussion of organization design in Section 10. If the parties have a choice of the functions $z_t^i(\cdot)$, it is advantageous to choose them to make (41) as weak as possible when due account is taken of the effect on the \tilde{q}_t^i through (37). In general, the $z_t^i(\cdot)$ functions that do this depend on the (q_t^1, q_t^2) agreed under the relational contract, with a trade-off between increasing q_t^1 and increasing q_t^2 . It has become conventional in the literature on relational contracts to use for comparison purposes the efficient qualities (q^{1*}, q^{2*}) defined by

$$q^{i*} = \arg \max_{q^i} y(q_t^1, q_t^2) - c^i(q^i), \quad i = 1, 2. \quad (42)$$

Then (41) with $q_t^i = q^{i*}$ for $i = 1, 2$ and all t becomes an incentive compatibility condition for achieving efficient qualities and the value of the discount factor δ that makes it hold with equality is the lowest for which efficient qualities can be achieved for given $z_t^i(\cdot)$. The feasible $z_t^i(\cdot)$ that results in the lowest such δ is said to be *relational efficient*.

Radner (1986) was concerned with how close the parties can get to a fully efficient outcome in a model of this type when (a) $b_t^i(\cdot)$ and $z_t^i(\cdot)$ are given exogenously, (b) there are no payments p_t^i , (c) there is no discounting ($\delta = 1$), and (d) stochastic terms that are imperfectly monitored enter the payoffs. Radner et al. (1986) give an example with discounting in which equilibria are bounded away from full efficiency even as the discount factor goes to one. Neither Radner (1986) nor Radner et al. (1986) considers a role for

contracts to alter $b_t^i(\cdot)$ or $z_t^i(\cdot)$. Garvey (1995) does just that in a model with no stochastic elements by assuming joint output is verifiable so that the parties can write a legally enforceable contract specifying $z_t^i(\cdot)$ as a function of $y(q_t^1, q_t^2)$. To simplify, he restricts contracts to constant shares of output and also imposes the restriction $b_t^i(\cdot) = z_t^i(\cdot)$. Since the $b_t^i(\cdot)$ then take only non-negative values, (41) does not need to be supplemented by conditions that make it incentive compatible to pay negative ones. Garvey (1995) uses an example with (potentially different) quadratic cost of quality functions $c^i(q_t^i)$ and $y(q_t^1, q_t^2) = q_t^1 + q_t^2$ to show that the optimal shares are closer to a half in a relational contract than in a spot market contract. Garvey (1995) relates these shares to ownership patterns but that is not central to the analysis.

In McAdams (2010), the joint output at t is affected not only by the qualities (q_t^1, q_t^2) but also by a partnership-specific “state of the world” that is stochastic. The state is persistent, in the sense of depending positively (in a first-order stochastic dominance sense) on the state and qualities delivered in the previous period. In his general model, monetary payments are unrestricted except that it must be incentive compatible for the partners to pay them. This structure yields potentially rich patterns of partnership development that McAdams (2010) describes in terms of transitions between (1) “dating” partnerships that last just one period with zero quality, because they are immediately seen to be unpromising; (2) “honeymoon” partnerships that survive the dating stage, which have higher current payoffs because they are better matches and, as a result, induce higher quality; (3) “hard time” partnerships that yield low current payoffs because the state is less propitious but may, nevertheless, be kept going because of the probability of improvement; and (4) “golden years” partnerships in which the state has become so favourable, and the quality delivered so high, that the partnership ends only with the “death” off one of the partners.

In Doornik (2006), one party’s chosen quality is not observed by the other. Instead, there is for each i an observable (but not verifiable) stochastic individual performance measure x^i with *iid* probability distribution conditional on quality $f(x^i | q^i)$ that satisfies MLRP (Assumption 1). Party i receives return y_t^i , which is a deterministic function of both the realized performance measures (x_t^1, x_t^2) but is not verifiable, so there cannot be a legally enforceable contract to share these. The parties can, however, agree a relational contract involving voluntary bonuses. Both parties are risk neutral. Doornik (2006)

shows that, if a relationship is feasible, there is always an optimal relational contract that is *invariant*, in the sense of being the same every period apart from possibly calling for termination after some outcomes even if continuation is feasible and efficient. When CDFC (Assumption 2) holds and the two parties are identical, an optimal invariant contract has the following properties: (a) the parties terminate the relationship with probability one when both performance levels are below some thresholds; (b) otherwise, the bonus payment takes one of two levels depending essentially on the likelihood ratios of the observable performance measures in a way that corresponds to the party with lower relative performance paying the other. An optimal invariant contract does not necessarily require termination when continuation is feasible and efficient. An example is when each party's return is independent of the other's performance measure so that returns are separable. In that case, each party can act in the role of the purchaser in the moral hazard model in Levin (2003) with respect to the quality decision by the other. But, more generally, inefficient termination in some circumstances ("burning money") is required to keep the relationship incentive compatible, as in the moral hazard model in Levin (2003) with subjective performance measures. For some values of the discount factor, there is no fixed payment that is self-enforcing in the absence of bonuses but efficient quality is implementable with an optimal relational contract. Thus price (fixed payment plus bonus) varies with costs, as observed with Japanese automobile subcontracting. Price, however, varies less than if it were determined by period-by-period bargaining in the absence of a relational contract.

Rayo (2007) considers n risk-neutral parties choosing qualities $q_t = (q_t^1, \dots, q_t^n)$ whose joint output y is verifiable, as in Garvey (1995), but a random variable with expected value $E[y | q]$. (There is also a party $n + 1$ who plays no role in production and serves merely to balance the payments between the parties but whose payments still have to be incentive compatible.) In addition, there is a signal x^i of party i 's performance with density $f_i(x^i | q^i)$ that is independent across parties, observed by all parties, but not verifiable by outsiders. Denote by I the set of those that participate in a relational contract. Each participant receives a constant share $\alpha^i \geq 0$ of the realised output y_t , where the α^i are specified in a legally enforceable agreement and sum to one over all participants. Participants can also receive a legally enforceable fixed payment p_t^i and voluntary bonus payments $b_t^i(x_t)$ conditioned on the vector of realized signals x_t for all

participants — by assumption x_t is a sufficient statistic for (x_t, y_t) . Rayo (2007) assumes that each participant i 's net bonus plus discounted payoff in the continuation equilibrium at $t + 1$ conditional on x_t is additively separable in her own performance signal x_t^i and the performance signals of the other participants, denoted x_t^{-i} , and so can be written $V^i(x_t^i) + W^i(x_t^{-i})$ for some functions V^i and W^i . (Without that assumption, the optimal contract would take the form of a “winner take all” tournament which, it is argued, would have practical problems of implementation.) He then shows that the constraints for no party to default on the required bonus payments can, as in Levin (2003), be pooled into the single aggregate *dynamic enforcement* (DE) constraint

$$\delta \left(S - \tilde{S} \right) \geq \sum_{i \in I} \left\{ \sup_{x_t^i} V^i(x_t^i) - \inf_{x_t^i} V^i(x_t^i) \right\}. \quad (43)$$

Rayo (2007) calls the term in braces with superscript i in this i 's *implicit incentive*. It is the difference, under the relational contract and measured at the point bonuses are to be paid, between participant i 's payoffs with the best and the worst possible values of x_t^i . These implicit incentives arise from the gain from having the relationship continue in the future, so their total over all participants cannot exceed the total gain, which is what (43) requires. That condition is effectively an extension of (24) to the case of many participants with the continuation equilibrium depending on the stochastic signals. If the continuation equilibrium did not depend on those signals, payoffs from $t + 1$ on would cancel on the right-hand side of (43), leaving only the bonus terms as in (24).

