Trust and Discretion in Agency Contracts*

Nabil I. Al-Najjar[†] Ramon Casadesus-Masanell[‡]

June 2001
This version: April 10, 2002

Abstract. We extend the standard agency framework to allow for two aspects largely overlooked in the literature, namely complex information and trust. We show that the introduction of these considerations leads to agency contracts that are incomplete. In addition to endogenously generating contractual incompleteness, our model illustrates that complex information and trust lead to sharply different conclusions compared to the standard agency model. In particular, not all the relevant, valuable information about the agent's action choice is incorporated in the equilibrium contract; and, even when inference is perfect, the principal may only be able to induce the agent to undertake low effort. We conclude that one main function of agency contracts is to protect the agent from possible opportunistic behavior of the principal.

* We thank Emre Ozdenoren, participants at the Summer School in Economic Theory 'Contracts and Markets with Boundedly Rational Agents' held in Venice, and seminar participants at University of Michigan and Harvard Business School for insightful comments. The second author gratefully acknowledges the support of the Harvard Business School division of research.

† Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. <al-najjar@nwu.edu>

‡ Harvard Business School, Boston MA 02163. <rmasanell@hbs.edu>

1. Introduction

The economic literature on agency studies the design of optimal contracts in settings where the parties are opportunistic and the informational content of the observable and verifiable signals on the agent's choice of action can be captured by *complete* contracts, in the sense that wage schedules specify one and only one payment for each conceivable outcome.[1] In this setting, the agent does not need to trust the principal because the contract confers no discretion to choose compensation after the outcome has been realized. In standard agency, trust and contractual incompleteness are worthless.

Nevertheless, real agency contracts are characterized by a substantial degree of incompleteness. The Law of Agency, for example, contains provisions informing how the agent should be paid when the contract leaves compensation unspecified.[2] According to the Restatement of the Law, when there is no written agreement for a definite amount, the agent should receive 'a fair value for his services.' In the standard model of principal-agent relationships this is a vacuous statement because the wage schedule *is* the contract.

We propose three variations to the standard model and show that in our more general framework, contractual incompleteness arises endogenously: we introduce a notion of complex information; we require that contracts are written down; and we allow for some degree of trustworthiness on the part of the principal.

Most of the existing literature assumes that the environment is sufficiently stable so that complete contracts can capture all information on the agent's choices

[1] See, for example, Berhold (1971), Ross (1973), Holmstrom (1979), Shavell (1979), Grossman and Hart (1983).

[2] See American Law Institute (1957) *Restatement of the Law, Second, Agency*, pg. 343: "If the amount of compensation is not otherwise agreed upon, as where no specific amount is stated and there is no customary rate for the services, it is inferred that, in a transaction in which some compensation is due, the parties have agreed that the agent is to receive the reasonable value of his services. In determining this, evidence of what other agents receive for similar services is competent, together with other factors, including the reputation of the agent, the skill with which the work is done, and the difficulty or danger of the task."

embedded in the observable signals: the information contained in the outcomes varies sufficiently smoothly from outcome to outcome so that contracts can be made fully sensitive to every minor detail of the signals.

We depart from this 'smoothness' assumption and allow information to vary drastically from outcome to outcome. We define complexity in these terms: roughly, information is *complex* if 'close by' outcomes lead to very different inferences on the agent's choice of action. Intuitively, an environment is complex if 'details matter.' Complexity arises not because of the number of possible outcomes per se is large, but because of the number of independent pieces of information that must be taken into account in writing optimal contracts. In contrast, in a *simple* environment the informational content of outcomes does not vary much for close by outcomes, just as in the standard model.

In order for contracts to be enforceable, they must be written down unambiguously. To formally model this idea, we describe the outcomes in terms of their objective features. We let the outcome space to have a product structure where each level of 'depth' represents a different feature. Contracts specify payments contingent on sets of features. Since contracts have to be of finite length, they will necessarily condition on *finitely many* features.

One main consequence of the coupling of complexity and finite definability is that no complete contract can fully capture the richness and variability of the information contained in the outcomes. Under complexity and finite definability, complete contracts fall short from being second best because useful information must necessarily be left out from the contract.

We also depart from previous work on agency in that we do not focus exclusively on contracts as the only possible method of governance. We allow for *trustworthiness* on the part of the principal. The principal is trustworthy if she can commit not to take advantage of the agent when the contract grants her discretion. A trustworthy principal compensates the agent for the full cost of effort, for the risk derived from the stochastic nature of outcomes, and, when applicable, for the risk ensuing from the belief that the principal may in fact be opportunistic.

In contrast, an opportunistic principal pays the agent as little as possible.

The presence of trust in a world of complex information and finite contracts, allows for the emergence of *incomplete* contracts. A contract is incomplete if it is set-valued. When the contract is incomplete, the principal has discretion to choose the wage after the outcome has been realized, within the limits set by the contract. We show that equilibrium contracts consist of a simple lower bound below which the principal is not allowed to pay and are open above. Trust restores flexibility when written contracts cannot capture the full variability of the environment.

We show that trust is necessary for the working of incomplete contracts and that there is a monotone relationship between the principal's level of trustworthiness and her expected profit.[3] Trust reduces the agent's risk bearing and, thus, it results in larger total surplus of the relationship. Since the principal makes a take-it-or-leave-it offer, she appropriates the gains.

In our setting, Holmstrom (1979)'s sufficient statistic result breaks down: when information is complex, some informative signals have to be left out of the the written contract. In addition, the standard model predicts that whenever the densities associated with each level of effort have different supports, the first best will be achieved by sufficiently punishing the agent if the realized outcome is impossible under high effort. In our framework, this result also breaks down because it may be too hard to write such a contract. If contracts could be of infinite length or, alternatively, if there was sufficient trust, both these results would be restored.

In conventional agency models, contracts trade off incentives and risk. In our model, one main function of written contracts is protecting the agent from the possibility that the principal may be opportunistic. If the principal was completely trustworthy, then the optimal contract would be no contract at all. It is the principal's imperfect trustworthiness what makes written contracts necessary in the first place.

[3] Throughout the paper we will refer to the principal as *she* and to the agent as *he*.

1.1. Literature review

The first work formalizing complexity considerations in contracting is Dye (1985). He considers a contracting problem with exogenously given cost of contracting on each contingency. This approach was more recently taken by Battigalli and Maggi (2000). One difficulty with such approach is that conclusions are sensitive to the way the cost of contracting is modelled. It is also hard using this approach to account for the possibility that initially complicated contingent actions/compensations can be codified in standard-form contracts with minimal writing cost. In other words, if the contingent actions are sufficiently routine in nature, one would expect dramatic reduction in writing costs using, for instance, standard-form contracts. One would thus like to capture incompleteness that *persists* after contracting parties exhausted all reasonable possibilities of reducing contracting cost by hard-wiring repetitive aspects that may initially appear as a complicated contingent action.

In response to these difficulties, Anderlini and Felli (1994) proposed a different approach where contracts are restricted only to be finitely defined. This formulation, which we largely follow in this paper, may be viewed as a way to capture the limit of positive but vanishingly small cost of contracting, without making any specific assumptions about the exact form of this cost.[4]

A problem with Anderlini and Felli's model is that they obtain an approximation result: in their setting the first best contract can be approximated arbitrarily closely by finitely-defined contracts. This clearly undermines the potential of such model to understand incomplete contracting.

An alternative approach has been proposed by Al-Najjar, Casadesus-Masanell, and Ozdenoren (2002). These authors introduce complexity in a model of the continuum resembling that of Anderlini and Felli (1994). In a somewhat different contracting setting they show that in a complex environment, the ex post optimal

[4]In Section 9 we present a formal argument along these lines.

course of action cannot be approximated arbitrarily closely by finitely-defined ex ante complete contracts. Their model of complexity has the attractive feature that, just as in the original Anderlini and Felli (1994), it is built directly on the continuum and does not need to resort to discrete (countable) versions of the state-space to introduce a meaningful notion of complexity that circumvents the approximation result.

Recently, Al-Najjar, Anderlini, and Felli (2000) have suggested a different approach to break Anderlini and Felli (1994)'s result, using a modelling device introduced in Al-Najjar (1999). The idea there, as in the present paper, is to use a discrete state space instead of the continuum. The discrete state space allows more concreteness in modelling complex objects such as functions and correspondences. Krasa and Williams (1999) provide yet another avenue to evade Anderlini and Felli's approximation theorem.

Al-Najjar, Casadesus-Masanell, and Ozdenoren (2002) present a behavioral foundation for decision making in a complex environment. This foundation can be easily extended to justify behavior under complexity in the discrete setting that we consider.

We adopt the feature structure of Anderlini and Felli (1994) and the discrete state space model introduced in Al-Najjar (1999). However, the economic questions we address are very different. Our focus is on trust and discretion in the classic principal-agent problem.[5] The central features of our model, namely moral hazard and delegation under asymmetric information, are absent in the works above which focus on co-insurance problems. In contrast to the endogenous model of agent's trustworthiness in Casadesus-Masanell (1999), we study the role of exogenous trust of the principal for the lay out of optimal incentive contracts in complex environments.

[5] For excellent surveys on the principal-agent literature, see Levinthal (1988), Sappington (1991), Gibbons (1998), and Prendergast (1999).

1.2. Outline of the paper

Section 2 reviews the standard agency model and characterizes optimal contracts. In Section 3 we introduce three variations to the standard model. First, we reinterpret the set of verifiable outcomes as collections of objective features; second, we allow for incomplete contracts; and third, we let the principal be trustworthy. We show that in the standard model, none of these variations affect the shape of the optimal contract. In Section 4 we construct the large, infinite outcome space as the limit of a sequence of finite outcome spaces. In Section 5 we introduce a model of complexity. In Section 6 we revisit the agency model, now with the richer large outcome space, and study optimal complete contracts in a setting that allows for complexity. We show that complete contracts are 'coarse' in that they do not vary as much with the signals as what standard theory would predict. In Section 7 we allow the principal to write incomplete contracts. We show that contractual incompleteness may be optimal when the outcome space is infinite. Section 8 presents a fully worked out example.

In Sections 6, 7, and 8 we examine optimal contracts without explicit consideration of contract-drafting costs. Thus, in our model contractual incompleteness does not arise because of some ad hoc cost of writing down contractual clauses. What generates incompleteness are drastic changes in the informational content of outcomes that agree on a large but finite number of features. In Section 9, however, we motivate our structure in terms of positive, but vanishingly small contract-drafting costs. Section 10 presents concluding remarks.

2. The Standard Agency Problem

Consider the classic contracting problem between a risk neutral principal and a risk averse agent. The agent's action choice (his effort level) determines, probabilistically, an outcome x from some finite outcome space

$$X_N = \{x_1, x_2, ..., x_{\#X_N}\}.$$

Outcomes are assumed to be observable and verifiable so that enforceable contracts can be written on them.

The agent takes one of two actions (or effort levels) $e \in \{H, L\}$. Effort is unobservable and costly. The cost of e is c_e, with $c_H > c_L$. We normalize $c_L = 0$. Each effort level induces a probability distribution on the set of outcomes, X_N. If the agent takes action e, then the outcomes are distributed according to density $\pi_e : X_N \to \mathbb{R}_+$. That is, $\pi_e(x_n)$ is the probability of x_n given e.

Incentives to work are provided by means of a take-it-or-leave-it contract $\alpha : X_N \to \mathbb{R}$ specifying a payment for each outcome, $x \in X_N$.

