American Economic Review 2009, 99:1, 265–294
http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.1.265

Cognition and Incomplete Contracts

By Jean Tirole*

Thinking about contingencies, designing covenants, and seeing through their implications is costly. Parties to a contract accordingly use heuristics and leave it incomplete. The paper develops a model of limited cognition and examines its consequences for contractual design. (*JEL* D23, D82, D86, L22)

In mainstream contract theory, thinking about contingencies, designing covenants, and seeing through their implications are costless. Contracts written by parties, accordingly, are efficient subject to explicit informational and participation constraints (complete contract theory) or within some given contracting set (incomplete contract theory). Contracts are never too detailed or too long.[1] Stylized facts, such as the benefits of economizing on contract completeness under relationship interactions (Stewart Macaulay 1963), or vertical integration (Oliver E. Williamson 1985), or the higher cost of negotiating long-term contracts (Timo Välilä 2005), live in a theoretical vacuum.

By contrast, the less formal bounded rationality approach (Herbert A. Simon 1955; Williamson 1975, 1985) recognizes the cost of gathering and processing information and emphasizes the use of heuristics in contract design. This paper aims to narrow the gap between these two strands. It follows the bounded rationality approach by accounting for cognitive limitations. But like mainstream theory, it takes a rational choice approach to contracting: parties are unaware, but aware that they are unaware.

Its general thrust goes as follows. The parties to a contract (buyer, seller) initially avail themselves of an available design, perhaps an industry standard. This design or contract is the best contract under existing knowledge. The parties are unaware, however, of the contract's implications, but they realize that something may go wrong with this contract; indeed, they may exert cognitive effort in order to find out about what may go wrong and how to draft the contract accordingly: put differently, a contingency is foreseeable (perhaps at a prohibitively high cost), but not necessarily foreseen. To take a trivial example, the possibility that the price of oil increases, implying that the contract should be indexed on it, is perfectly foreseeable, but this does not imply that parties will think about this possibility and index the contract price accordingly.

* Toulouse School of Economics, Manufacture des Tabacs, Aile Jean-Jacques Laffont, 21 allées de Brienne, 31000 Toulouse, France (e-mail: tirole@cict.fr). This paper was previously titled "Bounded Rationality and Incomplete Contracts." I am grateful to Patrick Bolton, Jacques Crémer, Mathias Dewatripont, Antoine Faure-Grimaud, Oliver Hart, Bentley MacLeod, David Martimort, Eric Maskin, Patrick Rey, Xiaojian Zhao, participants at the Melbourne Information Economics Revolution conference, the Contracts and Rationality London Business School conference, the 2007 Bavarian Graduate Program in Economics in Nuremberg, the Harvard-MIT Organizational Economics Seminar, a seminar at the Max Planck Institute for Research on Collective Goods, the editor, three referees, and especially Florian Ederer and Johannes Spinnewijn for helpful discussions and comments.

[1] In certain circumstances, optimal complete contracts take the form of "simple contracts" (e.g., Yeon-Koo Che and Donald B. Hausch 1999; Oliver Hart and John Moore 1999; Eric S. Maskin and Tirole 1999a; Georg Nöldeke and Klaus M. Schmidt 1995; Ilya R. Segal 1999). Similarly, short-term contracts may duplicate the outcome of optimal long-term contracts (e.g., Philippe Aghion and Patrick Bolton 1987; Douglas Diamond 1993; Benjamin E. Hermalin 2002; Patrick Rey and Bernard Salanié 1990, 1996).

This approach delivers two immediate implications that contrast with traditional contract theory. First, there are transaction costs of negotiating deals. Second, complete contracts may be wasteful contracts. Individual interests lead parties to fine-tune the contract whenever contract incompleteness could put them in a situation of being held up ex post. Completing contracts thus involves rent-seeking.

In this paper an incomplete contract is a contract specifying the available design, which is renegotiated whenever this design turns out not to be appropriate. To derive predictions as to when contracts are likely to follow the available design and later be renegotiated, I develop a specific illustration of this broader theme. A buyer and a seller can contract for delivery of a known good or design A. This specification may or may not be what will suit the buyer's need. In the latter case, a different, initially unknown specification A' delivers more surplus to the buyer provided that the seller collaborates. At the initial stage, though, the parties are aware only of A, although they know that it may not be the right design. Each may, before contracting, incur a cognitive cost to think about alternatives to A. A party who finds out that design A' is the appropriate one chooses whether to disclose the existence of (to describe) this design to the other party.

Then, following Kenneth J. Arrow's (1962) insight about the difficulty of licensing trade secrets, and that of Xavier Gabaix and David Laibson (2006) about shrouded attributes, I posit that the enunciation of A' is an "eye-opener." The very description of A' reveals that the state of nature is indeed such that A' rather than A is optimal. Put differently, a party's suggestion of contracting on A' gives the information away and prevents the knowledgeable party from fully benefiting from it, in the same way that Arrow's licensor cannot extract any royalties from the prospective licensee, or Gabaix and Laibson's supplier prefers to keep some contract attributes shrouded rather than mentioning the likely add-on and committing to its price.

Such "awareness-inducing information" is related to, but distinct from, the familiar concept of "hard information." In a standard hard-information model, the parties would be fully aware of the existence of designs A and A', but would be uncertain as to their respective payoffs (which one is appropriate); the contract can then be made contingent on one of the parties bringing hard information that A' is appropriate. No such contract is feasible here, since the description of A' simultaneously reveals the state of nature. My assumption implies that a party cannot charge the other party just for having come up with an alternative design.

This paper defines contract incompleteness in the following way. A contract is more incomplete if fewer resources are expended to identify the appropriate design; equivalently, contract incompleteness is measured by the probability that the design specified in the contract needs to be altered ex post. Transaction costs may be wastefully incurred, as it is in the parties' individual interest to know whether they are vulnerable to (or will benefit from) renegotiation. By contrast, adjustment costs vindicate some cognitive investments from a social point of view. Thus, two effects move the buyer away from optimal search. The first is the desire to avoid being held up ex post and forced to pay extra to adjust the contract. The second is the discount given at the outset by the seller, anticipating hold-up for the standard contract; this discount lessens the buyer's incentives to move away from it.

The paper is an inquiry into what drives equilibrium transaction costs. Its main insights can be summarized in the following way. As already pointed out:

(i) Cognition is a natural source of adverse selection in contractual relationships.

(ii) Contracts may be too complete, that is, there may be too much, not too little, information brought to bear on the design. Rent-seeking, and not only the avoidance of ex post contract adjustments, drives individual incentives for cognition.

I further show that:

(iii) Parties to a contract tend to specialize in identifying potential bad news for themselves/good news for the other party. For example, a seller may not have much incentive to reveal problems with the design; indeed, a central reason for a seller to look for problems is that if she does *not* find any, she might hesitate to go through with the transaction. Accordingly, I derive the conditions under which a party does not search for bad news for the other party.

(iv) Ex ante competition need not reduce transaction costs.

(v) Contracts are predicted to be *strictly* less complete under relational contracting or under vertical integration. Furthermore, long-term contracting may be *strictly* suboptimal. Thus, relational contracting, vertical integration, and short-term contracting generate (are not only responses to) contract incompleteness; this reverse causality has implications for empirical work.

The paper is organized as follows. Section I relates the paper to the literature. Section II develops the framework and analyzes its implications when only one party can engage in cognitive effort (Section IIA on one-sided cognition) and when both parties can (Section IIB on two-sided cognition). Section III analyzes the robustness of the conclusions (Section IIIA allows for multiple sellers and Section IIIB looks at mechanisms that can affect ex post bargaining powers). Section IV shows that contract incompleteness is implied by (and, as noted earlier, is not only a driver of) repeated relationships, the possibility of vertical integration, and short-term contracting. And Section V concludes.

I. Relationship to the Literature

The paper borrows unrestrainedly from, and brings together, several strands of the contract literature. With Kathryn Spier (1992) and Gabaix and Laibson (2006), it shares the view that contract incompleteness is related to asymmetric information. In Spier,[2] an informed party may not want to make the outcome contingent on certain known adverse contingencies for fear of signaling to the other party that these contingencies have high likelihood. Gabaix and Laibson's emphasis on the hold-up associated with add-ons that the buyer is unaware of is closely connected to the approach adopted here. On the other hand, neither of these two papers has cognitive costs, and contracts are too *incomplete*, not too complete, according to my definition.[3]

Ronald A. Dye (1985) modeled contract incompleteness by introducing a fixed cost per contingency included in the contract. This approach was criticized on the grounds that it may be cheap to include a well-formulated contingency in a contract and that it is unclear why writing costs should be proportional to the number of contingencies (Hart and Bengt Holmström 1987). In a sense, this paper follows Dye's impetus and attempts at opening Dye's black box of costs

[2] See also Aghion and Hermalin (1990) and Franklin Allen and Douglas Gale (1992). Luca Anderlini, Leonardo Felli, and Andrew Postlewaite (2006) extend this approach by letting a court rule out certain forms of contracts in order to maximize ex ante welfare.

[3] Contracts are incomplete in David Martimort and Salvatore Piccolo (2008) for a different reason: there, the parties to the contract are engaged in a game with another vertical structure. Committing to an incomplete contract modifies the latter's behavior and may turn out to be beneficial (see also Glenn Ellison 2005). That paper has no cognitive costs either. Another factor of contract renegotiation is their imperfect enforcement by courts (see J. Luis Guasch, Jean-Jacques Laffont, and Stéphane Straub (2006) for a model and empirical evidence).

of writing complete contracts.[4] And to use Benjamin Klein's (2002) terminology, it focuses on "search costs" (the costs of thinking through the contracts' implications), while Dye analyzes "ink costs" (the costs of actually writing contracts).

In Jacques Crémer, Fahad Khalil, and Jean-Charles Rochet (1998a, b), a party engages in rent-seeking information acquisition before contracting. Similarly, the paper shares with Jack Hirshleifer (1971) the insight that contracts may suffer from an excess provision of information. None of these contributions, though, deals with the renegotiation and hold-up associated with contracts built on imperfect cognition, and so the economic focus is accordingly different.

The paper borrows from the classic contributions of Sanford J. Grossman and Hart (1986), Hart and Moore (1990), and Williamson (1985), and from the subsequent literature on investment specificity the idea that contractual choices affect the extent of ex post hold-ups. The paper considers precontractual investments, while the literature focuses on postcontractual ones. Many contributions to incomplete contract theory posit foreseen, observable, but unverifiable contingencies or actions. In line with the general approach of Maskin (1999), the challenge is then, if possible, to make verifiable what is observable.[5] Renegotiation of a baseline contract (Bentley MacLeod and James Malcomson 1993; Aaron S. Edlin and Stefan Reichelstein 1996; Che and Hausch 1999) may do the trick; message spaces must not be too small (B. Douglas Bernheim and Michael D. Whinston 1998). And there always exists a contract with large enough message spaces that yields the optimal allocation if it is not renegotiated; renegotiation at best is innocuous or else lowers welfare (Hart and Moore 1988; Maskin and Moore 1999). In contrast with this literature, which ignores ex ante cognition, this paper focuses on the coarseness of the information on which the contract builds and links contract completeness to the transaction costs incurred in designing the contract.