Consider the implementation of some vector of qualities q with the vector of shares α . The share α^i provides an *explicit* incentive for participant i to deliver quality but any additional incentive has to be provided by the implicit incentive, that is via the term in braces with superscript i in (43). The minimum implicit incentive that will do this can be written $\Delta V^i(\alpha^i, q)$. Rayo (2007) first considers the case in which quality q_t^i is observed only by participant i and the $f_i(\cdot)$ satisfy MLRP and CDFC (Assumptions 1 and 2). Then, for implementing any optimal vector of qualities \hat{q} that is less than fully efficient, there exists an optimal allocation of shares $\hat{\alpha}$ such that one participant, call her l , has $\hat{\alpha}^l$ sufficiently high that $\Delta V^l(\hat{\alpha}^l, \hat{q}) = 0$. Participant l is the participant i for whom $-\partial \Delta V^i(\hat{\alpha}^i, \hat{q}) / \partial \alpha^i$ is largest. The intuition is as follows. (43) constrains the total implicit incentives that can be provided to all participants. Since incentives are in

short supply when efficient qualities cannot be attained, it is optimal to use the legally enforceable shares for that participant i for whom they are most effective at reducing the implicit incentive required, that is, for whom $-\partial \Delta V^i(\hat{\alpha}^i, \hat{q}) / \partial \alpha^i$ is largest. Intuitively, l is a relatively productive participant whose quality is hard to assess by means of x^l . Rayo (2007) shows that participant l can be interpreted as an endogenously chosen principal who, in the spirit of Alchian and Demsetz (1972), receives the total gain from continuation of the relationship but whose role as principal is just to make all the payments to other participants, not to monitor how they perform.

The other main case considered by Rayo (2007) is that in which all qualities are observable by all suppliers, that is, $x_t^i = q_t^i$. Then, provided some participants receive implicit incentives under an optimal contract, all do $(\Delta V^l(\hat{\alpha}^l, \hat{q}) > 0$ for all i) and profit shares are dispersed ($\hat{\alpha}^i < 1$ for all i). Rayo (2007) explains the intuition as follows. Suppose participant i receives no implicit incentive, so all her incentives come from $\hat{\alpha}^i > 0$. When q^i is chosen optimally, by the envelope theorem the change in i 's payoff from a small increase in q^i is of second order. Thus the cost of inducing slightly higher quality from agent i with a small implicit incentive is essentially negligible and hence always worthwhile. That $\hat{\alpha}^i < 1$ is the converse of this: if $\hat{\alpha}^i = 1$, i gets the full additional joint gain from additional quality and so would choose efficient quality without any implicit incentive. With an implicit incentive as well, i would choose quality above the efficient level, which cannot be optimal. Thus, given that there is *some* implicit incentive, it must be that $\hat{\alpha}^i < 1$.

The contrasting results for the two cases in Rayo (2007) illustrate the importance for incentive arrangements of the information quality of performance measures. When participants are able to monitor each other's performance only poorly, it is optimal to have one of them (one who is relatively productive but whose performance is particularly hard to assess) act as a principal who receives all the gain from continuation of the relationship and manages payments to all the others. When, on the other hand, participants are able to monitor each other's performance well, a partnership in which they all receive shares of current output and also some of the gain from continuing the relationship in the future is optimal.⁵

⁵Watson (1999), Watson (2002) and Sobel (2006) consider partnerships in which partners contribute to a joint output but do not make direct payments to each other. From a modelling perspective, this means that the dynamic enforcement constraints of the partners cannot be pooled into a single constraint. From a practical perspective, it makes the models less appropriate for studying contracts between firms.

10 Organization design

GM eventually took over Fisher Body because, on the interpretation in Klein (2000), efficiency requirements moved outside the zone of a self-enforcing relational contract. As the property rights literature stemming from Grossman and Hart (1986), Hart and Moore (1988) and Hart and Moore (1990) has emphasized, however, a takeover does not remove the underlying incentive issues, it just moves them inside the integrated firm. The basic models discussed in Section 3 above can be used to throw light on when it is more efficient to handle relational incentive issues within firms than between firms and, hence, make vertical integration worthwhile.

In Baker et al. (2002), non-integrated supply corresponds to the outsourcing model, integrated supply to the employment model. The difference arises because ownership of the assets used in supply is taken to confer ownership of the goods produced by those assets. Then, when a supplier is taken over, its previous owner becomes a manager (as the Fisher brothers did when GM took over Fisher Body) who still has to be induced to provide appropriate quality but can no longer withhold supply for use as a bargaining chip for payment because the output belongs to the downstream firm. In the event one party breaches a relational contract, the parties trade in the spot market thereafter (the so-called *grim trigger* strategies) but renegotiate ownership of the asset to provide the most efficient spot market operation. In the basic models of Section 3, the implication is that $\tilde{U}_{t+1} + \tilde{\Pi}_{t+1}$ in (8) for the employment model is the same as $\tilde{U}_{t+1} + \tilde{\Pi}_{t+1}$ in (15) for the outsourcing model because, whether or not the firms start out integrated, renegotiation leads to the same efficient spot market trading. Then, (8) with $\rho = 1$ is a weaker constraint than (15) for given q_t if $z_t(\tilde{q}_t) - c_t(\tilde{q}_t) > z_t(q_t)$, precisely the condition for use of a fixed price contract to be advantageous in the model of Itoh and Morita (2009) discussed in Section 5. Notice that every term in this condition involves

(In Sobel (2006), the probability that cheating is detectable by the courts can, in addition, be increased by incurring an enforcement cost.) Thomas and Worrall (2010) allow for transfers between agents via the equilibrium outcomes of a Nash demand game played each period after output has been produced. This has the effect of imposing credit constraints that can result in rewards to one party being postponed (backloaded) until a time at which the other party is unconstrained. The formal structure of cooperation in repeated partnerships is related to that of collusion in repeated oligopoly, of which there are recent models in Sannikov and Skrzypacz (2007) and Athey and Bagwell (2008). The former makes the relationship explicit in an application of its result that collusion cannot be sustained when the parties' actions are perfect substitutes, new noisy information arrives continuously, and the parties are able to respond to that information quickly. But, for reasons that are obvious in the context of illicit collusion, those models do not allow direct payments between the parties.

either the function $z_t(\cdot)$ or the quality \tilde{q}_t , both of which apply only when the relational contract is breached. It is only what happens in the event of breach that determines efficient organization design. Halonen (2002) considers another form of organization, *joint ownership*, which is discussed further below.