The contract is designed so that the agent is willing to enter the relationship: if $\bar{U}$ is the agent's reservation utility, then α must satisfy the *participation constraint*:

$$E_H[u(\alpha(x))] - c_H \geq \bar{U}. \tag{1}$$

The contract also needs to be *incentive compatible*; it must give incentives to take the action that the principal wants to implement. In particular, if the principal wants the agent to choose $e = H$, then α must be such that

$$E_H[u(\alpha(x))] - c_H \geq E_L[u(\alpha(x))] - c_L. \tag{2}$$

Given contract α and outcome x, the principal obtains benefit $B(x) - \alpha(x)$. We denote the expected benefit to the principal by V. To make the problem interesting, we assume that the function B is such that the principal would like to implement high effort. The expected cost to the principal of contract α that implements $e = H$ is

$$E_H \alpha(x). \tag{3}$$

In summary, the principal solves the problem:

$$\begin{array}{ll} & \min_\alpha E_H \alpha(x) \\ & \textit{subject to} \\ \text{(IC)} & H \in \arg\max_e E_e u(\alpha(x)) - c_e \\ \text{(IR)} & E_H u(\alpha(x)) - c_H \geq \bar{U} \end{array} \tag{P1}$$

The timing of the game is as follows: The principal makes a take-it-or-leave-it contract offer. The agent accepts or rejects the offer. If the contract is accepted, the agent exerts effort $e \in \{H, L\}$. If rejected, the game ends. An outcome x is realized according to π_e. Payments are made as prescribed by the contract α.

Optimal contracts are designed by minimizing (3) subject to (1) and (2). There is a trade-off between incentives and risk bearing. On the one hand, because risk bearing reduces the total surplus and the principal is risk neutral and the agent risk averse, the principal would ideally like to assume all risk by offering a constant wage schedule. On the other hand, if the wage schedule is constant, the agent chooses $e = L$ because $c_L < c_H$. Thus, the principal needs to carefully trade incentives and risk off when designing the incentives scheme.

Standard results by Holmstrom (1979), Grossman and Hart (1983), and others show that in equilibrium, constraints (1) and (2) are binding. The first order conditions characterizing the optimal contract are

$$\frac{1}{u'(\alpha(x_n))} = \lambda + \mu \left[1 - \frac{\pi_L(x_n)}{\pi_H(x_n)} \right] \text{ for all } x_n \in X_N, \tag{4}$$

where λ and μ are the Kuhn-Tucker multipliers associated to the Individual Rationality and Incentive Compatibility constraints, respectively. Because $\mu > 0$, the presence of moral hazard is costly: the principal would be strictly better off under symmetric information.

Condition (4) implies that the equilibrium wage varies as a function of the likelihood ratio, $\frac{\pi_L(x_n)}{\pi_H(x_n)}$. The likelihood ratio indicates the precision by which outcome x_n signals that $e = H$. The lower the likelihood ratio, the stronger the signal that the agent chose high effort.

3. Verifiability and incompleteness in in standard agency

In the standard agency model, contracts are costlessly enforceable, reflecting the view that the underlying outcomes of the agency relationship can be unambiguously described to the contracting parties and to outside enforcement mechanisms, such as a court. In this section we model explicitly the sense in which outcomes are verifiable. We then show that trust, discretion, and contractual incompleteness play no role in standard agency. The purpose of this exercise is to highlight the role of complex information which we introduce in later sections, and in particular, how complexity gives rise to incomplete agency contracts where trust and discretion are valuable.

3.1. Outcomes and verifiability

To model verifiability, we adapt the framework of Anderlini and Felli (1994) to agency relationships. Formally, we assume that there is an infinite number of binary *features*. Outcome x's j^{th} feature, denoted x^j, takes values in set $\{0, 1\}$. Thus, an outcome may be viewed as a point in the product space $\{0, 1\}^{\mathbb{N}}$.[6]

Intuitively, verifying an outcome is a process of checking combinations of elementary features. Each such feature is unambiguous, in the sense that there is no room for disagreement as to whether a given outcome does or does not have this feature. For example, one feature may represent a profit threshold, so we may use 1 to indicate that the threshold has been exceeded, and 0 otherwise. Another feature may be 'change in market share,' and so on. Define $A_m = \{x \in \{0,1\}^{\mathbb{N}} : x^m = 1\}$. That is, A_m is the set of outcomes whose m^{th} feature is 1. Let $\mathcal{A}$ be the *algebra* of subsets of X generated by $A_m,\ m = 1, 2, \ldots$ We refer to an event $A \in \mathcal{A}$ as *finitely defined* or *finitely describable*, because each such set can be fully pinned down by verifying a fixed finite number of features.

To understand the model, the reader may find it helpful to (somewhat artificially) distinguish between two types of features: those that are payoff relevant, and those that have a purely informational content. Suppose, for concreteness, that the principal's payoff at outcome x is $\sum_{k=1}^{K} 2^{-k} x^k$, for some integer $K > 0$. In this case, only features $\{1, \ldots, K\}$ are payoff relevant. Indeed, these features induce a simple linear ordering on outcomes that can be used to introduce monotonicity assumptions (such as the Monotone Likelihood Ratio Condition - MLRC). Assume that the remaining features $\{K+1, \ldots\}$ are payoff irrelevant, in the sense that two outcomes that agree on features $\{1, \ldots, K\}$ generate precisely the same payoff to the principal regardless of the values taken by features $\{K+1, \ldots\}$.

As is well-known, under optimal agency contracts, what the agent is paid at each outcome is not a function of the payoff-relevance of that outcome, but rather it is a function of the informational content of each outcome. The realization of a particular outcome is nothing but an (imperfect) signal of what the agent has done. The signalling value of an outcome may have little or nothing to do with the payoff implications of that outcome. Even if one assumes that the informational content

[6]For convenience, one may view the set of outcomes as the interval [0,1] where one identifies real numbers with their binary expansions. There are, however, important differences. First, $\{0,1\}^{\mathbb{N}}$ involves duplication as (countable many) real numbers may have different equivalent binary expansions. Second, the interval [0,1] includes a metric structure which implicitly assumes a particular ordering of the features. Thus, in [0,1] two outcomes that agree on the *first* K features are very close when K is large. $\{0,1\}^{\mathbb{N}}$ presumes no such ordering by importance.

of payoff-relevant features changes ‘smoothly’ (in the sense of MLRP, for instance), there is no reason to expect that payoff irrelevant features convey information in any simple manner (say, by inducing an implied linear ordering on outcomes). Our analysis hinges on the possibility that these pieces of information may interact in a complex way to determine the inference that can be drawn about the agent’s actions.

The above discussion is meant to illustrate our ideas under a sharp, artificial dichotomy between payoff-relevant and -irrelevant features. Our model is general in that it does not presume such dichotomy.

3.2. Coarse vs. incomplete contracts

There is no generally accepted definition of what an ‘incomplete’ contract means. In the literature, the informal idea of an ‘incomplete contract’ is sometimes introduced as contracts whose provisions do not vary finely with the outcome of the relationship. In other words, these contracts are ‘coarse’ in the sense that they lump together contingencies that a more detailed contract would have treated differently.

On the other hand, one may reasonably argue that coarse contracts are indeed complete. After all, they do specify definitive payment for each and every outcome of the relationship. Intuitively, an important characteristic of incomplete agency contracts is that (perhaps in addition to being coarse) they do not specify exactly what to do in every contingency. Under such contracts, although there may be restrictions on what the parties can and cannot do at each contingency, these restrictions do not completely pin down what the parties have to do. By leaving a *range* of possible outcomes, incomplete contracts are vague about how the relationship is to be resolved at each contingency. This is the definition implicit in the property rights literature, where the ownership of an asset conveys to the owner ‘residual rights of control’ to do whatever he pleases provided doing so is not in violation of other contractual provisions. See Grossman and Hart (1986).

In this paper we adopt a flexible framework that allows the introduction of both coarse and incomplete contracts. To do so, we generalize the wage schedule α to be a correspondence:

$$\alpha : X_N \twoheadrightarrow \mathbb{R}.$$

The interpretation is that for every outcome x, the contract specifies a *range* of possible payments for the principal to choose from *after* x has been realized. Formally,

DEFINITION 1: *A contract is* complete *if it prescribes a single wage for each outcome x. A contract is* incomplete *if it is multiple-valued at some x.*

To formalize the principal's ex post choice, given a contract α we let s_α be a selection of α; that is, $s_\alpha : X \to I\!R$ such that $s_\alpha(x) \in \alpha(x)$ for every x.

EXAMPLE 1: *An example of a contract is $\alpha(x) = [m_1, m_2]$ for all $x \in X_N$ with $m_1 < m_2$. In such a contract, after the agent has chosen the action and outcome x has been realized, the principal offers payment $s_\alpha(x) \in [m_1, m_2]$. This is an incomplete (or discretionary) contract because ex post the principal has some discretion to decide how much to pay the agent.*

EXAMPLE 2: *A second example is*

$$\alpha(x) = \begin{cases} m_1 & \text{if } 0 \le x < \frac{1}{2} \\ m_2 & \text{otherwise.} \end{cases}$$

This contract is complete; it gives no discretion to the principal as it specifies an exact payment for each possible outcome.

3.3. Trust

When the contract α is incomplete (not single-valued), it no longer determines the principal's behavior in the relationship. We allow the principal to be of one of two types *when time comes to compensate the agent*: she may be opportunistic or trustworthy.[7]

- The opportunistic principal pays the agent the lowest allowed under the contract. The selection rule under opportunism is denoted by $\underline{s}_\alpha(x)$.
- The trustworthy principal covers the full cost of effort; she also covers the risk ensuing from the stochastic relation between actions and outcomes; in addition, the trustworthy principal pays a premium to compensate for the ex ante risk that she may turn out to be opportunistic when time comes to remunerate the agent. The selection rule under trust is denoted by $\bar{s}_\alpha(x)$.

To abstract from signaling considerations, we assume that the contract is written before the principal's type is realized. With probability τ she will act trustworthily and with complementary probability she will behave opportunistically.[8]

The expected cost to the principal of contract α that implements $e = H$ is

$$E_H[\tau\bar{s}_\alpha(x) + (1-\tau)\underline{s}_\alpha(x)]. \tag{5}$$

In this model, the agent bears two types of risk. First, for any given level of effort, e, the exact realization of x is random. Second, the type of the principal is unknown: with probability τ she is trustworthy (*i.e.* unwilling to take advantage of the agent) and with probability $1-\tau$ she is opportunistic (*i.e.* self-interest seeker with guile). The modified individual rationality and incentive compatibility constraints are

[7] See Casadesus-Masanell and Spulber (2001) for a detailed analysis on the sources of trust in agency relationships.

