challenged agent and the challenger have to pay Δt, they could pay it to a third agent; finally (3), for the mechanism to work, for each agent i, we need to identify only one other agent beyond i that has observed para-meter θ_i.

Note, however, that the preceding mechanism relies to a great extent on the uncompromising faith in rationality of all players that underlies subgame-perfect equilibrium: If in step (a) of stage 1, agent 1 deviates from truth telling, agent 2 decides to challenge in step (b) *because he or she has faith in the fact that agent 1 will make a payoff-maximizing decision in step* (c). This reasoning is the key behind equilibrium truth telling in step (a). However, under this reasoning, a deviation from truth telling is sure to cost agent 1 a very large amount Δt. It then takes quite some confidence for agent 2 to think that agent 1 will "come back to his or her senses" in step (c) and optimize over a stake that is typically much smaller. This is in fact a general problem with subgame-perfect equilibrium in games where players take actions repeatedly: Deviations are always considered to be "one-shot deviations from rationality" that do not shatter the faith players have in the subsequent rationality of their opponents. However, here, in contrast to standard game theory, we have specifically *designed* a game instead of analyzing one whose rules are derived from economic stylized facts. And we have specifically chosen the cost of the initial deviation to be very large in case a challenge occurs, and as a result, it is really the move in step (a) that is the crucial one for agent 1. Therefore, the preceding criticism of subgame perfection is particularly relevant here.

A second criticism of the preceding mechanism is that, once an agent has been challenged, the continuation equilibrium may involve ex post inefficient outcomes, so that it is subject to the renegotiation critique. In the next section we present models that address this critique. We do it in the context of a specific bilateral investment setting.

12.3 The Holdup Problem

The notion of "hold-up problem" has first been defined and addressed in the seminal article by Goldberg (1976). It has been developed further by Klein, Crawford, and Alchian (1978) and Williamson (1975, 1985). Here we consider the formulation of the contracting problem giving rise to a hold-up problem due to Hart and Moore (1988): Two contracting parties, a prospective buyer and a prospective seller, can enter a relationship in which they can end up trading a quantity $q \in [0, 1]$ at a price P. The utility they obtain from trading depends on the buyer's valuation v and the seller's production cost c. These utilities are uncertain at the time of contracting and can be influenced by specific investments made by each party at an earlier date. Specifically, we make the following assumptions:

$$v \in \{v_L, v_H\}, \quad \text{with } v_L < v_H \quad \text{and} \quad \Pr(v_H) = j$$

where investment j costs the buyer $\psi(j)$, and

$$c \in \{c_L, c_H\}, \quad \text{with } c_L < c_H \quad \text{and} \quad \Pr(c_L) = i$$

where investment i costs the seller $\phi(i)$. Assume that the two investment-cost functions are increasing and convex, and that they are *sunk* whatever the ex post level of trade. The ex post payoff levels are thus

$$vq - P - \psi(j)$$

for the buyer and

$$P - cq - \phi(i)$$

for the seller. The timing is as follows: First, the parties contract; second, they simultaneously choose their investment levels i and j; third, they *both* learn the state of nature $\theta = (v, c)$; fourth, they execute the contract.

What is the first-best outcome? Assume for simplicity that

$$c_H > v_H > c_L > v_L$$

Under this assumption, the *ex post efficient* level of trade is $q = 1$ if $\theta = (v_H, c_L)$ and 0 otherwise. As for ex ante efficiency, since the parties are assumed to be risk neutral, it is equivalent to investment efficiency; that is, i and j must result from

$$\max_{i,j} \{ij(v_H - c_L) - \psi(j) - \phi(i)\}$$

Assuming an interior solution, the first-order conditions give us the optimal investment levels i^* and j^*:

$$i^*(v_H - c_L) = \psi'(j^*)$$

and

$$j^*(v_H - c_L) = \phi'(i^*)$$

The contracting problem the literature has analyzed is one where the state of nature $\theta = (v, c)$ and the investment levels i and j are not contractable, although θ is observable to both contracting parties ex post. If there is spot contracting ex post, after θ is realized and investments i and j are sunk, and if the gains from trade at that point are evenly divided between buyer and seller, there will be underinvestment in equilibrium as we have already noted and as Figure 12.1 below illustrates. The solid curves represent the optimal investment functions i^* and j^* while the dashed curves represent the best response functions under spot contracting:

$$\frac{1}{2}i(v_H - c_L) = \psi'(j^{br})$$

and

$$\frac{1}{2}j(v_H - c_L) = \phi'(i^{br})$$

The difficulty faced by the contracting parties ex ante is how to formulate an optimal long-term contract that is independent of θ, which mitigates

