Robust Incentive Contracts

by

BIRGER WERNERFELT*

Considering a principal–agent model in which the difficulty of the agent's action is better known *ex interim* than *ex ante*, we compare two contracting regimes; one with commitment to an *ex ante* negotiated contract, and one with an *ex interim* negotiated contract. The *ex ante* contract cannot have too strong incentives, but attempts to negotiate a stronger *ex interim* contract may result in bargaining failure. The relative efficiency of the two contracting regimes therefore depends on parameter values. The argument can be interpreted as an analysis of the trade-off between weak incentives in the firm and the possibility of unsuccessful negotiations in the market. (JEL: D 2, L 2)

1 Introduction

The paper compares more or less frequently negotiated incentive contracts in a dynamic environment. Given a series of tasks of varying difficulty, a single long-term contract is burdened by extra risk, but if the players try to negotiate a sequence of short-term contracts, they will occasionally fail to reach agreement. The efficiency and sustainability of the two contracting regimes is found to depend on how the tasks differ and how often they change. The comparison can be interpreted as the trade-off between weak incentives in the firm and the costs of bargaining in the market.

The first step of the argument is made in the context of a single-period principal–agent model in which the agent has to perform an *ex ante* unknown "ideal" task. The difficulty of the ideal task is *ex ante* unknown, but as it is identified, both players get private and public information about its difficulty. We compare the most efficient contracts from two different regimes: in the *spot contracting* regime the players try to negotiate a contract after the ideal task has been identified (*ex interim*), and in the *robust contracting* regime the parties are constrained to an *ex ante* negotiated contract. The advantage of spot contracting is that more information is known at the time of contracting, so that the incentives can be tailored to the task. The problem is that the players negotiate the contract under asymmetric information and thus risk

* The paper has benefited from Steven Tadelis's very detailed reading of an earlier draft, as well as comments by two anonymous referees and seminar participants at the University of California at Berkeley.

suffering bargaining failures. The pros and cons of robust contracting are the reverse. Negotiation takes place *ex ante*, before any informational asymmetries are realized, but the contract cannot depend on the difficulty of the ideal task. Comparing the two regimes, we therefore find that incentives are strongest under spot contracting, while robust contracting is more likely to implement the ideal task. In terms of efficiency, the result is, roughly, that robust contracting is more attractive when the players have more private and less common information about the difficulty of specific tasks.

We proceed to imbed the model in a dynamic setting in which the agent has to perform a sequence of different ideal tasks. In this context, we find conditions under which the players can sustain a commitment to the robust contracting regime, as well as conditions under which the prospect of future bargaining will enhance the efficiency of spot contracting. Both contracting regimes are found to be asymptotically efficient. We go on to suggest that the dynamic version of the model speaks to the theory of the firm by illuminating the trade-off between weak incentives in settings where one player "follows orders" and the costs of bargaining in the market.

Contracts between manufacturers and salespeople illustrate the basic trade-offs in the model. Over time, the manufacturer will often want to realign salesperson territories and change the set of products sold. If the salesperson is an employee, these adjustments are normally handled without any change in the compensation contract; the salesperson still gets a salary and a percentage of sales as commission. In contrast, if the salesperson is an independent representative, all changes in territory and many changes in products will trigger renegotiation of the contract. These renegotiations occasionally fail, leading to dissolution of the relationship and/or expensive legal action (NOVICK [1988, chapters 11f.]). On the other hand, independent representatives will normally have higher commission rates than employees (KOTLER [2000, p. 498]).

The paper has some relation to the literature on commitment. In particular, FUDENBERG, HOLMSTROM, AND MILGROM [1990] also compare the performance of a single long-term agency contract with that of a sequence of short-term contracts. They show that the former can be better if there is asymmetric information at the time of recontracting, exactly the case we are looking at. However, the possibility of failed negotiations does not play a role in their analysis.