[8] We offer two interpretations that can be used to justify eliminating the role of signaling through contract choice. First, contracts may be set by convention, and so do not reflect the type of any particular principal. Second, we may consider a model where the contract may signal the principal's type, but where we focus on a pooling equilibrium of the signalling game. Such equilibrium may be supported by a system of beliefs in which the agent interprets any deviation from the pooling contract as an indication that the principal is opportunistic.

$$\tau E_H[u(\overline{s}_\alpha(x))] + (1-\tau)E_H[u(\underline{s}_\alpha(x))] - c_H \geq \bar{U} \tag{6}$$

and

$$H \in \arg\max_e \{\tau E_e[u(\overline{s}_\alpha(x))] + (1-\tau)E_e[u(\underline{s}_\alpha(x))] - c_e\}. \tag{7}$$

In summary, the principal solves:

$$\begin{array}{ll} & \min_\alpha \tau E_H[\overline{s}_\alpha(x)] + (1-\tau)E_H[\underline{s}_\alpha(x)] \\ & \textit{subject to} \\ \text{(IC)} & H \in \arg\max_e \{\tau E_e[u(\overline{s}_\alpha(x))] + (1-\tau)E_e[u(\underline{s}_\alpha(x))] - c_e\} \\ \text{(IR)} & \tau E_H[u(\overline{s}_\alpha(x))] + (1-\tau)E_H[u(\underline{s}_\alpha(x))] - c_H \geq \bar{U} \end{array} \tag{P1-τ}$$

With these variations, the timing of the game is as follows: The principal makes a take-it-or-leave-it contract offer (which may allow for discretion). The agent accepts or rejects the offer. If the contract is accepted, the agent exerts effort $e \in \{H, L\}$. If rejected, the relationship ends. An outcome x is realized according to the distribution induced by π_e. Nature chooses the type of the principal. If the contract allows for discretion, then the principal chooses a payment in the allowed set, $\alpha(x)$. If she is opportunistic, she chooses $\underline{s}_\alpha(x)$, and if she is trustworthy, she chooses $\bar{s}_\alpha(x)$.

3.4. Incomplete contracts and trust in the standard model

The following proposition shows that in the standard model of agency there is no role for contractual incompleteness, no matter how trustworthy the principal is. In other words, given any feasible incomplete contract, there always is a complete contract that does just as well.

PROPOSITION 1*: In the standard model of agency incompleteness and trust have no value.*

Formally, let V_{P1} be the value of the program (P1) and, for any given value $\tau \in (0,1]$, let V^{τ}_{P1} be the value of the program (P1-τ). Then any α^* that solves (P1) must also solve (P1-τ).

PROOF: Suppose α^* solves (P1). Note that α^* is in the feasible set of (P1-τ) and that $V_{P1}(\alpha^*) = V^{\tau}_{P1}(\alpha^*)$. Towards a contradiction suppose that there is another contract $\hat{\alpha}$ such that $V^{\tau}_{P1}(\hat{\alpha}) \geq V^{\tau}_{P1}(\alpha)$ for all α in the feasible set of (P1-τ) and $V^{\tau}_{P1}(\hat{\alpha}) > V^{\tau}_{P1}(\alpha^*)$. The complete randomized contract $\hat{\alpha}'$ for (P1) that assigns wage $\bar{s}_{\hat{\alpha}}(x)$ with probability τ and $\underline{s}_{\hat{\alpha}}(x)$ with probability $1-\tau$ after outcome x is realized, gives the same payments with the same probabilities than contract $\hat{\alpha}$ for (P1-τ) when the principal is of type τ. Therefore, we must have that $\hat{\alpha}'$ is feasible for (P1) and that $V_{P1}(\hat{\alpha}') = V^{\tau}_{P1}(\hat{\alpha})$. But then,

$$V_{P1}(\hat{\alpha}') > V_{P1}(\alpha^*)$$

but this is impossible because a randomized contract can never achieve higher value to the principal than a non-randomized contract (see Grossman and Hart (1983), Remark 8 following Proposition 13). ■

As a simple corollary notice that the principal's expected payoff is always the same, no matter how trustworthy she is. The reason is that the optimal contract is *complete* and thus trust and discretion are never used.

In the following sections we show that contractual incompleteness emerges endogenously in the limit of a sequence of finite models like the one we just presented.

4. The Agency Problem with a Large Outcome Space

A common intuition in the literature is that as the outcome space becomes large, the contracting problem becomes more complex, and this complexity may lead to some form of incompleteness. For example, Posner (1986) writes:

"[S]ome contingencies, even though foreseeable in the strong sense that both parties are fully aware that they may occur, are so unlikely to occur that the costs of careful drafting to deal with them might exceed the benefits, when those benefits are discounted by the (low) probability that the contingency will actually occur."

In a similar vein, Hart and Moore (1999) observe that:

"In an ideal world the parties would write a contingent contract specifying exactly which good to be delivered in each state. However, if the number of states is very large, such a contract would be prohibitively expensive. So instead the parties will write an incomplete contract."

Although compelling, this suggested link between the size of the state space and the complexity of the contracting problem is not so obvious. For example, agency problems with a continuum of outcomes may admit simple optimal contracts that can be characterized quite easily. See, for example, Holmstrom (1979). The goal of this section is to introduce a model where the large outcome space leads to complexity and incompleteness in the contractual problem.

4.1. Large outcome spaces

Consider a sequence of finite outcome spaces $\{X_N\}_{N=1}^{\infty}$, where the N^{th} outcome space is the standard one introduced in Section 2:

$$X_N = \{x_1, \ldots, x_{\#X_N}\} \subset \{0,1\}^{\mathbb{N}}.$$

The uniform distribution λ_N on X_N assigns to each subset $B \subset \{0,1\}^{\mathbb{N}}$ a measure equal to its relative frequency:

$$\lambda_N(B) \equiv \frac{\#(B \cap X_N)}{\# X_N}.$$

We are interested in the agency problem when N is large. We gain a better insight into the problem by looking at the limiting model of $\{X_N\}$ as N goes to infinity. We denote the outcome space of the limiting model by X and define it to be $X = \cup_{N=1}^{\infty} X_N$. We would also like to have a counterpart on X for the uniform distribution λ_N in the finite models. The following assumption ensures that the limiting model X is 'well behaved' in the sense that it resembles the finite models along the sequence.

Assumption 1*: The sequence of outcome spaces $\{X_N\}_{N=1}^{\infty}$ satisfies:*

1. *$X_N \subset X_{N+1}$ for every N;*

2. *$\lim_{N\to\infty} \lambda_N(A)$ converges for every $A \in \mathcal{A}$;*

3. *There is $\epsilon > 0$ such that $\epsilon < \lim_{N\to\infty} \lambda_N(A_m) < 1-\epsilon$, for all $m \in \{1,2,3,...\}$.*

Example 3*: Sequences of outcome spaces satisfying Assumption 1 can be generated by taking 'typical random draws' from $\{0,1\}^{\mathbb{N}}$. Let P be the probability distribution on $\{0,1\}^{\mathbb{N}}$ such that $P(A_m) = \frac{1}{2}$ for every $m \in \mathbb{N}$ and features are independent. Consider independent random sampling from $\{0,1\}^{\mathbb{N}}$ according to P. Denote the draws by $\{z_1, z_2, ...\}$. Define the N^{th} model as $X_N = \{z_1, \ldots, z_N\}$. The $N+1^{th}$ model, X_{N+1}, is the N^{th} model plus one extra draw: $X_{N+1} = \{z_1, \ldots, z_N, z_{N+1}\}$. Then with probability 1, the sequence of finite outcome spaces thus constructed satisfies Assumption 1.*[9]

[9] That conditions 1 and 3 in the definition are satisfied is obvious. To establish condition 2, fix a set $A \in \mathcal{A}$, and consider the function $1_A : \{0,1\}^{\mathbb{N}} \to \{0,1\}$ that is the indicator function of the set A. Since the draws are independent, by the Strong Law of Large Numbers there is probability

Note that given $A \in \mathcal{A}$, we have, by definition of λ_N

$$\lambda_N(A) \equiv \lambda_N(X_N \cap A) \equiv \frac{\#(X_N \cap A)}{\#X_N}$$

Then, condition 2 in the definition simply states that the relative frequency of a set $A \in \mathcal{A}$ in the N^{th} model settles down to some limiting value. The idea here is that the sequence of outcome spaces, although increasing in size, resemble each other in terms of the distribution of the features in them. Condition 3 is a non-degeneracy condition whose role is to eliminate superfluous features that have zero limiting mass.

LEMMA 1*: For any outcome space X satisfying Assumption 1, there is a finitely additive probability measure λ on 2^X such that:*

$$\lambda(B) = \lim_{N\to\infty} \lambda_N(B) \tag{8}$$

for every $B \subset X$ for which the limit exists.

PROOF: See Appendix. ■

For the rest of the paper we fix an infinite outcome space X constructed as the limit of a sequence of finite outcome spaces satisfying Assumption 1 and we let λ be a finitely additive probability measure on X given by Lemma 1.

4.2. Actions and probabilities

Given X and λ, functions of the form $f : X \to I\!R$ can be integrated in a straightforward way (see Section A.1 in the Appendix). All we need in this paper is the

1 that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} 1_A(z_n) = P(A) \tag{$*$}$$

Since $\mathcal{A}$ is countable, there is probability 1 that (*) is satisfied simultaneously for all $A \in \mathcal{A}$. Taking $X = \{z_1, z_2, z_3 \ldots\}$ to be any typical draw completes the proof.

integral of a *simple* function f (*i.e.*, a function with finite range $\{y_1, \ldots, y_K\}$):

$$\int_X f\, d\lambda \equiv \sum_{k=1}^{K} y_k \ \lambda(\{x : f(x) = y_k\}). \tag{9}$$

That is, the integral of a simple function is the average of its values weighted by their frequencies, $\lambda(\{x : f(x) = y_k\})$. Note that this integral is well-defined for *any* simple function, since every subset of X is measurable.

In the agency context, each action induces a probability distribution on the set of outcomes X. It is convenient to represent this distribution (as we did in the finite case) using a density function π, where π will, of course, depend on the action the agent takes.

The following is a natural extension of the usual definition of a density function:

DEFINITION 2: *A function $\pi : X \to I\!R$ is a density function if $\pi(x) \geq 0$ for every x, and $\int_X \pi(x)\, d\lambda = 1$.*

For example, if

$$\pi(x) = \begin{cases} 2 & \text{if } \frac{1}{2} \leq x \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

then, the probability of the event $A = [0\,, \frac{3}{4}]$ is $\int_0^{\frac{3}{4}} \pi(x)\, d\lambda = \int_{\frac{1}{2}}^{\frac{3}{4}} 2 d\lambda = \frac{1}{2}$.

The central point of our model would be to show that not all densities have to take this simple form. Our model allows a density π as in the following example: π, takes values in $\{0, 2\}$ such that for each set $A \in \mathcal{A}$ with $\lambda(A) > 0$, $\frac{1}{\lambda(A)}\lambda\{x : \pi(x) = 0\} = \frac{1}{2}$.

4.3. Expected payoffs

Given a bounded function $f : X \to I\!\!R$ representing a state-contingent payoff, say, its expected payoff is given by

$$E_\pi f \equiv \int_X f(x)\pi(x)\, d\lambda.$$

4.4. The agency problem restated

The agency problem with a large outcome space may be written just as the standard problem of Section 2:

$$\begin{array}{ll} & \min_\alpha E_H \alpha(x) \\ & \textit{subject to} \\ \text{(IC)} & H \in \arg\max_e E_e u(\alpha(x)) - c_e \\ \text{(IR)} & E_H u(\alpha(x)) - c_H \geq \bar{U} \end{array} \tag{P1}$$

We solve for the optimal contract by minimizing pointwise the expected wage subject to the (IC) and (IR) constraints. As in the finite case, the optimal schedule is characterized by the likelihood ratios (eq. 4).

Let α^* be the solution to Problem (P1).

5. COMPLEXITY IN AGENCY CONTRACTS

In this section we present our notion of complexity. We study the implications of complexity for contract design in Sections 6 and 7.

