RAND Journal of Economics
Vol. 39, No. 2, Summer 2008
pp. 352–377

Costly contracting in a long-term relationship

Pierpaolo Battigalli*

and

Giovanni Maggi**

We examine a model of contracting where parties interact repeatedly and can contract at any point in time, but writing formal contracts is costly. A contract can describe the external environment and the parties' behavior in a more or less detailed way, and the cost of writing a contract is proportional to the amount of detail. We consider both formal (externally enforced) and informal (self-enforcing) contracts. The presence of writing costs has important implications both for the optimal structure of formal contracts, particularly the tradeoff between contingent and spot contracting, and for the interaction between formal and informal contracting. Our model sheds light on these implications and generates a rich set of predictions about the determinants of the optimal mode of contracting.

1. Introduction

■ Contracting in a long-term relationship may come in a variety of modes. Contracts may be formal (i.e., externally enforced), informal (i.e., self-enforcing), or a combination of the two. Formal contracting in turn may take the form of a long-term contingent contract, a sequence of "spot" contracts, or a mixture of the two approaches.[1] Empirical examples of these

*Bocconi University; pierpaolo.battigalli@unibocconi.it.
**Yale University; giovanni.maggi@yale.edu.
We thank two anonymous referees for very useful comments and suggestions. We also thank Luca Anderlini, Mariagiovanna Baccara, Alvaro Bustos, Arnaud Costinot, Avinash Dixit, Kfir Eliaz, Leonardo Felli, Astrid Gamba, Edoardo Grillo, Bentley MacLeod, Eric Maskin, Fabio Michelucci, Fausto Panunzi, David Pearce, Ludovic Renou, Joel Watson, the participants in conferences at Siena, Venezia, Ischia, Gertzensee, and Lousanne, and the participants in seminars at Northwestern University, PUC Rio de Janeiro, European University Institute, Bocconi University, DELTA, INSEAD, University of Arizona, University of Helsinki, Hebrew University of Jerusalem, and Tel Aviv University for helpful comments and discussions. Battigalli thanks Bocconi University for financial support. Maggi received financial support from the National Science Foundation. Both authors thank the NYU Stern School of Business for its hospitality during part of this project.

[1] By spot contracting, we mean that the contract is noncontingent but gets modified as circumstances change. The difference between contingent and spot contracting is nicely explained by Williamson (1985), who states that the writing of contracts "can be done with a great deal of care, in which case a complex document is drafted in which numerous contingencies are recognized, and appropriate adaptations by the parties are stipulated and agreed to in advance. Or the document can be very incomplete, the gaps to be filled by the parties as the contingencies arise. Rather, therefore, than contemplate all conceivable bridge crossings in advance, which is a very ambitious undertaking, actual bridge-crossing choices are only addressed as events unfold."

different contracting modes abound. The classic paper by Macaulay (1963) and a number of subsequent studies have highlighted that formal and informal contracting often coexist in long-term relationships, such as the ones between firms and suppliers or between employers and employees (see, for example, Baker et al., 1994, for a discussion of this literature). Also, it is not hard to find cases in which formal contracting exhibits a mix of spot and contingent contracting.[2] Further, there is anecdotal evidence of long-term relationships where contracts become gradually more contingent over time: for example, Meihuizen and Wiggins (2000) study supply contracts in the U.S. natural gas industry and find that many of the observed contracts tend to become richer and to incorporate a growing number of contingent clauses over time.

In this article, we present a dynamic principal-agent model in which the contracting mode is endogenously determined. A key feature of the model is that writing formal contracts is costly. We argue that writing costs have strong implications for the optimal mode of contracting, even if their magnitude is small. The model highlights that it may be optimal to use different contracting modes for different aspects of the relationship, and yields rich predictions on the impact of fundamental parameters on the optimal mode of contracting.

Next we give a more detailed preview of the model and of the main results.

We consider a multitask, principal-agent setting with verifiable contingencies and actions, where parties interact repeatedly over an infinite horizon. Principal and agent can contract at any point in time; importantly, this includes the possibility of "spot" contracting, meaning that in each period the contract is written after observing the state of nature. A contract can describe the external environment and the agent's behavior in a more or less detailed way, and the cost of writing a contract is increasing in the amount of detail. In each period, parties can save on writing costs by modifying the previous contract rather than drafting a whole new contract. We also allow for a form of dynamic writing economies: the cost of describing a given action/contingency for the first time may be higher than for subsequent times.

In our model a contract is binding only for the current period, but we could equivalently consider open-ended contracts, in the sense that contracts are automatically renewed unless modified or rejected. For this reason, if a contract is signed at $t = 1$ and is never modified, we can interpret it as a *long-term* contract.

In the first part of the article, we focus on situations where the parties rely entirely on formal contracting. A simple way to capture these situations is to consider the Markov perfect equilibrium of the game, which excludes the possibility of enforcing informal obligations through history-dependent punishments. We examine the optimal structure of formal contracts, and in particular the choice between contingent and spot contracts. At the intuitive level, it is not obvious whether the presence of writing costs should favor contingent or spot contracting: on the one hand, spot contracting avoids the cost of describing contingencies; on the other hand, spot contracts must describe the agent's *behavior* repeatedly, and this may push in favor of a contingent contract.[3]

We emphasize that spot contracting can induce efficient state-dependent behavior in spite of the noncontingent nature of the contracts, because we assume that in each period, parties can contract after observing the state realization. This assumption "stacks the deck" against contingent contracting, because it removes its most obvious advantage relative to a sequence of spot contracts, but nonetheless important tradeoffs remain, because, as we hinted above, contingent contracting saves on the costs of describing the agent's behavior repeatedly over time.[4]

[2] For example, in the construction industry, the relationship between contractors and subcontractors often exhibits a mix of spot and contingent contracting. The project specifications are replaced from project to project (spot contracting), but there are typically also some contingent clauses that apply to all projects (e.g., specifying circumstances under which a party can make changes to the project or terminate it). See Eccles (1981) for an analysis of contractor/subcontractor relationships in the U.S. construction industry.

[3] The idea that transaction costs might favor a long-term contingent contract has already been expressed informally, for example by Hart and Holmstrom (1987).

[4] If we assumed that in each period the contract must be written before observing the state realization, then the parameter region under which contingent contracting is optimal would expand, and this expansion would be more pronounced if the state of nature is less persistent.

Absent writing costs, the model has no predictive power, because there is a plethora of optimal contracting plans, including a long-term contingent contract, a sequence of spot contracts, and a host of intermediate solutions. But in the presence of (even very small) writing costs, the model yields a unique optimum. Each task is optimally governed in one of three ways: (i) a contingent clause, written once and for all; (ii) a sequence of spot clauses; or (iii) a noncontingent clause that is replaced at some point in time by a contingent clause (we refer to this as the "enrichment" approach). A contingent clause is optimal when the cost of describing contingencies is low relative to the cost of (re-)describing actions; a spot approach is optimal when this relative cost is high; and an enrichment approach may be optimal when this relative cost takes intermediate values.

If the enrichment approach is optimal for at least some of the tasks, the model predicts that the number of contingent clauses in the contract will increase over time. Thus our model can explain why sometimes contracts become gradually more contingent over time (as in the case documented by Meihuizen and Wiggins, 2000, for the natural gas industry).

Two important determinants of the optimal contracting mode are the degree of uncertainty and the durability of the relationship. We find that tasks characterized by a higher degree of uncertainty are more likely to be governed by contingent clauses, whereas lower-uncertainty tasks are more likely to be governed by a spot approach. Similarly, the model predicts that the fraction of tasks governed by spot contracting should be higher when the external environment is more stable. On the other hand, we find that the relative importance of spot contracting is lower when the relationship is more durable. Interestingly, then, the stability of the relationship and the stability of the external environment have opposite effects on the choice between spot and contingent contracts.

In the second part of the article, we examine the interaction between formal and informal contracting. As we mentioned at the outset, there is ample anecdotal evidence that these two modes of contracting often go hand in hand. The presence of writing costs provides a theoretical explanation for the coexistence of formal and informal contracting. Intuitively, the advantage of a self-enforcing contract is that it can be communicated informally, rather than being written formally, and this saves on writing costs. On the other hand, if parties are not very patient, they will not be able to govern all tasks by informal contract. The question is then under what conditions the parties will combine formal and informal contracting and, if so, how changes in the contracting environment affect the optimal mix of formal and informal contracting.

In the model, we allow for the possibility of informal contracting by focusing on the efficient subgame perfect equilibria of the game. Whereas formal rules are enforced by external courts, informal rules are enforced by the threat of reverting to an equilibrium that is worse for the deviator (punishment strategy). We show that, if the discount factor is above a certain threshold (which is close to $1/2$ if writing costs are close to zero), there exists a credible punishment strategy that keeps a deviator (principal or agent) at his maxmin payoff. This punishment essentially involves the deviator paying a fine to the other player, and then the relationship is brought back to a "normal" state.