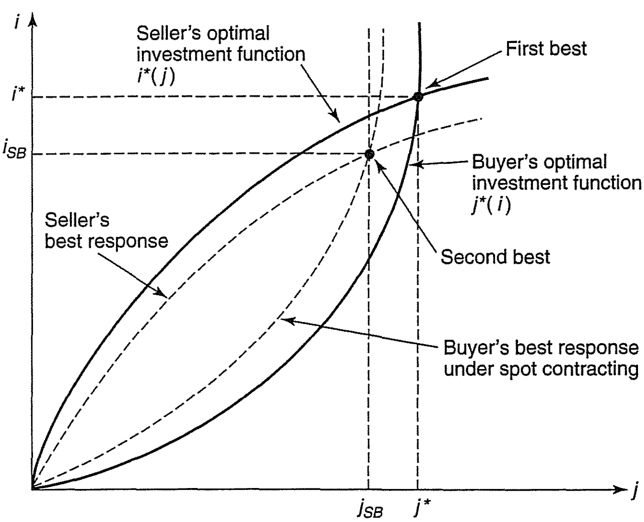

Figure 12.1
Underinvestment with a Holdup Problem

this underinvestment problem. As explained in the next subsections, the contributions in the literature differ in the following respects:

- First, they make different assumptions on the extent to which the level of trade is contractable: Chung (1991), Aghion, Dewatripont, and Rey (1994, hereafter ADR), and Noldeke and Schmidt (1995) allow for "specific-performance contracts," where the contract can specify a given level of trade that parties can request ex post, whether it is efficient or not. Instead, Hart and Moore (1988) focus on "contracting at will," where courts only enforce price schedules contingent on the level of trade, without being able to identify who was responsible for the possible failure to trade. This issue will prove crucial for the ability of the contract to achieve first-best outcomes or not.

- Second, while all contributions assume that the parties cannot commit not to undertake ex post Pareto-improving renegotiations, they differ in their assumptions about the ability to contractually influence the renegotiation process. ADR go furthest in contractual "renegotiation design," by assuming that relative bargaining powers in renegotiation can be contractually chosen. While they focus on specific exogenous bargaining games, Chung, Noldeke-Schmidt, and Hart-Moore end up with the same (one-sided) distribution of bargaining powers as ADR. Instead, Edlin and Reichelstein (1996) assume simple Nash bargaining powers in the renegotiation process. It turns out, however, that assumptions about bargaining powers are not as important as the distinction between specific performance and contracting at will.

- Finally, whereas all the preceding contributions rule out direct externalities, that is, any direct effect of the buyer's investment on the seller's cost or of the seller's investment on the buyer's valuation, Che and Hausch (1999) allow for such externalities. They show that not only is the first best not generally reachable anymore, but *the null contract* may be the optimal contract. And Segal (1999a) obtains a similar result without such direct externalities but for environments that are very "complex."

12.3.1 Specific Performance Contracts and Renegotiation Design

Let us first assume that the contract can specify "default options" that parties can request whenever trade is possible. In this case, one can define the level of trade $\tilde{q}$ such that

$$\tilde{q}(c_H - c_L) = \phi'(i^*)$$

and one can consider the following contractual mechanism: Once θ has been realized, the parties play the following game: in stage 1, the buyer can make an offer (P, q) to the seller; in stage 2, the seller accepts the offer (and trade takes place at these terms), or rejects it, in which case $\tilde{q}$ is traded, at a prespecified price $\tilde{P}$ designed to share the ex ante surplus according to initial bargaining strengths.

This mechanism implements the first best. Indeed, note first that the buyer has full bargaining power in the two-stage game. She will thus offer to trade the ex post efficient quantity while leaving the seller indifferent between this trade and his default-option payoff. While ex post efficiency is guaranteed, what about investment efficiency? The seller will anticipate obtaining his default option payoff whatever the ex post level of trade, so that he will solve

$$\max\{\tilde{P} - ic_L\tilde{q} - (1-i)c_H\tilde{q} - \phi(i)\}$$

By the construction of $\tilde{q}$, investment level i^* is the seller's optimal choice, whatever the buyer's investment may be. Finally, since the buyer has full bargaining power, she is residual claimant on her investment and solves

$$\max\{i^* j(v_H - c_L) - [\tilde{P} - i^* c_L\tilde{q} - (1-i^*)c_H\tilde{q}] - \psi(j)\}$$

She thus maximizes total surplus minus the payoff of the seller (which does not depend on her investment) and minus her cost of investment. Consequently, she chooses $j = j^*$ if the seller chooses $i = i^*$.

The preceding mechanism thus induces efficient *bilateral* investment and circumvents the "moral-hazard-in-teams" problem à la Holmström (1982b), that we discussed in Chapter 8. For the buyer, efficient investment is achieved simply by making her a residual claimant. More intriguing is the case of the seller, who has the appropriate incentive to invest *despite having no bargaining power at all*. His incentive to invest comes from his being able to request the default option, whose attractiveness rises when his production cost goes down. This option makes the seller's payoff sensitive to his investment, and, through this second instrument, both parties can have proper incentives to invest.