Is this framework useful for discussing GM's takeover of Fisher Body? Suppose it were initially the case that $z_t(\tilde{q}_t) - c_t(\tilde{q}_t) < z_t(q_t)$ so that (15) is weaker than (8). Then a given q_t is sustainable at a lower discount factor δ_t with the supplier a separate company, so the relational efficient organization is outsourcing. For vertical integration to become efficient, that inequality would have to be reversed. Transferring Buick body production to Flint was presumably jointly advantageous to GM and Fisher Body because the saving on the costs of transporting completed bodies outweighed any increased costs of inputs. It is not, however, obvious that a reduction in transportation cost alone would have changed the relational efficient organization. Suppose, for concreteness, the cost $c_t(q_t)$ at the body plant gate was unaffected by the change in location, the advantage being measured by an increase in $y_t(q_t)$ resulting from the saving in transportation cost of k per body for each of q_t bodies delivered. Under the *split the difference* bargaining outcome assumed by Baker et al. (2002), z_t is determined by the parties sharing equally the gain from the purchasing firm receiving the output over it being sold to a third party. Formally, with v_t the value for which the output can be sold to a third party,

$$z_t(q_t) = v_t(q_t) + \frac{1}{2} [y_t(q_t) - v_t(q_t)] = \frac{1}{2} [y_t(q_t) + v_t(q_t)]. \quad (44)$$

So, if the cost of transporting bodies to an alternative automobile manufacturer increased by the same k per body, $z_t(q_t)$ would remain unchanged and the relocation to Flint would leave the relational efficient ownership unchanged. This example is, of course, intended only to illustrate use of the theory, not to provide a full analysis of the situation faced by GM and Fisher Body. Some careful research is still needed to see whether the model can really be applied to that case and, if so, how.

Baker et al. (2002) show that there is another effect to consider. Suppose, as in Levin (2003), the outcome y is stochastic given q_t and so is z , with both *iid* given q_t so that the expected future payoff to the relationship does not depend on the current realization. Also as in Levin (2003), suppose the purchaser observes neither q_t nor the realization of the stochastic terms but does observe the realizations y_t and z_t . Then, in the notation

of Section 6, the jointly observed outcome is $\phi_t = \{y_t, z_t\}$. Because the purchaser does not observe q_t (and so cannot punish directly breach of an agreement about quality), the incentive compatibility condition for the supplier to deliver quality becomes (in a stationary environment)

$$q_t \in \arg \max_q E_y [P(\phi) | q] - c(q). \quad (45)$$

This is just like (21) for the model in Levin (2003) except that there is no argument θ because the agent can be of only one type. In the employment model, the incentive compatibility constraints (22) and (23) for the purchaser and the supplier to pay agreed bonuses apply just as in Levin (2003) and hence so does the *dynamic enforcement* (DE) constraint (24). In the outsourcing model, (45) continues to apply but, in line with (13) and (14) for the outsourcing model, (22) and (23) become

$$\delta (\Pi - \tilde{\Pi}) \geq \sup_{\phi} [b(\phi) - z]; \quad (46)$$

$$\delta (U - \tilde{U}) \geq \inf_{\phi} [z - b(\phi)]. \quad (47)$$

Hence, the interaction between the variation in bonuses and whether the firms are integrated affects the incentives to breach a relational contract, so the variation in payoffs, not just their expected values, can be important in the choice between organizational forms. Baker et al. (2002) make use of specific functional forms to show that the relational efficient organization does indeed depend on the variation in output for given q_t and also to derive further results.

Several papers consider developments of the model in Baker et al. (2002). Baker et al. (2001) extends it to allow supply to have some verifiable characteristics, as in the models in Section 5. Then the contractual payment p can be conditioned on those verifiable characteristics. The main result is that it is not possible to replicate the payoffs of spot outsourcing with a relational employment contract. Kvaløy (2007) considers an alternative to grim trigger strategies, specifically that the parties trade in the spot market for one period only and then revert to their relational contract. That affects \tilde{U} and $\tilde{\Pi}$ and hence the relational contracts that are sustainable. Ruzzier (2009) adopts a different specification of how $z_t(\cdot)$ is determined. Specifically, instead of assuming *split*

the *difference* bargaining as in (44), he assumes v_t is an outside option that can be taken up by the supplier only once bargaining with the purchaser has been broken off, giving the bargaining outcome

$$z_t = \begin{cases} \frac{1}{2}y_t, & \text{if } \frac{1}{2}y_t \geq v_t; \\ v_t, & \text{if } \frac{1}{2}y_t < v_t. \end{cases}$$

Then the distribution of v_t does not affect the choice between organizational forms if the top line applies (the case Ruzzier (2009) calls *high specificity*) and the distribution of y_t does not affect it if the bottom line applies (*low specificity*).

The underlying framework has also been extended to the case in which both parties take actions that contribute to the value of their joint output, the partnership models of Section 9. In Halonen (2002), both parties make their quality decisions simultaneously, these are observed by the other party, and either or both can be applied to an asset. If both are applied to the asset, the joint output is $y(q^1, q^2) = q^1 + q^2$; if only party i 's, party i 's output is λq^i , with $\lambda \in [0, 1]$, and party j 's output for $j \neq i$ is zero whatever q^j . (λ is known so there is no stochastic element.) Neither quality nor output are verifiable, so no sharing contract is legally enforceable. In a one-period relationship, and after one party has breached a relational contract, the parties still cooperate to apply both qualities to the asset because that gives a higher total output, which they share by bargaining that splits equally the difference between their joint output and what their individual outputs would be in the absence of cooperation. If the asset is owned by party i , these individual outputs are λq^i and 0. But the asset can also be owned jointly, in which case each party can prevent the other using it, so both have individual outputs of zero. With ownership by party 1, the rewards in a one-period relationship that result from this bargaining are, in the notation of Section 9,

$$z^1(q^1, q^2) = \lambda q^1 + \frac{1}{2} [(1 - \lambda) q^1 + q^2] = \frac{1}{2} [(1 + \lambda) q^1 + q^2] \quad (48)$$

$$z^2(q^1, q^2) = \frac{1}{2} [(1 - \lambda) q^1 + q^2], \quad (49)$$

so the optimal qualities defined in (37) are

$$\tilde{q}^1 = \arg \max_{q^1 \geq \underline{q}} \frac{1}{2} [(1 + \lambda) q^1 + q^2] - c(q^1) \quad (50)$$

$$\tilde{q}^2 = \arg \max_{q^2 \geq \underline{q}} \frac{1}{2} [(1 - \lambda) q^1 + q^2] - c(q^2). \quad (51)$$

With joint ownership, both parties' individual outputs are zero, so the rewards in a one-period relationship are given by (48) and (49) with $\lambda = 0$. Thus, party 2 chooses the same quality \tilde{q}^2 under both forms of ownership and, provided $\lambda > 0$, party 1 chooses higher quality \tilde{q}^1 under single ownership (but still less than the efficient level because $\lambda < 1$). Thus in a one-period relationship single ownership is better than joint ownership for $\lambda > 0$.