We then examine the conditions under which formal and informal contracting coexist. We find that it is optimal to combine formal and informal contracting if the discount factor lies in some intermediate range. Also, under a relatively mild condition—namely that there is at least one task with relatively low disutility—the optimal contract is partly or fully informal. Interestingly, this is true even if writing costs are very small. In fact, reducing writing costs may even *facilitate* informal contracting (more on this below).

When formal and informal contracting coexist, tasks characterized (other things equal) by a lower disutility for the agent are governed by informal contracting, whereas higher-disutility tasks are governed by formal contracting. This is because the agent has a stronger incentive to shirk from higher-disutility tasks.

The model offers interesting insights regarding the effects of changes in the fundamental parameters of the model on the optimal mix of formal and informal contracting. In order to make the comparative-statics analysis tractable, we consider a version of the model with a continuum of tasks. The first observation is that it is harder to give players incentives to abide by the informal rules when: (i) the available surplus from the relationship is lower; (ii) the discount factor is lower; (iii) uncertainty is higher; and (iv) writing costs are higher. Whereas the first two statements are intuitive, the role of uncertainty and writing costs is more subtle: the costs of spot contracting are incurred repeatedly over time, and these costs are higher when the environment is more uncertain and when amending the contract is more costly; as a consequence, increasing these costs strains incentives (if some tasks are governed by spot contracting). Based on these considerations, one might expect that any of the parameter changes mentioned above ([i]–[iv]) would lead to less informal contracting and more formal contracting. This indeed turns out to be true if, at the initial optimum, all formal tasks are governed in the same way (spot or contingent). Note that in this case, a reduction in writing costs facilitates informal contracting (as we hinted above). However, results may be quite different if the optimum is interior, in the sense that some tasks are governed by spot contracting and some by contingent contracting. For example, in this case, a change in the available surplus has *no* effect on the fraction of tasks governed informally, and an increase in writing costs or uncertainty leads to an *increase* in the fraction of tasks governed informally. These surprising results arise from a nontrivial three-way interaction between informal contracting, spot contracting, and contingent contracting.

Our analysis focuses on the case in which writing costs are relatively small, and more specifically, small enough that it is optimal to implement the first-best outcome. In other words, we focus on the case in which contracting is "complete."[5] The emphasis on this case is useful for two reasons. First, all the main insights can be brought out with small writing costs, and the exposition is considerably simpler than in the case of large writing costs. Second, this helps to clarify that contractual incompleteness per se is not essential to our results. In the concluding section, we briefly discuss the case in which writing costs are large.

In what follows, we discuss how our article relates to the previous literature.

The way we model contracting costs is similar to our previous paper, Battigalli and Maggi (2002). In particular, the cost of writing a contract is proportional to the amount of detail contained in the contract. The main difference is that the present article explores the implications of writing costs for the dynamics of (formal and informal) contracting, whereas our previous paper focuses on a static setting, and therefore is silent on issues of dynamic contracting.[6]

The interaction between formal and informal contracting has also been analyzed by Baker et al. (1994), Pearce and Stacchetti (1998), and Schmidt and Schnitzer (1995).[7] These papers propose a different explanation for the combined use of the two types of contract. They consider a repeated principal-agent game where formal contracts can be based only on verifiable signals of the agent's action, whereas informal contracts can be based on unverifiable signals. In these models, it may be optimal to combine a formal wage with an informal "bonus." Our model, on the other hand, explains why it may be efficient to regulate some tasks formally and some others informally (in our model there is no need for bonuses). Perhaps more importantly, the rationale for mixing formal and informal contracting in our model is not the presence of verifiable and nonverifiable signals, but the interaction between writing costs and self-enforcement constraints.

Relative to this literature, our analysis also yields different predictions on the interplay between formal and informal contracting. For example, one key result in Baker et al. (1994) is that, if the imperfection in formal contracting is sufficiently small, informal contracts cannot

[5] By complete contracting, we mean simply that the first-best outcome is implemented, so this may include contingent formal, spot formal, and informal contracting.

[6] We should mention that there are other papers on complexity costs as a cause of contractual incompleteness. In a static setting, see, for example, Dye (1985a), Anderlini and Felli (1994, 1999), MacLeod (2000), Krasa and Williams (2007), and Al Najjar et al. (2006). In a dynamic setting, see the theoretical papers by Tirole (2007) and Bolton and Faure-Grimaud (2007), and the empirical paper by Kahan and Klausner (1997).

[7] There is also a vast literature on purely self-enforcing contracts. Bull (1987), MacLeod and Malcomson (1989), and Levin (2003) are prominent examples of this literature.

be sustained. Therefore their model predicts that, if the verifiable signal becomes more precise, informal contracting may be undermined as a consequence. According to our analysis, even if formal contracting is close to perfect (i.e., if writing costs are close to zero), the optimum typically involves a mix of formal and informal contracting or a fully informal contract. Thus, our model suggests that informal contracting need not disappear as the formal-contracting system becomes more efficient.[8]

Finally, our model yields interesting predictions on the choice between spot and contingent contracts. The models by Baker et al. (1994) and Pearce and Stacchetti (1998) are silent about this issue because the nature of the contractual imperfection is different. In these models, the imperfection is given by the noisiness of verifiable signals about the agent's performance. This imperfection is exogenous and is present in every period, whereas in our model the imperfection (writing costs) is endogenous and need not be incurred in every period.

The article is organized as follows. In Section 2, we describe the contracting environment. In Section 3, we focus on the case in which only formal contracting is available. In Section 4, we extend the analysis to the possibility of informal contracting. In Section 5, we offer some concluding remarks.

2. The contracting environment

■ We analyze a repeated multitask principal-agent game where parties can contract in each period. Payoffs depend on actions and external contingencies, which are both verifiable in court but costly to describe in a written contract. Thus, the only contractual imperfection is the presence of writing costs. We model payoffs and writing costs in a similar way as in Battigalli and Maggi (2002). We start by describing the language used to write contracts.

The set $L^e = \{e_1, e_2, \ldots, e_N\}$ is a finite collection of primitive sentences, each of which describes an *elementary event* concerning the external environment. For example, e_1: "the passenger has a moustache," e_2: "the passenger's bag is red." The set $L^a = \{a_1, a_2, \ldots, a_N\}$ is a finite collection of primitive sentences describing *elementary actions*, for example, a_1 : "check the passenger's passport," a_2 : "search the passenger's bag."

With a slight abuse of terminology, we will use the notation e_k (respectively, a_k) to indicate both an elementary event (respectively, action) and the primitive sentence that describes it. We assume that this language is the (only) common-knowledge language for the parties and the courts. This ensures that there are no problems of ambiguous interpretation of the contract.

A *state* is a complete description of the exogenous environment, represented by a valuation function $s : L^e \to \{0, 1\}$, where $s(e_k) = 1$ means that primitive sentence e_k is true at state s and $s(e_k) = 0$ means that primitive sentence e_k is false at state s.[9] In other words, $s(e_k)$ is a dummy variable that takes the value one if elementary event e_k occurs and zero otherwise, and a state is a realization of the vector of dummy variables $(s(e_1), s(e_2), \ldots)$. Similarly, a *behavior* is a complete description of all elementary actions, represented by a valuation function $b : L^a \to \{0, 1\}$, where $b(a_k) = 1$ means that elementary action a_k is executed, and $b(a_k) = 0$ that a_k is not executed.

We assume a very simple payoff structure. There is a one-to-one correspondence between elementary actions and elementary events. The principal wants action a_k to be performed if and only if elementary event e_k occurs; in our airport example, the principal wants the agent to check the passenger's passport if and only if the passenger has a moustache, and to search his bag if and only if the bag is red. The principal and the agent are risk neutral. The principal gets incremental benefit π_k from "matching" $s(e_k)$ with $b(a_k)$, whereas he gets zero incremental benefit if there is a "mismatch." Formally, the principal's per-period payoff gross of writing costs is

[8] In Section 4, we will be more specific about the reasons for this divergence in results.

[9] To simplify the exposition, we describe the basic notation omitting time subscripts. We will introduce time subscripts later in this section, when we describe the dynamic aspects of the game.

$$\Pi(s, b, m) = \sum_{k=1}^{N} \pi_k I_k(s, b) - m, \tag{1}$$

where m is the (possibly negative) payment to the agent and $I_k(s, b) = s(e_k)b(a_k) + (1 - s(e_k))(1 - b(a_k))$ is a dummy variable that takes the value one if there is a match between event e_k and action a_k, and zero if there is a mismatch.