The preceding mechanism is in the spirit of subgame-perfect implementation, as it is a multistage mechanism with a unique subgame-perfect equilibrium (where stage 1 could be reinterpreted as having the buyer "announce θ"). Moreover, in this game each party acts only once, so that, in comparison with the Moore-Repullo example, the reliance on backward induction is less objectionable. Still, the mechanism relies on the ability of the parties to commit to ex post inefficient outcomes: If in stage 1 the buyer has made a "crazy offer," the seller may face two ex post inefficient possibilities and no way out of this suboptimal choice.

ADR, however, provide a reinterpretation of this mechanism that involves a much-weakened ability of the parties to commit not to engage in Pareto-improving renegotiations: Assume that, in the absence of a contract, the parties bargain—starting at a date t after θ has been observed—about the terms of trade in an alternating-offer bargaining game à la Rubinstein (1982). As is well known, in state $\theta = (v_H, c_L)$ there is a unique stationary subgame-perfect equilibrium of this game for any pair of discount factors $\delta_B < 1$ for the buyer and $\delta_S < 1$ for the seller. Indeed, when it is her turn to make an offer P^B, the buyer solves

$$\min P^B \text{ such that } P^B - c_L \geq \delta_S(P^S - c_L)$$

Similarly, when it is his turn to make an offer P^S, the seller solves

$$\max P^S \text{ such that } v_H - P^S \geq \delta_B(v_H - P^B)$$

Since at the optimum these two inequalities are binding, uniqueness follows (with trade taking place immediately at time t). If the seller makes the first offer, the price is

$$P^S = \frac{1-\delta_B}{1-\delta_B\delta_S} v_H + \delta_B \frac{1-\delta_S}{1-\delta_B\delta_S} c_L$$

If instead the buyer can make the first offer, the price is

$$P^B = \delta_S \frac{1-\delta_B}{1-\delta_B\delta_S} v_H + \frac{1-\delta_S}{1-\delta_B\delta_S} c_L$$

Since discount factors are less than 1 and $v_H > c_L$, the price offered by the buyer is lower than the price offered by the seller. If, however, both discount factors are equal and tend to 1, then both prices tend to $(v_H + c_L)/2$, so that gains from trade are shared equally. But if one party becomes very patient relative to the other ($\delta_k \to 1$ while δ_l remains bounded away from 1), he or she obtains the entire surplus from trade.

How can the contract influence the bargaining process? First assume that the contract can specify a default option $(\tilde{P}, \tilde{q})$ that each party can enforce when it is his or her turn to react to the offer by the other party: Beyond accepting or waiting one period to be able to make an offer, it can request the default option. This action turns the bargaining game into an alternating-offer game with an "outside" option, which has been studied by Binmore, Rubinstein, and Wolinsky (1986), who have defined the so-called *outside-option principle:* Call $(\tilde{U}_B, \tilde{U}_S)$ the parties' payoffs associated with the outside option. Call (U_B^*, U_S^*) the parties' (ex post efficient) payoffs in the bargaining game *without* the outside option. Then, if (U_B^*, U_S^*) Pareto-dominates $(\tilde{U}_B, \tilde{U}_S)$, the equilibrium payoffs of the game with the outside option are (U_B^*, U_S^*). If (U_B^*, U_S^*) does not Pareto-dominate $(\tilde{U}_B, \tilde{U}_S)$, the outcome of the game with outside option is ex post efficient and gives party k for whom $\tilde{U}_k > U_k^*$ his or her outside option payoff.

In order to limit the seller to his default-option payoff, one has to make sure this payoff is better for him than the outcome of bargaining without the default option. This outcome can be ensured by introducing a *penalty for delayed trade* that the seller would have to pay the buyer. If this penalty is big enough, the seller will immediately accept any offer (P, q) that is better for him than $(\tilde{P}, \tilde{q})$, in order to avoid delay.

The first-best outcome can thus be implemented. Moreover, it can be achieved in a "light" way, that is, through a relatively simple contract, consisting in a default option and a penalty for delayed trade. Simplicity is in fact achieved because the contract allows for *equilibrium renegotiation* and only *supplements* the underlying bargaining structure the agents have at their disposal in the absence of contracting.

One important assumption, however, is that renegotiation stops whenever the default option has been chosen by one party. This is reasonable if there is, for example, a fixed cost to be incurred whenever the seller produces, so that multiple trades over time would be excessively costly. In such a case, "requesting the default option" simply means that the seller unilaterally decides to produce. In some environments, however, the technology may allow the parties, after the default level of trade has occurred, to keep bargaining if the ex post efficient trade is higher than the default trade. This possibility would undermine the preceding results, as section 12.3.2.2 will show.