Finally, the paper is most closely related to interesting and independent work by Bolton and Antoine Faure-Grimaud (2007) on the relationship between information acquisition and contracts. That paper builds on their 2005 paper, in which a single decision maker may choose to incur delay costs (the counterparts of our cognitive costs) in order to take better decisions. The 2005 paper depicts individual decision making as a bandit problem, in which thinking ahead to define a complete action-plan enables the individual to react more promptly to a new event, but delays initial decisions. There is no direct transaction cost, but delay is incurred while waiting for information to accrue. The 2007 paper applies these ideas to two-party contracting. Actions are known from the start but parties are initially uncertain about the payoffs attached to a risky action and must choose between this action and a safe one with known consequences. The parties may have a disagreement as to when to select an action (even if they happen to rank the actions similarly). Due to nontransferable utility, the more impatient party cannot compensate the other for acting quickly. Interestingly, Bolton and Faure-Grimaud show that the impatient party may deliberately transfer control to the more patient/cautious one. Their paper and this one share the view that the pursuit of individual interests may make contracts too complete;[6] the two papers are complementary as they use different modeling techniques and stress rather different themes.

II. Cognitive Limitations and Equilibrium Contracts

Let us first describe the model.

[4] Other papers that have built and improved on Dye's approach are Anderlini and Felli (1994, 1999) and Pierpaolo Battigalli and Giovanni Maggi (2002), which take a rather different approach from that followed here.

[5] Early models of observable but unverifiable performances include Holmström (1982) and Crémer (1995).

[6] In Bolton and Faure-Grimaud, the contract is incomplete if it does not specify which action is taken once investment is sunk and the state of nature is revealed; that is, there is scope for further learning before the actual decision is selected.

Designs.—A buyer (B) and a seller (S) contract on the delivery of a good. Initially they can avail themselves of common-knowledge design A. This design (or for that matter an alternative design) costs c for the seller to produce and deliver to the buyer. I assume that c is large enough that the seller will insist on a contract before delivering any design: the seller would not want to invest c in the absence of a contract by fear of a hold-up.

With probability $1 - \rho$, design A is the appropriate design and delivers utility $v > c$ to the buyer. For example, in an industry where there is a standard form contract from which the parties can depart, $1 - \rho$ would capture how well the standard form contract fits a typical transaction.

With probability ρ, however, some other, initially indescribable,[7] design A' delivers utility v to the buyer while A delivers only $v - \Delta$, where $\Delta > 0$. Converting A into A' requires the seller's collaboration; the seller's supplemental cost of this conversion or "adjustment cost" is equal to $a \in [0, \Delta)$. That is, gains from renegotiation are $\Delta - a$. By contrast, if design A' is identified at the contracting stage, the total cost of production is c: there is no adjustment cost.[8]

The adjustment cost a can be interpreted in several ways. If A and A' refer to physical designs such as computer codes or characteristics of an engine, the move from A to A' may involve rewriting lines of code or making the already developed engine consistent with the new specs, involving costs that would have been avoided, had the specs being chosen right initially. If A and A' refer to pure contract design (covenants, indexation, etc.), then the adjustment cost may be thought of as the real cost brought about by surprises in liquidity positions or by suboptimal risk sharing. In either case, a can also stand for the transaction costs or the risk of negotiation breakdown involved in drafting a new contract.

For example, in a *technology licensing* interpretation, the seller licenses a technology to the buyer. To implement the technology to full functionality for his own use, though, the buyer may or may not need a license to another patent owned by the seller that he was not aware of. In a *procurement* interpretation, the authority may discover later on that the specified design does not fit his need and that further work may be demanded from the contractor. Finally, the model has an interpretation in terms of the *implications of a given contract*. Suppose that the contract describes a course of action, but that it may be feasible for the seller to achieve the same contractual requirement in a way that is both cheaper for her and less attractive to the buyer. Assuming that the seller's cost savings associated with perfunctory provision of the output is smaller than the buyer's loss in surplus, the buyer then will renegotiate with the seller to adhere to the spirit, if not the letter, of the contract.

We assume that design A generates gains from trade even in the absence of cognition:

$$v - c - \rho a > 0.$$

Contract Renegotiation.—If the contract specifies the trade of design A' at some price p, then it is implemented without any renegotiation: the seller incurs cost c, delivers the good, and the buyer obtains utility v.

By contrast, when the contract specifies design A, the seller incurs cost c, delivers A, and the buyer takes possession of it. The buyer then learns whether the design is appropriate. If not, he renegotiates with the seller in order to obtain the adjustment to A'.

[7] The model is thus in the spirit of Maskin and Tirole (1999b), in that the von Neumann–Morgenstern utilities are known, but actions or contingencies may not be describable (they are just "numbers").

[8] We locate the uncertainty on the buyer's value. Symmetrically, one could assume that the value is always v, and that the seller's cost of producing is c if A is appropriate and $c + \Delta$ (which can be reduced back to c at cost a for the buyer by switching to design A') if A is not appropriate. The roles of the seller and the buyer would then be reversed. Section IIB combines the two forms of hold-up.

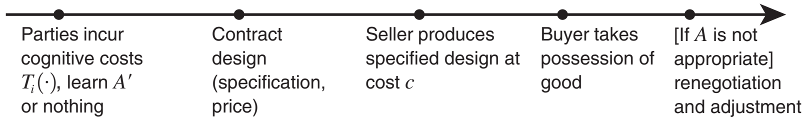

Figure 1. Timeline

We will assume that the buyer and the seller have bargaining power β and σ, respectively $(\beta + \sigma = 1)$, where a party's bargaining power measures the share of the gains from trade that the party can secure in a negotiation. That is, we apply the generalized Nash bargaining solution with weights β and σ. For notational simplicity, the bargaining powers are the same ex ante and ex post. Obviously, this assumption can no longer be sustained when we consider ex ante competition.

Because in most of the paper we will make assumptions that guarantee the existence of a pure-strategy equilibrium and therefore symmetric information on the equilibrium path, the Nash bargaining solution is always well defined. A possible interpretation of the Nash bargaining solution is that the two parties engage in sequential bargaining (with frequencies of offers β and σ, say) and do not update their beliefs when made off-the-equilibrium-path offers (see footnote 13).

Transaction Costs.—Before contracting, the buyer and the seller can incur thinking or cognitive costs $T_B(b)$ and $T_S(s)$, respectively.[9] If A is the appropriate design, they learn nothing from their investigation; if A' is the appropriate design, they learn A' with probability b and s, respectively; and they learn nothing with probability $1 - b$ and $1 - s$. The choices b and s are unobserved by the other party and are individually rational. The functions T_i (for $i \in \{B, S\}$) are smooth, increasing, and convex functions such that $T_i(0) = 0$, $T_i'(0) = 0$, and $T_i(1) = +\infty$. To guarantee that the solution of (4) below is unique and thereby shorten the analysis, we make the following, maintained assumption:

ASSUMPTION 1: $T_B''(b) > \rho^2(1 - \rho)\sigma\Delta/(1 - \rho b)^2$ *for all* $b \in [0, 1]$.

Furthermore, and as discussed in the introduction, the enunciation of A' by i fully reveals to $j \neq i$ that the proper design is A'.

Cognitive costs have a broad range of interpretations (including the managers' psychic cost of focusing on issues they are unfamiliar with, their opportunity cost of not devoting time to other important activities, or the fees paid to lawyers and consultants for advice on contracting).[10] The magnitude of cognitive costs is also revealed indirectly by the substantial incompleteness of many contracts and by the costs of this incompleteness. The timing is summarized in Figure 1.

The timing assumption that the buyer first takes delivery and possession of the good and then discovers whether the design is appropriate (he gets v if A is appropriate and $v - \Delta$ if it is not and the contract is not renegotiated) rules out complex schemes in which the buyer's willingness to pay for the service is elicited prior to actual delivery. One can, for example, imagine that the buyer appropriates the know-how when consuming the good or that the buyer at that stage really needs the good and cannot credibly dispense with it.

[9] As in Mathias Dewatripont and Tirole (2005), for example, the cost of cognition is depicted as a sampling cost. The idea is that, through a combination of cognitive attention and the usual cognitive mechanisms (inference, associativeness, etc.), the agents may stumble on (become aware of) implications of the current design and an alternative to it.

[10] For example, a water concession contract may be a few thousand pages long.

We will say that the contract is *more incomplete* if the ex ante probability that it specifies the available design A is higher.

A. *One-Sided Cognition*

Let us assume in a first step that only the buyer has the ability to learn about the appropriate design. For example, the buyer may have an easier access to information about what he really needs. More fundamentally, we will later observe that there is a basic asymmetry between the buyer's and the seller's incentives to search for a potential hold-up by the seller.

The *socially efficient* level of cognition $\hat{b}$ equalizes the marginal cost of thinking and its marginal benefit (the avoidance of the adjustment cost a when A' is the appropriate design):

$$T_B'(\hat{b}) = \rho a.$$

When the adjustment from A to A' is costless ($a = 0$), then any investment in cognition is pure rent-seeking in this model; that is, the socially optimal levels of cognitive efforts are equal to 0. We will more generally investigate whether the provision of cognitive effort is subject to free riding (has a public good flavor).

(*a*) *Deterministic Cognition Region*

Let us first look for a pure-strategy equilibrium. Let b^* denote the equilibrium probability that the buyer discovers that A is not appropriate when this is indeed the case.

Suppose that the buyer learns nothing. The contract then specifies delivery of design A. Let

$$\hat{\rho}(b) \equiv \frac{\rho(1-b)}{1-\rho b}$$

denote the posterior probability that A is not appropriate conditional on cognitive intensity b and unawareness. If A' is the appropriate design, the seller captures a fraction σ of the renegotiation gain, creating a hold-up:[11]

$$h \equiv \sigma(\Delta - a).$$

Thus, everything is as if the buyer were able to make the adjustment himself, and so receive $v - a$, but had to pay a "tax" h to the seller.

On the *equilibrium path*, $b = b^*$ and the buyer and the seller both expect a hold-up benefit (for the seller) or cost (for the buyer) $\hat{\rho}(b^*)h$. The ex ante price $p(b^*)$ for design A accounts for the possible hold-up, and so[12]

[11] As we will see, this model exhibits another hold-up: the expropriation of the buyer's cognitive investment by the seller at the time of contract writing. By "hold-up," we will henceforth mean the postcontractual hold-up by the seller.