The most closely related work is that of BAJARI AND TADELIS [2001] and TADELIS [2002], who also show that robust incentives cannot be too strong when there is possibility of bargaining failure. In the former paper, the central endogenous variable is the probability of renegotiation as determined by the extent to which an incentive contract is complete. This is then compared with an alternative contract that is completely flat and therefore can be costlessly renegotiated. Based on these results, TADELIS [2002] assumes that stronger contracts cause renegotiation costs to go up. In both cases, the idea is that incentive strength affects bargaining costs. In the present paper, the causality goes the other way. Bargaining costs are incurred on a per-renegotiation basis, but a more frequently renegotiated contract can be stronger because more relevant information is known at the time of writing.

We next derive the central result of the paper in a very simple one-period setting; the dynamic extension is discussed in section 3. Sections 4 and 5 contain a discussion of limitations and an interpretation of the results as speaking to the theory of the firm.

2 *Static Model*

A seller may create value for a buyer by exerting effort on one of a large number of possible tasks. Only the ideal task has value, and the identity of this task is *ex ante* unknown. We use the subscript j to indicate a specific task, and introduce *ex interim* information in a very simple way by assuming that output is given by

$$x_j = e_j + \varepsilon_{js} + \varepsilon_{jp} + \varepsilon_{jb} + \varepsilon_t, \tag{1}$$

where e_j is effort, ε_{js}, ε_{jp}, and ε_{jb} are task-specific difficulty parameters, and ε_t is noise.[1] *Ex ante*, when the players negotiate in the robust contracting regime, neither ε_{js}, ε_{jp}, ε_{jb}, nor ε_t is known. However, *ex interim*, when they negotiate in the spot contracting regime, the seller knows ε_{js}, the buyer knows ε_{jb}, and both know ε_{jp}. The noise is only realized *ex post*, after the seller has chosen and expended effort.

We assume that ε_{js}, ε_{jp}, ε_{jb}, and ε_t are independently distributed as $N(0, \sigma_s^2)$, $N(0, \sigma_p^2)$, $N(0, \sigma_b^2)$, and $N(0, \sigma_t^2)$, respectively. The seller's cost of effort is $e_j^2/2$, and if he gets payments w, his utility is $-\exp[-\gamma(w - e_j^2/2)]$, where $\gamma > 0$ is the seller's risk attitude. The buyer is risk-neutral, and each unit of the ideal task is worth 1 to her, while nonideal units are worth nothing. We can therefore invoke the usual arguments to focus on linear contracts of the form $w_j(x_j) = \alpha_j x_j + \beta_j$ in the spot contracting regime and $w(x_j) = \alpha x_j + \beta$ in the robust contracting regime (Holmstrom and Milgrom [1987]). The parameters α and β have very distinct roles in this contract. Since α determines how much the seller's payment is affected by output, it measures the strength of incentives in the model. Furthermore, because the utility function allows us to summarize the utility of the risk-averse player by its monetary certainty equivalent, any split of surplus can be implemented by an appropriate value of β. We therefore assume that the buyer selects $\alpha_j(\alpha)$, while the players negotiate over $\beta_j(\beta)$ to determine the amount of surplus the seller can expect.

We still need to specify how the players negotiate over the β's. This is not problematic in the robust contracting regime, since negotiations there take place *ex ante* when the players have symmetric information. Things are more difficult in the spot-market case, since we know from Myerson and Satterthwaite [1983] that two-sided incomplete information implies that there has to be some expected loss of surplus. A sea of different extensive forms could possibly govern the spot-market negotiations, and one modeling choice would be to specify a particular

[1] An alternative, perhaps more appealing formulation is one in which the buyer's valuation and the seller's costs are imperfectly known *ex ante*. Such a formulation yields similar results.

negotiation game form and express the outcome as a probabilistic function of the realization of the players' private information. However, in order to help focus the attention on the main trade-offs, we will abstract from these complications and assume that the distribution of outcomes is independent of the players' private information. That is, *we assume that spot-market negotiations fail with exogenous probability $\lambda > 0$, and otherwise result in a β that is independent of the realizations of ε_{js} and ε_{jb}*. While this in effect is equivalent to selecting a less-than-second-best reduced form, it does not confound our comparison of the two regimes. First, the expected losses from bargaining failure in the spot contracting regime remain positive, while they are zero in the robust contracting regime. Secondly, if negotiated contracts do depend on the players' private information, the uncertainty associated with spot contracts is reduced, allowing the players to use stronger incentives. But our simplification is defensible because we find that the incentives resulting from it still are stronger than those in the robust contracting regime. In the Appendix, we nevertheless offer a more general comparison of the two regimes without the simplifying assumptions.