12.3.2 Option Contracts and Contracting at Will

The original holdup model of Hart and Moore (1988) did not introduce renegotiation design, but instead considered the following bargaining game (which we present in simplified form, following Noldeke and Schmidt, 1995): After θ has been realized, the parties can simultaneously send one another new written trading offers. Assume now that trade can take only the values 0 and 1, so that an offer consists in at most two prices, since it can concern only "trade" and "no trade." Assume the initial contract was a pair (P_0, P_1) (for no trade and trade, respectively) and call the new offers by the buyer and seller (P_0^B, P_1^B) and (P_0^S, P_1^S). Once the stage where the parties exchange new offers is over, trade does or does not take place (in a way that we will specify) and a payment is made, possibly after the intervention of courts. These are assumed to be able to observe whether trade has taken place or not and to enforce the corresponding payment. This payment is the original one, that is, P_k if a quantity k was exchanged $(k = 0, 1)$, unless a party finds it in his or her interest to show the court a new written offer made by the other party. These assumptions guarantee that the initial contract protects each party against unilateral violations while allowing for Pareto-improving renegotiations.

In equilibrium, the only offers ever shown to the court are those sent by the buyer to the seller saying she accepts a *higher* price and those sent by the seller to the buyer saying he accepts a *lower* price. Why would such offers ever be written? Because they may be the only way to ensure that the other party accepts the ex post efficient trade.

The outcome of this game depends on the ability to enforce default options, as we now show.

12.3.2.1 Option Contracts

Noldeke and Schmidt (1995) allow for specific performance contracts, as in ADR. They consider option contracts, where the seller receives a price P_0 if the good is not delivered, and has the option to deliver the good and receive an additional payment K (so that $P_1 = P_0 + K$).

How does renegotiation proceed in this setting? Three cases have to be distinguished:

- First, consider the case $K < c_L$. Barring renegotiation, the seller never has an incentive to deliver the good in this case. This is ex post efficient whenever the buyer's valuation is low or the seller's cost is high (or both), in which case no trade takes place and the equilibrium payment made by the seller is P_0. But what happens in state $\theta = (v_H, c_L)$, where trade is efficient? Achieving trade requires raising the premium the seller receives for delivery at least up to c_L. In fact, the buyer can make sure not to have to raise it further: By sending a letter offering $P_1^B = P_0 + c_L$, the buyer induces the seller to deliver the good and knows the seller has the incentive to show this letter to the court. Indeed, otherwise the enforced price would be the lower P_1. Moreover, any letter sent by the seller requiring a higher price would simply not be shown to the court by the buyer. Consequently, just as in ADR, one party has full bargaining power: In this game, the party who behaves in equilibrium in a way that is suboptimal *at the initial contract prices* has no bargaining power at all. And since here it is only the seller that makes the trade decision (unlike in the next subsection), it is always the buyer who has full bargaining power.

• Second, consider the case $c_L < K < c_H$. Barring renegotiation, ex post inefficiency occurs when both the seller's cost and the buyer's valuation are low, since the seller would find it profitable to deliver the good although $v_L < c_L$. Once again, the buyer can extract the full surplus from renegotiation, by sending a letter agreeing to a higher price for no trade, that is, $P_0^B = P_1 - c_L$.

• Finally, when $c_H < K$, the seller always wants to deliver the good, and this case is inefficient when the buyer's valuation is low or the seller's cost is high (or both). As before, the buyer can extract the full surplus from renegotiation by sending a letter agreeing to a higher price for no trade, that is, $P_0^B = P_1 - c$, where c is the realized cost of the seller.

This mechanism is similar to the one presented in the previous subsection: The buyer has full bargaining power, while the seller receives his initial contract payoff, which depends on the value of his cost even in cases where trade may not take place ex post. Let us focus on the case where $c_L < K < c_H$, where the seller obtains P_0 when his cost is high, and $P_1 - c_L = P_0 + K - c_L$ when his cost is low. He thus chooses his investment to solve

$$\max \{i(K - c_L) - \phi(i)\}$$

To ensure an appropriate investment choice, K has to be chosen so that

$$K - c_L = \phi'(i^*)$$

This implies $K < c_H$, since from the definition of the first-best outcome, we have

$$\phi'(i^*) = j^*(v_H - c_L) < v_H - c_L < c_H - c_L$$

Given the choice of K, it is optimal for the seller to make the first-best investment choice, whatever the investment decision of the buyer. As for the buyer, as in the previous subsection, she is residual claimant with respect to her investment choice—having full bargaining power in the renegotiation—so she chooses $j = j^*$ when the seller chooses $i = i^*$. The first-best outcome is implemented again, this time with a simple option contract $(P_0, P_1 = P_0 + K)$.

12.3.2.2 Contracting at Will

Assume now, as in the original Hart-Moore (1988) model, that specific performance contracts cannot be enforced by courts. Think, for example, that, were the seller to deliver the good, the buyer could always claim it is not of "appropriate" quality. If quality is unverifiable, the court can only observe whether trade took place and enforce quantity-contingent price schedules, but cannot distinguish who is responsible for the lack of trade. In this setup, trade takes place only if *both* parties want it to happen. Hart and Moore use the following model: After θ has been realized, the parties can exchange written messages with new price offers. Then both parties simultaneously decide whether they want to go ahead with trade. Only if both agree does trade take place, followed by the same associated payments as before.