[12] Because the ex ante and ex post bargaining powers have been assumed to be identical, this condition is equivalent to the condition that the ex ante price splits the default (prior to renegotiation) surplus:

$$\sigma(v - c - \hat{\rho}(b^*)\Delta) = p(b^*) - c.$$

The analogous (and equivalent) condition for the buyer is:

$$\beta[v - c - \hat{\rho}(b^*)a] = [v - \hat{\rho}(b^*)(a + h)] - p(b^*).$$

$$\sigma[v - c - \hat{\rho}(b^*)a] = p(b^*) - [c - \hat{\rho}(b^*)h], \tag{1}$$

or

$$p(b^*) = c + \sigma[v - c - \hat{\rho}(b^*)\Delta].$$

The left-hand side of (1) is the share, accruing to the seller, of the total surplus as perceived when bargaining over design A, that is, $v - c - \hat{\rho}(b^*)a$. The right-hand side of (1) is the seller's profit. Her opportunity cost is $c - \hat{\rho}(b^*)h$, since with conditional probability $\hat{\rho}(b^*)$, she will have the opportunity to hold the buyer up for an amount h. The price $p(b^*)$ is such that the seller indeed obtains a fraction σ of the ex ante total surplus. The term $\hat{\rho}(b^*)h$ can be interpreted as a *hold-up discount*.[13]

If the buyer finds out that A' is appropriate, then he will rationally insist that the contract specify the delivery of A', as he knows that under design A he will be held up with probability 1: by disclosing A', the buyer receives fraction β of the gains from trade, i.e., $\beta(v - c)$. By concealing design A', she gets $v - (a + h) - p(b^*)$, where as earlier $p(b^*)$ shares the gains from trade:

$$v - \hat{\rho}(b^*)(a + h) - p(b^*) = \beta(v - c - \hat{\rho}(b^*)a).$$

Thus, disclosing A' increases the buyer's utility by

$$\Delta U_B = \beta a + (1 - \hat{\rho}(b^*))\sigma\Delta > 0.$$

The first term in ΔU_B is a true efficiency gain (of which the buyer appropriates a fraction β), while the second represents the unanticipated expropriation by the seller that the buyer cannot charge for ex ante if he does not disclose design A'.

Gross of cognitive costs, the payoffs for design A' are

$$\beta(v - c) \text{ for the buyer and } \sigma(v - c) \text{ for the seller.}$$

Finally, consider the buyer's choice of cognitive effort b. Assuming for the moment that the buyer trades even when having learned nothing, the latter solves

$$\max_{\{b\}}\{-T_B(b) + \rho b\beta(v - c) + \rho(1 - b)[v - a - h - p(b^*)] + (1 - \rho)[v - p(b^*)]\} \tag{2}$$

$$\Leftrightarrow \max_{\{b\}}\{-T_B(b) + \beta(v - c) + (1 - \rho b)\hat{\rho}(b^*)\sigma\Delta - \rho(1 - b)(a + h)\}.$$

To obtain (2), note that with probability ρb, the buyer proposes design A' and obtains share β of the joint surplus $v - c$. With probability $1 - \rho b$, the two parties contract on design A at price

[13] The generalized Nash bargaining solution allows us to abstract from the issue of inferences that are drawn by the seller in a sequential bargaining game when the buyer makes an off-the-equilibrium-path price offer. Consider an alternating-move bargaining game delivering fraction (β, σ) of the surplus in a complete information setup. Then the solution $p(b^*)$ still prevails under (potentially) asymmetric information as long as the seller's beliefs are passive (that is, the seller still believes that the buyer has chosen $b = b^*$ when the buyer makes an off-the-equilibrium-path offer). By contrast, the other "natural" foundation for the generalized Nash bargaining solution (seller (buyer) makes take-it-or-leave-it offer with probability $\sigma(\beta)$) requires some modification to the analysis. Because the buyer can refuse to trade when the seller makes a take-it-or-leave-it offer, his payoff function is not smooth at $b = b^*$ in a hypothetical pure-strategy equilibrium. The analysis is then similar to that of the case $\beta \leq \beta_0$ below. We opt for simplicity by making an assumption that guarantees the existence of a pure-strategy equilibrium.

$p(b^*)$. With probability $\rho(1 - b)$, the appropriate design is not identified and the buyer must bear the adjustment cost a augmented by the hold-up h. Differentiating (2) (together with the equilibrium condition $b = b^*$) yields first-order condition:

$$T_B'(b^*) = \rho a + \rho[h - \hat{\rho}(b^*)\sigma\Delta]. \tag{3}$$

The left-hand side of (3) is the buyer's marginal cost of cognition. His marginal benefit, the right-hand side of (3), is composed of three terms. The first, ρa, is the social benefit. The second ρh, is the buyer's marginal benefit of avoiding a hold-up. The third, $-\rho\hat{\rho}(b^*)\sigma\Delta$, corresponds to the reduction in the bargained price when specifying design A' rather than design A (the bargained price for design A accounts for the possibility of hold-up and is therefore lower than that for design A').

Uniqueness of b^* and Comparative Statics.—Using the Bayesian updating condition, $\hat{\rho}(b) = \rho(1 - b)/(1 - \rho b)$,

$$T_B'(b^*) = \rho a + \rho\left[h - \frac{\rho(1 - b^*)}{1 - \rho b^*}\sigma\Delta\right]. \tag{4}$$

In particular, in the absence of adjustment cost ($a = 0$, i.e., $h = \sigma\Delta$), the equilibrium level of effort is given by

$$T_B'(b^*) = \frac{\rho(1 - \rho)}{1 - \rho b^*}h.$$

Under Assumption 1 (and the fact that $T_B'(0) = 0$ and $T_B(1) = +\infty$), condition (4) has a unique solution, which lies in $(0, 1)$. One can interpret Assumption 1 in terms of an interesting *strategic complementarity* between the buyer's cognitive choice, b, and the seller's anticipation thereof, b^*. A seller who anticipates a high level of cognition will accept only a low hold-up discount for design A. In turn, a high price for design A makes design A' relatively more attractive and therefore encourages the buyer to engage in more cognition. Assumption 1 puts a bound on this strategic complementarity so as to ensure uniqueness of the deterministic cognition equilibrium.

Condition (4) implies that b^*, and therefore the equilibrium transaction costs increase with the adjustment cost a and with the utility loss Δ.[14]

Nonlocal Deviations.—To check that the cognitive intensity b^* given by (4) is indeed an equilibrium, we need to check that the buyer does not want to deviate and not trade when remaining unaware (the maximand in (2) implicitly assumes that trade always occurs). This is indeed the case if $b \geq b^*$ (as the buyer is then at least as confident as the seller about the absence of hold-up opportunity) or for b a bit below b^*. But for b much lower than b^*, the buyer puts much more weight than the seller on the design not being appropriate, and a bargaining breakdown may occur at the ex ante stage.

[14] As long as the transaction costs exceed the socially optimal level ($b^* > \hat{b}$), b^* also increases with the seller's bargaining power σ (see below for the corresponding comparison).

Expression (2) indeed describes the buyer's payoff as long as the buyer's expected utility when contracting on A (i.e., when being unaware) is positive.[15] Let b' denote the buyer's optimal cognitive intensity when planning not to write a contract when unaware:

$$b' = \underset{\{b\}}{\operatorname{argmax}}\{-T_B(b) + \rho b \beta (v - c)\}.$$

A necessary and sufficient condition for b^* to be the equilibrium cognitive strategy is thus

$$\beta[v - c - \rho(1 - b^*)a] - T_B(b^*) \geq \rho b' \beta(v - c) - T_B(b'), \tag{5}$$

where we make use of the fact that parties are symmetrically informed on the equilibrium path and share the gains from trade $v - c - \rho(1 - b^*)a$ in proportions (β, σ).

Appendix 1 shows that (5) is satisfied if and only if $\beta \geq \beta_0$ for some $\beta_0 \in (0, 1)$.

Are Transaction Costs Too High or Too Low?—The issue of whether the buyer's information acquisition is excessive or insufficient is simply a question of the external effect of specifying A' on the seller. When the buyer specifies A', the effect on the seller is

$$\Delta U_S = \sigma a - [1 - \hat{\rho}(b^*)]\sigma\Delta.$$

The first term in ΔU_S is the share of the true efficiency gain borne by the seller. The second term is (the negative of) the unanticipated expropriation by the seller. Condition (3) can be rewritten as

$$T_B'(b^*) = \rho(a - \Delta U_S).$$

There is excessive cognition if and only if $\Delta U_S < 0$, or

$$\left[\frac{1 - \rho}{1 - \rho b^*}\right]\Delta > a. \tag{6}$$

Fixing a, condition (6) is satisfied[16] if and only if $\Delta \geq \Delta^*(a)$ for some $\Delta^*(a) \geq a$. Note that $\Delta^*(0) = 0$: in the absence of adjustment cost, the contract is always too complete. Conversely, (6) is violated and so the contract is too incomplete when there is no gain to renegotiation ($\Delta \equiv a$).

These results fit with our earlier intuition. Cognition is socially excessive if it is just meant to avoid hold-ups. By contrast, for large adjustment costs, which end up being shared between the two parties, free riding is paramount and cognition is socially suboptimal.

[15] That is, $v - p(b^*) - \hat{\rho}(b)(a + h) \geq 0$, or $\beta(v - c) \geq \hat{\rho}(b)(h + a) - \hat{\rho}(b^*)(h + \sigma a)$. Because $\hat{\rho}(b) \leq \rho$, a sufficient condition for this inequality to be satisfied for all b when $a = 0$ is that the buyer's bargaining power be sufficiently large in relation to the relative hold-up stake:

$$\frac{\beta}{1 - \beta} \geq \frac{\rho\Delta}{v - c}.$$

[16] Recall that b^* is an increasing function of Δ and a.

Does It Pay to Be Bright/Experienced?—One may wonder whether brightness or experience benefits the contracting parties.[17] Intelligence or experience with this type of contract can be described in this framework as a reduction in the marginal cost of cognition, $T_B(b;\nu)$; that a higher ν corresponds to a higher intelligence, or experience corresponds to the condition $\partial\big(T_B'(b;\nu)\big)/\partial\nu < 0$. Condition (4) implies that a brighter buyer is more likely to find out a contractual shortcoming $(\partial b^*/\partial\nu > 0)$. This implies that with positive adjustment costs, a bright buyer is a more attractive trading partner for the seller (who receives $\sigma[v - c - \rho(1 - b^*)a]$). As for the buyer, condition (2) implies that being *perceived* as stupid/inexperienced, by increasing the hold-up discount, benefits the buyer. The ideal situation for the buyer is therefore to be bright and not to be perceived so.

"Something's Fishy."—Condition (4) may still hold even if A' is not necessarily the final design when A is not the appropriate one. Namely, suppose that when the buyer learns A' ex ante, there is positive probability that the appropriate design is some A'' rather than A'. As long as the buyer cannot further investigate whether A' or A'' obtains, then assuming there is no adjustment cost from A' to A'', condition (4) holds: the buyer under design A' receives expected gross payoff $\beta(v - c)$ *on and off* the equilibrium path (i.e., for any b, and not only for $b = b^*$). Put differently, the important feature behind (4) is that the identification of A' creates symmetric information, not that it is the final design or that it eliminates ex post hold-ups by the seller.