Under a relational contract, that is not necessarily the case. The formal analysis can be done by inserting the appropriate functional forms into (41). In contrast to Baker et al. (2002), Halonen (2002) considers primarily the assumption that ownership is not renegotiated in the event of breach by either party, which corresponds to making \tilde{S}_{t+1} a function of λ . When $\lambda = 0$, single ownership and joint ownership are again equivalent. With joint ownership, nothing changes as λ changes (because λ is then irrelevant). However, with a single owner, \tilde{S}_{t+1} is an increasing function of λ for $\lambda \in [0, 1)$ because, from (50), \tilde{q}^1 is increasing in λ and, from (51), \tilde{q}^2 is independent of λ . So the long-term cost of breaching a relational contract for given (q^1, q^2) is decreasing in λ . The short-term gain from doing so is also decreasing in λ . Halonen (2002) shows that, for $c(q) = q^\gamma$ with $\gamma > 1$ and an optimal payoff structure when efficient qualities are achievable, these two effects exactly balance when $\gamma = 2$. For $\gamma < 2$, single ownership is relational efficient; for $\gamma > 2$, joint ownership is relational efficient. Halonen (2002) also considers the assumption that ownership is renegotiated in the event of default (so that \tilde{S}_{t+1} is independent of λ) but only for a specific numerical example.⁶

In Bragelien (2007) too, both parties make quality decisions but, as in Baker et al. (2002) but unlike in Halonen (2002), output is an *iid* random variable for given quality and quality decisions are unobserved by the other party. Neither output nor quality is verifiable. The benefit to each party if they subsequently cooperate is their own quality

⁶Applied to the model in Baker et al. (2002), joint ownership might be interpreted as implying that the supplier cannot sell the supply to a third party without the agreement of the purchaser. Then, in the formulation in (44), $v_t(q_t) = 0$ for all q_t and hence $z_t(q_t) = y_t(q_t)/2$.

plus an additive stochastic term; it is non-contractible but observed by both. If, after observing the benefit, either decides not to cooperate, their benefits are lower and, in addition to depending on what the benefits would have been had they cooperated, also depend on the ownership structure. If they cooperate, they share the gains from doing so. The form payments take under an optimal relational contract consists of each party receiving the payoff it would receive in the absence of cooperation, together with a fixed payment independent of the benefits from cooperation and a payment depending on the realized values of those benefits that takes one of just two values (a one-step payment schedule). With sufficient symmetry, the last of these payments depends on which agent has the higher realized gain from cooperation and so resembles a tournament. This extends results in Levin (2003) for the case in which just one party decides quality to the case in which both do.

The main focus of Bragelien (2007) concerns the choice of organizational structure. Non-cooperation results in the parties operating as if in a spot market ever after. The paper considers both the case discussed by Halonen (2002), in which the parties are committed to the ownership structure for ever once it has been determined, and that discussed by Baker et al. (2002), in which the parties can renegotiate their ownership structure if they move to a spot market relationship. For the first case, Bragelien (2007) gives conditions that favour an ownership structure that gives weaker (and, conversely, stronger) incentives in spot market interaction. They differ from those in Halonen (2002) because of the uncertainty. The generalization of Baker et al. (2002) is that the potential benefits are continuous, not just binary. For this case, the general result is that the ownership structure should never give worse incentives in the spot market for one party without giving better incentives for the other. (“Better” here may be stronger or weaker depending on whether that party would otherwise choose quality below or above the efficient level.) There are also results for correlated performance measures. Bragelien (2007) illustrates the results with linear contracts, which may be optimal (if error terms are uniformly distributed and functions linear) and for which explicit solutions can be derived.

Blonski and Spagnolo (2007) question whether, when modelled properly, there is actually a difference between the case in which ownership can be renegotiated following breach and the case in which it cannot. The essence of their argument, formalized in a

model with both parties choosing quality but no stochastic elements, is as follows. Trading in a spot market environment after one party breaches is inefficient, so the parties would gain by renegotiating back to continued cooperation. Doing that does not destroy incentive compatibility if the parties use *restitution strategies*, based on an idea in Farrell and Maskin (1989), in which the payoff to the breaching party in the continuation equilibrium played following breach is reduced to the level that party would have obtained by trading in the spot market ever after. (In the absence of wealth constraints, that can be achieved by a single monetary payment for restitution but this is not the only possibility.) The reason restitution strategies do not upset incentive compatibility is that, if defaulting at t , the supplier still gets continuation payoff \tilde{U}_{t+1} and, if defaulting at t , the purchaser still gets continuation payoff $\tilde{\Pi}_{t+1}$ — it is only the continuation payoff of the non-breaching party that is affected by the restitution. But, if the parties were using the relational efficient organization before breach, they will have no reason to renegotiate the organizational structure after breach. Thus it makes no difference whether they can renegotiate that structure.

Baker et al. (2010) consider a different issue of organization design: how to allocate decision rights in an organisation in order to adapt decisions to changing circumstances. To illustrate, consider two parties and a single decision about quality q_t that can, in principle, be taken by either of them. Which has the right to do so has to be determined before all uncertainty, represented by a random state of the world s , is resolved and the party given that right takes the decision in its own interest once the uncertain state of the world has been revealed.⁷ So, in a one-period relationship, it is optimal to give the decision right to whichever party's short-term self-interested decisions for each state s yield the highest expected joint payoff. But in an on-going relationship in which actual decisions are observed by all parties, voluntary bonus payments and the payoffs in the continuation equilibrium can be conditioned on the state and on the decision made, so the decision-maker can be induced to take decisions not solely in a short-term self-interested way. Then the party to whom the decision right is given becomes just like the supplier in the employment model of Section 3.1 except that there is uncertainty at the time the

⁷Baker et al. (2010) contrast their approach with the incomplete contracts approach of Grossman-Hart-Moore in the following way. In the latter, a legally enforceable agreement about the decision to be taken can always be negotiated between the parties after any uncertainty is resolved no matter who has the formal right to take it. In the former, there is insufficient time to negotiate such an agreement, a situation referred to by Williamson (2000) as “maladaptation in the contract execution interval”.

purchaser commits to p_t that is resolved by the time the supplier makes the decision q_t . (Baker et al. (2010) also allow for payments to be made between the state being revealed and the decision q_t taken.) Assume for simplicity that the state s is *iid* and denote its realized value at t by s_t . Then, in an obvious analogy with equation (1), we can write party 1's payoff conditional on being the decision-maker in period t and on the relational contract continuing as

$$U_t^1(s_t) = p_t + b_t(q_t, s_t) - c_t^1(q_t, s_t) + \delta_t E_s U_{t+1}^1(s), \quad \text{for all } t \geq \tau, \quad (52)$$

where $c_t^1(q_t, s_t)$ is now to be interpreted as the net cost to party 1 of making decision q_t in state s_t instead of the decision that maximizes her short-term interest. One can make exactly the same notational changes to the purchaser's payoffs for the non-decision-making party 2 and so derive the incentive compatibility condition conditional on s_t , corresponding to (8) with $\rho = 1$,

$$\begin{aligned} \delta_t E_s [S_{t+1}(s) - \tilde{S}_{t+1}(s)] &\equiv \delta_t E_s [U_{t+1}^1(s) - \tilde{U}_{t+1}^1(s) + \Pi_{t+1}^2(s) - \tilde{\Pi}_{t+1}^2(s)] \\ &\geq c_t^1(q_t, s_t), \quad \text{for all } t \geq \tau. \end{aligned}$$

Since the decision right could, in principle, be allocated to different parties in successive periods, the joint future gain $E_s [S_{t+1}(s) - \tilde{S}_{t+1}(s)]$ is independent of who has the decision right in period t , as well as of the decision taken at t . The parties can thus implement the decision rule $q_t(s)$ for all s with the decision allocated to party 1 in period t only if

$$\delta_t E_s [S_{t+1}(s) - \tilde{S}_{t+1}(s)] \geq \max_s c_t^1(q_t(s), s).$$

One can derive a similar condition for when the decision is allocated to party 2. The incentive compatibility requirement is obviously weaker if the decision right is allocated to the party i for which $\max_s c_t^i(q_t(s), s)$ is smaller. The decision rule $q_t(s)$ can thus be implemented for all s by some allocation of the decision right only if

$$\delta_t E_s [S_{t+1}(s) - \tilde{S}_{t+1}(s)] \geq \min_{i \in \{1,2\}} \max_s c_t^i(q_t(s), s), \quad \text{for all } t \geq \tau.$$

This example illustrates the underlying idea. Baker et al. (2010) allow for richer structures

with many parties and many decisions for which the right to decide can be allocated, not necessarily all to the same party. There are then many potential governance structures that specify who has the right to make which decision.