Consider a simple contract (P_0, P_1). Under these terms, trade takes place if and only if $v \geq P_1 - P_0 \geq c$. If this condition is not met, at least one party prefers not to trade, and no trade is the outcome. This is ex post efficient unless $\theta = (v_H, c_L)$, in which case we have the following two possibilities:

- If $P_1 - P_0 > v_H > c_L$, although trade is efficient, the buyer finds it too expensive. In this case, it is the seller who has full bargaining power, since he can ensure trade by sending the buyer a letter agreeing to $P_1^S = P_0 + v_H$. With this offer, the buyer is ready to trade and is unable to induce the seller to agree to any lower price.
- If $v_H > c_L > P_1 - P_0$, it is now the seller who finds trade too costly for the price difference. The buyer now has full bargaining power, being able to induce trade by agreeing to $P_1^B = P_0 + c_L$. As before, the full bargaining power goes to the party who would benefit from trading at the initial contracting prices.

The preceding reasoning indicates that in this example there is no loss of generality in choosing P_0 and P_1 such that $v_H \geq P_1 - P_0 \geq c_L$, since price differences that do not satisfy these inequalities are renegotiated so as to (just) satisfy them. How much investment does such a renegotiation-proof contract generate? The key idea with at-will contracting is that, in case of disagreement, the starting point of renegotiation is always no trade, so that the parties benefit from their own investment only when it is efficient to trade. The payoffs are therefore

$$ij[v_H - (P_1 - P_0)] - P_0 - \psi(j)$$

for the buyer, and

$$P_0 + ij[(P_1 - P_0) - c_L] - \phi(i)$$

for the seller. Consequently, the classical moral-hazard-in-teams problem arises, with underinvestment as the outcome. Hart and Moore thus provide foundations to the arguments of Goldberg, Klein, Crawford, and Alchian, and Williamson that underinvestment is likely when long-term contracts are incomplete.

Note finally that the preceding result assumes that the parties sign a simple contract (P_0, P_1). Could they improve upon it using message games? Hart and Moore show that they cannot, so that, just as in subsections 12.3.1 and 12.3.2.1, the optimum is achieved through a simple contract that relies on equilibrium renegotiation.

12.3.3 Direct Externalities

Let us now introduce direct investment externalities, as done by Che and Hausch (1999). For simplicity, assume that only the seller can profitably invest in the relationship. Assume, moreover, that the seller's investment has an impact not only on his production cost but also on the quality of the product. Specifically, the seller's investment influences not only his cost in that $\Pr(c_L) = \beta i$ but also the buyer's valuation in that $\Pr(v_H) = \gamma i$. The first best is then the result of

$$\max\{\beta\gamma i^2(v_H - c_L) - \phi(i)\}$$

Assuming an interior solution, the first-order condition implies

$$2\beta\gamma i^*(v_H - c_L) = \phi'(i^*)$$

As shown by Che and Hausch, direct externalities can dramatically affect the efficiency properties of contracts. To illustrate, assume the parties can initially sign a simple specific-performance contract $(\tilde{P}, \tilde{q})$. Following Edlin and Reichelstein (1996), assume also that the renegotiation that follows the investment choice and the realization of θ can be represented by generalized Nash bargaining, leading to ex post efficiency with a share

α of the surplus from renegotiation going to the seller and a share $(1 - \alpha)$ going to the buyer. The seller's expected payoff from the contract $(\tilde{P}, \tilde{q})$ is his default payoff plus a share α of the surplus from renegotiation, that is,

$$E_{(v,c)|i}\{\tilde{P} - c\tilde{q} + \alpha[(v-c)q^* - (v-c)\tilde{q}]\} - \phi(i)$$

where q^* is the first-best level of trade [that is, 0 unless $\theta = (v_H, c_L)$]. This expression is therefore equal to

$$\tilde{P} + \alpha\beta\gamma i^2(v_H - c_L) - [(1-\alpha)E_{(v,c)|i}(c\tilde{q}) + \alpha E_{(v,c)|i}(v\tilde{q})] - \phi(i)$$

or, equivalently,

$$\tilde{P} + \alpha\beta\gamma i^2(v_H - c_L) - \tilde{q}\,\{(1-\alpha)[c_H - \beta i(c_H - c_L)] + \alpha[v_L + \gamma i(v_H - v_L)]\} - \phi(i)$$

This latter expression identifies three effects from an increase in investment i on the seller's payoff:

- The first effect $[\alpha\beta\gamma i^2(v_H - c_L)]$ refers to the fact that the seller captures a share α from the surplus generated by investing. For $\alpha < 1$, this in itself is insufficient to avoid underinvestment.
- The second effect $[\tilde{q}(1 - \alpha)\beta i(c_H - c_L)]$ arises because, by investing, the seller improves his own default payoff. Just as in the earlier cases, this provides additional incentives to invest and the more so the higher the default option $\tilde{q}$.
- The third, countervailing, effect $[-\tilde{q}\alpha\gamma i(v_H - v_L)]$, is due to the fact that, by investing, the seller improves the default option *of the buyer*. This externality results in a *disincentive* to invest. Once again, the effect is stronger the higher the default option $\tilde{q}$.