The agent's interests are always in conflict with the principal's, in the sense that his preferred actions are always opposite the principal's preferred actions. Formally, the agent's one-period utility is

$$U(s, b, m) = m - \sum_{k=1}^{N} d_k I_k(s, b). \tag{2}$$

We will often refer to the job of matching action k with event k as "task" k. The parameter d_k thus captures the agent's disutility from performing task k. We let $\mathbf{N} = \{1, \ldots, N\}$ denote the set of tasks and we use bold capital letters to denote subsets of tasks as in $\mathbf{K} \subseteq \mathbf{N}$.[10]

Each party's reservation payoff is zero. Assuming $0 < d_k < \pi_k$ for all k, the parties' joint surplus (gross of writing costs) is maximized when the agent performs all tasks $k \in \mathbf{N}$. We will refer to this as the *first-best* outcome. Payoffs are common knowledge to the contracting parties, and the state and the parties' behavior are verifiable in court. Thus, there are no issues of moral hazard or adverse selection. We assume that preferences and realized payoff levels are not verifiable in court, and that the principal cannot "sell the activity" to the agent (i.e., the agent cannot be made the recipient of the revenue).[11]

Next we define a contract. A contract is a pair $(\mathbf{g}, m)$, where $\mathbf{g} = (g_k)_{k\in\mathbf{N}}$ is a set of N clauses and m is a transfer from the principal to the agent (wage). Each clause g_k regulates a task. Given our simple matching structure between actions and events, we can restrict our attention to four types of clause: (i) a contingent clause, constraining the agent to do a_k if and only if e_k occurs, $C_k : [a_k \leftrightarrow e_k]$; (ii) a noncontingent positive clause, constraining the agent to do a_k whatever happens, $R_k : [a_k]$; (iii) a noncontingent negative clause, constraining the agent to do *not* a_k whatever happens, $\bar{R}_k : [\neg a_k]$; and (iv) the empty clause, D_k (for discretion), that imposes no constraint on the agent for task k. For example, if $N = 3$, the set of clauses (R_1, D_2, C_3) constrains the agent to do a_1 whatever happens and to do a_3 if and only if elementary event e_3 occurs, leaving the agent free with regard to task 2. Note that, because we include the empty clause among the possible clauses, there is no loss of generality in assuming that the number of clauses in the contract is N.[12]

The parties interact repeatedly over an infinite horizon and have common discount factor $\delta \in (0, 1)$. The parameter δ can also be interpreted as capturing the stability of the relationship.[13] Within each period t, the timing is the following: the state of nature s_t is observed; then the principal offers a contract $(\mathbf{g}_t, m_t)$ to the agent, incurring the associated writing costs; if the contract is accepted, the principal makes payment m_t and then the agent acts

[10] We could consider more general payoff structures where the interests of principal and agent may not always be in conflict. For example, we could allow for the possibility that the agent incurs a cost d_k if and only if he takes action a_k, and the principal gets a benefit π_k if and only if the agent takes action a_k when e_k occurs. With more general payoffs the notation would be more complicated, but nothing substantial would change in the results.

[11] If preferences were verifiable, the first-best outcome could trivially be achieved by a contract of the form "the agent's behavior must maximize the sum of the parties' utilities." On the other hand, if realized payoff levels were verifiable, the first-best outcome could be achieved by offering the agent a transfer that increases one-for-one with the principal's realized payoff level. And selling the activity to the agent would be equivalent to specifying a contingent transfer as in the previous point.

[12] We note that in this model, there is nothing to gain from making wages contingent on the state or on the agent's behavior.

[13] The parameter δ can be interpreted as the composition of two parameters, $\delta = \rho\delta'$, where ρ is the probability that the game will continue and δ' is the true discount factor.

(being constrained by the contract). The contract at time t is binding only for the current period.[14]

We assume that elementary events are independent of each other, and that each elementary event is governed by a Markov process. Let the transition probabilities for event e_k be $\Pr[s_t(e_k) = 1 \mid s_{t-1}(e_k) = 1] = p_k$ and $\Pr[s_t(e_k) = 0 \mid s_{t-1}(e_k) = 0] = q_k$. We define labels so that $p_k \geq q_k$ for all k. We assume that the process exhibits nonnegative persistence, that is, $q_k \geq 1/2$. Note that this allows for an *i.i.d.* process ($p_k = q_k = \frac{1}{2}$). If we allowed for $q_k < 1/2$ the analysis would be more cumbersome, but the main qualitative insights would not change (see footnote 19 for a brief discussion of this extension). Finally, we assume that the initial state is $s_1(e_k) = 1$ for all k; given our definitions, this means that the initial state is the one with highest degree of persistence. The extension to the more general case in which also the initial state is random is straightforward but tedious.

If at time t the principal wants to offer a different contract from at time $t-1$, he can save on writing costs—whose exact nature will be specified below—by *modifying* the existing contract rather than writing a whole new contract. To capture this idea, we assume that the current contract is obtained from the previous contract by amending some (or possibly none) of its clauses. In other words, each clause from the previous contract is still effective in the current contract unless it is amended. More formally, the set of clauses at time t is given by

$$\mathbf{g}_t = f\left(\mathbf{g}_{t-1}, \mathbf{g}_t^A\right) = \left((g_{k,t-1})_{k \in \mathbf{N} \setminus \mathbf{K}_t^A}, (\alpha_{k,t})_{k \in \mathbf{K}_t^A}\right),$$

where $\alpha_{k,t} \in \{C_k, R_k, \overline{R}_k, D_k\}$ ($\alpha_{k,t} \neq g_{k,t-1}$) is the amendment for task k and $\mathbf{g}_t^A = (\alpha_{k,t})_{k \in \mathbf{K}_t^A}$ is the set of amendments. It is natural to assume that at time $t = 0$ there is no contract in place, and therefore we can write the initial condition as $g_{k,0} = D_k$ for all k.

We can now describe the costs of writing contracts. Writing primitive sentences is costly. We allow for a simple form of dynamic writing economies: writing a given primitive sentence $\xi \in L^e \cup L^a$ for the first time is (weakly) more costly than writing it the subsequent times. If $c(\xi)$ denotes the cost of writing ξ for the first time, we assume that the cost of writing ξ each subsequent time is $r \cdot c(\xi)$, where $r \in (0, 1]$. Specifying the wage and writing logical connectives (such as $\neg$ or $\rightarrow$) in the contract is costless.

In each period t, the principal incurs the costs of modifying the previous contract, that is, the writing costs of the amendments $\mathbf{g}_t^A$.[15] These costs (which are history dependent, due to the recalling economies) can be derived using the assumptions above. Focusing on task k, there are only a few relevant possibilities that we need to consider. (i) Writing a contingent clause C_k at time $t = 1$ costs $c(a_k) + c(e_k)$. (ii) Writing a noncontingent clause (R_k or $\bar{R}_k$) at time $t = 1$ costs $c(a_k)$. (iii) Replacing clause R_k with clause $\bar{R}_k$ or vice versa costs $r \cdot c(a_k)$, because this involves recalling an already-described action. (iv) Replacing a noncontingent clause (R_k or $\bar{R}_k$) with a contingent clause (C_k) costs $r \cdot c(a_k) + c(e_k)$, because this involves describing a new elementary event and recalling an already-described action. (v) The empty clause D_k, of course, involves no cost. Also, removing a clause (i.e., replacing it with the empty clause D_k) is costless.[16]

We emphasize that in each period, the contract is written *after* the state is observed. Thus, writing a contingent contract would not make sense if the parties interacted for just one period. But with repeated interaction, as we will see shortly, a contingent contract may be efficient.

[14] Note that we do not allow for multiperiod contracts (i.e., contracts that constrain players for future periods), and we implicitly assume that the principal cannot pay bonuses (payments in excess of the wage specified in the formal contract), but in the Conclusion we will argue that both of these restrictions are without loss of generality in this model. We will also argue that our results would not change if the payment (possibly including a bonus) were made after the agent acts.

[15] Note that writing a clause at $t = 1$ can be seen as a modification of the empty clause D_k, which is the initial clause at $t = 0$.

[16] Introducing a cost of removing clauses would make our notation heavier without changing our results.

3. Formal contracting

■ In this section, we focus on situations where parties rely entirely on formal contracting. A simple way to capture these situations is to look at the Markov perfect equilibrium (MPE) of the game. This excludes the possibility of punishing the agent if he "shirks" on tasks that are not covered by the formal contract, and hence informal contracting is ruled out.

The analysis of this section can be interpreted as applying to long-lasting relationships where the parties for some reason cannot coordinate on efficient equilibria (that is, equilibria where informal contracting is sustained by the threat of reverting to an equilibrium that is worse for the deviator). Moreover, we note that the same qualitative results would obtain if we considered the backward-induction equilibrium of the finite-horizon version of our game with at least three periods, and thus our results can also be broadly interpreted as applying to relationships that have a relatively short duration.[17]

An additional reason for analyzing the MPE of the game is that this allows us to focus sharply on the tradeoffs between different modes of formal contracting (e.g., spot vs. contingent) without the confounding effects due to the presence of informal contracting.