11 Rent and relational contracts

Crucial to sustaining a relational contract is that the sum of the future payoffs to the two parties from sticking to it is larger than the sum of their future payoffs from breaching it. If that is not the case, $S_{t+1} = \tilde{S}_{t+1}$ in (8) and no relational contract can be sustained. In other words, the parties must have a rent from continuing their relationship. The rent can go to either party — (8) is a condition on the total rent, not its distribution and the payments p_t and b_t can be used to distribute that rent in any proportions. But that rent must exist. The same applies to all the other models described in this chapter.

In much of the literature, the required rent comes from the assumption that, if one of the parties breaches the relational contract, they either trade in the spot market thereafter or, if spot market trade cannot be sustained, they cease to trade completely. Even when, as in Blonski and Spagnolo (2007), they do not *actually* do that in the event of breach, the *possibility* of doing so is still important in ensuring that it is incentive compatible for the breaching party to pay restitution to the non-breaching party. In practice, of course, each party may have other potential trading partners. This is not necessarily a problem if the potential partners observe which party has breached and treat the breaching party as a pariah, or even if other potential partners observe only that a relationship has come to an end and treat both of the parties to that relationship as pariahs, see MacLeod and Malcomson (1989). But, in practice, there are many reasons other than breach for a relational contract to come to an end that third parties may not be able to distinguish from breach. As Bewley (1999, p. 298) emphasizes in the case of employment, many employers find it difficult to get information about potential employees from former employers. Dixit (2003) discusses market solutions motivated by the Sicilian Mafia in which intermediaries sell, at a profitable price, either information about the past behaviour of potential partners or enforcement services. But a third party has no reason to pay that price unless it would reveal something inherent about the potential partner that is relevant to the productivity of a future relationship.

Without some way to generate the required rent despite outside markets, relational contracts will break down, as Bulow and Rogoff (1989) note in the case of sovereign debt. There are, however, ways in which that rent can arise even if potential partners do not discriminate against those whose previous relationships have ended. Specific investments upfront are an obvious one. But, in general, there is no reason why the efficient level of specific investment, even if positive, should be large enough to sustain a relational contract, let alone one that delivers efficient quality. An alternative widely discussed in the literature is the possibility that at least one party is not in a position to form another relationship straightaway. This may be because of exogenous market frictions due to search and matching, as in Kranton (1996*b*), Ramey and Watson (1997), Ramey and Watson (2001) and Sobel (2006). Or it may be generated endogenously through unemployment on one side of the market, as originally modelled in Shapiro and Stiglitz (1984) in the context of efficiency wages. In the latter case, it does not matter which of the parties has difficulty rematching. It is an *assumption* of Shapiro and Stiglitz (1984) that the rent must go to employees in the form of efficiency wages, rather than to firms as the result of excess vacancies, as MacLeod and Malcomson (1989) and MacLeod and Malcomson (1998) show. But there must be a rent to one party or the other. For more on this issue in the context of employment, see Malcomson (1999).

Another possibility is explored by Kranton (1996*a*): starting the relationship with low quality so that there is little to lose if the other party defaults and having quality build up over time. Recall that, in the basic employment model, (4), (7) and (8) are necessary and sufficient conditions for a quality sequence to be an equilibrium and that these are all inequalities. Moreover, only future gains (that is, from the second period of the relationship $\tau + 1$ on) affect the left-hand side of (8) and lowering q_τ reduces the right-hand side. Thus, for any sequence of qualities $(q_\tau, q_{\tau+1}, \dots)$ that can be sustained as an equilibrium, a sequence $(q'_\tau, q_{\tau+1}, \dots)$ with $q'_\tau < q_\tau$ can also be sustained as an equilibrium provided the individual rationality conditions (4) and (7) are still satisfied. As long as q_τ is not above the efficient level, it also yields less gain to starting a new relationship and, hence, less temptation to breach an existing one. Moreover, setting $q'_\tau < q_t$ creates some slack in (8) even if there were none before, so reducing $q_{\tau+1}$ (and hence the left-hand side of (8) for $t = \tau$) by a small amount will not upset incentive compatibility for period τ , creates some slack in the incentive compatibility constraint for $\tau + 1$, and does not affect

incentive compatibility from $\tau + 2$ on. Thus, in the limit, one can always set $q_t = \underline{q}$ for as many periods after τ as needed to reduce U_τ to \bar{U}_τ and Π_τ to $\bar{\Pi}_\tau$, a way suggested by Murphy and Topel (1990) to remove the need for involuntary unemployment in the model of Shapiro and Stiglitz (1984). Watson (1999), Watson (2002), and Sobel (2006) similarly explore starting relationships in a small or inefficient way.

Essentially, this is a way of dissipating the potential gains from new relationships in order to make starting a new relationship less attractive than continuing an existing one. There are other ways this can be done. Klein and Leffler (1981) explore advertising. Carmichael and MacLeod (1997) explore gifts that cost the giver more than they are worth to the recipient. There are no doubt other possibilities.

12 Courts and empirical evidence

This chapter is primarily about theory. Empirical evidence on contracts between firms is discussed in the chapter by Lafontaine and Slade (2011), including the issues of rent and reputation that arise in relational contracts. But one issue of particular relevance to relational contracts is not covered there: the quality of the legal system.

Relational contracts are concerned with agreements that can be enforced without resort to courts. The spirit of much of the theory discussed here is that, although an effective legal system exists, important elements of the relationship cannot be enforced legally because courts do not have the information to do that. Relational contracts are, however, also valuable for enforcing agreements that courts could, in principle, enforce but cannot, in practice, be relied upon to do so. The parties may, for example, be in different legal jurisdictions, or courts may be too corrupt, cumbersome or otherwise ineffective. Djankov et al. (2003) document how slow and ineffective courts can be, even in countries with highly-developed legal systems, over such seemingly simple things as the collection of a check returned for non-payment. Such circumstances provide good reason for parties to use relational contracts to avoid recourse to courts whenever they are effective substitutes, not just when crucial information is unverifiable (though the relational contracts may then not be able to rely on even non-contingent payments being legally enforceable). Greif (1994) documents the way in which 11th-century Maghribi traders used reputations to enforce relational contracts when trading between countries.

Fafchamps (1996) and Bigsten et al. (2000) document the widespread use of relational contracts among African businesses, and their less widespread use where legal institutions best support business. McMillan and Woodruff (1999) find that private sector firms in Vietnam provide customers with more credit when alternative suppliers are hard to locate and when customers are identified through business networks, both of which may provide sources of sanction for default. Banerjee and Duflo (2000) demonstrate the importance of reputation effects in Indian software. But while relational contracts can clearly act as a substitute for well-functioning courts, better formal institutions can also foster contracting, as Johnson et al. (2002) find in post-communist countries.