The last two effects are linear in $\tilde{q}$, which does not appear in the first effect. Raising $\tilde{q}$ thus raises incentives to invest if and only if

$$\alpha\gamma(v_H - v_L) \leq (1-\alpha)\beta(c_H - c_L)$$

If this condition is not satisfied—which happens if α or γ is large or β is small, for example—then setting $\tilde{q} = 0$ is optimal: The null contract is the optimal initial contract. One case that is very intuitive is $\beta = 0$: When the seller's investment only improves the valuation of the buyer and not the seller's cost, positive default options only improve the buyer's bargaining position when the seller invests more. Consequently, they act as a disincentive to invest and are counterproductive in comparison with the null contract.[2]

Che and Hausch, moreover, have shown that no message-contingent contract can improve on simple contracts of the form $(\tilde{P}, \tilde{q})$. This is thus a general lesson from the entire holdup literature we have considered here, whether the optimal contract achieves the first-best outcome (as in ADR or Noldeke-Schmidt) or not (as in Hart-Moore or Che-Hausch): By relying on equilibrium renegotiation, it is possible to derive simple optimal contracts.

12.3.4 Complexity

The contracts we have discussed so far in this chapter assume unverifiability of the state of nature but not of trades, a type of unverifiability stressed in the Grossman-Hart-Moore incomplete-contract paradigm. One reason trades may not be contractable, however, is the excessive *complexity* involved in specifying ex ante the nature of transactions: The exact specifications of ex post transactions may not be known yet—for example, if we are talking of new products and if investment concerns R&D. The mechanism-design question is then, What can contracts achieve when actions are contractable ex post but not ex ante, especially in "complex" environments?

Maskin and Tirole (1999a, 1999b) have identified conditions under which ex ante noncontractability of actions is irrelevant, thereby questioning the incomplete-contract methodology à la Grossman-Hart-Moore. Their key observation is that, while Grossman, Hart, and Moore assume noncontractability of actions ex ante, they assume that the *payoff consequences* of the various actions that could be taken ex post can be *foreseen*. This foresight is indeed necessary in order to be able to make rational investment choices before uncertainty is realized. But, remarkably, it also allows Maskin and Tirole to construct a mechanism (which builds upon subgame-perfect-implementation results) that does not require specifying actions ex ante. Instead, the mechanism specifies transfers and the right to make offers ex post contingent on announcements about the state of nature (and its payoff consequences).

Maskin and Tirole make an important methodological contribution. Their approach, however, is rather abstract, and we shall only discuss it in this

2. Bernheim and Whinston (1998b) expand on this point in a more general setting.

chapter in the context of a specific holdup model developed by Segal (1999a). Segal's main goal is in fact to define a notion of complexity of the trading environment and to relate it to the effectiveness of contracting. In his setting, when complexity grows without bounds, contracting loses its power and we are left with the null contract, just as in Che and Hausch (1999). Segal thus provides a foundation for contractual "incompleteness" connected to the difficulty of specifying in advance the realization of uncertainty. In this section, we develop Segal's insight while relying on the formulation, results, and proofs contained in Hart and Moore (1999), who simplify the analysis and illustrate and qualify the results of Maskin and Tirole (1999a).

Assume a contracting problem between two risk-neutral agents, a buyer and a seller. Only the seller can invest to raise the surplus from trade. The problem has three stages: in stage 1 contracting takes place, in stage 2 the seller invests, and in stage 3 the state of nature is observed by both parties and trade takes place. Assume it is always efficient ex post to trade one unit of a good, or "widget," but there is uncertainty ex ante about which "type" of widget should be traded. There are initially N types of widgets. Types are contractible ex post. One type only should be traded. Call it the "special" widget, generating constant valuation v for the buyer and random cost c for the seller. Ex ante, $c \in \{c_L, c_H\}$, with $c_L < c_H < v$ and $\Pr(c_L) = i$, where $\phi(i)$ is, as before, the seller's investment cost. There are thus no direct externalities. The other widgets are called "generic," with production cost for the seller equal to[3]

$$c_n^g = c_L + \frac{n}{N}(c_H - c_L), \quad \text{for } n = 1, \dots, N-1$$

The problem ex post is that of recognizing which is the special widget among the N possible widgets. The parameter N is thus a measure of complexity. As will become clear, the key fact will be that, as N becomes large, c_n^g "fills" the interval $[c_L, c_H]$.

Assume complete symmetry ex ante across widgets, both in terms of the probability of being the special one and in terms of production cost when generic. A state of nature is a cost realization of the special widget and a permutation of the $N - 1$ generic widgets plus the special one. Each state

3. The seller's investment has no impact on the cost of generic widgets. This assumption is made only for simplicity.

of nature is thus equally likely. Finally, assume that c and the c_n^g's are observable but unverifiable at stage 3.