Formally, an MPE is a subgame perfect equilibrium where strategies depend only on the payoff-relevant state variable (see Fudenberg and Tirole, 1991). The payoff-relevant state variable has four components: the current state of the environment s_t, the set of clauses used in the previous period, $\mathbf{g}_{t-1}$, the set of tasks $\mathbf{M}^a_{t-1}$ for which the corresponding elementary action (a_k) has been described in the past, and the set of tasks $\mathbf{M}^e_{t-1}$ for which the corresponding elementary event (e_k) has been described in the past. We will denote the overall state variable by $X_t = (\mathbf{g}_{t-1}, \mathbf{M}^a_{t-1}, \mathbf{M}^e_{t-1}, s_t)$. Notice that, because the agent's current choice has no impact on future states, Markov strategies do not depend on past actions of the agent. This implies that in the MPE the agent minimizes his disutility in every period, taking the inefficient action for every task not covered by the contract.

In the MPE, the wage m_t is set at the minimum level that induces the agent to accept the proposed contract. Because the determination of the wage is a trivial aspect of the analysis, we will focus on the set of clauses. Solving for the MPE boils down to maximizing the expected discounted value of the surplus net of writing costs.

To state the problem formally, define a *contracting policy* as a function of the form $X_t \mapsto \mathbf{g}^A_t = \psi(X_t)$. This function induces, for each (t, τ, X_t) $(t, \tau \geq 0)$, a random value for the surplus net of writing costs at date $t + \tau$, which we denote $\boldsymbol{\sigma}^{\psi}_{t+\tau} \mid X_t$. The problem can then be stated as

$$\forall X_t, \forall \tau \geq 0, \quad \max_{\psi} \mathrm{E}\left[\sum_{\tau=0}^{\infty} \delta^{\tau} \boldsymbol{\sigma}^{\psi}_{t+\tau} \mid X_t\right]. \tag{MP}$$

The *optimal* contracting policy is the solution to problem (MP).

Most of the interesting points can be brought out by considering strictly positive but relatively small writing costs. We assume that writing costs are sufficiently small that the optimal contracting policy implements the first-best outcome, or in other words, "complete" contracting is optimal. A simple sufficient condition for this is

$$\forall k, \ c(e_k) + c(a_k) < \pi_k - d_k. \tag{SC}$$

In the concluding section, we will discuss how results change when writing costs are large, so that "incomplete" contracting may be optimal.

Given our assumptions, we can derive the optimal contracting policy by looking separately at each task k. For this reason, to keep the exposition simple, we drop the task subscript k for the remainder of this section. We will reintroduce it when we examine informal contracting, as there will be nontrivial interactions among tasks.

[17] The MPE of the infinite-horizon game can be viewed as the limit of the backward-induction equilibrium of the finite-horizon game as the horizon approaches infinity.

We show that, for a given task, there are only three *rules* that can be optimal (see proof of Proposition 1):

(i) At time $t = 1$ a C clause is written, and it is never modified. We will refer to this as a *contingent* rule, and denote it $\mathcal{C}$.
(ii) At time $t = 1$ a noncontingent clause R is written, and it is amended every time the realization of s_t changes. We refer to this as a *spot* rule, and denote it $\mathcal{S}$.
(iii) At $t = 1$ a clause R is written, and it is permanently replaced by a clause C the first time $s_t(e) = 0$ occurs. We refer to this as an *enrichment* rule, and denote it $\mathcal{E}$.[18]

If we dropped the assumption $q \geq 1/2$ other rules could be optimal, but the essence of the results would not change.[19]

Whereas it is intuitive that the contingent rule and the spot rule are candidates for an optimum, it is less obvious why the enrichment rule might be optimal. Broadly speaking, the reason the $\mathcal{E}$ rule might be preferred to the $\mathcal{C}$ rule is that it allows postponing the cost of describing contingencies, and the reason it might be preferred to the $\mathcal{S}$ rule is that it requires changing the contract only once rather than multiple times. Below we will be able to give some intuition about the conditions under which $\mathcal{E}$ is indeed preferred to $\mathcal{C}$ and $\mathcal{S}$.

Before proceeding, we add a note of interpretation regarding the contingent rule $\mathcal{C}$. Technically, this rule involves renewing the same contingent contract every period. But, as we noted in Introduction, we can interpret this as an open-ended contract (i.e., automatically renewed unless modified). Thus, we can interpret it as a contract that is signed once and for all at $t = 1$, or a *long-term* contingent contract.

Because all three rules ($\mathcal{C}$, $\mathcal{S}$, and $\mathcal{E}$) implement the first best, we only need to compare the present discounted value (PDV) of writing costs for each rule, denoted $\widehat{c}(\cdot)$. Simple calculations show that

$$\begin{aligned}
\widehat{c}(\mathcal{C}) &= c(a) + c(e), \\
\widehat{c}(\mathcal{S}) &= c(a) + r \cdot c(a) \sum_{t=2}^{\infty} \delta^{t-1} \Pr[s_t(e) \neq s_{t-1}(e)] \\
&= c(a) + r \cdot c(a) \cdot \frac{\delta(1-p)[1-\delta(2q-1)]}{(1-\delta)[1-\delta(p+q-1)]}, \text{ and} \\
\widehat{c}(\mathcal{E}) &= c(a) + [r \cdot c(a) + c(e)] \sum_{t=2}^{\infty} \delta^{t-1} p^{t-2}(1-p) \\
&= c(a) + [r \cdot c(a) + c(e)] \frac{\delta(1-p)}{1-\delta p}.
\end{aligned}$$

Letting $\gamma \equiv \frac{c(e)}{r \cdot c(a)}$ denote the cost of writing an elementary event relative to the cost of recalling an elementary action, it follows from the above expressions that (i) $\mathcal{C}$ is preferable to $\mathcal{S}$ if and only if $\gamma < \overline{\gamma}_{\mathcal{C}/\mathcal{S}} \equiv \frac{\delta(1-p)[1-\delta(2q-1)]}{(1-\delta)[1-\delta(p+q-1)]}$; (ii) $\mathcal{E}$ is preferable to $\mathcal{C}$ if and only if $\gamma > \overline{\gamma}_{\mathcal{C}/\mathcal{E}} \equiv \frac{\delta(1-p)}{1-\delta}$; and (iii) $\mathcal{E}$ is preferable to $\mathcal{S}$ if and only if $\gamma < \overline{\gamma}_{\mathcal{E}/\mathcal{S}} \equiv \frac{\delta(1-q)[1-\delta(2p-1)]}{(1-\delta)[1-\delta(p+q-1)]}$.

Note that $\overline{\gamma}_{\mathcal{C}/\mathcal{E}} \leq \overline{\gamma}_{\mathcal{C}/\mathcal{S}} \leq \overline{\gamma}_{\mathcal{E}/\mathcal{S}}$ if and only if $p \geq q$, which holds by assumption. Thus, rule $\mathcal{C}$ is optimal if and only if $\gamma \leq \overline{\gamma}_{\mathcal{C}/\mathcal{E}}$, rule $\mathcal{S}$ is optimal if and only if $\gamma \geq \overline{\gamma}_{\mathcal{E}/\mathcal{S}}$, and rule $\mathcal{E}$ is optimal if and only if $\gamma \in [\overline{\gamma}_{\mathcal{C}/\mathcal{E}}, \overline{\gamma}_{\mathcal{E}/\mathcal{S}}]$. Moreover, $\mathcal{E}$ is strictly optimal for some value of γ if and only if $\overline{\gamma}_{\mathcal{C}/\mathcal{E}} < \overline{\gamma}_{\mathcal{E}/\mathcal{S}}$, which in turn holds if and only if $p > q$. We have just shown the following result.

[18] Note that a *rule* as specified above does not completely determine a *policy*, because it specifies amendments only for reachable states. For example, the rule saying that clause C has to be written at time $t = 1$ and kept thereafter does not specify what to do after C is removed. This partial description of the optimal policy is sufficient for our purposes.

[19] In the working paper version (Battigalli and Maggi, 2003), we consider a more general stochastic process. There we show that, if q is allowed to be lower than $1/2$, it might be optimal to use *exceptions* rather than amendments. But aside from this point, the main qualitative results of the model do not change.

Proposition 1. Under assumption (SC), the optimal choice of contracting mode depends on the contracting cost parameters only through the ratio $\frac{c(e)}{r \cdot c(a)} \equiv \gamma$. The optimal rule is $\mathcal{C}$ if γ is small, $\mathcal{S}$ if γ is high, and $\mathcal{E}$ if γ is intermediate. The $\mathcal{E}$ rule is strictly optimal for some value of γ if and only if $p > q$.

This proposition highlights a key tradeoff involved in the choice of formal contracts. There are three contracting modes that can be optimal. The first approach is to write a long-term contingent contract. This involves a higher upfront cost, but it does not involve any subsequent costs. The second is to write a noncontingent contract and modify it as the environment changes; this is what we call "spot contracting." The advantage of this contracting mode is that there is no need to describe contingencies, but only the agent's actions; the disadvantage is that writing costs must be incurred repeatedly over time. The third possibility (the "enrichment" rule) is one that does not fall squarely into either category of spot or contingent contracting: the contract is initially noncontingent, but is replaced later on by a contingent contract.