In the context of this model, the two contracting regimes can be more specifically defined and analyzed as follows.

Spot Contracting. At the start of the game, the ε_{js} and ε_{jp} associated with the ideal task are revealed to the seller, while the buyer learns the realizations of ε_{jp} and ε_{jb}. The players then proceed to negotiate over $w_j(x_j)$. Per our simplifying assumption, these negotiations fail with probability λ. After agreement on a contract, the seller chooses a level of effort, output is observed, and payments are made. Neither gets any payoff without a contract.

In this regime, the seller's certainty equivalent payoff is given by

$$\alpha_j(e_j + \varepsilon_{js} + \varepsilon_{jp}) + \beta_j - e_j^2/2 - \gamma\alpha_j^2\left(\sigma_b^2 + \sigma_t^2\right)/2. \tag{2}$$

So he sets $e_j^s = \alpha_j$, and if bargaining succeeds, the negotiated fixed payment is

$$\beta_j^s = -\alpha_j^2/2 - \alpha_j\varepsilon_{jp} + \gamma\alpha_j^2\left(\sigma_b^2 + \sigma_t^2\right)/2 + \pi, \tag{3}$$

where π is the seller's expected surplus and we rely on our simplifying assumption. Given this, the buyer's expected payoff is

$$\alpha_j + \varepsilon_{jp} + \varepsilon_{jb} - \alpha_j^2/2 - \gamma\alpha_j^2\left(\sigma_b^2 + \sigma_t^2\right)/2 - \pi, \tag{4}$$

and she therefore sets

$$\alpha_j^s = 1/\left(1 + \gamma\left[\sigma_b^2 + \sigma_t^2\right]\right). \tag{5}$$

Because negotiations may fail, no task is implemented with probability λ.

Robust Contracting. Before the ideal task is identified, the players negotiate a contract $w(x_j)$. Because this contract is negotiated before any asymmetric information is revealed, the negotiation succeeds with probability one. After negotiations, the ε_{js} and ε_{jp} associated with the new ideal task are revealed to the seller, while the buyer learns the realizations of ε_{jp} and ε_{jb}. At this point the players can neither renegotiate the contract nor exit the relationship. Instead, the buyer asks the seller to work on the ideal task, the latter chooses a level of effort, output is observed, and payments are made.

A robust contract gives the seller an expected certainty-equivalent payoff of

$$\alpha e_j + \beta - e_j^2/2 - \gamma\alpha^2\left(\sigma_s^2 + \sigma_p^2 + \sigma_b^2 + \sigma_t^2\right)/2. \tag{6}$$

So he will set $e_j^r = \alpha$, and the negotiated fixed payment is

$$\beta^r = -\alpha^2/2 + \gamma\alpha^2\left(\sigma_s^2 + \sigma_p^2 + \sigma_b^2 + \sigma_t^2\right)/2 + \pi. \tag{7}$$

Given this, the buyer's expected payoff is

$$\alpha - \alpha^2/2 - \gamma\alpha^2\left(\sigma_s^2 + \sigma_p^2 + \sigma_b^2 + \sigma_t^2\right)/2 - \pi. \tag{8}$$

She therefore sets

$$\alpha^r = 1/\left(1 + \gamma\left[\sigma_s^2 + \sigma_p^2 + \sigma_b^2 + \sigma_t^2\right]\right), \tag{9}$$

and the regime implements the ideal task with probability one.

Comparing the two contracting regimes, we see that the incentives are stronger under spot contracting, while robust contracting implements more of the ideal tasks. Because the robust contracts are written when the seller faces more unresolved uncertainty, incentives are riskier and thus more costly.