The first best involves trading the special widget in all states of nature and setting investment so as to solve

$$\max_i \{i(v - c_L) + (1 - i)(v - c_H) - \phi(i)\}$$

What happens under noncontractibility of i and unverifiability of c and the c_n^g's? Intuitively, the answer to this question may depend on whether the parties can or cannot commit not to renegotiate, and on whether widget types can or cannot be described ex ante.

12.3.4.1 No Renegotiation

Even if the widgets cannot be described at stage 1 (but remember they can be at stage 3), the first best can be achieved provided the parties can commit not to renegotiate: Just set up a mechanism in which the seller can make the buyer a take-it-or-leave-it offer at stage 3. Since he is endowed with full bargaining power, the seller will choose the first-best investment level.

This result is consistent with Maskin and Tirole's claim that ex ante non-describability may not matter: Clearly, since the first best can be achieved without specifying ex ante any explicit message-contingent outcome,[4] being able to perform such ex ante description adds nothing.

While the no-renegotiation benchmark is exceedingly simple, things become much more involved when the parties cannot commit not to engage in subsequent Pareto-improving renegotiations.

12.3.4.2 Renegotiation

Assume for simplicity that the buyer has full bargaining power in renegotiation. Call P_i the expected price obtained by the seller if his cost is c_i. The seller chooses his investment i to solve

$$\max_i \{i(P_L - c_L) + (1 - i)(P_H - c_H) - \phi(i)\}$$

First-best investment obtains if $P_L = P_H$, but investment decreases when $P_H - P_L$ increases. In particular, in the absence of a prior contract, the buyer

4. With one exception: "No trade" has to be contractable ex ante, since the buyer must be allowed to decline the seller offer. Maskin and Tirole derive their results under the assumption that there is at least one level of trade (e.g., no trade) that is contractable ex ante.

offers the seller $P_i = c_i$ ex post, which leads the seller to choose $i = 0$.

The striking result from this model is that, even if the widgets can be described at stage 1, because of the parties' inability to commit not to engage in Pareto-improving renegotiations, the optimal contract implies $i \to 0$ when $N \to \infty$. The gains from contracting thus vanish when com-plexity grows without bounds and the parties might as well not bother to sign any initial contract. Note that this result generalizes to any distribution of bargaining powers in the renegotiation process, as well as to bilateral investments.

Let us now establish this result. The underlying logic is to observe how incentive-compatibility requirements translate into post-renegotiation prices that depend almost one for one on the cost realization of the seller, just as under the null contract.

Specifically, take any mechanism M. Define a state of nature (L, τ) as one where the special widget costs c_L and where the $N - 1$ generic widgets are arranged according to a permutation τ. Without loss of generality, take the special widget to be widget 1, and have widgets $2, 3, \ldots, N$ cost

$$c_L + \frac{1}{N}(c_H - c_L), \quad c_L + \frac{2}{N}(c_H - c_L), \ldots, \quad c_L + \frac{N-1}{N}(c_H - c_L)$$

respectively. Denote the equilibrium strategies of the two parties when playing mechanism M in state of nature (L, τ) as $m_B(L, \tau)$ and $m_S(L, \tau)$ (that is, their announcements, truthful or not, of the state of nature that is mutually revealed to them). Define the price $P(L, \tau)$ as the equilibrium price at which the special widget is traded in state (L, τ), possibly after renegotiation. This yields surpluses $v - P(L, \tau)$ for the buyer and $P(L, \tau) - c_L$ for the seller.

Consider now state of nature (H, τ'), with cost c_H for the special widget and a new permutation of the widgets: The special widget becomes widget N, and widgets $1, 2, 3, \ldots, N - 1$ now cost

$$c_L + \frac{1}{N}(c_H - c_L), \quad c_L + \frac{2}{N}(c_H - c_L), \ldots, \quad c_L + \frac{N-1}{N}(c_H - c_L)$$

Denote the equilibrium strategies of the two parties when playing mechanism M in state of nature (H, τ') as $m_B(H, \tau')$ and $m_S(H, \tau')$, and define the price $P(H, \tau')$ as the equilibrium price at which the special widget is traded in state (H, τ'), possibly after renegotiation.

What does incentive compatibility of truthful reporting of the state of nature imply for prices $P(L, \tau)$ and $P(H, \tau')$? Let us focus on the following two incentive constraints:

1. In state of nature (H, τ'), the seller should not play as if the state of nature were (L, τ).
2. In state of nature (L, τ), the buyer should not play as if the state of nature were (H, τ').