The parameter γ captures the cost of describing contingencies relative to the cost of (re-) describing actions. It is intuitive that, if γ is sufficiently low, a contingent rule is optimal, and if γ is sufficiently high, a spot rule is optimal. As a consequence, the enrichment rule $\mathcal{E}$ can be optimal only for intermediate values of γ. Proposition 1 also states that $\mathcal{E}$ is strictly optimal for a range of γ if and only if the degree of persistence is asymmetric across the two states $s(e) = 0$ and $s(e)$, that is, $p > q$. To gain some intuition for this result, let us first note that $\mathcal{E}$ is better than $\mathcal{S}$ only if $c(e)$ is smaller than the expected future costs of amending a rigid clause every time it is needed (the cost of describing the appropriate action at $t = 1$ is the same under both approaches and hence cancels out). When $p = q$ the two states $s(e) = 0$ and $s(e) = 1$ are symmetric and hence the present value of future costs under rule $\mathcal{S}$ is independent of the current state. If this present value is larger than $c(e)$, then it must be optimal to use the contingent clause from the first period, that is, it is optimal to use rule $\mathcal{C}$. This explains intuitively why rule $\mathcal{E}$ can be optimal for a range of γ *only if* $p > q$. Proposition 1 goes beyond this intuition, and establishes also that rule $\mathcal{E}$ is indeed optimal for a range of γ if $p > q$.

Notice an interesting implication of the model: if the enrichment rule $\mathcal{E}$ is optimal for a nonempty subset of the tasks, *the contract becomes gradually more complex over time*, as noncontingent clauses get replaced by contingent clauses.

We emphasize that, if there were no writing costs, the model would have little predictive power, because there would be a vast multiplicity of optimal contracting plans: any contracting plan that implements the first best would be optimal, including a complete long-term contingent contract, a sequence of complete spot contracts, and a whole host of intermediate solutions. However, as the analysis makes clear, the introduction of (even very small) writing costs is sufficient to pin down a unique optimum and deliver strong predictions on the optimal structure of formal contracts.

The model also generates interesting predictions on the impact of changes in the contractual environment on the choice between spot and contingent contracting. The relative importance of spot contracting can be captured by the fractions of tasks regulated by a pure spot rule. Recall that, even though we dropped the task index k from the notation, tasks may be heterogeneous with respect to any of the relevant parameters, so the optimal contracting plan may regulate tasks in different ways.

We start by examining the impact of a change in the degree of uncertainty in the external environment. A simple way to parameterize the degree of uncertainty in this model is to consider the case of symmetric persistence, that is, $p = q$. Then, a decrease in p increases the degree of uncertainty (recalling the restriction $p \geq 1/2$).[20] If the p parameter varies across tasks, one can decrease (weakly) all the p_k's.

[20] There is also another way to increase uncertainty: one can simply decrease p_k holding q_k constant for each k. Our result holds also in this case.

We find that an increase in uncertainty leads to a decrease in the fraction of tasks regulated by spot contracting. The intuition is simple: if uncertainty is higher, the external state is expected to fluctuate more over time, and this increases the expected cost of using a spot approach, whereas the cost of a contingent rule is not affected by uncertainty.

Another parameter that has an important impact on the tradeoff between contingent and spot contracting is the discount factor δ. As usual, this parameter can be interpreted as capturing not only the players' pure time discounting but also the probability at a given period that the game will continue in the next period. For this reason, we will often think about δ as capturing the *durability* of the relationship.

We find that the fraction of tasks regulated by spot contracting is decreasing in δ. The intuitive reason is that a spot approach saves on the cost of describing external events in the first period at the price of paying the cost of describing actions in future periods. The following proposition, whose proof can be found in the Appendix together with the proofs of all remaining results, states these results.

Proposition 2. Under assumption (SC):

(i) an increase in uncertainty leads to a decrease in the fraction of tasks regulated by spot contracting.
(ii) the fraction of tasks regulated by spot contracting is decreasing in δ.

This proposition suggests an interesting empirical prediction: spot contracting should be relatively more prevalent when the external environment is less uncertain and when the relationship is less durable. Thus, the stability of the external environment and the internal stability of the relationship have opposite effects on the tradeoff between spot and contingent contracting.

It is also worth highlighting a simple corollary of Proposition 2: if uncertainty varies across tasks, then tasks characterized by a higher degree of uncertainty are more likely to be regulated by contingent clauses, whereas lower-uncertainty tasks are more likely to be regulated by a spot approach.

Finally, one can examine the effects of changes in contracting costs, using the result of Proposition 1. One might take the view that, if the legal system improves over time, contracting costs tend to go down for all tasks. Then Proposition 1 tells us that the effect of this change depends on exactly how the various contracting cost parameters go down, and in particular on how the ratio $\frac{c(e)}{r \cdot c(a)}$ changes. Thus, for example, if the costs of describing events and actions ($c(e)$ and $c(a)$) go down proportionally, then the tradeoff between spot and contingent contracting is not affected. But one might think that also the recalling cost r tends to go down, because this parameter in a broad sense captures the cost of modifying an existing contractual provision, and this is probably lower when the legal system is more efficient. In this case, then, a general reduction of contracting costs would favor spot contracting over contingent contracting.

□ **Digression 1. The role of language.** Next, we argue that our language-based approach to writing costs plays an important role for the predictions of the model. To make this point, we consider a more "traditional" specification of writing costs, similar in spirit to Dye (1985a).

Let $\{s^1, \ldots, s^M\}$ be the set of states and $\{b^1, \ldots, b^M\}$ the set of behaviors (where $M = 2^N$), and assign indices so that it is efficient to do b^j if and only if the state is s^j, for all j. Now assume that, unlike our model, it is not possible to break down the description of a state or behavior into its elementary constituents. Let c^s be the cost of describing a state and c^b the cost of describing a behavior, and suppose that c^s and c^b are small, so that it is optimal to implement the first-best outcome. Keep all other assumptions of our model unchanged.

In this version of the model, if N is not too small, a sequence of spot contracts is optimal, and in particular it dominates a contingent contract. To see this, note that a complete contingent contract must specify the efficient behavior b^j for each state s^j and therefore its cost (paid once and for all) is $2^N(c^s + c^b)$, whereas a sequence of spot contracts costs at most c^b in each period, and therefore its discounted cost does not exceed $\frac{c^b}{1-\delta}$.

Thus, this alternative specification of writing costs implies that spot contracting is typically optimal. This is in stark contrast with our model, where a contingent contract may well be optimal even if N is large; indeed, with our specification of writing costs, the optimal Markovian policy is essentially independent of N (recall that the problem is separable in the N tasks). For example, suppose $N \geq \log_2(\frac{1}{1-\delta})$. In this case, under the specification *à la* Dye, contingent contracting is dominated for any c^s and c^b, whereas under our specification, contingent contracting is optimal for a whole region of parameters (in particular, if (SC) holds and $\frac{c(e_k)}{r \cdot c(a_k)}$ is sufficiently small for all k).

This should clarify our statement that the nature of language matters for the predictions of the theory. The question of which type of language is more relevant is an open one, but we have argued elsewhere (Battigalli and Maggi, 2002) that a language of the type considered in this article is likely to be more efficient than a Dye-type language, and that it is closer to the languages that we observe in reality.

Another remark about language is in order. We assumed that the language described at the outset is the only common-knowledge language. In principle, the parties could construct a new language, for example by attaching a new primitive sentence to each state and to each behavior, and write a contract with the new language. Note that the parties would have to attach a vocabulary that translates the new language into the original one, in order for the courts to be able to interpret the contract. If the relationship is one-shot, the new language cannot be more efficient than the original one, because the cost of writing the vocabulary in the contract is at least as large as its benefits. In a repeated relationship, however, this approach might in principle be efficient.[21] A more general model would allow for this kind of recoding of the language, but we conjecture that the main qualitative results would not be affected. We already allow for the possibility of "recalling" the description of an action a at low cost; the effect of coding would probably be very similar.

□ **Digression 2. The Maskin-Tirole argument.** In a well-known 1999 paper, Maskin and Tirole have argued that unforeseen contingencies (or, by a straightforward extension of their argument, the costs of describing contingencies) do not imply inefficiencies in contracting, provided that an appropriate message-based mechanism is played after contingencies are observed and before actions are taken. Here we argue that the results of our model would not change if such mechanisms were available, and hence the Maskin-Tirole irrelevance argument does not apply to our setting.

The key observation is that a message-based mechanism would offer no advantage relative to a spot contract. To see this, first recall an important aspect of our model that distinguishes it from Maskin and Tirole's model: here there are no *ex ante* investment actions. For this reason, in our model, spot contracts are sufficient to induce the agent to take the efficient actions. In principle, a message-based mechanism might be useful if it allowed parties to avoid the costs of describing the *actions* to be taken. But can this be the case? The answer is no. Even if a message-based mechanism is played after contingencies are observed, the actions to be taken must still be described formally, that is, in such a way that they can be enforced in court. But if the formal description of actions is subject to the costs that we have assumed in this article, then a message-based mechanism can never be preferable to a spot contract. It follows immediately that these mechanisms are redundant in our setting.