To compare total surplus under the two regimes, we first maintain the (unreasonable) assumption that the probability of bargaining failure (λ) is independent of the extent of asymmetric information (σ_s^2, σ_b^2). In this case (5) and (9) suggest that spot contracting is more efficient when the seller has more important private information about task-specific difficulty (σ_s^2), when there is more public information about task-specific difficulty (σ_p^2), and when the buyer has less important private information about task-specific difficulty (σ_b^2). Under the realistic assumption that the probability of bargaining failure is larger when there is more asymmetric information, the effect of increases in the seller's private information (σ_s^2) becomes ambiguous. So we can conclude that *robust contracting is more efficient when the difficulty of tasks appears less similar to the buyer, but more similar to the public.* Intuitively, spot contracting does not relieve the seller of the risk coming from the buyer's private information, but does relieve him of the risk coming from the *ex interim* public information.

We will now look at a dynamic version of the model to discuss the players' incentives to break the robust contract as well as the possibility that they treat spot-contract bargaining as a repeated game.

3 *Dynamic Extension*

The static analysis of the robust contracting regime was based on the assumption that the players refrain from renegotiating the contract. When robust contracting is most efficient in a static setting, this absence of renegotiation can possibly be justified in a dynamic version by appeal to an implicit contract in the form of a subgame-perfect equilibrium of the repeated game. The idea is that the players aim to maximize their discounted payoffs from infinitely repeated play of a stage game that at time τ looks as follows:

(i) Ex Interim Revelation of Information. The $\varepsilon_{js\tau}$ and $\varepsilon_{jp\tau}$ associated with the ideal task are revealed to the seller, while the buyer learns the realizations of $\varepsilon_{jp\tau}$ and $\varepsilon_{jb\tau}$.

(ii) Either Player May Insist on Renegotiation. The seller selects $S_\tau \in \{0, 1\}$, and the buyer selects $B_\tau \in \{0, 1\}$. If both select 0, play continues with the contract $w_{\tau-1}(x)$ used in the previous stage, and the players will have the option of using this contract again in stage $\tau + 1$. If either or both select 1, they proceed to negotiate over a new contract $w_\tau(x)$. If these negotiations succeed, the players will have the option of using the new contract again in stage $\tau + 1$. If the negotiations fail, there is no further activity in stage τ, but the players have the option of reverting to the contract $w_{\tau-1}(x)$ in stage $\tau + 1$.

(iii) Production and Payoff. The seller chooses a level of effort, output is observed, and payments are made.

If robust contracting is most efficient in a static setting, it is not hard to show that there exist equilibria in which the players aim to sustain a commitment not to renegotiate and thus stay in the robust contracting regime. It is enough to look at a simple case in which they employ trigger strategies prescribing permanent reversal to spot contracting after any attempt at renegotiation. With these strategies the players will be able to stay in the robust contracting regime until one of them is faced with a sufficiently extreme realization of task difficulty in the form of a low $\varepsilon_{js\tau} + \varepsilon_{jp\tau}$ for the seller or a high $\varepsilon_{jp\tau} + \varepsilon_{jb\tau}$ for the buyer. As we know from the folk theorem, deviations from the commitment not to renegotiate become increasingly rare as the between-stage discount rate decreases.

The folk theorem also tells us that there are many other equilibria of the repeated game, including some in which the seller commits to a particular level of effort. We will now perform an asymptotic analysis to show that *the first best can be achieved in either contracting regime, as the frequency of change grows very high.* To model situations in which the ideal task changes more or less frequently, we hold constant the mean and variance of "yearly" output and look at the effects of having $n > 1$ periods per "year." In this case the interperiod discount factor is the nth root of the "year-to-year" discount factor.

Expressed in these terms, the slope of the robust contract is

$$\alpha^r = 1/\big(1 + \gamma\big[\sigma_s^2 + \sigma_p^2 + \sigma_b^2 + \sigma_t^2\big]/n\big). \tag{10}$$

Since the limit of this is 1, we conclude that robust contracting becomes asymptotically sustainable and first best as the frequency of adjustments grows. Three forces help enhance the efficiency and sustainability of robust contracting as adjustments occur more often. Low interadjustment discount rates make it easier to uphold an implicit contract not to renegotiate; the reduced standard deviation of each adjustment makes extreme realizations less likely and allows the use of a closer to first-best contract.