Both unilateral deviations amount to the players choosing a pair of strategies $[m_B(H, \tau'), m_S(L, \tau)]$. Without loss of generality, assume that mechanism M prescribes, upon such a pair of announcements, a starting point of renegotiation where the buyer has to pay the seller an amount $\tilde{P}$ and where widget n is traded with probability x_n, while no trade happens with probability

$$1-\sum_{n=1}^{N} x_n \geq 0$$

Since we assumed full bargaining power for the buyer in renegotiation, the incentive constraint in state of nature (H, τ') to avoid deviation 1 by the seller is

$$\tilde{P}-\sum_{n=1}^{N} x_n\left[c_L+\frac{n}{N}(c_H-c_L)\right] \leq P(H, \tau')-c_H$$

Similarly, in order to avoid deviation 2 by the buyer in state of nature (L, τ), and keeping in mind that the ex post efficient trade will result anyway, what we need is for the seller not to lose from this deviation, or

$$\tilde{P}-\sum_{n=1}^{N} x_n\left[c_L+\frac{n-1}{N}(c_H-c_L)\right] \geq P(L, \tau)-c_L$$

These two inequalities imply

$$P(H, \tau')-P(L, \tau) \geq c_H-c_L-\sum_{n=1}^{N} \frac{x_n}{N}(c_H-c_L) \geq \frac{N-1}{N}(c_H-c_L)$$

This condition has to be satisfied for any pair (τ, τ'), and therefore, since all permutations are equally likely, the expected price the seller receives when his cost is c_H, minus the expected price he receives when his cost is c_L, is at least

$$\frac{N-1}{N}(c_H-c_L)$$

which tends to $c_H - c_L$ when $N \to \infty$. Consequently, the seller obtains none of the gains from investing, just as without any prior contract. We have thus established the result.

The lessons of this result are twofold. First, as in Maskin and Tirole (1999), we have a case here where ex ante describability of widgets once again does not matter: dropping it cannot hurt, since in the limit there is already no value of contracting even with ex ante describability. Second, what matters here is not whether widgets are describable ex ante but whether the parties can commit not to engage in ex post Pareto-improving renegotiations.

Let us end with an important point concerning the interpretation of this result. The proof has focused on incentive compatibility for states of nature (L, τ) and (H, τ'). These states differ in the cost level of the special widget: widget 1, the special one in state (L, τ), costs c_L, while widget N, the special one in state (H, τ'), costs c_H. More importantly, the two states of nature are "extreme" in the cost differences of the these two widgets when they are not special: widget N is the *most expensive* generic widget in state (L, τ), while widget 1 is the *cheapest* generic widget in state (H, τ'). As a result, it is particularly difficult to prevent the seller from claiming that the state of nature is (L, τ) when it is in fact (H, τ') (so that widget 1 rather than N should be produced) and conversely for the buyer. In this model the cost difference between the most expensive and cheapest generic widgets increases with N, the "complexity" of the environment. Reiche (2003a) makes the point, however, that the value of contracting in this kind of setting can also go to zero if the seller's investment is "ambiguous," that is, has a value that can be negative if the wrong ex post action is taken (because of technological complementarity between the investment and the specific widget to be produced). Ambiguity works, just like complexity, because it generates the same negative correlation between the cost of the ex post efficient widget and the cost of the "associated" ex post inefficient widget as in states (H, τ) and (H, τ').

12.3.4.3 Describability

Although describability did not matter in the preceding analysis, Hart and Moore also provide an example where it does. This is a variation on the previous analysis that has *no* uncertainty. Assume then, without loss of generality, that widget 1 is always the special widget, while widgets 2 to N are the generic ones, costing, respectively,

$$c_L + \frac{1}{N}(c_H - c_L), c_L + \frac{2}{N}(c_H - c_L), \ldots, \quad c_L + \frac{N-1}{N}(c_H - c_L)$$

Here, with describability, it is easy to achieve the first-best outcome: write a specific-performance contract where the parties agree to trade at stage 3 one unit of widget 1 at a fixed price, for example, v.

Instead, without describability at stage 1, we are in the same setup as in the previous subsection, with no benefit of contracting when complexity grows without bound, that is, when N becomes large. Indeed, assume there are N "names" at stage 1, each of which is equally likely to describe the special widget at stage 3, and these are the only ways to describe widgets at stage 1. So, as Hart and Moore (1999, p. 125) write, "Even though the buyer and the seller know at stage 1 which widget is the special one, they have no words to describe it, other than the N names, any one of which may turn out to be appropriate at stage 3." This problem is thus the same as the one considered in the previous subsection where, without commitment against Pareto-improving renegotiation, there is asymptotically no value of contracting when complexity grows.

In contrast to the Maskin-Tirole result, describability matters here because of the combined effect of *renegotiation and risk neutrality:* When renegotiation can be assumed away, first-best implementation can be obtained despite undescribability, as in section 12.3.4.1. When the parties are risk averse, renegotiation-proofness can to some extent be circumvented by "creating risk," and thus ex post inefficiency, following some messages sent by the parties. In such a case, Maskin and Tirole manage to construct mechanisms that still achieve first-best implementation despite undescribability. As the preceding example shows, this result is not possible anymore under risk neutrality.

12.4 Ex Post Unverifiable Actions

The previous section focused first on ex ante and ex post contractable actions, and then, in the last subsection, on ex ante noncontractable but ex post contractable actions. We now focus on ex ante *and* ex post noncontractable actions. This type of actions has in fact already been considered earlier: Think of the *effort choice of moral-hazard models,* a choice that is