The methodology developed here also extends to a situation in which a modification to the standard design may trigger further cognition. To illustrate this point in the simplest manner, consider the following "stationary environment." Let $A_0 \equiv A$. If the buyer learns about design A_k, then there is probability $1 - \rho$ that this design is appropriate and probability ρ that it is not. By exerting cognitive effort $T_B(b_k)$, the buyer discovers design A_{k+1} with probability ρb_k (and can then continue his cognitive process). Suppose further that there is no adjustment cost ($a = 0$) so as to simplify expressions. One can search for an equilibrium in which, conditional on knowing design A_k, (i) the buyer exerts cognitive effort $b_k \equiv b^{**}$, (ii) if he does not discover anything wrong with A_k, he discloses it and trades at price p, and (iii) the value function is stationary. Let V_k denote the value function:

$$V_k \equiv \max_{\{b_k\}} \{-T_B(b_k) + \rho b_k V_{k+1} + (1 - \rho)(v - p) + \rho(1 - b_k)(v - a - h - p)\},$$

$$V_k = V \quad \text{for all } k, \qquad b_k = b^{**} \quad \text{for all } k,$$

and

$$p = c + \sigma[v - c - \hat{\rho}(b^{**})\Delta].$$

After some manipulation, one obtains

$$T_B'(b^{**}) = \rho\left[\frac{1 - \rho}{1 - \rho b^{**}}\, h - \frac{T_B(b^{**})}{1 - \rho b^{**}}\right]$$

[17] I am grateful to Joel Sobel for suggesting this question.

as compared to

$$T_B'(b^*) = \rho\left[\frac{1-\rho}{1-\rho b^*}\ h\right].$$

The final design has a higher probability of being inappropriate in the multistage cognition problem: the buyer realizes that finding a fault with the candidate design will not alleviate the endogenous adverse selection concern and will lead him to incur future expected cognition cost $T_B(b^{**})[1 + \rho b^{**} + (\rho b^{**})^2 + \cdots] = T_B(b^{**})/[1 - \rho b^{**}]$.

Postcontractual Cognition?—Does the buyer ever have an incentive to look for problems after the contact has been signed but before adjustment costs must be incurred (i.e., in time to adjust the design without having to incur the adjustment cost)? The answer is no. Intuitively, the hold-up by the seller can no longer be avoided once a contract has been signed and so the incentives to engage in cognition are diminished. To show this formally, let $T_B(b + \hat{b})$ denote the cognition cost, with b and $\hat{b}$ denoting the pre- and postcontractual efforts.[18] To see that $b = b^*$ and $\hat{b} = 0$ is an equilibrium, note that a postcontract identification of design A' yields an extra βa to the buyer, and so the marginal gain of cognition at $\hat{b} = 0$ is

$$-T_B'(b^*) + \rho\beta a \leq 0$$

(with strict inequality if $\sigma > 0$).

(*b*) *Randomized-Cognition Region*

Conversely, if $\beta < \beta_0$, the buyer randomizes over his cognition choice. Furthermore, the ex ante contract negotiation necessarily breaks down with strictly positive probability; for, in the absence of breakdown, the equilibrium price for design A must be independent of the value b chosen by the buyer (since the buyer prefers the lowest possible price among those consistent with a deterministic agreement regardless of his cognitive choice). Consequently, (2) and (4) must hold, a contradiction.

Appendix B provides an explicit solution when the seller makes a take-it-or-leave-it offer ($\beta = 0$) and adjustments are costless ($a = 0$). The equilibrium involves the buyer randomizing over some interval $[0, \bar{b}]$ and the seller randomizing over the price interval $[v - \rho\Delta, v - \hat{\rho}(\bar{b})\Delta]$ for design A (the seller charges v for design A'). The buyer has a zero expected utility, as one would expect, given that the seller has full bargaining power. The lower bound $\underline{b}$ on the buyer's cognitive intensity is necessarily equal to zero since the seller will never charge less than the willingness to pay, $v - \hat{\rho}(\underline{b})\Delta$, of the most pessimistic "type" $\underline{b}$. Thus the buyer's utility is $-T_B(\underline{b})$, and so $\underline{b} = 0$ for this utility to be nonnegative.

To illustrate this randomization in a simpler manner, let us restrict the analysis to the case of two feasible levels of cognition.

Two Levels of Cognition.—Suppose that $\beta = a = 0$ and that the buyer chooses between cognition levels 0 (at no cost) and b (at cost T_B: we omit the argument for notational simplicity). Let ρ and $\hat{\rho} < \rho$ denote the corresponding posteriors, and $p_0 \equiv v - \rho\Delta$ and $p_b \equiv v - \hat{\rho}\Delta$, the buyer's associated willingness to pay for design A.

[18] One may have in mind that the buyer (rationally) investigates first the most obvious/cheapest routes to find flaws in the design, and then tries harder ones.

Let the buyer choose cognition level 0 (respectively b) with probability g (respectively $1 - g$). Given that it is suboptimal for the seller to charge prices differing from p_0 and p_b, let the seller charge p_0 (respectively p_b) with probability f (respectively $1 - f$).

To make things interesting, assume that

$$(1 - \rho b)(\rho - \hat{\rho})\Delta > T_B.$$

This condition ensures that the deterministic choice of no cognition is not an equilibrium;[19] for, if it were, then $f = 1$, and by deviating to cognition level b, the buyer would receive[20]

$$(1 - \rho b)(v - p_0 - \hat{\rho}\Delta) - T_B = (1 - \rho b)(\rho - \hat{\rho})\Delta - T_B > 0.$$

The buyer cannot choose cognition level b for certain either, since if he did the seller would charge p_b for sure $(f = 0)$ and so the buyer's utility would be equal to $-T_B < 0$. In equilibrium therefore, the buyer must play a mixed strategy and be indifferent between the two cognitive choices:

$$f(1 - \rho b)(\rho - \hat{\rho})\Delta - T_B = 0.$$

Similarly, the seller must be indifferent between p_0 and p_b:[21]

$$p_0 - c + \frac{g\rho + (1 - g)(1 - \rho b)\hat{\rho}}{g + (1 - g)(1 - \rho b)}\Delta = \frac{(1 - g)(1 - \rho b)}{g + (1 - g)(1 - \rho b)}[p_b - c + \hat{\rho}\Delta],$$

or, after some manipulation:[22]

$$\frac{g}{1 - g} = (1 - \rho b)\frac{(\rho - \hat{\rho})\Delta}{v - c}.$$

Note that when the gains from trade, $v - c$, increase, the seller's pricing behavior (f) remains unchanged (the prices p_0 and p_b increase with v, of course). By contrast, the buyer exerts more cognition (g decreases); otherwise the seller would be too eager to guarantee a sure agreement, i.e., to charge the lower price p_0.

[19] In the continuous-cognition analysis of Appendix B, this is ensured by the assumption that $T_B'(0) = 0$.

[20] The buyer is charged v and has a zero gross utility when bargaining over design A'. Yet the buyer prefers to disclose design A' when aware of it since he receives a negative utility of design A even for the lower price p_0: $v - p_0 - \Delta = -(1 - \rho)\Delta < 0$.

[21] Note that the probability of "type" 0 *conditional* on design A is higher than g and is equal to $g/[g + (1 - g)(1 - \rho b)]$.

[22] To interpret this condition, rewrite it as $g(v - c) = [(1 - g)(1 - \rho b)](\rho - \hat{\rho})\Delta$. The left-hand side is the expected loss of charging p_b rather than p_0 (with probability g, the buyer chooses no cognition and would have accepted p_0, yielding surplus $v - c$ to the seller). The right-hand side is the expected gain associated with rent extraction when the buyer engages in cognition b.

(*c*) *Summary*

PROPOSITION 1 (One-Sided Cognition):

(*i*) *When* $\beta \geq \beta_0$*, the buyer incurs cognitive cost* $T_B(b^*)$ *where* b^* *is given by*

$$T_B'(b^*) = \rho a + \rho[h - \hat{\rho}(b^*)\sigma\Delta].$$

The buyer's cognitive cost increases with the severity of the hold-up problem (increases with Δ*) and with the adjustment cost* a*.*

The buyer fully bears the deadweight loss associated with cognition and has utility $\beta[v - c - \rho(1 - b^*)a] - T_B(b^*)$*, while the seller has utility* $\sigma[v - c - \rho(1 - b^*)a]$*.*

The contract is too complete (there is too much cognition, i.e., $b^* > \hat{b}$*) if and only if*

$$\left[\frac{1 - \rho}{1 - \rho b^*}\right]\Delta > a.$$

(*ii*) *When* $\beta < \beta_0$*, contract negotiations break down with positive probability.*

B. *Two-Sided Cognition*

Suppose that the seller's cognition is no longer "prohibitively" expensive. The seller can learn about design A' (when relevant) with probability s at cost $T_S(s)$. We look for conditions under which the one-sided cognition equilibrium $(b = b^*, s = 0)$ is also a two-sided cognition equilibrium. To this purpose, we investigate two questions. First, does the seller have an incentive to disclose the existence of design A' when she becomes aware of it? Second, does the seller refuse to trade when not learning about a potential hold-up? If the answers to these two questions are negative, then the seller gains nothing from becoming informed and, a fortiori, does not want to incur costs to become informed.[23]

Disclosure of Design A'.—Suppose that the buyer selects b^*. If the buyer finds out that the design should be A', then any cognitive effort by the seller is wasted. Suppose therefore that the seller, but not the buyer, becomes aware that design A' is the appropriate one. The seller obtains

$$\sigma(v - c)$$

by revealing A'. By concealing A', she obtains, instead,

$$p(b^*) - c + h = \sigma(v - c) + [h - \hat{\rho}(b^*)\sigma\Delta].$$

Concealing design A' therefore dominates disclosure if and only if

$$h \geq \hat{\rho}(b^*)\sigma\Delta \Leftrightarrow \Delta - a \geq \hat{\rho}(b^*)\Delta.$$

This is nothing but the condition (condition (6)) under which transaction costs are too high. Put differently, the excessive incentive for the buyer to engage in cognition is linked to the seller's

[23] This holds regardless of the degree of correlation between the two cognitive activities.

incentive not to disclose. This is natural: excessive buyer cognition arises when the disclosure of A' by the buyer hurts the seller, and therefore also corresponds to the condition under which the seller does not want to disclose design A' herself. The seller does not disclose if the hold-up benefit h that she will enjoy ex post by not disclosing exceeds the hold-up discount, $\hat{\rho}(b^*)\sigma\Delta$, on the bargained price under design A; similarly, the buyer has excessive incentives to invest in cognition if the seller's hold-up benefit exceeds the ex ante price reduction associated with design A'.

Incentive for Not Trading.—If the seller does not benefit from disclosing design A', the only potential incentive for the seller to invest in cognition is to avoid trade when she does not discover an opportunity for a hold-up. Consider a candidate equilibrium in which only the seller does not disclose when learning design A' and trades even if unaware of a hold-up opportunity. The price is then $p(b^*)$ for design A. If $p(b^*) \geq c$, or

$$v - c \geq \hat{\rho}(b^*)\Delta, \tag{7}$$

however, the seller always wants to trade even if sure that design A is appropriate.