13 Concluding remarks

Extensive as it is, this chapter does not cover anything like all the literature on relational contracts. As noted in the Introduction, it focuses on theory as applied to supply relationships between firms, and does not consider applications to purely financial contracts, employment or relationships involving sovereign bodies. The inability to commit that underlies relational contracts is inherent to relationships involving sovereign bodies because there is nobody to enforce a contract on such a body even when the relevant information is publicly available. Thus, relational contracts have been applied to issues of sovereign debt, to treaties between sovereign states, and to the relationships between rulers and ruled within states. Acemoglu (2003), for example, develops incentive compatibility conditions similar to those used here in analysing taxation under dictatorship and democracy.

This chapter has, moreover, not discussed relational contracts concerned only with risk sharing. Indeed, for tractability, most of the literature surveyed here assumes all parties are risk neutral, though there are notable exceptions such as Pearce and Stacchetti (1998) and MacLeod (2003). Issues of risk aversion perhaps deserve greater attention, So (in addition to other issues noted above) do issues of which of the multiple equilibria will be played in practice and of the implications of non-stationary environments.

There are also foundational issues that deserve greater attention. The findings of Macaulay (1963) provided an important motivation for economic research on relational contracts. However, models of the type presented here certainly do not capture all those

findings. One role for relational contracts discussed by Macaulay (1963) is as a substitute for planning exchange relationships completely, a framework within which the relationship can be adjusted in the face of either unforeseen contingencies or inconsistent perceptions by the parties about what should happen in certain eventualities. That is not a feature of the models discussed above. Formally, equilibria in those models are in terms of strategies that specify what actions the parties are to take for every history that can occur, and thus for every possible contingency. Moreover, equilibrium strategies are consistent in the sense that each party's strategy is a best response to the other's. Relational contracts in these models are a substitute for enforcement by courts, not a substitute for careful planning. There is certainly no modelling of unforeseen contingencies. As in so much of economics, incorporating unforeseen contingencies remains a major challenge. Kreps (1996) reflects on these issues and possible ways forward.

Some earlier writers about relational contracts, such as Dore (1983), contrasted them with standard economic analysis and, in particular with economists' perceptions of allocative efficiency. As Dore (1983, p. 472) put it: "Any economist, at least any economist worth his neo-classical salt, would be likely to scoff at the idea. Just think, he would say, of the market imperfections, of the misallocation and loss of efficiency involved." That was, however, before most of the literature discussed above was written. It should be apparent from this survey that research by economists on relational contracts is precisely about how the parties may achieve *more* efficient outcomes by making use of relational contracts in circumstances in which courts cannot be relied upon to enforce contractual agreements.

There remain important differences of interpretation, however, about the notion of *trust*. Trust is seen by many non-economists as central to relational contracts. Some of the literature surveyed above interprets an equilibrium with quality above the level sustainable in a one-period relationship as involving trust because at least one party leaves itself vulnerable to exploitation by the other. As MacLeod (2007*b*, p. 609) puts it: "In a relational contract, one party trusts the other when the value from future trade is greater than the one period gain from defection." In this spirit, Kvaløy and Olsen (2009) interpret an increase in the discount factor δ as an increase in trust because it increases the stake that the parties are prepared to risk. That interpretation of trust is not, however, universally accepted. In the context of exchanges in which the gains from future dealings

are sufficiently highly valued to induce cooperation, Sabel (1993, p. 1135) writes that: “it would be wrong to associate cooperation with trust at all, because cooperation results from continuous calculation of self-interest rather than a mutually recognized suspension, however circumscribed, of such calculation.” Sako (1992) develops different concepts of trust along these lines applicable to inter-firm relations in Britain and Japan. These writers view relational contracts as being about more than self-interested cooperation.

Whether or not they capture all the elements of relational contracts, however, the economic models discussed here have certainly been invaluable in illuminating important features of economic relationships. Further research will no doubt illuminate yet more.

References

- Abreu, D., Milgrom, P. and Pearce, D. (1991), ‘Information and timing in repeated partnerships’, *Econometrica* **59**(6), 1713–1733.
- Acemoglu, D. (2003), ‘Why not a political coase theorem? social conflict, commitment, and politics’, *Journal of Comparative Economics* **31**(4), 620–652.
- Alchian, A. A. and Demsetz, H. (1972), ‘Production, information costs, and economic organization’, *American Economic Review* **62**(5), 777–795.
- Asanuma, B. (1989), ‘Manufacturer-supplier relationships in Japan and the concept of relation-specific skill’, *Journal of the Japanese and International Economies* **3**(1), 1–30.
- Asanuma, B. and Kikutani, T. (1992), ‘Risk absorption in Japanese subcontracting: A microeconomic study of the automobile industry’, *Journal of the Japanese and International Economies* **6**(1), 1–29.
- Athey, S. and Bagwell, K. (2008), ‘Collusion with persistent cost shocks’, *Econometrica* **76**(3), 493–540.
- Baker, G., Gibbons, R. and Murphy, K. J. (1994), ‘Subjective performance measures in optimal incentive contracts’, *Quarterly Journal of Economics* **109**(4), 1125–1156.
- Baker, G., Gibbons, R. and Murphy, K. J. (2001), ‘Bringing the market inside the firm?’, *American Economic Review* **91**(2), 212–218.

- Baker, G., Gibbons, R. and Murphy, K. J. (2002), ‘Relational contracts and the theory of the firm’, *Quarterly Journal of Economics* **117**(1), 39–84.
- Baker, G., Gibbons, R. and Murphy, K. J. (2010), s happens: Relational adaptation in contracts, firms, and other governance structures. MIT, Department of Economics.
- Baker, G. P. (1992), ‘Incentive contracts and performance measurement’, *Journal of Political Economy* **100**(3), 598–614.
- Banerjee, A. V. and Duflo, E. (2000), ‘Reputation effects and the limits of contracting: A study of the Indian software industry’, *Quarterly Journal of Economics* **115**(3), 989–1017.
- Baron, D. P. (1989), Design of regulatory mechanisms and institutions, in R. Schmalensee and R. D. Willig, eds, ‘Handbook of Industrial Organization’, Vol. II, Elsevier Science, Amsterdam, chapter 24, pp. 1347–1447.
- Battigalli, P. and Maggi, G. (2002), ‘Rigidity, discretion, and the costs of writing contracts’, *American Economic Review* **92**(4), 798–817.
- Battigalli, P. and Maggi, G. (2008), ‘Costly contracting in a long-term relationship’, *RAND Journal of Economics* **39**(2), 352–377.
- Bernheim, B. D. and Whinston, M. D. (1990), ‘Multimarket contact and collusive behavior’, *RAND Journal of Economics* **21**(1), 1–26.
- Bernheim, B. D. and Whinston, M. D. (1998), ‘Incomplete contracts and strategic ambiguity’, *American Economic Review* **88**(4), 902–932.
- Bewley, T. F. (1999), *Why Wages Don’t Fall During a Recession*, Harvard University Press, Cambridge, MA.
- Bigsten, A., Collier, P., Dercon, S., Fafchamps, M., Gauthier, B., Gunning, J. W., Oduro, A., Oostendorp, R., Patillo, C., Soderbom, M., Teal, F. and Zeufack, A. (2000), ‘Contract flexibility and dispute resolution in African countries’, *Journal of Development Studies* **36**(4), 1–37.