5.1. Finitely defined contracts

Consider a function $f : X \to \mathbb{R}$ that represents a contract between two parties. It is natural to require that any such contract can be written down in terms of the elementary, objective features. An example of a function that can be written down in terms of elementary features is

$$f(x) = \begin{cases} 0 & \textit{if first feature} = 0 \textit{ and second feature} = 0, \\ 1 & \textit{if first feature} = 0 \textit{ and second feature} = 1, \\ 2 & \textit{if first feature} = 1 \textit{ and second feature} = 0, \\ 3 & \textit{if first feature} = 1 \textit{ and second feature} = 1. \end{cases} \tag{10}$$

More generally, $\mathcal{A}$-measurable functions are the only functions that can unambiguously be written down in terms of elementary features.

For a contract to be enforceable, it must be written down. A difficulty one faces in formalizing the idea that 'contracts must be written down' is where to draw the line between what can and cannot be written down. Are 10 features (2^{10} outcomes) too many? What about 100 features (2^{100} outcomes)? Rather than resort to ad hoc criteria based on the number of features or writing cost, we simply allow the principal to write *any* contract she wants as long as it conditions on a *finite* number of features. Thus, our only requirement is that contracts are *finitely defined*.

DEFINITION 3: *A contract α is* finitely-defined *if it is $\mathcal{A}$-measurable.*

Finitely-defined contracts can unambiguously be written down, communicated, and reproduced. Finitely-defined contracts wind up conditioning on intervals: if J is the index of the last feature included in the contract, then the schedule is effectively dividing the outcome space into 2^J intervals and can possibly assign a different compensation to each one of these intervals.

In the standard agency model of Section 2 with finitely many outcomes, all contracts are finitely defined; thus, they can be written down. However, the solution to the agency problem with a large outcome space (Section 4.4), the optimal contract, α^*, is not necessarily finitely defined.[10]

Thus, we need to impose one additional constraint to problem (P1); namely, that contracts have to be written down. Let $\mathcal{C}$ be the set of all finitely-defined contracts. The finite definability constraint (FD) is

$$\alpha \in \mathcal{C}.$$

The following example illustrates finite definability.

EXAMPLE 4: *Suppose that for $e \in \{H, L\}$, π_e takes values in $\{0, 2\}$ so that for each set $A \in \mathcal{A}$ with $\lambda(A) > 0$, $\frac{1}{\lambda(A)}\lambda\{x : \pi_e(x) = 0\} = \frac{1}{2}$. Further, suppose that for each $A \in \mathcal{A}$ and $k \in \{0, 2\}$, $\{x \in A : \pi_H(x) = k\} \cap \{x \in A : \pi_L(x) = k\} = \emptyset$. Clearly, after observing the outcome, the principal knows with certainty whether the agent chose high or low effort. Contract α^* (as characterized in Section 4.4) gives the agent his reservation utility plus c_H if the observed outcome belongs to the outcome subset generated by $e = H$ and less than his reservation utility plus c_L if the observed outcome comes from $e = L$.*

In this example, contract α^* cannot be written down because on every $A \in \mathcal{A}$ (no matter how tiny A is), there is a 50/50 distribution of outcomes from H and L. In this example, the likelihood ratios characterizing α^* vary more than what

[10] In fact, α^* need not even be the limit of a sequence of finitely defined contracts.

the most finely grained finitely-defined contract can capture. In fact, in this case the optimal finitely-defined contract is a constant payment and only $e = L$ can be implemented.

5.2. Complexity

Complex functions vary more than what finitely defined sets can capture. Formally, let $\mathcal{F}$ be the set of all finitely defined functions from X to $\mathbb{R}$, then

DEFINITION 4: *A function* $g : X \to \mathbb{R}$ *is* complex over $A \in \mathcal{A}$ *if*

$$\inf_{f \in \mathcal{F}} \int_A |f(x) - g(x)| \, d\lambda > 0.$$

A function $g : X \to \mathbb{R}$ *is* simple over A *if it is not complex over* A.

For a function to be complex it should not be possible to approximate it arbitrarily closely by a sequence of finitely-defined functions.[11]

We say that *information is complex* if the likelihood ratio (as a function of x) is complex over some finitely defined set A. Clearly, complex information implies that at least one of the density functions π_e is a complex function itself, but the converse is not necessarily true.

By definition of complexity, the likelihood ratios characterizing the optimal contract are (are not) finitely defined when the environment is not (is) complex. Therefore,

REMARK 1: $\alpha^* \in \mathcal{C}$ if and only if information is not complex over the entire outcome space.

[11] Note that the limit of a sequence of finitely defined functions need not be finitely defined but it is not complex either. This captures the intuition that there is nothing complex about continuous functions such as $g(x) = x$ which are limits of sequences of finitely defined functions. (Continuity of functions on X is defined analogously to continuity of functions on the continuum, $[0, 1]$. That is, g is continuous at x if and only if given $\epsilon > 0$ there exists $\delta > 0$ such that $|g(y) - g(x)| < \epsilon$ for all y satisfying $|y - x| < \delta$.)

6. THE AGENCY PROBLEM REVISITED (I) - COARSE CONTRACTS

We now show that when information is complex, the equilibrium complete contract is coarser than α^*, the optimal complete contract as derived in Section 4.4. Intuitively, even though α^* is in some sense the best possible complete contract, the principal may have trouble writing it down.

We are interested on the solution to Problem (P1) when the additional finite definability constraint is imposed:

$$\begin{array}{ll} & \min_\alpha E_H\alpha(x) \\ & \textit{subject to} \\ \text{(IC)} & H \in \arg\max_e E_e u(\alpha(x)) - c_e \\ \text{(IR)} & E_H u(\alpha(x)) - c_H \geq \bar{U} \\ \text{(FD)} & \alpha \in \mathcal{C} \end{array} \tag{P2}$$

Let $\widetilde{\alpha}$ be the solution to Problem (P2). In this section we characterize $\widetilde{\alpha}$. As we just pointed out, ideally the principal would like to set $\widetilde{\alpha} = \alpha^*$ (see Section 4.4). By Remark 1, when information is simple α^* is the solution to (P2). However, when information is complex α^* does not satisfy (FD). In what follows, we show that under complexity, the optimal finitely defined contract can be obtained by solving a 'modified' version of (P1).

For tractability reasons, we make two technical assumptions on the densities π_e. First, even as we allow for very general density functions, we simplify the analysis by considering only densities that take on finitely many values:

ASSUMPTION 2*: Densities π_e, $e \in \{H, L\}$ take on finitely many values.*

In addition, we restrict attention to densities with averages that settle down on finitely-defined sets:

ASSUMPTION 3*: There is an $\mathcal{A}$-measurable finite partition of X, $\mathcal{P}(X) = \{A^1, ..., A^N\}$, such that if $A^n \in \mathcal{P}(X)$ and $\{A_1^n, A_2^n\}$ is a finitely defined partition of A^n with $\lambda(A_i^n) > 0,\ i \in \{1,2\}$, then*

$$\frac{1}{\lambda(A_1^n)} \int_{A_1^n} \pi_e(x)\, d\lambda = \frac{1}{\lambda(A_2^n)} \int_{A_2^n} \pi_e(x)\, d\lambda$$

for $e \in \{H, L\}$.

Assumption 3 is a regularity condition requiring the existence of a finitely defined partition of the set of outcomes X such that each density function takes values in some constant proportions on the elements of the partition, no matter how finely one keeps further subdividing the members of that partition into smaller finitely defined sets.[12]

We define $\varphi_e(x)$ to be the irreducible averages of $\pi_e(x)$. That is, for $x \in A^n$ with $A^n \in \mathcal{P}(X)$,

$$\varphi_e(x) = \frac{1}{\lambda(A^n)} \int_{A^n} \pi_e(x)\, d\lambda.$$

Clearly, since π is a density, so is φ.[13] We will sometimes refer to φ_e as the 'smoothed' density of π_e. Likewise, $\frac{\varphi_L(x)}{\varphi_H(x)}$ is the 'smoothed' likelihood ratio corresponding to $\frac{\pi_L(x)}{\pi_H(x)}$. Finally, notice that if f is finitely defined, then

$$E_{\pi_e} f(x) = E_{\varphi_e} f(x).$$

EXAMPLE 5*: Suppose that $\mathcal{P}(X)$ consists of two sets $\{A^1, A^2\}$. Suppose that the only difference between the elements of A^1 and A^2 is that on all $x \in A^1$ the first feature is 0 and on A^2 it is 1. On A^1, $\pi_H(x)$ is a 50/50 distribution on values $\frac{1}{4}$ and*

[12] In Al-Najjar, Casadesus-Masanell, and Ozdenoren (2002), we provide a derivation of Assumption 3 from more primitive conditions. The argument there (which is performed in the continuum) relies on the martingale convergence theorem (or the Radon-Nikodym derivative). Unfortunately, with finitely-defined probability measures both the martingale convergence theorem and the Radon-Nikodym derivative may fail and thus we need to assume the settling down of each density π rather than derive it from more basic principles.

[13] $\pi \geq 0 \Rightarrow \varphi \geq 0$ and $1 = \int_X \pi\, d\lambda = \int_{\mathcal{P}(X)} \pi\, d\lambda = \sum_{A \in \mathcal{P}(X)} \int_A \pi\, d\lambda = \sum_{A \in \mathcal{P}(X)} \lambda(A)\varphi(x) = \int_X \varphi\, d\lambda$.

$\frac{1}{2}$ and $\pi_L(x) = 1$ for all x. On A^2, $\pi_H(x)$ is a 50/50 distribution on values $\frac{1}{4}$ and 3 and $\pi_L(x)$ is again 1 for all x. Assume π_H and π_L are uncorrelated. It is easy to check that $\pi_H(x)$ and $\pi_L(x)$ are well defined densities satisfying Assumptions 2 and 3.

In this example, the likelihood ratios are

$$\frac{\pi_L(x)}{\pi_H(x)} = \begin{cases} 50/50 \textit{ distribution on } 4 \textit{ and } 2, \textit{ for } x \in A^1 \\ 50/50 \textit{ distribution on } 4 \textit{ and } \frac{1}{3}, \textit{ for } x \in A^2. \end{cases}$$

Therefore, information is complex (see Section 5.2).

It is easy to derive the 'smoothed' densities φ. On A^1, $\varphi_H(x) = \frac{3}{8}$ and on A^2, $\varphi_H(x) = \frac{13}{8}$. $\varphi_L(x) = 1$ for all $x \in X$. The smoothed *likelihood ratios are*

$$\frac{\varphi_L(x)}{\varphi_H(x)} = \begin{cases} \frac{8}{3}, \textit{ for } x \in A^1 \\ \frac{8}{13}, \textit{ for } x \in A^2. \end{cases}$$

Let $f : X \to \mathbb{R}$ be as in equation (10), a finitely defined function. Then,

$$E_{\pi_H} f(x) = E_{\varphi_H} f(x) = \tfrac{17}{8},$$
$$E_{\pi_L} f(x) = E_{\varphi_L} f(x) = \tfrac{3}{2}.$$

We are now ready to characterize the solution to problem (P2):

PROPOSITION 2*: The solution to (P2) (where expectations are taken using π), coincides with the solution to*

$$\min_{\alpha}\{E\alpha(x) \textit{ subject to (IC) and (IR)}\} \qquad \textit{(P2')}$$

where all expectations are taken using φ.