PROPOSITION 2 (Two-Sided Cognition): *The one-sided cognition equilibrium for $\beta \geq \beta_0$ is also a two-sided-cognition equilibrium ($b = b^*, s = 0$) provided that:*

(*i*) *Contracts are too complete in the one-sided cognition equilibrium:*

$$\Delta - a \geq \hat{\rho}(b^*)\Delta, \tag{8}$$

and

(*ii*)

$$v - c \geq \hat{\rho}(b^*)\Delta. \tag{9}$$

Proposition 2 captures a basic asymmetry between the buyer's and the seller's gains from learning about an opportunity to hold up the seller. The buyer gains from being informed about the probability of a hold-up as this enables him to avoid it. By contrast, when aware of an opportunity for a hold-up, the seller may not want to disclose it; her only potential use of this information is then to refuse to trade when she is pessimistic about the possibility of an additional income ex post. If the scope for hold-up is, however, limited (note that a sufficient condition for (9) to be satisfied is that $v - c \geq \rho\Delta$), then it is not worth forgoing gains from trade in order to economize on the hold-up discount.[24]

[24] Condition (9) is stronger than needed for the property that the seller does not incur any cognitive cost. For example, suppose that (i) the seller chooses cognitive cost $T_S(s)$ after being offered to negotiate on design A (so the seller knows that the buyer is unaware of an alternative design) and (ii) the probabilities of discovering A' if appropriate are independent for the buyer and the seller. Then the seller exerts no cognitive effort when bargaining over A if

$$\sigma(v - c) \geq \max_{\{s\}}\{\hat{\rho}(b^*)s[p(b^*) + h - c] - T_S(s)\}.$$

In the extreme case of free cognition ($T_S \equiv 0$ and so $s = 1$), this condition boils down to $v - c \geq \hat{\rho}(b^*)\Delta$, that is, to (9); but it is in general weaker than (9).

When Does the Seller Engage in Cognitive Effort?—While Proposition 2 captures a fundamental asymmetry in the incentives to engage in cognition about the potential for hold-up by the seller, more generally the seller may exert cognitive effort for one of three reasons:

(i) *Severe hold-up discount.* First, if the hold-up discount $(\hat{\rho}(b^*)\sigma\Delta)$ is large (either (8) or (9) fails), the seller may want to learn about the inappropriateness of design A, either to avoid the adjustment cost or not to trade when unaware of a hold-up opportunity.

(ii) *Good news for the buyer.* We have so far assumed that cognition may unveil *bad news* (a hold-up opportunity) for the buyer. We could more symmetrically assume that cognition may also unveil *good news* for the buyer (some unexpected benefit delivered by the seller's technology). We then have a natural *specialization principle*: the buyer exerts cognitive effort in order to identify potential hold-ups while the seller tries to push the price up by finding reasons why her technology delivers value to the buyer.

(iii) *Bad news for the seller.* The seller may attempt to learn about bad news for herself. Indeed, consider the more symmetric model with three states of nature. With probability ρ, the appropriate design is A', which delivers v at cost c, while A delivers $v - \Delta$ at the same cost. With probability μ, the appropriate design is A'', which delivers v at cost c, while A delivers v at cost $c + \Delta$ for the seller. For example, contract A fails to identify a loophole in a noncompete clause; due to Bertrand competition in an adjacent market, this loophole brings little or no revenue to the buyer, but costs Δ to the seller. Alternatively, design or contract A may fail to specify some marketing, technological, or standardization effort by the buyer, costing Δ to the seller. When the appropriate design is A' or A'', but the initial contract specifies A, adjustment cost $a \in [0, \Delta)$ needs to be incurred (by the party who does not lose Δ) to move to the appropriate design.

Finally, with probability $1 - \rho - \mu$, A is the appropriate design and delivers value v at cost c.

Suppose that the two parties have a comparative advantage in investigating their own payoff, i.e., here in searching for bad news for themselves: the buyer can learn design A' (if relevant) with probability b at cost $T_B(b)$, and the seller can learn design A'' (if relevant) with probability z at cost $\mathcal{T}_S(z)$. The symmetry of the situation suggests that the seller indeed engages in cognition so as to identify bad news for herself. Appendix C checks that this is indeed the case and provides the equilibrium conditions for the cognitive efforts.

III. Robustness

A. *Does Competition Reduce Transaction Costs?*

This section provides some preliminary insights on the impact of ex ante competition among sellers. The bargaining weights β and σ (and relatedly the hold-up stake h) now refer exclusively to the ex post bargaining strengths. For the sake of simplicity and unless otherwise specified, we assume one-sided cognition and no adjustment cost. We generalize the analysis of Section IIA to the case of Bertrand competition among (at least two) sellers.

Competition for design A' yields price equal to cost (c) and buyer gross surplus $v - c$. Suppose next that the proposed design is A and that the equilibrium level of cognition is b^* (as it turns out, this level will be the same as in Section IIA, so we do not introduce new notation). The competitive price is then

$$p = c - \hat{\rho}(b^*)h.$$

Underpricing (low-ball bidding) thus reflects a hold-up discount.

And so the buyer solves

$$\max_{\{b\}}\{-T_B(b) + \rho b(v - c) + \rho(1 - b)[v - c - [1 - \hat{\rho}(b^*)]h] + (1 - \rho)[v - c + \hat{\rho}(b^*)h]\},$$

or, letting $H(b,b^*) \equiv [(1 - \rho)\hat{\rho}(b^*) - \rho(1 - b)[1 - \hat{\rho}(b^*)]]h$,

$$\max_{\{b\}}\{-T_B(b) + (v - c) + H(b,b^*)\}.$$

By contrast, under bilateral monopoly (Section IIA), the buyer solved

$$\max_{\{b\}}\{-T_B(b) + \beta(v - c) + H(b,b^*)\}.$$

Thus, competition allows the buyer to appropriate an extra $\sigma(v - c)$ (the seller's share of the total gross surplus), but does not affect his incentive to exert cognition.

Next, consider the buyer's option not to contract on design A. The buyer then solves

$$\max_{\{b\}}\{-T_B(b) + \rho b(v - c)\}.$$

By comparison, under bilateral ex ante monopoly, he solved

$$\max_{\{b\}}\{-T_B(b) + \rho b\beta(v - c)\}.$$

This analysis yields:

PROPOSITION 3: *Under Assumptions 1 and 2 and ex ante seller competition, there exists* $\beta_0^c < \beta_0$ *such that a pure-strategy equilibrium exists if and only if* $\beta \geq \beta_0^c$. *The cognitive intensity* b^* *is then the same as under ex ante bilateral monopoly.*

Proposition 3 shows that competition in general does not reduce the buyer's transaction costs. Intuitively, the buyer is worried about the occurrence of an ex post hold-up, and this concern is the same regardless of the extent of ex ante competition.

By contrast, Proposition 2 suggests why the analysis of two-sided cognition (Section IIB) must be amended. We saw that in an ex ante bilateral-monopoly situation, the seller may not want to sink cognitive effort because she may enjoy a positive surplus even when she is unaware of hold-up opportunity. Ex ante competition destroys this surplus. Indeed, the seller loses $\hat{\rho}(b^*)h$ in the absence of hold-up. Thus, provided that $T_S'(0) = 0$, the sellers sink cognitive effort. The situation then resembles one of common values, with positive affiliation among sellers, and negative affiliation between the sellers and the buyer. Interestingly, the winner's curse faced by the sellers depends not only on the auction's design, but also on the other sellers' and the buyer's cognitive investments. And the buyer might try to limit sellers' access to information in order to avoid costly rent seeking. We leave the fascinating topic of auction design with transaction costs for future research.

B. *Mechanisms to Modify Ex Post Bargaining Power*

We have focused on contracts that specify a design and a price. However, the fact that other designs (in the current model, A') are not initially describable may not prevent parties from writing contracts that elicit these designs when they later enter the parties' awareness; what can be achieved then depends on the specification of the stochastic structure for von Neumann–Morgenstern payoff functions.[25] On the other hand, the robustness of the cognitive rent-seeking point and the implications derived in this paper hinge only on the buyer's concern about being held up because of a faulty contract design. Even if one were able to control through a contract the sharing pattern for the gain of adjustment from design A to design A', it is in general unwise not to let the seller receive any benefit from this adjustment. Let us give three reasons for this.

Finding Ex Post Solutions.—Suppose that restoring design A''s utility under faulty design A requires ex post innovativeness by the seller and the buyer. Ex post sharing of the gain from trade then determines the ex post incentives to find solutions. Any mechanism is characterized by the shares $\hat{\beta}$ and $\hat{\sigma}$ of the gains from trade $\Delta - a$ realized when a design is adjusted with gain Δ for the buyer and cost a to the seller.[26] Because $\hat{\beta} + \hat{\sigma} = 1$, it may not be feasible to make both parties choose efficient levels of cognition *and* the proper ex post effort to find a solution.

Even under one-sided cognition, discouraging cognitive investments will generally entail other costs. Consider, for example, the buyer's cognitive investment. From (4), b^* increases with $\hat{\sigma}$ and reducing excess cognition requires lowering $\hat{\sigma}$. It is easy to think of reasons why the seller must receive a minimum fraction of the surplus when the design is adjusted. Let us depict two such setups.

Degradation.—In the first, when A is the right design, the buyer can degrade it, thereby appropriating some private benefit ξ and creating a cost d to the seller (for example the buyer may fail to invest in the technology, and thereby to generate spillovers for the seller); the value of A to the buyer is then only $v - \Delta$. Assume a cost $a \in (\xi, \Delta)$ of adjusting it back, to a design that is payoff equivalent to A; then the states of nature in which A must be adjusted to A' because of degradation and because A was inappropriate to start with are indistinguishable from a von Neumann–Morgenstern perspective, and so any mechanism must deliver the same surpluses to the buyer and the seller. To avoid a degradation (which is undesirable if $a + d > \xi$), the share of the ex post surplus, $\hat{\sigma}(\Delta - a)$, received by the seller in case of renegotiation must be sufficiently large. This puts a lower bound on the seller's hold-up, and therefore on the buyer's incentive to engage in cognitive rent-seeking.

[25] On this, see Maskin and Tirole (1999b).

[26] As in Maskin and Moore (1999) and Segal and Whinston (2002), we rule out third parties ("budget breakers"). Recall also that we are in a risk-neutral world, which further limits what can be implemented (see Maskin and Tirole 1999b).

Seller Investment.—Another standard reason why the seller must get some ex post surplus is that some specific investment by the seller must be encouraged. Suppose, for example, costless adjustments ($a = 0$) and that, conditional on A being the appropriate design, a postcontractual investment by the seller raises the buyer's utility by Δ after some adjustment. Again, assume that the adjustments from A to A', whether justified by an inappropriate initial design or by seller investment, are payoff-equivalent (they raise the buyer's utility by Δ). Encouraging specific investment by the seller requires providing her with enough surplus from adjustments, thus creating incentives for the buyer to invest in cognition.

Appendix D goes into the analysis of this section more formally and shows that even in the most favorable case, in which ex post bargaining powers can be fine-tuned through an initial contract, the possibility of cognitive rent seeking remains.

IV. Determinants of Contract Incompleteness and Reverse Causality

This section draws some implications of the model to highlight the two-way relationship between reputational concerns, control rights, and specific investments on the one hand, and contract incompleteness on the other.

Relational Contracting and Incompleteness.—The literature has emphasized that firms will search for repeated relationships when contracts are (for an exogenous reason) incomplete. Following Macaulay (1963), who argued that a key virtue of relational contracting is that parties can count on each other to abide by the spirit of the contract, and can therefore economize on the cost of specifying its letter, I point out that causality runs in both directions.

A dynamic model of relationship contracting can be found in Tirole (2007). The gist of the analysis can be apprehended in reduced form as follows. Suppose that, when faced with the opportunity to hold up the buyer, the seller exploits it only with probability x. With probability $1 - x$, the seller abides by the spirit of the contract rather than by its letter,[27] and makes the adjustment at cost (a). In a full-fledged reputation model, x decreases with the sellers' patience and increases with the fraction of a priori opportunistic sellers.