The analysis of repeated spot contracting is different because of the possibility that bargaining strategies can depend on actions in past bargains. This means that the players can enhance the performance of spot contracting. In particular, it may be possible to reduce the probability of bargaining failure by playing strategies that allow the players to pool some of the incentive constraints over several bargaining occasions (LEVIN [2003]). So, as the interperiod discount rate goes to zero, bargaining failures should vanish so that $\lambda \to 0$. Furthermore, there is a direct effect on static equilibrium contracts mirroring that in (10), such that with n "yearly" adjustments

$$\alpha_j^s = 1/\left(1 + \gamma\left[\sigma_b^2 + \sigma_t^2\right]/n\right). \tag{11}$$

For spot contracting, there are then two forces helping to enhance its efficiency as adjustments occur more often: Low interadjustment discount rates make bargaining more efficient, and the reduced standard deviation of each adjustment allows the use of a closer to first-best contract.

Since both regimes asymptotically can implement the first best, the arguments in the present paper do not help us choose one or the other. However, this does mean that frequent adjustments render the choice irrelevant. When changes occur with very high frequency, it may be necessary to worry about the additional communication–bargaining activity demanded by spot contracting. If the players anyway are going to agree on a contract that differs very little from the most recent one, is it hardly worth spending time discussing it (WERNERFELT [1997]).

4 *Limitations*

The two contracting regimes compared in the previous sections are obviously not the only candidates. A particularly interesting alternative is a mechanism in which the seller can select from an *ex ante* designed menu of contracts after he has received his private information. The advantage of such an arrangement is that it allows the players to avoid negotiating under asymmetric information and thus the risk of bargaining failure. However, because the scheme cannot offer the seller less risk than spot contracting and the contracts have to differ in terms beyond the intercept, the truth-telling constraints force some distortions on the incentives. So while a menu-based contracting regime may be more efficient in some regions of the parameter space, it does not dominate either of the two we look at. (It is also interesting that we see so few real-life examples from this class of contracting regimes.)

Consistent with the sales-force example from the Introduction, the model assumes the availability of a single scale on which all (possibly ideal) tasks can be measured. In many cases this seems like a fair assumption, because agents often perform a rather narrow range of tasks. For example, the tasks could consist of sewing different models of clothes or washing different windows. On the other hand, there are clearly other examples, such as secretarial work, in which the natural units of different tasks are very heterogeneous. In such cases, the only feasible robust contract is flat. While such contracts obviously have poor incentive properties, they may still be preferred to spot contracts, especially when combined with some subjective measures.

5 *Interpretation*

This paper has identified a new relationship between incentives and bargaining costs by showing how more frequent negotiation allows stronger incentives. If we interpret the robust contracting regime as an employment relationship, the paper becomes part of a small but growing literature that defines the employment relationship by the absence of bargaining over adjustments (WERNERFELT [1997], [2002]; BOLTON AND RAJAN [2001]; SIMESTER AND KNEZ [2002]). In contrast to models in which the players would like to, but cannot, make contracts depend on minutiae (e.g., MACLEOD AND MALCOMSON [1989]), this literature defines the firm as a self-enforcing agreement in which contracts *by choice* are not adjusted on a case-by-case basis. Within this definition of the firm, the paper explains that employees face weaker incentives than independent contractors because their contracts carry additional risk by virtue of not reflecting all available information. This contrasts with other explanations based on the definition that the employee does not own the productive assets. This definition implies that he cannot successfully bargain for a large share of the surplus (GROSSMAN AND HART [1986]), cannot be compensated for hard-to-measure additions to residual claims (HOLMSTROM AND MILGROM [1994]), and should not be tempted to abuse the assets too much (WILLIAMSON [1985, p. 132]).