- Blonski, M., Ockenfels, P. and Spagnolo, G. (2009), Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. J. W. Goethe University, Department of Economics.
- Blonski, M. and Spagnolo, G. (2007), Relational efficient property rights. University of Frankfurt, Department of Economics.
- Board, S. (2008), Relational contracts and the value of loyalty. UCLA, Department of Economics.
- Bragelien, I. (2007), Asset ownership and relational contracts. Norwegian School of Economics and Business Administration.
- Brown, M., Falk, A. and Fehr, E. (2004), 'Relational contracts and the nature of market interactions', *Econometrica* **72**(3), 747–780.
- Bull, C. (1987), 'The existence of self-enforcing implicit contracts', *Quarterly Journal of Economics* **102**(1), 147–159.
- Bulow, J. and Rogoff, K. (1989), 'Sovereign debt: Is to forgive to forget?', *American Economic Review* **79**(1), 43–50.
- Calzolari, G. and Spagnolo, G. (2010), Relational contracts and competitive screening, Discussion Paper 7434, CEPR.
- Carlton, D. W. (1986), 'The rigidity of prices', *American Economic Review* **76**(4), 637–658.
- Carmichael, H. L. and MacLeod, W. B. (1997), 'Gift giving and the evolution of cooperation', *International Economic Review* **38**(3), 485–509.
- Chassang, S. (2010), 'Building routines: Learning, cooperation and the dynamics of incomplete relational contracts', *American Economic Review* **100**(1), 448–465.
- Che, Y.-K. and Yoo, S.-W. (2001), 'Optimal incentives for teams', *American Economic Review* **91**(3), 525–541.
- Cr  mer, J. (1986), 'Cooperation in ongoing organizations', *Quarterly Journal of Economics* **101**(1), 33–49.

- Daido, K. (2006), ‘Formal and relational incentives in a multitask model’, *International Review of Law and Economics* **26**(3), 380–394.
- Di Tella, R. and MacCulloch, R. (2002), ‘Informal family insurance and the design of the welfare state’, *Economic Journal* **112**, 481–503.
- Dixit, A. (2003), ‘On modes of economic governance’, *Econometrica* **71**(2), 449–481.
- Djankov, S., La Porta, R., Lopez-de Silanes, F. and Shleifer, A. (2003), ‘Courts’, *Quarterly Journal of Economics* **118**(2), 453–517.
- Doornik, K. (2006), ‘Relational contracting in partnerships’, *Journal of Economics and Management Strategy* **15**(2), 517–548.
- Dore, R. (1983), ‘Goodwill and the spirit of market capitalism’, *British Journal of Sociology* **34**(4), 459–482.
- Evans, R. (2008), ‘Simple efficient contracts in complex environments’, *Econometrica* **76**(3), 459–491.
- Fafchamps, M. (1996), ‘The enforcement of commercial contracts in Ghana’, *World Development* **24**(3), 427–448.
- Fairburn, J. A. and Malcomson, J. M. (1994), ‘Rewarding performance by promotion to a different job’, *European Economic Review* **38**(3/4), 683–690.
- Fairburn, J. A. and Malcomson, J. M. (2001), ‘Performance, promotion, and the Peter principle’, *Review of Economic Studies* **68**(1), 45–66.
- Farrell, J. and Maskin, E. (1989), ‘Renegotiation in repeated games’, *Games and Economic Behavior* **1**(4), 327–360.
- Fehr, E. and Falk, A. (2002), ‘Psychological foundations of incentives’, *European Economic Review* **46**(4/5), 687–724.
- Fong, Y.-F. and Li, J. (2010a), Information revelation in relational contracts. Northwestern University, Kellogg School of Management.
- Fong, Y.-F. and Li, J. (2010b), Relational contracts, efficiency wages, and employment dynamics. Northwestern University, Kellogg School of Management.

- Fuchs, W. (2007), ‘Contracting with repeated moral hazard and private evaluations’, *American Economic Review* **97**(4), 1432–1448.
- Garvey, G. T. (1995), ‘Why reputation favors joint ventures over vertical and horizontal integration: A simple model’, *Journal of Economic Behavior and Organization* **28**(3), 387–397.
- Gietzmann, M. B. and Larsen, J. G. (1998), ‘Motivating subcontractors to perform development and design tasks’, *Management Accounting Research* **9**(3), 285–309.
- Greif, A. (1994), ‘Cultural beliefs and the organization of society: A historical and theoretical reflection on collectivist and individualist societies’, *Journal of Political Economy* **102**(5), 912–950.
- Grossman, S. J. and Hart, O. D. (1986), ‘The costs and benefits of ownership: A theory of vertical and lateral integration’, *Journal of Political Economy* **94**(4), 691–719.
- Halac, M. (2009), Relational contracts and the value of relationships. Columbia University.
- Halonen, M. (2002), ‘Reputation and the allocation of ownership’, *Economic Journal* **112**(481), 539–558.
- Hart, O. D. and Moore, J. (1988), ‘Incomplete contracts and renegotiation’, *Econometrica* **56**(4), 755–785.
- Hart, O. and Moore, J. (1990), ‘Property rights and the nature of the firm’, *Journal of Political Economy* **98**(6), 1119–1158.
- Holmström, B. (1999), ‘Managerial incentive problems: A dynamic perspective’, *Review of Economic Studies* **66**(1), 169–182.
- Holmstrom, B. and Milgrom, P. (1991), ‘Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design’, *Journal of Law, Economics, and Organization* **7**(Special Issue), 24–52.
- Hörner, J. (2002), ‘Reputation and competition’, *American Economic Review* **92**(3), 644–663.

- Iossa, E. and Spagnolo, G. (2009), Contracts as threats: On a rationale for rewarding A while hoping for B. Brunel University, Department of Economics.
- Itoh, H. and Morita, H. (2009), Formal contracts, relational contracts, and the holdup problem. Hitotsubashi University.
- Johnson, S., McMillan, J. and Woodruff, C. (2002), ‘Courts and relational contracts’, *Journal of Law, Economics, and Organization* **18**(1), 221–277.
- Kawasaki, S. and McMillan, J. (1987), ‘The design of contracts: Evidence from Japanese subcontracting’, *Journal of the Japanese and International Economies* **1**(1), 327–349.
- Klein, B. (2000), ‘Fisher–General Motors and the nature of the firm’, *Journal of Law and Economics* **43**, 105–141.
- Klein, B., Crawford, R. G. and Alchian, A. A. (1978), ‘Vertical integration, appropriable rents, and the competitive contracting process’, *Journal of Law and Economics* **21**, 297–326.
- Klein, B. and Leffler, K. B. (1981), ‘The role of market forces in assuring contractual performance’, *Journal of Political Economy* **89**(4), 615–641.
- Kranton, R. E. (1996a), ‘The formation of cooperative relationships’, *Journal of Law, Economics, and Organization* **12**(1), 214–233.
- Kranton, R. E. (1996b), ‘Reciprocal exchange: A self-sustaining system’, *American Economic Review* **86**(4), 830–851.
- Kranz, S. and Ohlendorf, S. (2009), Renegotiation-proof relational contracts with side payments. University of Bonn, Department of Economics.
- Kreps, D. M. (1996), ‘Markets and hierarchies and (mathematical) economic theory’, *Industrial and Corporate Change* **5**(2), 561–595.
- Kvaløy, O. (2007), ‘Asset specificity and vertical integration’, *Scandinavian Journal of Economics* **109**(3), 551–572.
- Kvaløy, O. and Olsen, T. E. (2006), ‘Team incentives in relational employment contracts’, *Journal of Labor Economics* **24**(1), 139–169.