The generalization of condition (4) is then

$$T_B'(b^*(x)) = \rho a + \rho[xh - \sigma\hat{\rho}(b^*(x))a].$$

For example, in the absence of adjustment cost $(a = 0)$ and when opportunistic sellers always pool with honest ones ($x = 0$), $b^*(0) = 0$: the buyer need not engage in cognition as he knows that he can rely on the seller to behave fairly (in the spirit of the contract) ex post.

Contract Length and Vertical Integration.—Suppose now that the parties choose between the "basic contracting mode" (on A or A') as described above or an "alternative contracting mode"; in order to encompass several applications, let us leave this alternative mode abstract and just assume that it generates an extra cost K, and it does not leave scope for hold-ups. For simplicity,

[27] As Macaulay (1963) noted long ago: "There is a hesitancy to speak of legal rights or to threaten to sue in these negotiations." This view is based on interviews with businessmen and lawyers. Representative narratives are: "If something comes up, you get the other man on the telephone and deal with the problem. You don't read legalistic contract clauses at each other if you ever want to do business again. One doesn't run to lawyers if he wants to stay in business because one must behave decently;" and "You can settle any dispute if you keep the lawyers and accountants out of it. They just do not understand the give-and-take needed in business."

there is no adjustment cost $(a = 0)$, and β is sufficiently large for the pure-strategy equilibrium to exist in Proposition 1 $(\beta \geq \beta_0)$. The buyer will not opt for the alternative mode if

$$\beta(v - c) - T_B(b^*) \geq \beta(v - c - K)$$

or

$$\beta K \geq T_B(b^*). \tag{10}$$

Note that the buyer then exerts a positive externality on the seller who receives $\sigma(v - c)$ on average rather than $\sigma(v - c - K)$.

If (10) is violated however, then the alternative contracting mode must be chosen with positive probability.[28] The analysis is then similar to the analysis of mixed-strategy equilibria in Section IIA.

Let us give two applications of this idea:

(i) *Delays and contract length.* The "alternative contracting mode" could refer to contracting and incurring the cost c of manufacturing the design tomorrow rather than today. This makes the project less attractive (for example the buyer's utility is $v - K$ as opposed to v if things get started today),[29] but has the benefit that the relevant design (A or A') will be costlessly learned by both parties, economizing on transaction costs.

This very stylized model offers a metaphor for the trade-off involved in writing detailed long-term contracts: filling out the details early on, before the uncertainty resolves itself, gives rise to transaction and adverse selection costs. A slightly more sophisticated version of the same idea can be found in Tirole (2007), in which K stands for the cost of not protecting another investment through a long-term contract rather than for a cost of delaying the start of a project. For example, there could be two periods 1 and 2 (at which the uncertainty about the design resolves exogenously). Delivery occurs at date 2 regardless of the choice of contract. The "basic mode" involves contracting at date 1 on A or A'. The "alternative mode" consists in waiting and signing a spot contract at date 2. This interpretation fits, for example, with Paul L. Joskow's (1987) empirical evidence on coal contracts, which shows that such contracts are much shorter when ex post competition makes the prospect of hold-up more remote.[30]

[28] But not with probability 1. Otherwise $b = 0$ and the two parties are better off contracting on A, with probability ρ of a hold-up.

[29] See also Bolton and Faure-Grimaud (2007).

[30] The relationship between transaction costs, incompleteness, and contract length is illustrated, for example, by the contrast between traditional procurement contracts and public-private partnerships (PPPs) in which the conception and construction stage is bundled with longer-term operations and maintenance. PPPs (Private Finance Initiative (PFI) deals in the United Kingdom) are known to be prone to both contractual incompleteness and high transaction costs (Välilä 2005). The UK National Audit Office (2003) argues that "Private Finance Initiative deals remain very costly to negotiate and these costs need to be factored into the assessment." For example, the mere cost of negotiations for the London Underground deal was £180 million (and this number excludes bidders' transaction costs); the average cost of National Health Service PFI deals for external advisors only is about 3.7 percent of the capital value of the projects.

(ii) *Ownership.* The endogeneity of the extent of completeness provides a rationale for the conventional wisdom that vertical integration economizes on contracting costs, i.e., that a key benefit of firms is that they allow less explicit contractual specifications.[31]

Suppose that the buyer ex post can "fix" (costlessly modify by himself) a wrong design A, provided that he acquires ownership of the seller's technology and not only of the seller's good.[32] Such "buyer ownership"[33] thus prevents ex post hold-ups.

The transfer of the seller's technology to the buyer, however, creates a deadweight loss. The literature has analyzed a number of reasons why it may be so, and we will therefore not be interested in formalizing a particular cost; for example, the transfer may create competition in the R&D or some downstream product market, generating cost K. Alternatively, transferring the bargaining power through ownership may involve the costs discussed in Section IIIB.

This idea is consistent with, and provides a potential rationale for, available empirical evidence. Jonathan Levin and Steven Tadelis (2007) document that local governments tend to provide public services in-house when contracting on these services is complex. To the extent that transaction costs increase with the complexity, it is natural that local governments use vertical integration to economize on them for highly complex or unknown operations.

Sharon Novak and Scott Stern (2007) investigate the relationship between make-or-buy choice and quality improvements over the life cycle in the automobile industry. That paper shows that quality increases over time for in-house components and remains roughly constant for outsourced components. My analysis provides one possible rationale for this observation. Flaws in design (the inappropriateness of A in our model) are unveiled and corrected only over time; thus the "adjustment cost" refers not only to the technical cost, but also to the buyer's lost time incurred in producing a lower-quality output.[34] Because the probability of adjustment, $\rho(1 - b)$, is higher under vertical integration than under outsourcing, quality improvements are more likely under vertical integration.[35]

V. Alleys for Future Research

The introduction listed the main insights. Rather than restating them, let us conclude with a couple of alleys for research.

The paper has stressed *precontractual* transaction costs. Similar techniques could be used to give content to transaction costs that are incurred at the *implementation* stage. For example, while in Maskin and Tirole (1999b) actions and contingencies may be indescribable at the contracting stage, they are costlessly describable ex post. Thinking through alternatives ex post,

[31] For example, Klein (2002) argues: "If [the pioneering Grossman and Hart model of integration] has the advantage of taking the incompleteness of contracts seriously, it does not consider the key aspect of the contractual arrangement we identify with the firm, namely that it involves less explicit contractual specification and more flexibility." Note, though, that in recent interesting work, Hart and Moore (2008) do endogenize the degree of contractual incompleteness by introducing the behavioral feature of feelings of entitlement that make economic agents want to shade on performance when they feel shortchanged.

[32] This is one of many ways of apprehending the transfer of a control right from the seller to the buyer. For example, a transferable control model à la Aghion, Dewatripont, and Rey (2004) could alternatively be envisioned.

[33] "Buyer ownership" may not be exclusive. For example, the seller may have granted a block license for her entire intellectual property, and still be able to use the intellectual property herself.

[34] To formalize this idea in the context of this paper's model, one can assume that the buyer's value is $v - \Delta$ for some time after the product is introduced and v once the seller has adjusted the design.

[35] An alternative explanation, noted by Novak and Stern, builds on Patrick Bajari and Tadelis (2001), which shows that adjustments are less likely to be made under high-powered incentive schemes (outsourcing here). The idea in Bajari and Tadelis is that accounting structures are permanently different under vertical integration and separation, and that fixed-price contracts are fraught with asymmetric information ex post, and therefore harder to renegotiate.

however, involves transaction costs that are similar to those formalized here. It is certainly worth considering such ex post transaction costs.

Second, while the paper has stressed that a specialization in the cognitive investigations is natural, interesting patterns of strategic complementarities/substitutabilities may emerge. This would arise in particular if both parties' "thinking deeply enough about the issues" were needed for the realization of gains from trade.

Third, the standard designs (A) and the cognitive cost functions (T_i) are influenced by past contractual experimentation within and outside the industry. Familiarity with designs and their implications is path-dependent. The dynamics of contractual incompleteness and heterogeneity across relationships is a fascinating topic for research.

Fourth, this line of thought may have implications for the separation of tasks of designing and building.[36] An architect or designer may be rewarded for getting the design right at the outset. An integrated designer-builder may economize on costs, but may require a strong reputation mechanism, akin to that envisioned in Section IV.

Finally, we have assumed that the contract designers are residual claimants. In practice, they often are agents for their respective organizations. Their incentives to sink transaction costs (opportunity cost of their time, legal and consulting fees, etc.) depend on the design of their incentive package and on the internal monitoring setup. As observed by Macaulay (1963), managers are usually tempted to write incomplete contracts and count on relational contracting to discipline their counterpart. Their hierarchy and legal councils, by contrast, are advocates of more rigor and try to avoid the off–balance sheet liabilities associated with adjustment costs and hold-ups. Put differently, the extent of contract incompleteness depends on the firms' internal organizations.

Appendix A: Existence of a Cut-Off β_0

Note, first, that for $\beta = 1$ and for any a, the buyer is always better off trading:

$$\max_{\{b\}}\{-T_B(b) + (v - c) - \rho(1 - b)a\} > \max_{\{b\}}\{-T_B(b) + \rho b(v - c)\}$$
$$= \max_{\{b\}}\{-T_B(b) + v - c - (1 - \rho b)(v - c)\},$$

where we make use of the assumption that there are gains from trade even in the absence of cognition ($v - c - \rho a > 0$). Note also that for $\beta = 0$, the left-hand side of (5) is equal to $-T_B(b^*) < 0$ and the right-hand side to zero. Let

$$U^T(\beta) \equiv -T_B\big(b^*(\beta)\big) + \rho b^*(\beta)\beta(v - c) + \big[1 - \rho b^*(\beta)\big]\beta\big[v - c - \hat{\rho}\big(b^*(\beta)\big)a\big]$$

and

$$U^{NT}(\beta) \equiv -T_B\big(b'(\beta)\big) + \rho b'(\beta)\beta(v - c).$$

Consider a β such that the buyer wants to contract on design A when the seller expects $b^*(\beta)$ and, furthermore:

$$U^T(\beta) = U^{NT}(\beta).$$

[36] I'm grateful to a referee for this suggestion.

It must be the case that when choosing $b'(\beta)$ and the seller still expects $b^*(\beta)$, the buyer does not want to contract on design A. Furthermore, for the buyer not to be willing to trade at price $p(b^*(\beta))$, it must be the case that

$$b'(\beta) < b^*(\beta).$$

And so:

$$\begin{aligned}\frac{dU^T}{d\beta} &= v - c - \rho(1 - b^*)a + (1 - \rho b^*)\frac{d\hat{\rho}}{db^*}\frac{\partial b^*}{\partial \beta}\sigma\Delta \\ &> v - c - \rho(1 - b^*)a \\ &> \rho b'(v - c) = \frac{dU^{NT}}{d\beta}.\end{aligned}$$

This establishes the existence of a unique cut-off β_0.