4. Formal and informal contracting

■ In reality, long-term relationships are often governed by informal contracts, or by a combination of formal and informal contracts. Informal contracts have an important advantage

[21] We thank Leonardo Felli and Luca Anderlini for bringing this point to our attention.

over formal contracts, namely that they do not require a written description of actions and contingencies, so they save on writing costs.[22] The shortcoming of informal contracts, on the other hand, is the absence of an external enforcement mechanism. Because an informal contract must be self-enforcing, if players are not sufficiently patient it may not be possible to implement the first-best outcome with a fully informal contract. In what follows, we will examine more closely this tradeoff between formal and informal contracting. Our main objective in this section is to understand under what conditions it is efficient to govern the relationship by formal contracting, by informal contracting, or by a combination of the two, and how changes in the contracting environment affect the optimal contracting mode.

□ **Optimal contracting plans.** Consider the game of Section 2. The way we allow for informal contracts is by focusing on efficient subgame perfect equilibria, of the game. In such equilibria, there may be tasks for which the agent always takes the efficient action even if not constrained by a formal clause, in which case we say that the task is governed by an *informal* (contingent) rule. Informal rules are enforced by the threat of reverting to an equilibrium that is worse for the deviator, rather than by external courts.

For reasons of tractability and to facilitate comparison with the case of formal contracting, we focus on subgame perfect equilibria of a simple kind: we assume that on the equilibrium path, each task is governed either by an informal rule or by one of the formal rules considered in the previous section ($\mathcal{C}$, $\mathcal{E}$, or $\mathcal{S}$). In analogy with the standard terminology of self-enforcing contracts, we label equilibria of this kind *self-enforcing contracting plans.* A self-enforcing contracting plan is *optimal* if there is no other self-enforcing plan that gives a higher payoff to both players. By risk neutrality and transferable utility, optimal contracting plans are those that maximize the present value of the total surplus net of writing costs.[23] The assumption that players focus on optimal contracting plans seems reasonable for situations in which there is sufficient mutual understanding between players, so that they are able to coordinate on a "good" equilibrium.

Optimal behavior is supported by the threat of punishment. In this type of game, the harshest punishment that can be inflicted on a player is to give him an overall payoff (PDV) of zero, that is, the player's *maxmin*. The way we will structure the presentation is the following: first we explain how tasks are optimally governed on the equilibrium path, assuming provisionally that maxmin punishments are credible (i.e., subgame perfect); and then we consider the conditions under which maxmin punishments are indeed credible. As will become clear, maxmin punishments are credible if the discount factor is above a certain threshold.

Such a contracting plan is described by a tuple $(\mathbf{I}, \mathbf{C}, \mathbf{E}, \mathbf{S}, (m_t)_{t\geq 1})$, where bold capital letters denote subsets of tasks and m_t is a stochastic wage, that is, $m_t : H_t \to \mathbb{R}$, where H_t is the set of histories of shocks $h_t = (s_1, \ldots, s_t)$.[24] $(\mathbf{I}, \mathbf{C}, \mathbf{E}, \mathbf{S})$ is an ordered partition of the set of tasks $\mathbf{N} = \{1, \ldots, N\}$ specifying that tasks $k \in \mathbf{I}$ are agreed-upon informally and tasks $k \in \mathbf{C}$ (respectively, $\mathbf{E}, \mathbf{S}$) are regulated formally by a $\mathcal{C}_k$ (respectively, $\mathcal{E}_k$, $\mathcal{S}_k$) rule. (Recall that a "rule" specifies how the clause for a given task evolves over time as a function of the state realization.) We will refer to a tuple $(\mathbf{I}, \mathbf{C}, \mathbf{E}, \mathbf{S})$ as a *task partition.*

We say that a task partition is *incentive compatible* if it is part of a self-enforcing contracting plan. Next we derive a necessary condition for the incentive compatibility of a task partition, assuming that maxmin punishments are credible. If some tasks are regulated by informal rules, the agent has the opportunity to shirk on those tasks (i.e., increase his current utility by not taking the

[22] Arguably, the cost of writing an enforceable contract is higher than the cost of communicating an informal agreement, because for the contract to be enforced by courts it must be written according to the commonly accepted legal standards, which may be quite cumbersome to meet. In particular, it is not sufficient that the language used in the formal contract be common knowledge to the contracting parties; it has to be common knowledge to the parties *and* the courts, and this may require effort and skills (or lawyers).

[23] The reader can refer to Levin (2003) for a proof of an analogous statement.

[24] It should be clear that, even if the wage is state dependent, it is not contractually specified as a contingent wage in the formal contract, but it is written period by period after observing the state s_t, so it involves no writing costs.

efficient action). The most effective way to prevent the agent from shirking on the informal tasks is to give him all the net surplus from the second period onward if he does not shirk and give him a PDV of zero if he shirks. Then the agent will not shirk only if the PDV of future net surpluses is at least as large as the current benefit from shirking. To express this condition formally, it is convenient to introduce some more notation. Let $\Pi(\mathbf{K}) = \sum_{k\in\mathbf{K}} \pi_k$ and $D(\mathbf{K}) = \sum_{k\in\mathbf{K}} d_k$ where $\mathbf{K} \subseteq \mathbf{N}$. Also, let $\widehat{c}_t(\mathbf{C}, \mathbf{E}, \mathbf{S} \mid h_\tau)$ denote the expected present value of writing costs from time t conditional on history $h_\tau (\tau \leq t)$.[25] Then it is clear that a task partition is incentive compatible only if the following set of aggregate incentive constraints hold:

$$\forall t \geq 1, \forall h_t \in H_t, \qquad \frac{\delta}{1-\delta}[\Pi(\mathbf{N}) - D(\mathbf{N})] - \delta \widehat{c}_{t+1}(\mathbf{E}, \mathbf{S} \mid h_t) \geq D(\mathbf{I}). \tag{IC}$$

The left-hand side of (IC) is the expected present value of future net surpluses conditional on history h_t. The right-hand side of (IC) is the agent's disutility from performing the informal tasks, which is also the benefit from shirking.[26]

We prove in the Appendix that, under the same condition that makes maxmin punishments credible, condition (IC) is also sufficient for incentive compatibility. It follows that an optimal contracting plan minimizes the PDV of writing costs subject to the incentive constraint (IC).

We now argue that, if δ is above a certain threshold, there exists a credible punishment strategy that keeps the deviator at his maxmin. Moreover, we show that the punishment can be constructed in such a way that the contracting plan after a deviation is optimal in the sense defined above. To convey the basic intuition, we focus on the case in which writing costs are negligible and we consider only two simple deviations: "stiffing" by the principal, that is, offering a lower-than-equilibrium wage, and "shirking" by the agent, that is, not performing some informally governed tasks. The idea is to punish the deviator by shifting all the surplus to the non-deviator through a change in the transfer, while keeping the task partition and the agent's actions as before the deviation.

Suppose first that the agent shirks at time $t-1$. Then from period t onward, the agent goes back to the same behavior as before the deviation, but the transfer changes: in period t, the transfer is such that the agent gets a PDV of zero (this may require a negative transfer); and from period $t+1$ onward, the transfer is such that the agent is the recipient of all the net surplus. This kind of stick-and-carrot punishment ensures that the agent has an incentive to keep working hard also in the punishment phase, in spite of getting a PDV of zero from period t on. Next consider a deviation by the principal. Suppose he offers a lower-than-equilibrium wage in period $t-1$. Then the agent rejects the offer (so that there is no surplus at time $t-1$), and from period t onward the players revert to the same behavior as before the deviation, except that the transfer is such that the principal gets a PDV of zero, and hence the agent becomes the recipient of all the net surplus. Note that for this punishment to be credible, the agent must have an incentive to reject the off-the-path offers by the principal. This incentive can be given in a simple way: if the agent accepts the off-the-path offer, he becomes the punished player in period t, and he gets punished just as described above.

Let us check that the proposed punishment strategy is credible. By construction, the principal has no incentive to deviate: this is because the agent rejects the offer and there is no surplus at $t-1$, hence there is not even a one-shot gain from cheating. Also, the agent clearly has no incentive to shirk. We only need to check if the agent has an incentive to reject off-the-path offers. We can focus

[25] Note that, for $t \geq 2$, $\widehat{c}_t(\mathbf{C}, \mathbf{E}, \mathbf{S} \mid h_\tau)$ depends only on $\mathbf{E}$ and $\mathbf{S}$ because the cost of contingent clauses is entirely paid up front.