PROOF: See Appendix. ■

Intuitively, the optimal complete contract will condition on as many finitely defined sets as it is useful to. Assumption 3 guarantees that there is a last feature

j that it is valuable conditioning on. Since the expectations of finitely defined functions are the same under π and φ, a contract satisfying (IC) and (IR) in Problem (P2'), will also satisfy (IC) and (IR) in Problem (P2). Further, since φ_e are finitely defined, the solution to (P2') will also satisfy (FD). The argument is complete by noticing that the objective functions are the same in both problems.

Let $\tilde{\alpha}$ be the solution to (P2) and recall that V is the principal's expected utility. Clearly, if α^* could be written down, $V(\alpha^*) \geq V(\tilde{\alpha})$ with strict inequality when information is complex. This is so because (P2) is (P1) with an additional constraint (finite definability) that is binding whenever the problem is complex.

6.1. Complexity implies coarse complete agency contracts

Under complexity, the agency problem looks very differently if assessed from an ex ante rather than an ex post point of view. Ex post, after a specific outcome has been realized, the principal has a 'clearer' idea as to what has been the agent's action.[14] It is useful to think of Problem (P2) as the 'ex ante problem' and of Problem (P1) as the 'ex post problem.' If information is complex, the optimal ex ante complete contract is coarse: the principal would have liked to offer a much finer incentive schedule but the finite-definability constraint precluded her from doing so.

Writing the contract ex ante, before the outcome is known, requires evaluating *all* possible contingencies in advance, and this is much harder than evaluating ex post the specific contingency that has been realized. Ex ante, the principal would like to introspect and write a complete contract that captures her full knowledge about the agent's actions at each instance. But the contract needs be *written down.* Effectively, having to come up with a 'writable' contract means that many different

[14] In some cases, as in Example 4, the principal may know *exactly* the agent's action after the fact. Of course, our model allows for intermediate cases where the principal, after observing the outcome, has a better understanding as to what has been de agent's action, but she still does not know for sure.

contingencies will have to be bunched together: many different contingencies will have to have the same ex post payments.

The finite definability restriction translates into noise: the contract cannot be as finely grained as the principal would like it to be. The contract ends up being based on the ex ante likelihood ratios that reflect the extra noise added by complexity. The best a complete contract can do is to prescribe payments that are efficient *on average* over the range of contingencies covered by its clauses. However and as we show below, finitely-defined incomplete contracts allow for finer adjustment to the specific contingencies that arise ex post. We turn now to the study of optimal incomplete contracts.

7. The Agency Problem Revisited (II) - Incomplete Contracts

In this section we show that if the principal is trustworthy, then, depending on the complexity of the environment, the optimal contract is incomplete. In contrast to Proposition 1, trust and contractual incompleteness are valuable in a complex world.

We begin by adapting Definition 1 to the infinite outcome space.

Definition 5*: A contract is* complete *if it is complete at almost every x. A contract is* incomplete *if it is incomplete on a set $B \in X$ with $\lambda(B) > 0$.*[15]

[15] It is important not to confuse finite-definability with incompleteness. A contract can be finitely defined but complete. Also, a contract can be not-finitely-defined but incomplete. Finite definability means that contracts need be written down. Incompleteness, on the other hand, confers discretion to the principal to choose ex post, after outcome x has been realized.

The extended model where the principal may be trustworthy is:

$$
\begin{array}{lll}
 & \min_\alpha \tau E_H[\overline{s}_\alpha(x)] + (1-\tau)E_H[\underline{s}_\alpha(x)] & \\
 & \textit{subject to} & \\
\text{(IC)} & H \in \arg\max_e\{\tau E_e[u(\overline{s}_\alpha(x))] + (1-\tau)E_e[u(\underline{s}_\alpha(x))] - c_e\} & \text{(P3)} \\
\text{(IR)} & \tau E_H[u(\overline{s}_\alpha(x))] + (1-\tau)E_H[u(\underline{s}_\alpha(x))] - c_H \geq \bar{U} & \\
\text{(FD)} & \alpha \in \mathcal{C} &
\end{array}
$$

Let α^τ be the solution to (P3). By an argument mimicking that in the proof of Proposition 1, it can be easily shown that if we did not require the contract to be finitely defined, there would be no role for trust and contractual incompleteness. Thus, absent (FD),

$$\alpha^\tau = \alpha^*.$$

By Remark 1, complete contract $\tilde{\alpha}$ (the solution to (P2)) is optimal for (P3) when information is not complex.[16]

7.1. Optimal incomplete contracts

Note that if there is no trust ($\tau = 0$), then, regardless of the complexity of the problem, $\tilde{\alpha}$ is optimal.[17] When the problem is complex and $\tau > 0$, $\tilde{\alpha}$ is generally not optimal for (P3). To see this, suppose that the principal is completely trustworthy ($\tau = 1$), then the ex post optimal schedule α^* can be achieved *de facto* by a fully incomplete contract (or *no contract at all*): give discretion to the principal to choose any payment she wishes ex post, after the outcome has been realized. Because $\tau = 1$, the agent knows that he will be paid according to the ex post likelihood ratios characterizing α^*, $\frac{\pi_L(x)}{\pi_H(x)}$. By construction, this schedule satisfies (IC), (IR), and (FD) in (P3). Because under complexity $V(\alpha^*) > V(\tilde{\alpha})$,

[16] This can be shown by an argument similar to that in the proof of Proposition 2.

[17] As we show below, when $\tau > 0$ but small, the optimal contract may still be $\tilde{\alpha}$.

the principal will be strictly better off with the incomplete contract than with the contract that solves (P2). Thus,

REMARK 2: When $\tau = 1$, the fully incomplete contract achieves the highest expected payoff to the principal.

The following proposition characterizes the solution to (P3) for intermediate levels of trustworthiness ($0 < \tau < 1$).

PROPOSITION 3: *When τ is sufficiently large and information is complex, the contract α^{τ} that solves Problem (P3) is partially incomplete: it gives the principal discretion to choose payment ex post from a set of allowed wages. Moreover, if there is a cost $\delta > 0$ to write contractual clauses, then α^{τ} consists of a lowest bound under which the principal is not allowed to pay and is left open above.*

PROOF: See Appendix. ■

Optimal contracts consist of a finitely defined lower bound and are open above. The lower bound protects the agent from the possibility that the principal may be opportunistic. The openness above allows the trustworthy principal to use her ex post knowledge to compensate the agent. Trust acts as a substitute for written contracts.

Written agency contracts exist to protect the agent from the possible opportunism of the principal as well as to provide incentives for performance. If the principal was completely trustworthy, there would be no point in having written contracts. The agent would feel compelled to work because he would know that the wage would be fair. It is because of imperfect trust that contracts need be written down. When τ is close to 1, incentives come mainly from the expectation that the wage will be fair ex post. However, the agent needs some protection (the lowest bound) because the principal may turn out to be opportunistic. At the other extreme, when τ is close to 0, incentives cannot come from the expectation of a fair wage. In this case, incentives have to be written down.

It is interesting to notice that the optimal contract in Proposition 3 looks very similar to the incomplete contracts *assumed* in recent papers on the interaction between implicit and explicit incentives (see Pearce and Stacchetti (1998) and Baker, Gibbons, and Murphy (1994)). In this sense, our work provides a foundation for the kind of incomplete contracts assumed in those papers. In both Pearce and Stacchetti (1998) and Baker, Gibbons, and Murphy (1994), the agent's incentives derive from payments based on verifiable signals (this corresponds to our finitely defined lower bound) and from implicit promises by the principal of bonuses for good behavior (this corresponds to the payments above the lowest bound that the trustworthy principal grants when there is enough evidence that the agent has worked hard).

7.2. Comparative statics

We now analyze the relationship between the principal's trustworthiness and her expected gains from the agency relationship. The following proposition shows that the principal is better off the more trustworthy she is. Therefore, if the principal could credibly commit to acting in a trustworthy manner, she would.

PROPOSITION 4: *Suppose $\tau > \tau'$, that τ' is large enough so that the construction in Proposition 3 is valid, and that information is complex. Then,*

$$V(\alpha^{\tau}) > V(\alpha^{\tau'}).$$

PROOF: See Appendix. ■

The intuition is clear. As the principal's trustworthiness increases, the agent faces less risk because it is less likely that the principal will use her discretion to take advantage of the agent. As risk bearing diminishes, the principal needs to offer less expected compensation to induce the agent to accept the contract and the principal can capture more of the total surplus generated.

Clearly,

Corollary 1*: Under complex information,*

$$V(\alpha^*) > V(\alpha^\tau) > V(\widetilde{\alpha}).$$

Proof: See Appendix. ■

8. Numerical Example

It is interesting to provide a fully worked out example to illustrate Propositions 3 and 4.

8.1. Setup

Let the agent's vNM utility function be $u(\alpha(x)) = \sqrt{\alpha(x)}$ and

$$\mathcal{P}(X) = \begin{cases} A^1 = \left\{x \in X : x^1 = 0\right\} & (\textit{roughly } [0, \tfrac{1}{2}]) \\ A^2 = \left\{x \in X : x^1 = 1\right\} & (\textit{roughly } [\tfrac{1}{2}, 1]) \end{cases}$$

Assume further that the agent's reservation utility is $\bar{U} = 50$ and that $c_H = 10$ and $c_L = 0$.

Let the densities be as in Example 5: On A^1, $\pi_H(x)$ is a 50/50 distribution on values $\frac{1}{4}$ and $\frac{1}{2}$ and $\pi_L(x) = 1$ for all x. On A^2, $\pi_H(x)$ is a 50/50 distribution on values $\frac{1}{4}$ and 3 and $\pi_L(x)$ is again 1 for all x. Assume π_H and π_L are uncorrelated.

8.2. Ex post optimal complete contract

By using the likelihood ratios computed in Example 5, one can easily compute the ex post optimal complete contract (the solution to (P1)):

$$\alpha^*(x) = \begin{cases} 1,685.66, & \textit{for all } x \in A^1 \textit{ with } \frac{\pi_L(x)}{\pi_H(x)} = 4 \\ 2,881.01, & \textit{for all } x \in A^1 \textit{ with } \frac{\pi_L(x)}{\pi_H(x)} = 2 \\ 1,685.66, & \textit{for all } x \in A^2 \textit{ with } \frac{\pi_L(x)}{\pi_H(x)} = 4 \\ 4,123.10, & \textit{for all } x \in A^2 \textit{ with } \frac{\pi_L(x)}{\pi_H(x)} = \frac{1}{3}. \end{cases}$$

If α^* was implementable, the expected cost to the principal would be $3,663.16$. The problem is, of course, that this contract is not enforceable because it does not describe the outcomes in terms of their features.

8.3. Optimal finitely defined contract in the absence of trust

On A^1, $\varphi_H(x) = .375$ and on A^2, $\varphi_H(x) = 1.625$. $\varphi_L(x) = 1$ on all $x \in X$. Thus, the optimal finitely-defined, complete contract in the absence of trust is

$$\tilde{\alpha}(x) = \begin{cases} 1,156, & \textit{for all } x \in A^1 \\ 4,356, & \textit{for all } x \in A^2. \end{cases}$$

This results in a cost to the principal of $3,756$. Therefore, $V(\alpha^*) > V(\tilde{\alpha})$.

8.4. Optimal incompleteness

We now allow the principal to be trustworthy ($\tau > 0$) and solve (P3).