- Kvaløy, O. and Olsen, T. E. (2009), ‘Endogenous verifiability and relational contracting’, *American Economic Review* **99**(5), 2193–2208.
- Laffont, J.-J. and Tirole, J. (1988), ‘The dynamics of incentive contracts’, *Econometrica* **56**(5), 1153–1175.
- Lafontaine, F. and Slade, M. E. (2011), Inter-firm contracts: Evidence, *in* R. Gibbons and J. Roberts, eds, ‘Handbook of Organizational Economics’, Princeton University Press.
- Levin, J. (2002), ‘Multilateral contracting and the employment relationship’, *Quarterly Journal of Economics* **117**(3), 1075–1103.
- Levin, J. (2003), ‘Relational incentive contracts’, *American Economic Review* **93**(3), 835–857.
- Macaulay, S. (1963), ‘Non-contractual relations in business: A preliminary study’, *American Sociological Review* **28**, 55–67.
- MacLeod, W. B. (2003), ‘Optimal contracting with subjective evaluation’, *American Economic Review* **93**(1), 216–240.
- MacLeod, W. B. (2007a), ‘Can contract theory explain social preferences?’, *American Economic Review* **97**(2), 187–192.
- MacLeod, W. B. (2007b), ‘Reputations, relationships, and contract enforcement’, *Journal of Economic Literature* **45**, 595–628.
- MacLeod, W. B. and Malcomson, J. M. (1988), ‘Reputation and hierarchy in dynamic models of employment’, *Journal of Political Economy* **96**(4), 832–854.
- MacLeod, W. B. and Malcomson, J. M. (1989), ‘Implicit contracts, incentive compatibility, and involuntary unemployment’, *Econometrica* **57**(2), 447–480.
- MacLeod, W. B. and Malcomson, J. M. (1998), ‘Motivation and markets’, *American Economic Review* **88**(3), 388–411.
- Macneil, I. R. (1978), ‘Contracts: Adjustment of long-term economic relations under classical, neoclassical, and relational contract law’, *Northwestern University Law Review* **72**(6), 854–904.

- Mailath, G. J. and Samuelson, L. (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press, Oxford, UK.
- Malcomson, J. M. (1999), Individual employment contracts, *in* O. Ashenfelter and D. Card, eds, ‘Handbook of Labor Economics’, Vol. 3B, Elsevier, Amsterdam, chapter 35, pp. 2291–2372.
- Malcomson, J. M. (2011a), Relational incentive contracts with persistent private information. University of Oxford, Department of Economics.
- Malcomson, J. M. (2011b), Relational incentive contracts with private information. University of Oxford, Department of Economics.
- Martimort, D. and Stole, L. (2009), ‘Market participation in delegated and intrinsic common-agency games’, *RAND Journal of Economics* **40**(1), 78–102.
- McAdams, D. (2010), Performance and turnover in a stochastic partnership, ERID Working Paper 59, Duke University.
URL: <http://ssrn.com/abstract=1677568>
- McMillan, J. and Woodruff, C. (1999), ‘Interfirm relationships and informal credit in Vietnam’, *Quarterly Journal of Economics* **114**(4), 1285–1320.
- Milgrom, P. R. and Segal, I. (2002), ‘Envelope theorems for arbitrary choice sets’, *Econometrica* **70**(2), 583–601.
- Mukherjee, A. (2008), ‘Sustaining implicit contracts when agents have career concerns: The role of information disclosure’, *RAND Journal of Economics* **39**(2), 469–490.
- Murphy, K. M. and Topel, R. H. (1990), Efficiency wages reconsidered: Theory and evidence, *in* Y. Weiss and G. Fishelson, eds, ‘Advances in the Theory and Measurement of Unemployment’, Macmillan, London.
- Obstfeld, M. and Rogoff, K. S. (1996), *Foundations of International Macroeconomics*, MIT Press, Cambridge, MA.
- Pearce, D. G. and Stacchetti, E. (1998), ‘The interaction of implicit and explicit contracts in repeated agency’, *Games and Economic Behavior* **23**(1), 75–96.

- Prendergast, C. (1993), ‘The role of promotion in inducing specific human capital acquisition’, *Quarterly Journal of Economics* **108**(2), 523–534.
- Radner, R. (1985), ‘Repeated principal-agent games with discounting’, *Econometrica* **53**(3), 1173–1198.
- Radner, R. (1986), ‘Repeated partnership games with imperfect monitoring and no discounting’, *Review of Economic Studies* **53**(1), 43–57.
- Radner, R., Myerson, R. and Maskin, E. (1986), ‘An example of a repeated partnership game with discounting and with uniformly inefficient equilibria’, *Review of Economic Studies* **53**(1), 59–69.
- Ramey, G. and Watson, J. (1997), ‘Contractual fragility, job destruction, and business cycles’, *Quarterly Journal of Economics* **112**(3), 873–911.
- Ramey, G. and Watson, J. (2001), ‘Bilateral trade and opportunism in a matching market’, *Contributions to Theoretical Economics* **1**(1), article 3. <http://www.bepress.com/bejte/contributions/vol1/iss1/art3>.
- Ray, D. (2002), ‘The time structure of self-enforcing agreements’, *Econometrica* **70**(2), 547–582.
- Rayo, L. (2007), ‘Relational incentives and moral hazard in teams’, *Review of Economic Studies* **74**(3), 937–963.
- Ruzzier, C. (2009), Asset specificity and vertical integration: Williamson’s hypothesis reconsidered, Working Paper 09-119, Harvard Business School.
- Sabel, C. F. (1993), ‘Studied trust: Building new forms of cooperation in a volatile economy’, *Human Relations* **46**(9), 1133–1170.
- Sako, M. (1992), *Prices, Quality and Trust: Inter-Firm Relations in Britain and Japan*, Cambridge University Press, Cambridge, UK.
- Sannikov, Y. and Skrzypacz, A. (2007), ‘Impossibility of collusion under imperfect monitoring with flexible production’, *American Economic Review* **97**(5), 1794–1823.

- Schmidt, K. M. and Schnitzer, M. (1995), ‘The interaction of explicit and implicit contracts’, *Economics Letters* **48**(2), 193–199.
- Schöttner, A. (2008), ‘Relational contracts, multitasking, and job design’, *Journal of Law, Economics, and Organization* **24**(1), 138–162.
- Shapiro, C. and Stiglitz, J. E. (1984), ‘Equilibrium unemployment as a worker discipline device’, *American Economic Review* **74**(3), 433–444.
- Sobel, J. (2006), ‘For better or forever: Formal versus informal enforcement’, *Journal of Labor Economics* **24**(2), 271–297.
- Telser, L. G. (1980), ‘A theory of self-enforcing agreements’, *Journal of Business* **53**(1), 27–44.
- Thomas, J. P. and Worrall, T. (2010), Dynamic relational contracts with credit constraints. Available at SSRN: <http://ssrn.com/abstract=1599525>.
- Thomas, J. and Worrall, T. (1988), ‘Self-enforcing wage contracts’, *Review of Economic Studies* **55**(4), 541–554.
- Watson, J. (1999), ‘Starting small and renegotiation’, *Journal of Economic Theory* **85**(1), 52–90.
- Watson, J. (2002), ‘Starting small and commitment’, *Games and Economic Behavior* **38**(1), 176–199.
- Williamson, O. E. (2000), ‘The new institutional economics: Taking stock, looking ahead’, *Journal of Economic Literature* **38**(3), 595–613.
- Yang, H. (2009), Nonstationary relational contracts. Ohio State University.