[26] Another way to interpret (IC) is the following. Suppose the principal pays a stochastic "efficiency wage" $m_t(h_t)$ which in expectation is just enough to keep the agent from shirking. Then the expected wage must be $\overline{m} = D(\mathbf{N}) + \frac{1-\delta}{\delta} D(\mathbf{I})$. The principal would not renege on the promised wage after the first period if $\Pi(\mathbf{N}) - D(\mathbf{N}) - m_{t+1}(h_{t+1}) + \delta \frac{\Pi(\mathbf{N}) - D(\mathbf{N}) - \overline{m}}{1-\delta} - \widehat{c}_{t+1}(\mathbf{E}, \mathbf{S} \mid h_{t+1}) \geq 0$. Taking the expectation conditional on h_t and substituting the value of $\overline{m}$ we obtain (IC), which is therefore a necessary condition for the inequality above. Setting $m_{t+1}(h_{t+1}) = \overline{m} - [\widehat{c}_{t+1}(\mathbf{E}, \mathbf{S} \mid h_{t+1}) - \widehat{c}_{t+1}(\mathbf{E}, \mathbf{S} \mid h_t)]$ minimizes the strain on the binding constraint and makes (IC) fully equivalent to the inequality above.

on offers that give the agent less than the whole net surplus. According to the candidate equilibrium strategies, by accepting such an offer the agent would lose all the future net surpluses. Neglecting writing costs, a sufficient condition for rejection is $[\Pi(\mathbf{N}) - D(\mathbf{N})] < \frac{\delta}{1-\delta}[\Pi(\mathbf{N}) - D(\mathbf{N})]$. The left-hand side of this inequality is an upper bound to the benefit of accepting. The right-hand side is the present value of future surpluses, which is the opportunity cost of accepting (assuming that the candidate equilibrium strategies are followed from the next period). If $\delta > \frac{1}{2}$ this condition is satisfied, and the proposed punishment strategy profile is credible.

The punishment strategy we sketched above can be made precise and generalized to cover all possible deviations. This leads us to the following lemma, which incorporates both the result on punishment strategies and the characterization result illustrated above.

Lemma 1. Suppose that condition (SC) holds.

(i) There exists δ^* (function of other parameters) such that, for $\delta \geq \delta^*$, for every beginning-of-period subgame and every player i, there is an optimal contracting plan that keeps the continuation payoff of i at his maxmin (zero). The critical level δ^* approaches $\frac{1}{2}$ as writing costs become negligible.

(ii) If $\delta \geq \delta^*$, then a contracting plan is optimal if and only if the corresponding task partition solves the following problem:

$$\begin{aligned} &\min_{\mathbf{I},\mathbf{C},\mathbf{E},\mathbf{S}} \widehat{c}(\mathbf{C}, \mathbf{E}, \mathbf{S}) \\ &\text{s.t. } (IC), \end{aligned} \tag{P}$$

where $\widehat{c}(\mathbf{C}, \mathbf{E}, \mathbf{S})$ denotes the unconditional PDV of writing costs.

Note that the punishment strategies identified by Lemma 1(i) are credible in the strong sense that they are not Pareto dominated. Therefore, optimal contracting plans that employ such punishment strategies satisfy a form of renegotiation proofness. This appealing feature of our punishment strategies stems from the fact that the relationship is not broken following a deviation but reverts to a "normal" state, and the way a deviator is punished is essentially that he pays a fine to the other player (in the form of a lower wage if the deviator is the agent, or a higher wage if the deviator is the principal).[27] We also note that assumption (SC) is not necessary for part (i) of the lemma, although it allows for a simpler proof.[28]

We are now ready to analyze optimal contracting plans. In particular, we want to examine under what conditions the optimum is fully informal, fully formal, or involves a mix of formal and informal contracting. The following proposition provides an answer to this question.

Proposition 3. Let $d^{\min} \equiv \min_k d_k$. At an optimal contracting plan:

(i) if $\frac{\delta}{1-\delta} \geq \frac{D(\mathbf{N})}{\Pi(\mathbf{N})-D(\mathbf{N})}$, all tasks are regulated informally.

(ii) if $\frac{d^{\min}}{\Pi(\mathbf{N})-D(\mathbf{N})} < \frac{\delta}{1-\delta} < \frac{D(\mathbf{N})}{\Pi(\mathbf{N})-D(\mathbf{N})}$, and writing costs are sufficiently small, some tasks are regulated formally and some are regulated informally.

(iii) if $\frac{\delta}{1-\delta} < \frac{d^{\min}}{\Pi(\mathbf{N})-D(\mathbf{N})}$, all tasks are regulated formally.

Proposition 3 states that formal and informal contracting will coexist if the discount factor lies in some intermediate range. If the discount factor is sufficiently small, then introducing any

[27] As in Levin (2003), our punishment strategies are Pareto dominated by other equilibrium strategies in subgames starting with a response by the agent, whenever they prescribe a rejection. However, we think this is reasonable if the agent has to reply immediately, and hence there is no time for negotiations about a complex self-enforcing agreement.

[28] Unlike (SC), condition $\delta \geq \delta^*$ is essential for Lemma 1. (To see this, consider the extreme case of δ equal to zero: then there is a unique subgame perfect equilibrium (SPE)—the Markovian equilibrium—in which the principal offers a formal contract and makes a positive profit in every period.) However, we suspect this condition is not essential for our qualitative results. If it is not satisfied, it may not be possible to keep the principal at his maxmin payoff in the punishment phase, in which case the principal's incentive constraints will be more stringent, and this is likely to result in fewer tasks being regulated informally, but our qualitative results are still likely to hold.

of the tasks in the informal contract violates (IC), and therefore the optimum is fully formal. If, on the other hand, the discount factor is sufficiently high, a fully informal contract is incentive compatible, and hence optimal.[29]

To gain intuition for Proposition 3 recall that, in the equilibria we are considering, the agent anticipates getting the PDV of future net surpluses ($\frac{\delta}{1-\delta}[\Pi(\mathbf{N}) - D(\mathbf{N})]$) if he does not shirk, and note that $d^{\min}$ is a lower bound to the agent's one-time benefit from shirking, while $D(\mathbf{N})$ is an upper bound to this benefit. Clearly, then, if the PDV of future net surpluses is lower than $d^{\min}$, no task can be governed by informal contracting; this is case (iii). It should also be clear that, if the PDV of future net surpluses exceeds $D(\mathbf{N})$, all the tasks can be governed by informal contracting (case i). And if the PDV of future net surpluses lies strictly between $d^{\min}$ and $D(\mathbf{N})$, some but not all tasks can be governed informally, and hence the optimum involves a mix of formal and informal contracting (case ii).[30]

It is worth emphasizing that, if there is at least one task with relatively low disutility (i.e., if $d^{\min}$ is relatively low), then the parties can sustain at least *some* informal contracting, and this is true *even if writing costs are very small*. This is in sharp contrast with a result obtained by Baker et al. (1994), who show that if the imperfection in formal contracting is small, it is impossible to sustain any informal contracting. This is not due to the differences between their contracting game and ours, but rather to a restriction that they make on the punishment strategy: instead of focusing on optimal punishments, as we do here, they assume that deviations are punished by reversion to the static Nash equilibrium of the contracting game. In our model, the analogous restriction would be to assume that deviations are punished by reversion to the Markov perfect equilibrium: if writing costs are sufficiently small, this punishment strategy cannot deter deviations by the principal, hence it cannot support any informal contracting.[31]

□ **The impact of changes in the contracting environment.** We now examine how changes in the fundamental parameters affect the mode of contracting. We will focus on changes in the degree of uncertainty, the level of contracting costs, the amount of available surplus, and the durability of the relationship (discount factor).

To gain tractability, we analyze a version of the model with a continuum of tasks. This simplifies the analysis because it ensures that the incentive constraint is satisfied as an equality and allows the use of differential calculus. The model with a continuum of tasks should be thought of as an approximation of the discrete model with a large number of tasks.[32]

We also impose some symmetry assumptions. For each task, the cost of describing an elementary event is c^e, the cost of describing an elementary action is c^a, and the probability that the same event will occur in the next period is p (symmetric persistence).

[29] The reason we need sufficiently small costs at point (ii) is that if noncontingent formal rules are used, then the present value of future writing costs enters the incentive constraint. See the proof of Proposition 3 in the Appendix.

[30] Proposition 3 does not tell us *which* tasks are formally regulated and which tasks are informally regulated. Intuitively, in this model, there is a tendency for lower-disutility tasks to be regulated informally and for higher-disutility tasks to be regulated formally, because the former impose less strain on the incentive constraint than the latter. The reason this is only a "tendency" is because of the discrete nature of the optimization problem. This will become an exact statement in the next subsection, where we consider a version of the model with a continuum of tasks and impose some symmetry assumptions.

[31] To see this, note first that if some tasks are regulated informally, the principal must pay an efficiency wage in order to motivate the agent to work (see footnote 26). This creates a temptation for the principal to renege on the wage promised to the agent, and hence the principal has a strictly positive one-time gain from cheating. Second, note that the MPE gives all the surplus to the principal, and if writing costs are small, this is close to the maximum potential surplus. Thus the future loss from punishment for the principal, if any, is negligible, and cannot outweigh his one-time gain from cheating.

[32] The reason we did not consider a continuum of tasks from the start is that the symmetry assumptions we impose in this section for tractability reasons are not needed in the previous analysis and exclude interesting possibilities, such as the use of enrichment rules. Thus the continuum-of-tasks version of the model allows for clean comparative-statics analysis, but this comes at the price of losing a bit of generality.