Notice that within finitely defined set A^1, there are two informational equivalence classes: those outcomes with likelihood ratio 4 and those with likelihood ratio 2. Similarly for A^2. Therefore, the principal needs only compute four numbers: (a) a minimum bound in case the realized outcome falls in A^1, (b) a minimum bound in case x falls in A^2, (c) the wage to be paid if the outcome falls in A^1 and the likelihood ratio suggests that the agent took high effort, and, finally, (d) the wage if the outcome falls in A^2 and the likelihood ratio suggests that the agent took high effort.

This will be an incomplete contract that will give discretion to the principal: If x belongs to A^1, then the principal chooses between (a) and (c) and if it belongs to A^2, then she chooses between (b) and (d). The opportunistic principal will always choose either (a) or (c), but the trustworthy principal will choose 'fairly.'

Suppose that $\tau = \frac{1}{2}$, then the optimal incomplete contract is

$$\alpha^{\tau=\frac{1}{2}}(x) = \begin{cases} \{1,083.52, \quad 2,157.60\}, & \textit{if } x \in A^1 \\ \{3,833.63, \quad 4,767.94\}, & \textit{if } x \in A^2. \end{cases}$$

The expected cost to the principal is $3,735.48$. Thus,

$$V(\alpha^*) > V(\alpha^{\tau=\frac{1}{2}}) > V(\tilde{\alpha}).$$

Finally, note that if there is a cost $\delta > 0$ of writing each number in the contract, the optimal incomplete contract will specify two numbers only:

$$\alpha^{\tau=\frac{1}{2}}(x) = \begin{cases} [1,083.52, \ \infty), & \textit{if } x \in A^1 \\ [3,833.63, \ \infty), & \textit{if } x \in A^2. \end{cases}$$

Figure 1 illustrates Proposition 4. It shows that the expected cost to implement $e = H$ is a decreasing function of the principal's trustworthiness.

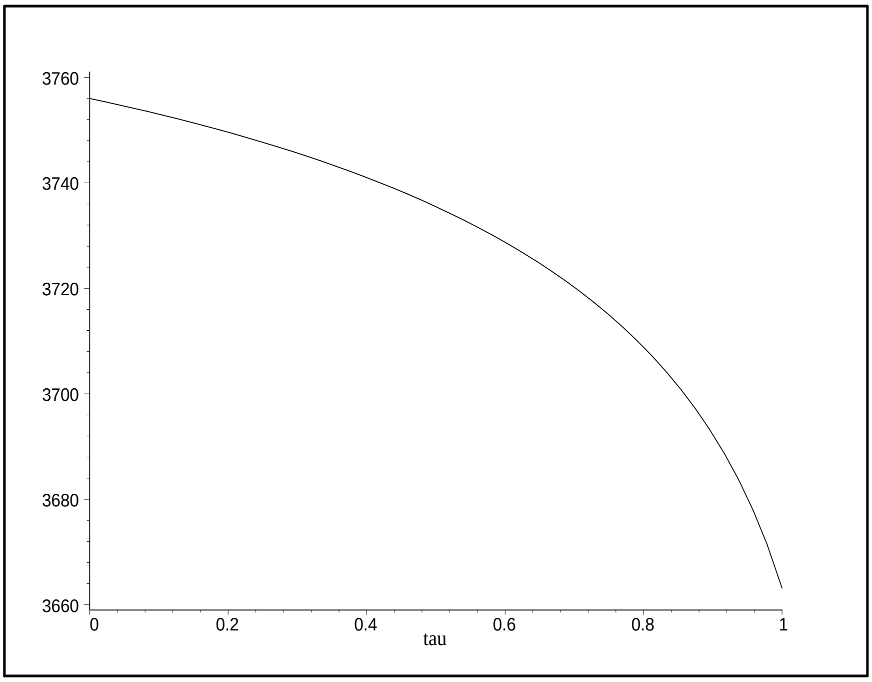

Figure 1: Expected cost of implementing $e = H$ as a function of τ

9. Drafting and Contracting Cost

In this section, we reinterpret the finite definability assumption by explicitly considering the costs of drafting contracts. It is uncontroversial that finite definability is not too restrictive an assumption: written contracts need be of finite length. A more interesting question is whether or not finite definability is too permissive. Is it too much to allow the principal to write contracts of *any* finite length?

In what follows we show that our model can be interpreted as the limit of a sequence of finite outcome agency problems where the principal can make the contract as finitely grained as she needs to but where there is a cost to writing contractual clauses that goes to zero as the size of the outcome space increases. Limits are taken in such a way that the dimension of the outcome space goes to infinity *faster* than drafting costs go to zero.

9.1. Modelling complexity cost

For any finite set $J = \{j_1, \ldots, j_l\}$ of indices, let $\mathcal{A}_J$ denote the algebra generated by features $j_1, \ldots, j_l$. To every finitely defined contract α corresponds an algebra Σ_α.[18] Define the *dimension* of α by:

$$l(\alpha) \equiv \min_J \{J : \Sigma_\alpha \subset \mathcal{A}_J\}.$$

Roughly, the contract α cannot be fully described using less than $l(\alpha)$ features. For example, the constant contract (specifying payment that does not depend on the outcome) has dimension 0, and a contract that conditions payment on a single feature has dimension 1.

Assumption 4*: A* contracting cost function *is any function that assigns to each contract α a value $C(\alpha) \geq 0$ such that:*

$$C(\alpha) \to \infty \text{ as } l(\alpha) \to \infty. \tag{11}$$

Example 6*: A simple example of a cost function is $C(\alpha) = c \cdot l(\alpha)$, with $c > 0$.*

In this example, adding an extra feature to a contract has cost $c > 0$ and as the length of the contract increases without bound so does the cost of writing the contract.

We now restate problem (P3) by removing the FD constraint and by adding complexity cost:

[18]That is, Σ_α is the smallest algebra of sets with respect to which α is measurable. Clearly, Σ_α is a finite algebra.

$$\begin{array}{ll} & \min_\alpha \tau E_H[\overline{s}_\alpha(x)] + (1-\tau)E_H[\underline{s}_\alpha(x)] + C(\alpha) \\ & \textit{subject to} \\ \text{(IC)} & H \in \arg\max_e\{\tau E_e[u(\overline{s}_\alpha(x))] + (1-\tau)E_e[u(\underline{s}_\alpha(x))] - c_e\} \\ \text{(IR)} & \tau E_H[u(\overline{s}_\alpha(x))] + (1-\tau)E_H[u(\underline{s}_\alpha(x))] - c_H \geq \bar{U} \end{array} \tag{P4}$$

Let α_C^τ be the solution to (P4). Assumption 4 implies that the only contracts relevant for problem (P4) are finitely defined contracts (as contracts that depend on infinitely many features have infinite cost). Thus, our modelling of cost of contracting provides a justification, derived from more primitive assumptions, as to why we limit attention to contracts that are only finitely defined.

It is easy to generate a constant contract as the optimal contract when cost is high enough. More interesting is the fact that cost creates a genuine trade-off between contract complexity and the principal's expected payoff.[19]

9.2. Optimal contracting with positive but vanishing cost

Why do we allow for *all* finitely defined contracts, even though these may have arbitrarily large, but finite, cost? The reason is that in many instances finite contracts, no matter how large they are, are inexpensive to write. For agency relationships that occur repeatedly or for those requiring routine tasks, standard form contracts will generally be used.[20] Thus, in this paper we aim to characterize optimal contracts after the parties have exhausted all reasonable possibilities

[19]Consider a problem where the optimal contract specifies one payment $\alpha(x_1)$ for outcomes where the first feature is equal to 0 or when the first feature is 1 and features 2 through 100 are equal to zero; otherwise the contract specifies a different payment $\alpha(x_2)$. In this case, if the cost of specifying features in the contract is large enough the principal may be better off by offering the constant contract.

[20]The quintessential example of a standard agency contract is the Law of Agency itself. The *Restatement of the Law, Second, Agency* (American Law Institute (1957)) is 5,308 pages long and acts as a 'default contract' for agency relationships in the US. The Restatement is a rather long (but finite) contract that effectively has zero cost for the marginal agency relationship.

of reducing contracting cost by hard-wiring repetitive aspects that may initially appear complicated.

The question is then: when finite (but long) contracts are inexpensive to write, should we expect the solution to (P4) coincide with that of (P3) (as derived in Section 7)? The following proposition shows that under Assumption 4, the solution to (P4) *is* the solution to (P3) as long as the cost of adding extra features to the contract is low.

PROPOSITION 5*: Let $\{C_m\}$ be a sequence of cost functions such that $C_m(\alpha) \to 0$ for every finitely defined contract α. Let α^τ be the optimal contract in problem (P3). Then there is $\bar{M} \in \mathbb{N}$ such that for every $m > \bar{M}$,*

$$\alpha^\tau_{C_m} = \alpha^\tau.$$

PROOF: See Appendix. ■

10. CONCLUDING REMARKS

10.1. Incomplete contracts may be essential to realizing gains from trade

The presence of trust allows for exchange to occur in situations where its absence would preclude trade.[21]

An illustration is Example 4. There, because the best complete contract is constant, the principal can only motivate the agent to undertake low effort. If the agent fully trusted the principal, then high effort is easily implemented: the principal pays the agent $\alpha(x) = \overline{U} + c_H$ if the observed outcome reflects $e = H$ and $\alpha(x) = \overline{U} - \delta$, for some $\delta > 0$ otherwise.

[21] For the same conclusion in a model where trust is placed on the agent, see Casadesus-Masanell (1999).

Nevertheless, trust does not work miracles. In order for certain actions to be implementable there must be sufficient trust. In Example 4, as $\tau \to 0$, we have that $V(\alpha^\tau) \to -\infty$. Thus, if trust is sufficiently low, the principal is better off staying out of the relationship: she cannot provide incentives to implement the high action at an economically viable cost.

10.2. Distinct support result breaks down

One main result from standard agency is that if the supports of π_H and π_L are different, then the first best can be obtained by punishing the agent sufficiently harshly if the realized outcome could only have accrued under $e = L$.

Example 4 shows how in our more general setting this result breaks down. Here, the *full support* of π_H and π_L are distinct, yet in the absence of trust, not only the first best is not achievable, but, as we just discussed, the principal can only implement $e = L$.

The problem, of course, is that when information is complex, the schedule that would implement the first best is not finitely defined. When the finite definability constraint is imposed, only constant contracts are 'optimal' and with these, only low effort is implementable.

10.3. Sufficient statistic result breaks down

Another landmark result in the standard model of agency is that the optimal contract should condition on all available *informative* signals.[22] (See Holmstrom (1979).) Yet this is at odds with the reality of agency contracts, which appear not be as finely sensitive to all information as theory suggests.

[22]Formally, the principal should condition the wage on a sufficient statistic for all the signals. An additional signal y is informative about the agent's action if and only if the original signal x is not a sufficient statistic for (x, y). An informative signal y conveys information about the agent's choice of action in addition to that which is already in x.

Example 4 shows that this intuition does not apply to agency problems with complex information. The most that a written contract can do is to condition on finitely many features, even if ex post the agent's choice of action is obvious.

There is valuable information in the signal x that is being obviated by the optimal finitely defined contract. The problem is that under complexity this valuable information varies more than what the most sophisticated finitely defined contract can capture.

10.4. Immunity to Maskin-Tirole's critique

In our model, contractual incompleteness arises because complexity prevents parties from writing detailed ex ante contracts. One critique to this modelling approach is that contracting parties may be able to eliminate the source of incompleteness by playing 'message games.' That is, they may be able to commit ex ante to play a game such that the equilibrium of that game leads to implement ex post the outcome they would have liked to contract on ex ante (but weren't able to because of complexity, or other impediments to contracting).