Appendix B: Contract Negotiation Breakdown when the Seller Has Full Bargaining Power

When the seller has full bargaining power, the buyer receives no surplus when unveiling design A'. The buyer's incentive to acquire information must then stem from a post-information-acquisition rent that he receives when having acquired information and being rather confident that there will be no hold-up. But such a rent can exist only if the seller is uncertain about how much information was acquired. Moreover, the lowest b on the equilibrium support must be equal to zero since no rent can accrue to this "type." The lowest possible offer for design A is therefore $p = v - \rho\Delta$. Consider the following candidate mixed-strategy equilibrium:

- The seller's offer p follows cumulative distribution $F(p)$ on $[v - \rho\Delta, \overline{p}]$;
- The buyer's cognitive effort b follows cumulative distribution $G(b)$ on $[0, \overline{b}]$.

The seller will never charge a price for design A that is refused with probability 1, and so $\overline{p} = v - \hat{\rho}(\overline{b})\Delta$, since "type" $\overline{b}$ is the type most optimistic about the absence of hold-up.

When exerting cognitive effort b, the buyer accepts all offers satisfying $v - p - \hat{\rho}(b)\Delta \geq 0$ and rejects others. His utility is then:

$$U(b) = -T_B(b) + (1 - \rho b)\int_{v-\rho\Delta}^{v-\hat{\rho}(b)\Delta}[v - p - \hat{\rho}(b)\Delta]\,dF(p) \equiv -T_B(b) + (1 - \rho b)R(b).$$

Note that $U(0) = 0$. For the buyer to play a mixed strategy, it must be the case that $U(b) = 0$ on $[0, \overline{b}]$; or

$$-T_B'(b) - \rho R(b) - (1 - \rho b)\frac{d\hat{\rho}}{db}\Delta F(v - \hat{\rho}(b)\Delta) = 0.$$

Using $d\hat{\rho}/db \equiv -\rho(1 - \rho)/(1 - \rho b)^2$, and letting $B(p)$, an increasing function, be defined by

$$p \equiv v - \hat{\rho}(B(p))\Delta,$$

and

$$K(b) = T_B'(b)(1 - \rho b) + T_B(b)\rho$$

(with $K' > 0$ and $K(0) = 0$),

$$F(p) = \frac{K(B(p))}{\rho(1-\rho)\Delta}. \tag{A1}$$

Using $F(\overline{p}) = 1$, the upper bound of the support, $\overline{b}$, is given by

$$[1 - \rho\overline{b}]T_B'(\overline{b}) + \rho T_B(\overline{b}) = \rho(1-\rho)\Delta. \tag{A2}$$

Note that $\overline{b} < b^*$. We must further show that $U'(b) \le 0$ for $b \ge \overline{b}$. For such values, the buyer always trades, at average price $p^e = E[p]$. Because, for $b > \overline{b}$,

$$U(b) = -T_B(b) + (1 - \rho b)[v - p^e - \hat{\rho}(b)\Delta],$$

$$U'(b) = -T_B'(b) - \rho[v - p^e] + \rho = -T_B'(b) + T_B'(\overline{b}) < 0.$$

Let us now turn to the determination of $G(b)$. Let $\pi(p)$ denote the seller's expected profit, conditional on design A:

$$H_0\pi(p) = \int_{\hat{\rho}^{-1}((v-p)/\Delta)}^{\overline{b}} [p - c + \hat{\rho}(b)\Delta](1 - \rho b)\, dG(b),$$

where $H_0 \equiv \int_0^{\overline{b}}(1 - \rho b)\, dG(b)$ is the probability of design A.

This can be rewritten as a function of the cut-off

$$\tilde{b} \equiv \hat{\rho}^{-1}\left(\frac{v-p}{\Delta}\right):$$

$$H_0\pi(\tilde{b}) = \int_{\tilde{b}}^{\overline{b}} \left[v - c - [\hat{\rho}(\tilde{b}) - \hat{\rho}(b)]\Delta\right](1 - \rho b)\, dG(b).$$

For π to be constant on $[0, \overline{b}]$, it must be the case that

$$-\frac{d\hat{\rho}}{d\tilde{b}}\Delta\left[\int_{\tilde{b}}^{\overline{b}}(1 - \rho b)\, dG(b)\right] - (v - c)(1 - \rho\tilde{b})g(\tilde{b}) = 0. \tag{A3}$$

Let $H(\tilde{b}) \equiv \int_{\tilde{b}}^{\overline{b}}(1 - \rho b)\, dG(b)$, and $x \equiv ((1 - \rho)\Delta)/(v - c)$.

The distribution G admits an atom at $b = \overline{b}$. Let y denote this atom. Then

$$H_0 = (1 - \rho\overline{b})y + \int_0^{\overline{b}^-} (1 - \rho b) g(b)\, db = (1 - \rho\overline{b})y + x\int_0^{\overline{b}^-} \frac{\rho}{(1 - \rho b)^2} H(b)\, db,$$

or

$$H_0 = (1 - \rho\overline{b})y + x\left[\frac{y}{1 - \rho\overline{b}} - H_0 + 1 - y\right]. \tag{A4}$$

Furthermore,

$$\ln\frac{H_0}{H(\overline{b})} = x\left[\frac{1}{1 - \rho\tilde{b}} - 1\right], \tag{A5}$$

yielding in particular (for $\tilde{b}$ converging to $\overline{b}$):

$$\ln\frac{H_0}{y} = x\left[\frac{1}{1 - \rho\overline{b}} - 1\right]. \tag{A6}$$

We thus obtain two equations (A4 and A6) with two unknowns (H_0 and y), yielding thereafter $H(\tilde{b})$ from (A5) and then $g(\tilde{b})$ from (A3). Like in the two-cognition-level case studied in the text, the distribution of "net prices" $p - v$ is independent of the gains from trade parameters v and c. By contrast, gains from trade $v - c$ affect the distribution G of cognition levels through the parameter x.

Appendix C: Symmetric Hold-Up

Let

$$h \equiv \sigma(\Delta - a) \quad \text{and} \quad k \equiv \beta(\Delta - a)$$

denote the hold-up stakes,

$$\theta_B(b,z) \equiv \frac{\rho(1 - b)}{\rho(1 - b) + \mu(1 - z) + (1 - \rho - \mu)},$$

$$\theta_S(b,z) \equiv \frac{\mu(1 - z)}{\rho(1 - b) + \mu(1 - z) + (1 - \rho - \mu)},$$

denote the posterior probabilities and

$$\theta(b,z) \equiv \theta_B(b,z) + \theta_S(b,z).$$

In a pure-strategy equilibrium (b^*, z^*), the transaction price $p = p(b^*, z^*)$ for design A is given by

$$\beta[v - c - \theta(b^*, z^*)a] = v - c - \theta_B(b^*, z^*)[a + h] + \theta_S(b^*, z^*)k, \tag{A7}$$

and the transaction costs by

$$T_B'(b^*) = \rho[a + h - \theta_B\sigma\Delta + \lambda_S\beta\Delta] = \rho[(1 - \theta_B\sigma + \theta_S\beta)a + (1 - \theta_B)h + \theta_S k] \tag{A8}$$

and

$$\mathcal{T}_S'(z^*) = \mu[(1 - \theta_S\beta + \theta_B\sigma)a + (1 - \theta_S)k + \theta_B h]. \tag{A9}$$

Appendix D: Formal Analysis of Section IIIB

This Appendix makes the heuristic analysis of Section IIIB more formal.

Degradation: Section IIIB argued that if the buyer can degrade the good delivered to him and later renegotiate to restore its full value, preventing degradation requires that the seller must get "enough" of the gains from renegotiation, thus inducing the buyer to engage in (possibly excessive) cognition.

Consider four ex post states of nature:

ω_1: The initial design A is appropriate and delivers the full utility v;

ω_2: Design A is appropriate, but yields only $v - \Delta$ as the buyer has degraded it and thereby enjoyed private benefit ξ. The seller can, at adjustment cost a, restore the value to v;

ω_3: Design A' is appropriate, but this was not identified ex ante. The seller can, at adjustment cost a, raise value from $v - \Delta$ to v;

ω_4: The appropriate design A' was contracted for ex ante and yields value v.

State ω_3 has probability $\rho(1 - b)$, state ω_4 probability ρb, and state ω_1 probability $1 - \rho$ in the absence of degradation and zero in case of degradation. Let $U_B(\omega)$ denote the buyer's ex post utility in state ω (the ex post utility is the utility obtained when the initial contract is implemented and perhaps renegotiated; it includes transfers, but does not include (sunk) cognition costs and benefit from degradation).

What can be implemented in state of nature ω in general depends on the description of feasible actions and payoffs in that state of nature. We will just assume that the states of nature ω_2 and ω_3 are identical in terms of von Neuman–Morgenstern (VNM) utility functions. We now show that *even if* the contract can ex post induce any levels of utility $U_B(\omega_i)$ it wants (but necessarily satisfying $U_B(\omega_2) = U_B(\omega_3)$), controlling cognition is inconsistent with preventing degradation (inducing seller investment in the second illustration).

Note that for any mechanism that does not induce degradation:

$$U_B(\omega_1) \geq U_B(\omega_2) + \xi,$$

(the buyer must be willing not to degrade design A when learning that it is appropriate),

$$U_B(\omega_4) = \beta(v - c),$$

(negotiation under symmetric awareness of A'),

$$U_B(\omega_2) = U_B(\omega_3),$$

(ex post VNM payoffs are the same in states ω_2 and ω_3, even though these states have different origins),

$$U_B(\omega_4) - U_B(\omega_3) = \frac{T_B'(b)}{\rho},$$

(buyer's incentive compatible cognition), and

$$[1 - \hat{\rho}(b)]U_B(\omega_1) + \hat{\rho}(b)U_B(\omega_3) = \beta[v - c - \hat{\rho}(b)a],$$

(ex ante bargaining over design A).

Combining these conditions, implementable levels of cognition consistent with the absence of degradation satisfy

$$\frac{T_B'(b)}{\rho} = [1 - \hat{\rho}(b)]\xi + \hat{\rho}(b)\beta a.$$

Let $\hat{b}$ be the efficient level of cognition:

$$\frac{T_B'(\hat{b})}{\rho} = a.$$

Cognition is necessarily excessive if

$$[1 - \hat{\rho}(\hat{b})]\xi > [1 - \hat{\rho}(\hat{b})\beta]a.$$

Finally, note that degradation is inefficient if $a + d > \xi$, where d is the negative externality of degradation on the seller (see the main text).

Seller investment: The analysis is very similar to that of degradation. For notational simplicity, we assume that $a = 0$, so efficient cognition is $\hat{b} = 0$. States ω_3 and ω_4 are as above. The other two ex post states are described as follows:

ω_1: Design A is the best design and yields $v - \Delta$ (there is nothing to be done);

ω_2: Provided that the seller has already sunk private cost $K > 0$, design A, which yields $v - \Delta$, can costlessly be turned (with the seller's complicity) into design A' and thereby deliver v to the buyer.

State ω_1 has probability μ and state ω_2 probability $1 - \rho - \mu$ (if the state is ω_2 but the seller fails to invest, the ex post state reverts to ω_1). Again, ex post payoffs $U_B(\omega_i)$ do not include sunk cognition and investment costs.