It is also interesting to compare our findings with SIMON's [1951] argument that employment is more attractive when the variance in the cost of tasks is smaller. He makes an implicit supergame argument and relies on the possibility that the employee may quit if faced with a very adverse cost realization. He does not allow this in the market. We are looking at the polar opposite case by assuming that the players always honor the robust (employment) contract, but may fail to reach agreement in the market. If the probability of failed negotiation depends very steeply on the extent of private information, we then get the opposite result, that the market is more efficient if there is less variance in the components of difficulty about which the players are privately informed. In the case of public information, however, our model agrees with Simon's, since less variance in the components of difficulty about which both players are informed adds to the relative efficiency of employment by limiting the associated decay of incentives.

Appendix: Proof that Spot Contracting May Be Inferior to Robust Contracting

Our simplifying assumptions make spot contracting less efficient in two ways: first, because bargaining failure in the most efficient game forms will be more likely when gains from trade are smaller, and secondly because negotiation will reveal some of the players' private information, allowing them to use stronger incentives. However, we know from MYERSON AND SATTERTHWAITE [1983] that there have to be some inefficient bargaining failures and thus some expected loss of surplus. We will denote the latter by $\phi > 0$. Furthermore, no matter how much information is revealed in negotiation, the seller will always face uncertainty about ε_t. The negotiated spot contract can therefore never have a slope that is steeper than $\alpha_j^{s\prime} = 1/(1 + \gamma\sigma_t^2)$. So the incentive advantages of spot contracting are at most the difference between this and $\alpha^r = 1/(1 + \gamma[\sigma_s^2 + \sigma_p^2 + \sigma_b^2 + \sigma_t^2])$, while the disadvantages from bargaining failures are $\phi > 0$. It follows that robust contracting is better for sufficiently small values of $\sigma_s^2 + \sigma_p^2 + \sigma_b^2$. *Q.E.D.*

References

BAJARI, P., AND S. TADELIS [2001], "Incentives versus Transaction Costs: A Theory of Procurement Contracts," *RAND Journal of Economics*, 32, 387–407.

BOLTON, P., AND A. RAJAN [2001], "The Employment Relation and the Theory of the Firm: Arm's Length Contracting vs. Authority," Working Paper, Department of Economics, University of Princeton, NJ.

FUDENBERG, D., B. HOLMSTROM, AND P. MILGROM [1990], "Short-Term Contracts and Long-Term Agency Relationships," *Journal of Economic Theory*, 51, 1–31.

GROSSMAN, S., AND O. D. HART [1986], "The Costs and Benefits of Ownership," *Journal of Political Economy*, 97, 691–719.

HOLMSTROM, B., AND P. MILGROM [1987], "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55, 303–328.

–– AND –– [1994], "The Firm as an Incentive System," *American Economic Review*, 84, 972–991.

KOTLER, P. [2000], *Marketing Management: Analysis, Planning, Implementation, and Control*, Millennium Edition, Prentice Hall: Upper Saddle River, NJ.

LEVIN, J. D. [2003], "Relational Incentive Contracts," *American Economic Review*, 93, 835–857.

MACLEOD, W. B., AND J. M. MALCOMSON [1989], "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica*, 57, 447–480.

MYERSON, R. B., AND M. A. SATTERTHWAITE [1983], "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

NOVICK, H. J. [1988], *Selling through Independent Reps*, Amacom: New York.

SIMESTER, D. I., AND M. KNEZ [2002], "Direct and Indirect Bargaining Costs and the Scope of the Firm," *Journal of Business*, 75, 283–302.

SIMON, H. A. [1951], "A Formal Theory of the Employment Relationship," *Econometrica*, 19, 293–305.

TADELIS, S. [2002], "Complexity, Flexibility, and the Make-or-Buy Decision," *American Economic Review*, Papers and Proceedings, 92, 433–437.

WERNERFELT, B. [1997], "On the Nature and Scope of the Firm: An Adjustment-Cost Theory," *Journal of Business*, 70, 489–514.

– – [2002], ”Why Should the Boss Own the Assets?” *Journal of Economics and Management Strategy*, 11, 473–485.

WILLIAMSON, O. E. [1985], *The Economic Institutions of Capitalism*, Free Press: New York.

Birger Wernerfelt
MIT Sloan School of Management
50 Memorial Drive
Cambridge, MA 02142
USA
E-mail:
bwerner@mit.edu