Maskin and Tirole (1999) formulated this critique, focusing mainly on the buyer-seller model. They also note that hidden-action principal-agent models, like the one we consider here, are not vulnerable to this critique. The reason, roughly, is that implementation mechanisms work by exploiting ex post differences in agent preferences across states. In the hidden-action agency model considered here, the principal and the agent have the same utility functions independent of the outcome realized, and so no sorting is possible.

APPENDIX

A.1. Integration with respect to a finitely additive measure

The integral of simple functions is given in Section 4.2. For more general functions, we have:

Definition 6: *The integral of a bounded function $f : X \to I\!R$ is*

$$\int_X f \, d\lambda = \lim_{n\to\infty} \int_X f_n(x) \, d\lambda,$$

where $f_n : X \to I\!R,\ n \in I\!N$, is any sequence of simple functions that converges uniformly to f.[23]

This integral is well-defined for all bounded functions and does not depend on the particular approximating sequence $\{f_n\}$.

Integration with respect to λ_N reduces to taking the simple average of its values over X_N:[24]

$$\int_X f \, d\lambda_N \equiv \frac{1}{\# X_N} \sum_{x \in X_N} f(x).$$

A.2. Proofs

Proof of Lemma 1: Let l_∞ denote the set of all bounded sequences with the supremum norm. That is, for $x = (x_1, \ldots)$, $\| x \| = \sup_n x_n$. We use the fact that there exists a function $T : l_\infty \to I\!R$ that is:

1. *linear:* $T(cx + dy) = cT(x) + dT(y)$;
2. *non-negative:* $T(x) \geq 0$ if $\inf_n x_n \geq 0$;
3. *preserves identity:* $T(1, 1, 1, \ldots) = 1$;
4. *translation invariant:* $T(x_1, x_2, \ldots) = T(x_2, x_3, \ldots)$.

---

[23] That is, for every $\epsilon > 0$ there is N such $n > N$ implies that $|f(x) - f_n(x)| < \epsilon$ for every $x \in X$.

[24] The integral f on a subset $A \in \mathcal{A}$ is covered by this notation, since $\int_A f \, d\lambda \equiv \int_X \chi_A f \, d\lambda$, where χ_A is the indicator function of A.

Any such function is known as a Banach limit (see Rao and Rao (1983, p.39-41)). One may think of T as a generalized limit. In fact, the above reference shows that:

$$\liminf_n x_n \leq T(x) \leq \limsup_n x_n;$$

so, in particular, $T(x) = \lim_{n\to\infty} x_n$ whenever the limit exists.

Fix one such function T. For an arbitrary set $A \subset X$, define

$$\lambda(A) = T(\lambda_1(A), \lambda_2(A), \ldots).$$

It is clear that λ thus defined is a finitely additive probability on X. Furthermore, Equation 8 must hold since T assigns to every convergent sequence its limit. ■

PROOF OF PROPOSITION 2: Let $\tilde{\alpha}$ be the solution to Problem (P2'). That is, $V_{P2'}(\tilde{\alpha}) \geq V_{P2'}(\alpha)$, for all contract α in the feasible set of (P2'). Because φ_e, $e \in \{H, L\}$ are finitely defined, so are the induced likelihood ratios and $\tilde{\alpha}$. Recall that the expectations of finitely defined functions with respect to π are equal to those taken using the corresponding φ. Therefore, $\tilde{\alpha}$ is in the feasible set of (P2). Towards a contradiction, suppose that there is another contract $\breve{\alpha}$ different from $\tilde{\alpha}$ that satisfies all the constraints in (P2) but yields lower expected cost to the principal, so that $V_{P2}(\tilde{\alpha}) < V_{P2}(\breve{\alpha})$. Notice that $\breve{\alpha}$ is in the feasible set of (P2'). Since $\breve{\alpha}$ is finitely defined, the expectations under φ must coincide with those under π. But then,

$$V_{P2'}(\breve{\alpha}) = V_{P2}(\breve{\alpha}) > V_{P2}(\tilde{\alpha}) = V_{P2'}(\tilde{\alpha})$$

a contradiction. ■

PROOF OF PROPOSITION 3: To characterize the optimal contract, we transform Problem (P3) into an equivalent problem that is tractable using standard methods.

First, we need some notation: Let $\lambda\left\{x \in A^n : \pi_e(x) = \pi_e^j(A^n)\right\} \equiv \lambda_e^j(A^n)$ and let $\pi_e^j(A^n)$ be π_e's j^{th} value on $A^n \in \mathcal{P}(X)$, $j = 1, ..., J_e^n$ is an arbitrary ordering. With this notation, there is a simple expression for the probability that finitely-defined event A^n will occur if the agent's effort is e:

$$\int_{A^n} \pi_e(x)\, d\lambda = \sum_{j=1}^{J_e^n} \lambda_e^j(A^n)\pi_e^j(A^n). \tag{A.1}$$

Now, by Assumption 2, $\pi_e(x)$ ($e \in \{H, L\}$) take finitely many values on each $A^n \in \mathcal{P}(X)$. Thus, the likelihood ratios will also take finitely many values on A^n, $n = 1, ..., N$. Because the likelihood ratios indicate the informational content of outcomes, there are finitely many *informational equivalence classes* of outcomes.

Formally, let $\hat{\pi}_L$ and $\hat{\pi}_H$ be two numbers. Consider the ratio

$$R(\hat{\pi}_H, \hat{\pi}_L) = \frac{\hat{\pi}_L \lambda \{x \in A^n : \pi_L(x) = \hat{\pi}_L\}}{\hat{\pi}_H \lambda \{x \in A^n : \pi_H(x) = \hat{\pi}_H\}}.$$

Because π_e may take only finitely many values, there are finitely many possible values of R. Let $J^n \in \mathbb{N}$ be the number of values that R takes on A^n. Each x in A^n has a value R^j for some $j \in \{1, ..., J^n\}$ associated with it. With some abuse of notation, we let this value be $R(x)$. We say that two outcomes y and z in A^n are informationally equivalent if $R(y) = R(z)$. Informationally equivalent outcomes will induce the same wage. As a consequence, we need only compute finitely many ex post payments.

Consider the subpartition of A^n into sets of informationally equivalent outcomes

$$\{x \in A^n : R(x) = R^j, j = 1, ..., J^n\}.$$

The components of the subpartition will not be finitely defined (a consequence of Assumption 3). Let, $\pi_H^j(A^n)$ and $\pi_L^j(A^n)$ be the values of π_H and π_L on the j^{th} element of the subpartition.

Let $\lambda^j(A^n) \equiv \lambda \left\{x \in A^n : R(x) = R^j\right\}$. Then, $\Pi_e^j(n) \equiv \pi_e^j(A^n)\lambda^j(A^n)$ is the probability that an outcome in equivalence class j in A^n will ensue if the agent chooses effort e. By definition (eq. A.1), $\sum_{j=1}^{J^n} \Pi_e^j(n) = \Pr\{x \in A^n | e\}$. Clearly, $\sum_{n=1}^{N} \sum_{j=1}^{J^n} \Pi_e^j(n) = 1,\ e \in \{H, L\}$. Notice that probability distributions Π_e, $e \in \{H, L\}$ reflect the risk associated with not knowing with certainty the equivalence class j and the finitely defined set A^n where the outcome x will fall after effort has been chosen. However, these probability distributions do not capture the risk associated with not knowing the principal's type; the fact that she may turn out to be opportunistic. Below, we construct probability distributions $\bar{\Pi}_e$, $e \in \{H, L\}$ that reflect such additional risk.

Within each A^n, we fix the ordering j so that

$$\frac{\Pi_L^j(n)}{\Pi_H^j(n)} < \frac{\Pi_L^{j+1}(n)}{\Pi_H^{j+1}(n)} \tag{A.2}$$

for all $j = 1, ..., J^n - 1$.

Let $\alpha^j(n)$ be the wage that the *trustworthy principal* pays the agent if the j^{th} informational equivalence class on A^n is realized and consider the following problem:

$$\begin{array}{ll} & \min_\alpha \sum_{n=1}^{N} \sum_{j=1}^{J^n} \bar{\Pi}_H^j(n)\alpha^j(n) \\ & \textit{subject to} \\ \text{(IC)} & H \in \arg\max_e \{\sum_{n=1}^{N} \sum_{j=1}^{J^n} \bar{\Pi}_e^j(n) u(\alpha^j(n)) - c_e\} \\ \text{(IR)} & \sum_{n=1}^{N} \sum_{j=1}^{J^n} \bar{\Pi}_H^j(n) u(\alpha^j(n)) - c_H \geq \bar{U} \end{array} \tag{P3'}$$

where

$$\bar{\Pi}_e^j(n) = \begin{cases} \tau\Pi_e^j(n) & \forall\ n \ and \ \ j = 1, ..., J^n - 1 \\ \sum_{j=1}^{J^n-1}(1-\tau)\Pi_e^j(n) + \Pi_e^{J^n}(n) & \forall\ n \ and \ \ j = J^n. \end{cases} \tag{A.3}$$

This is a standard agency problem that can be solved applying standard techniques. The likelihood ratios characterizing the optimal contract are for each n

$$\frac{\bar{\Pi}_L^j(n)}{\bar{\Pi}_H^j(n)} = \frac{\Pi_L^j(n)}{\Pi_H^j(n)} \tag{A.4}$$

for $j = 1, ..., J^n - 1$ and

$$\frac{\bar{\Pi}_L^{J^n}(n)}{\bar{\Pi}_H^{J^n}(n)} = \frac{\sum_{j=1}^{J^n-1}(1-\tau)\Pi_L^j(n) + \Pi_L^{J^n}(n)}{\sum_{j=1}^{J^n-1}(1-\tau)\Pi_H^j(n) + \Pi_H^{J^n}(n)}. \tag{A.5}$$

Assume that τ is sufficiently large (hypothesis in the Proposition) so that expression (A.5) is strictly larger than (A.4). With this assumption, the lowest equilibrium wage on A^n corresponds to informational equivalence class J^n.

Let $\alpha^{\tau\prime}$ be the solution to (P3'). Notice that $\alpha^{\tau\prime}$ is a standard complete contract.

Consider the following *incomplete* contract:

$$\alpha^\tau(x) = \begin{cases} \{\alpha^1(1),\ \alpha^2(1),\ ...,\ \alpha^{J^1}(1)\} & if\ x \in A^1 \\ \{\alpha^1(2),\ \alpha^2(2),\ ...,\ \alpha^{J^2}(2)\} & if\ x \in A^2 \\ \vdots & \vdots \\ \{\alpha^1(N),\ \alpha^2(N),\ ...,\ \alpha^{J^N}(N)\} & if\ x \in A^N \end{cases} \tag{A.6}$$

where $\alpha^j(n)$ are given by contract $\alpha^{\tau\prime}$. Contract α^τ is incomplete as it gives the principal freedom to choose compensation between the possible payments on each A^n: the principal places the realized x in the corresponding finitely defined set, A^n, and if trustworthy, she chooses

$$\bar{s}_{\alpha^\tau}(x) = \alpha^j(n)$$

where j is the realized informational equivalence class. If she is opportunistic, she chooses

$$\underline{s}_{\alpha^\tau}(x) = \alpha^{J^n}(n).$$

Claim 1*:* α^τ solves (P3).