One has:

$$U_B(\omega_2) = U_B(\omega_3),$$

(the VNM payoff functions are the same in states ω_2 and ω_3),

$$\frac{T_B'(b)}{\rho} = U_B(\omega_4) - U_B(\omega_3),$$

(privately optimal cognition, assuming $b \geq 0$),

$$U_B(\omega_4) = \beta(v - c),$$

(bargaining over design A'),

$$\left[\frac{1 - \rho - \mu}{1 - \rho b}\right][U_S(\omega_2) - U_S(\omega_1)] \geq K$$

$$\Leftrightarrow \left[\frac{1 - \rho - \mu}{1 - \rho b}\right][U_B(\omega_1) - U_B(\omega_2) + \Delta] \geq K,$$

(seller investment is incentive compatible), and

$$\frac{\mu}{1 - \rho b} U_B(\omega_1) + \frac{1 - \rho - \mu}{1 - \rho b} U_B(\omega_2) + \frac{\rho(1 - b)}{1 - \rho b} U_B(\omega_3) = \beta\left(v - c - \frac{\mu}{1 - \rho b}\Delta - K\right),$$

(ex ante bargaining over design A).

Combining these conditions yields

$$\frac{T_B'(b)}{\rho} = \left[\frac{\mu}{1 - \rho - \mu} + \frac{\beta}{1 - \rho b}\right]K - (1 - \beta)\frac{\mu\Delta}{1 - \rho b}.$$

In particular if μ is small, cognition is above the efficient level.

REFERENCES

Aghion, Philippe, and Patrick Bolton. 1987. "Contracts as a Barrier to Entry." *American Economic Review*, 77(3): 388–401.

Aghion, Philippe, Mathias Dewatripont, and Patrick Rey. 2004. "Transferable Control." *Journal of the European Economic Association*, 2(1): 115–38.

Aghion, Philippe, and Benjamin Hermalin. 1990. "Legal Restrictions on Private Contracts Can Enhance Efficiency." *Journal of Law, Economics, and Organization*, 6(2): 381–409.

Allen, Franklin, and Douglas Gale. 1992. "Measurement Distortion and Missing Contingencies in Optimal Contracts." *Economic Theory*, 2(1): 1–26.

Anderlini, Luca, and Leonardo Felli. 1994. "Incomplete Written Contracts: Undescribable States of Nature." *Quarterly Journal of Economics*, 109(4): 1085–1124.

Anderlini, Luca, and Leonardo Felli. 1999. "Incomplete Contracts and Complexity Costs." *Theory and Decision*, 46(1): 23–50.

Anderlini, Luca, Leonardo Felli, and Andrew Postlewaite. 2006. "Should Courts Always Enforce What Contracting Parties Write?" http://econ.lse.ac.uk/staff/lfelli/papers/Parametric.pdf.

Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Incentive Activity: Economic and Social Factors*, ed. Richard R. Nelson, 609–26. Princeton: Princeton University Press.

Bajari, Patrick, and Steven Tadelis. 2001. "Incentives versus Transaction Costs: A Theory of Procurement Contracts." *RAND Journal of Economics*, 32(3): 387–407.

Battigalli, Pierpaolo, and Giovanni Maggi. 2002. "Rigidity, Discretion, and the Costs of Writing Contracts." *American Economic Review*, 92(4): 798–817.

Bernheim, B. Douglas, and Michael D. Whinston. 1998. "Incomplete Contracts and Strategic Ambiguity." *American Economic Review*, 88(4): 902–32.

Bolton, Patrick, and Antoine Faure-Grimaud. 2005. "Thinking Ahead: The Decision Problem." National Bureau of Economic Research Working Paper 11867.

Bolton, Patrick, and Antoine Faure-Grimaud. 2007. "Satisficing Contracts." Unpublished.

Che, Yeon-Koo, and Donald B. Hausch. 1999. "Cooperative Investments and the Value of Contracting." *American Economic Review*, 89(1): 125–47.

Crémer, Jacques. 1995. "Arm's Length Relationships." *Quarterly Journal of Economics*, 104(2): 275–95.

Crémer, Jacques, Fahad Khalil, and Jean-Charles Rochet. 1998a. "Contracts and Productive Information Gathering." *Games and Economic Behavior*, 25(2): 174–93.

Crémer, Jacques, Fahad Khalil, and Jean-Charles Rochet. 1998b. "Strategic Information Gathering before a Contract Is Offered." *Journal of Economic Theory*, 81(1): 163–200.

Dewatripont, Mathias, and Jean Tirole. 2005. "Modes of Communication." *Journal of Political Economy*, 113(6): 1217–38.

Diamond, Douglas W. 1993. "Bank Loan Maturity and Priority when Borrowers Can Refinance." In *Capital Markets and Financial Intermediation*, ed. Colin Mayer and Xavier Vives, 46–68. Cambridge: Cambridge University Press.

Dye, Ronald A. 1985. "Costly Contract Contingencies." *International Economic Review*, 26(1): 233–50.

Edlin, Aaron S., and Stefan Reichelstein. 1996. "Holdups, Standard Breach Remedies, and Optimal Investment." *American Economic Review*, 86(3): 478–501.

Ellison, Glenn. 2005. "A Model of Add-on Pricing." *Quarterly Journal of Economics*, 120(2): 585–637.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, 121(2): 505–40.

Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy*, 94(4): 691–719.

Guasch, J. Luis, Jean-Jacques Laffont, and Stephane Straub. 2006. "Renegotiation of Concession Contracts: A Theoretical Approach." *Review of Industrial Organization*, 29(1–2): 55–73.

Hart, Oliver D., and Bengt Holmström. 1987. "The Theory of Contracts." In *Advances in Economic Theory: Fifth World Congress*, ed. Trumann F. Bewley, 71–155. Cambridge: Cambridge University Press.

Hart, Oliver D., and John Moore. 1988. "Incomplete Contracts and Renegotiation." *Econometrica*, 56(4): 755–85.

Hart, Oliver D., and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy*, 98(6): 1119–58.

Hart, Oliver D., and John Moore. 1999. "Foundations of Incomplete Contracts." *Review of Economic Studies*, 66(1): 115–38.

Hart, Oliver D., and John Moore. 2008. "Contracts as Reference Points." *Quarterly Journal of Economics*, 123(1): 1–48.

Hermalin, Benjamin E. 2002. "Adverse Selection, Short-Term Contracting, and the Underprovision of On-the-Job Training." http://www.bepress.com/bejeap/contributions/vol1/iss1/art5.

Hirshleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review*, 61(4): 561–74.

Holmström, Bengt. 1982. "Managerial Incentive Problems. A Dynamic Perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck*, 209–30. Helsinki: Swedish School of Business.

Joskow, Paul L. 1987. "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets." *American Economic Review*, 77(1): 168–85.

Klein, Benjamin. 2002. "The Role of Incomplete Contracts in Self-Enforcing Relationships." In *The Economics of Contracts: Theories and Applications*, ed. Eric Brousseau and Jean-Michel Glachant, 59–71. Cambridge: Cambridge University Press.

Levin, Jonathan, and Steven Tadelis. 2007. "Contracting for Government Services: Theory and Evidence from U.S. Cities." National Bureau of Economic Research Working Paper 13350.

Macaulay, Stewart. 1963. "Non Contractual Relations in Business." *American Sociological Review*, 28(1): 55–70.

MacLeod, Bentley and James Malcomson. 1993. "Investments, Holdup, and the form of Market Contracts." *American Economic Review*, 83(4): 811–37.

Martimort, David, and Salvatore Piccolo. 2008. "The Strategic Value of Incomplete Contracts for Competing Hierarchies." http://idei.fr/vitae.php?i=46&lang=fr#id4.

Maskin, Eric. 1999. "Nash Equilibrium and Welfare Optimality." *Review of Economic Studies*, 66(1): 23–38.

Maskin, Eric, and John Moore. 1999. "Implementation and Renegotiation." *Review of Economic Studies*, 66(1): 39–56.

Maskin, Eric, and Jean Tirole. 1999a. "Two Remarks on the Property-Rights Literature." *Review of Economic Studies*, 66(1): 139–49.

Maskin, Eric, and Jean Tirole. 1999b. "Unforeseen Contingencies and Incomplete Contracts." *Review of Economic Studies*, 66(1): 83–114.

Nöldeke, Georg, and Klaus M. Schmidt. 1995. "Option Contracts and Renegotiation: A Solution to the Hold-up Problem." *RAND Journal of Economics*, 26(2): 163–79.

Novak, Sharon, and Scott Stern. 2007. "How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry." National Bureau of Economic Research Working Paper 13235.

Rey, Patrick, and Bernard Salanié. 1990. "Long-Term, Short-Term and Renegotiation: On the Value of Commitment in Contracting." *Econometrica*, 58(3): 597–619.

Rey, Patrick, and Bernard Salanié. 1996. "Long Term, Short Term, and Renegotiation: On the Value of Commitment with Asymmetric Information." *Econometrica*, 64(6): 1395–1414.

Segal, Ilya. 1999. "Complexity and Renegotiation: A Foundation for Incomplete Contracts." *Review of Economic Studies*, 66(1): 57–82.

Segal, Ilya, and Michael D. Whinston. 2002. "The Mirrlees Approach to Mechanism Design with Renegotiation (with Applications to Hold-up and Risk Sharing)." *Econometrica*, 70(1): 1–45.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics*, 69(1): 99–118.

Spier, Kathryn E. 1992. "Incomplete Contracts and Signaling." *RAND Journal of Economics*, 23(3): 432–43.

Tirole, Jean. 2007. "Bounded Rationality and Incomplete Contracts." Unpublished.

UK National Audit Office. 2003. *Delivering Better Value For Money from the Private Finance Initiative*. London: The Stationary Office.

Välilä, Timo. 2005. "How Expensive Are Cost Savings? On the Economics of Public-Private Partnerships." *EIB Papers*, 10(1): 94–119.

Williamson Oliver E. 1975. *Markets and Hierarchies: Analysis of Antitrust Implications*. New York: Free Press.

Williamson Oliver E. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. New York: Free Press.

This article has been cited by:

1. PATRICK BOLTON, ANTOINE FAURE-GRIMAUD. 2010. Satisficing Contracts. *Review of Economic Studies* **77**:3, 937-971. [CrossRef]
2. Leon Yang Chu, David E. M. Sappington. 2010. Contracting with private knowledge of signal quality. *The RAND Journal of Economics* **41**:2, 244-269. [CrossRef]
3. Kvaløy Ola, Olsen Trond E.. 2009. Endogenous Verifiability and Relational ContractingEndogenous Verifiability and Relational Contracting. *American Economic Review* **99**:5, 2193-2208. [Abstract] [View PDF article] [PDF with links]
4. Antonio Estache, , Liam Wren-Lewis. 2009. Toward a Theory of Regulation for Developing Countries: Following Jean-Jacques Laffont's LeadToward a Theory of Regulation for Developing Countries: Following Jean-Jacques Laffont's Lead. *Journal of Economic Literature* **47**:3, 729-770. [Abstract] [View PDF article] [PDF with links]
5. Surajeet Chakravarty, W. Bentley MacLeod. 2009. Contracting in the shadow of the law. *The RAND Journal of Economics* **40**:3, 533-557. [CrossRef]
6. Xeni Dassiou, Jon Stern. 2009. Infrastructure Contracts: Trust and Institutional Updating. *Review of Industrial Organization* **35**:1-2, 171-216. [CrossRef]