To see this, note that by construction, α^τ is finitely defined and is in the feasible set of (P3). Towards a contradiction, suppose that there is another contract $\breve{\alpha}^\tau$ different from α^τ that satisfies the constraints in (P3) but yields lower expected cost to the principal, so that $V_{P3}(\breve{\alpha}^\tau) > V_{P3}(\alpha^\tau)$.

Without loss of generality we may assume that $\breve{\alpha}^\tau$ is constant on the components of $\mathcal{P}$ (this follows from Assumption 3). Let $\breve{\alpha}^{\tau\prime}$ be the complete contract (on informational equivalence classes) that when equivalence class j in finitely defined set A^n is realized, assigns with probability τ the wage paid by the trustworthy principal under $\breve{\alpha}^\tau$ when x belongs to informational equivalence class j in finitely defined set A^n and with complementary probability assigns the lowest allowed wage by $\breve{\alpha}^\tau$ on A^n. By construction, $\breve{\alpha}^{\tau\prime}$ is in the feasible set of (P3') and $V_{P3'}(\breve{\alpha}^{\tau\prime}) = V_{P3}(\breve{\alpha}^\tau)$. But then,

$$V_{P3'}(\breve{\alpha}^{\tau\prime}) = V_{P3}(\breve{\alpha}^\tau) > V_{P3}(\alpha^\tau) = V_{P3'}(\alpha^{\tau\prime})$$

a contradiction.

Thus, incomplete contract α^τ (eq. A.6) is the solution to Problem (P3): the agent is willing to accept the contract and feels compelled to choose high effort, the contract is finitely defined and it minimizes expected cost to the principal.

Notice that when the density functions π_e take many values, it may be time consuming to write contract α^τ (because there may be many informational equivalence classes). There is, though, a very simple way to write α^τ: just specify the lowest bound and leave it open above.

More formally, if there is a cost $\delta > 0$ to writing down in a piece of paper each contingency and/or payment, the 'cheapest' way to write down contract α^τ is to just specify the lower bound for each element of $\mathcal{P}$. That is,

$$\alpha^\tau(x) = \begin{cases} \left[\alpha^{J^1}(1),\ \infty\right) & \text{if } x \in A^1 \\ \vdots & \vdots \\ \left[\alpha^{J^N}(N),\ \infty\right) & \text{if } x \in A^N \end{cases}$$

The selection rule $\{\bar{s}_{\alpha^\tau}, \underline{s}_{\alpha^\tau}\}$ guarantees α^τ's optimality. ■

PROOF OF PROPOSITION 4: We want to show that given τ and τ' with $\tau > \tau'$, the likelihood ratio distribution under τ is a mean-preserving spread of that under τ'. Then, by Kim (1995)'s Proposition 1, the value to the principal of the τ problem is strictly larger than that of the τ' problem:

$$V(\alpha^\tau) > V(\alpha^{\tau'}).$$

Let

$$L_\tau(z) = \Pr\left[\frac{\bar{\Pi}_L}{\bar{\Pi}_H} = z \middle| e = H\right].$$

Notice that the expectation of the likelihood ratio distribution on A^n,

$$E\left[L_\tau(z)\right] = \frac{\Pi_L^1}{\Pi_H^1}\Pi_H^1\tau + \frac{\Pi_L^2}{\Pi_H^2}\Pi_H^2\tau + ... + \frac{\Pi_L^{J^n-1}}{\Pi_H^{J^n-1}}\Pi_H^{J^n-1}\tau +$$

$$+\frac{\sum_{j=1}^{J^n-1}(1-\tau)\Pi_L^j+\Pi_L^{J^n}}{\sum_{j=1}^{J^n-1}(1-\tau)\Pi_H^j+\Pi_H^{J^n}}\sum_{j=1}^{J^n-1}(1-\tau)\Pi_H^j+\Pi_H^{J^n}$$

$$=\tau\sum_{j=1}^{J^n-1}\Pi_L^j+\sum_{j=1}^{J^n-1}(1-\tau)\Pi_H^j+\Pi_H^{J^n}=\Pr\left[A^n\,|\,e=H\right],$$

is independent of τ. Thus, the mean is preserved as τ changes.

To see that the likelihood ratio distribution spreads out as τ increases, note that the set of informationally equivalent outcomes with likelihood ratio $\frac{\bar{\Pi}_L}{\bar{\Pi}_H}=\frac{\Pi_L}{\Pi_H}$, does not change as τ varies. Thus, the probability of the set containing the first k informationally equivalent outcomes ($k<J^n$) is $\tau\sum_{j=1}^{k}\Pi_H^j$, an increasing function of τ.

The only problem arises for the z such that the set of informationally equivalent outcomes with $\frac{\bar{\Pi}_L}{\bar{\Pi}_H}\leq z$ contains all outcomes $j=1,...,J^n$ (because the value of $\frac{\bar{\Pi}_L}{\bar{\Pi}_H}$ for the last outcome depends on τ).

Recall that

$$\frac{\bar{\Pi}_L^{J^n}}{\bar{\Pi}_H^{J^n}}=\frac{\sum_{j=1}^{J^n-1}(1-\tau)\Pi_L^j+\Pi_L^{J^n}}{\sum_{j=1}^{J^n-1}(1-\tau)\Pi_H^j+\Pi_H^{J^n}}.$$

Therefore,

$$\frac{\partial}{\partial\tau}\left(\frac{\bar{\Pi}_L^{J^n}}{\bar{\Pi}_H^{J^n}}\right)=\frac{-\sum_{j=1}^{J^n-1}\Pi_L^j\Pi_H^{J^n}+\sum_{j=1}^{J^n-1}\Pi_H^j\Pi_L^{J^n}}{\left(\sum_{j=1}^{J^n-1}(1-\tau)\Pi_H^j+\Pi_H^{J^n}\right)^2},$$

which is positive because of the construction in the proof of Proposition 3 (eq. A.2). Thus, as τ increases so does the largest value of the likelihood ratio (the likelihood ratio associated with the lowest payment). This takes care of the largest value of z and we can safely conclude that as τ increases, the likelihood ratio distribution under τ' is a mean-preserving spread of that under τ.

Now by Kim's Proposition 1,

$$V(\alpha^\tau)>V(\alpha^{\tau'}),$$

establishing the result. ■

PROOF OF COROLLARY 1: Follows from $\lim_{\tau\to1}\alpha^\tau=\alpha^*$ and $\lim_{\tau\to0}\alpha^\tau=\widetilde{\alpha}$. ■

PROOF OF PROPOSITION 5: Notice first that α^τ is in the feasible set of (P4). Suppose there is another contract $\dot{\alpha}^\tau$ such that $V_{P4}(\dot{\alpha}^\tau)>V_{P4}(\alpha^\tau)$. By Assumption 4, $\dot{\alpha}^\tau$ is finitely defined. Now, $V_{P4}(\dot{\alpha}^\tau)>V_{P4}(\alpha^\tau)$ implies $V_{P3}(\dot{\alpha}^\tau)-C_m(\dot{\alpha}^\tau)>V_{P3}(\alpha^\tau)-C_m(\alpha^\tau)$ or $V_{P3}(\dot{\alpha}^\tau)-V_{P3}(\alpha^\tau)>C_m(\dot{\alpha}^\tau)-C_m(\alpha^\tau)$. Clearly, as $m\to\infty$, the left hand side of this last inequality goes to zero, but then, for large enough m this contradicts the assumption that α^τ is optimal for (P3). Consider now the sequence $\{\alpha_{C_m}^\tau\}_{m=1}^\infty$ of optimal contracts for (P4). Notice that by Assumption 4, all these contracts are finitely defined and thus are in the feasible set of (P3). By Assumptions 2 and

3, there is a last feature after which, conditioning on finitely many additional features does not increase the value of the contract. For large enough m, contract $\alpha^{\tau}_{C_m}$ will condition on all the payoff-relevant features and will therefore also be optimal for (P3). ■

REFERENCES

AL-NAJJAR, N. (1999): "Complexity as a Barrier to Competitive Imitation," Mimeo - MEDS Department, Kellogg GSM, Northwestern University.

AL-NAJJAR, N., L. ANDERLINI, AND L. FELLI (2000): "Incomplete Contracts in a Complex World," Mimeo.

AL-NAJJAR, N., R. CASADESUS-MASANELL, AND E. OZDENOREN (2002): "Probabilistic Representation of Complexity," *Journal of Economic Theory*, Forthcoming.

AMERICAN LAW INSTITUTE (1957): *Restatement of the Law, Second, Agency.* St. Paul, Minnesota: American Law Institute Publishers.

ANDERLINI, L., AND L. FELLI (1994): "Incomplete Written Contracts: Undescribable States of Nature," *Quarterly Journal of Economics*, 109, 1085–1124.

BAKER, G., R. GIBBONS, AND K. MURPHY (1994): "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics*, 109, 1125–56.

BATTIGALLI, P., AND G. MAGGI (2000): "Contracting over Time when Writting is Costly," Mimeo - Bocconi University and Princeton University.

BERHOLD, M. (1971): "A Theory of Linear Profit-Sharing Incentives," *Quarterly Journnal of Economics*, 85, 460–478.

CASADESUS-MASANELL, R. (1999): "Trust in Agency," Mimeo - MEDS Department, Kellogg GSM, Northwestern University.

CASADESUS-MASANELL, R., AND D. F. SPULBER (2001): "Trust and Incentives in Agency," Mimeo - Harvard Business School and Northwestern University.

DYE, R. A. (1985): "Costly Contract Contingencies," *International Economic Review*, 26, 233–250.

Gibbons, R. (1998): "Incentives in Organizations," *Journal of Economic Perspectives*, 12, 115.

Grossman, S., and O. Hart (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*, 51, 7–46.

——— (1986): "The Costs and Benefits of Ownership - A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691–719.

Hart, O., and J. Moore (1999): "Foundations of Incomplete Contracts," *Review of Economic Studies*, 66, 115–138.

Holmstrom, B. (1979): "Moral Hazard and Observability," *Bell Journal of Economics*, 10, 74–91.

Kim, S. K. (1995): "Efficiency of an Information System in an Agency Model," *Econometrica*, 63, 89–102.

Krasa, S., and S. Williams (1999): "Incompleteness as a Constraint in Contract Design," Mimeo - Department of Economics, University of Illinois at Urbana-Champaign.

Levinthal, D. (1988): "A Survey of Agency Models in Organizations," *Journal of Economic Behavior and Organization*, 9, 153.

Maskin, E., and J. Tirole (1999): "Unforeseen Contingencies and Incomplete Contracts," *Review of Economic Studies*, 66, 83–114.

Pearce, D. G., and E. Stacchetti (1998): "The Interaction of Implicit and Explicit Contracts in Repeated Agency," *Games and Economic Behavior*, 23, 75–96.

Posner, R. (1986): *Economic Analysis of Law*. Boston, MA: Little Brown.

Prendergast, C. (1999): "The Provision of Incentives in Firms," *Journal of Economic Literature*, 37, 7.

Rao, K., and M. Rao (1983): *Theory of Charges*. New York: Academic Press.

Ross, S. A. (1973): "The Economic Theory of Agency: The Principal's Problem," *American Economic Review*, 63, 134–139.

Sappington, D. (1991): "Incentives in Principal-Agent Relationships," *Journal of Economic Perspectives*, 5, 45.

Shavell, S. (1979): "Risk Sharing and Incentives in the Principal and Agent Relationship," *Bell Journal of Economics*, 10, 55–73.