Without loss of generality, we can order tasks according to increasing disutility. For any positive fraction of tasks $x \in (0, 1]$, let $d(x)$ denote the disutility of the x-quantile of the distribution of tasks; for example, $d(\frac{1}{2})$ is the disutility of the median task. Similarly, $\pi(x)$ is the principal's profit from the x-quantile task if the agent takes the efficient action for this task. By construction, $d(x)$ is an increasing function. We let $D(x) = \int_{(0,x]} d$ denote the cumulative disutility of the set of tasks below the x-quantile. Similarly, we let $\Pi(x) = \int_{(0,x]} \pi$. Because d is increasing and strictly positive, D is strictly increasing and convex. We make the further regularity assumption that d is *strictly* increasing. Then $D(x)[\Pi(x)]$ is the cumulative disutility (profit) of the tasks with disutility lower than $d(x)$. We maintain the small-costs assumption: $\pi(x) - d(x) > c^e + c^a$ for all x. Finally, we assume that the continuum of tasks has size $N = 1$ (normalization).

Let us reconsider problem (P) by going to the continuum limit. First notice that, by symmetric persistence, enrichment rules are not optimal (cf. Proposition 1). Then a task partition is a measurable function $\boldsymbol{\Psi} : (0, 1] \to \{I, C, S\}$. We let x_I denote the *fraction* of tasks governed by informal rule, that is, the (Lebesgue) measure of the set $\boldsymbol{\Psi}^{-1}(I)$; similarly, x_C and x_S are, respectively, the fractions of tasks governed by contingent and spot rule.

Our assumptions imply that x_I must be strictly positive at an optimum: by the small-cost assumption the expected PDV of future net surpluses of a fully formal contracting plan is strictly positive, and therefore it is possible to save on writing costs by switching a small fraction of tasks to informal rule without violating the incentive constraint. Note that this feature is a consequence of the continuum of tasks: unlike case (iii) in Proposition 3, here it is always possible to govern informally at least a small measure of tasks.

It is clear that, if δ is sufficiently high, a fully informal contracting plan is self-enforcing, and hence it is optimal. A necessary and sufficient condition for this is $\frac{\delta}{1-\delta} > \frac{D(1)}{\Pi(1)-D(1)}$ (cf. Proposition 3). Next we focus on the more interesting case $\frac{\delta}{1-\delta} < \frac{D(1)}{\Pi(1)-D(1)}$, where the optimum involves a mix of formal and informal contracting. In this case, the incentive constraint must be satisfied as an equality at an optimum (suppose not, then it is possible to increase slightly the fraction x_I without violating the constraint, thus decreasing writing costs).

Next we make an important observation: at an optimum, *the tasks with disutility below a certain threshold d^* are governed by informal rule, and the tasks with disutility higher than d^* are governed by formal rule* (neglecting sets of tasks of measure zero). To see this, suppose the above statement is not true. Then there must be two subsets of tasks X_I and X_F with the same positive measure such that for all $x' \in X_I$, $x'' \in X_F$, x' is governed by informal rule, x'' is governed by formal rule, and $d(x') > d(x'')$. Now make the tasks in X_F informal and those in X_I formal: this does not change total expected writing costs (because tasks are homogeneous with respect to writing costs and probabilities of shocks), but creates slack in the incentive constraint, because it decreases the total disutility of tasks governed by informal rule; this slack can be used to increase X_I and lower writing costs.

Note that, under our symmetry assumptions, it does not matter which of the higher-disutility tasks are governed by contingent rule and which are governed by spot rule. Given the observations we just made, we can formulate the problem in terms of choosing the fractions of tasks governed by the three rules (x_I, x_C, and x_S), and we have the following preliminary result.

Lemma 2.

(i) If $\frac{\delta}{1-\delta} > \frac{D(1)}{\Pi(1)-D(1)}$, the optimum entails $x_I^* = 1$ (fully informal contracting).
(ii) If $\frac{\delta}{1-\delta} < \frac{D(1)}{\Pi(1)-D(1)}$, then at an optimum the incentive constraint is satisfied with equality, the tasks in $(0, x_I^*]$ are governed by informal rule, and the tasks in $(x_I^*, 1]$ are governed by formal rule, where $0 < x_I^* < 1$. Therefore the problem can be restated as

$$\min_{0 \leq x_I, x_C, x_S \leq 1} \left[(c^e + c^a) x_C + \left(c^a + \frac{\delta(1-p)}{1-\delta} r c^a \right) x_S \right] \tag{P'}$$

$$\text{s.t.} \quad x_I + x_C + x_S = 1$$

$$\text{and} \qquad \frac{\delta}{1-\delta}\left[\Pi(1) - D(1) - (1-p)rc^a x_S\right] - D(x_I) = 0. \tag{IC'}$$

The (IC') constraint has a similar interpretation as the (IC) constraint in the previous section, but it is important to emphasize one aspect that will play a key role in what follows. It is transparent from (IC') that the use of spot rules strains incentives, whereas the use of contingent rules does not. Formally, increasing x_S whereas holding x_I constant decreases the slack in (IC') ($\frac{\delta}{1-\delta}[\Pi(1) - D(1) - (1-p)rc^a x_S] - D(x_I)$), whereas x_C does not enter (IC') at all. Intuitively, this is a consequence of the fact that the cost of contingent rules is paid upfront, while the cost of spot rules is incurred repeatedly over time.

We are now ready to analyze the effects of parameter changes on the fractions x_I, x_C, and x_S. We consider small parameter changes, and focus on the more interesting case $\frac{\delta}{1-\delta} < \frac{D(1)}{\Pi(1)-D(1)}$. We will consider changes in the discount factor, the degree of uncertainty, the level of writing costs, and the available surplus. Because we treat the available surplus as a parameter, in what follows we use the shorter notation Π (rather than $\Pi(1)$).

One might conjecture that parameter changes that improve incentives (in the sense that they create slack in the incentive constraint (IC') given the initial task partition) would increase the extent of informal contracting. We show that this is indeed the case at non-interior solutions of (P'), where either $x_S = 0$ or $x_C = 0$, but surprising results may obtain if we start from an interior solution ($x_S > 0, x_C > 0$).

If spot contracting is not used ($x_S = 0$), (IC') yields $x_I = D^{-1}(\frac{\delta}{1-\delta}[\Pi - D(1)])$; therefore an increase in δ or Π yields an increase in $x_I (D^{-1}$ is increasing) compensated by a decrease in x_C. Other parameter changes do not have any effect because they do not affect (IC') in this case. If contingent contracting is not used ($x_C = 0$), then we have $D(x_I) = \frac{\delta}{1-\delta}[\Pi - D(1) - (1-p)rc^a(1-x_I)]$, and hence an increase in δ, Π, or p, or a decrease in writing costs, increases x_I.

The analysis is less straightforward if we start from an interior solution. The effects of small parameter changes can be gauged by inspection of the first-order conditions. Differentiating the Lagrangian of problem (P') and eliminating the multiplier, we obtain

$$\frac{d(x_I) - \frac{\delta}{1-\delta}(1-p)rc^a}{d(x_I)} = \frac{c^a + \frac{\delta}{1-\delta}(1-p)rc^a}{c^a + c^e}. \tag{FOC}$$

The right-hand side is the ratio between the cost of spot contracting ($c^a + \frac{\delta}{1-\delta}(1-p)rc^a$) and the cost of contingent contracting ($c^a + c^e$).[33] Because the left-hand side is less than 1, at an interior solution spot contracting must be cheaper than contingent contracting ($c^a + c^e > c^a + \frac{\delta}{1-\delta}(1-p)rc^a$). If the opposite inequality holds, then spot contracting is not used at all, because it is more costly and also hurts incentives.

First note that the surplus Π does not appear in (FOC), and therefore a change in Π does not affect x_I. Thus, interestingly, an increase in the available surplus does *not* lead to more informal contracting, even though it improves incentives.

Next consider an increase in the degree of uncertainty $(1-p)$. This increases the cost of spot contracting, and hence it strains incentives (as we remarked above). But surprisingly, this leads to an *increase* in informal contracting (x_I). To see this, note that, given the task partition, an increase in $(1-p)$ makes the left-hand side of (FOC) smaller than the right-hand side; to restore the equality it is necessary to increase $d(x_I)$, hence x_I, because this increases the numerator of the left-hand side proportionally more than the denominator. A way to understand this result intuitively is the following: an increase in uncertainty makes spot contracting less attractive, causing a replacement of spot rules with contingent rules, which in turn creates slack in the incentive constraint and allows an increase in x_I. An analogous result obtains for an equiproportional increase in writing costs: if c^a and c^e increase proportionally, the left-hand side

[33] The reader can check that the expressions for the costs of spot and contingent contracting developed in Section 3 reduce to the expressions above when the symmetry assumptions are imposed.

of (FOC) becomes smaller while the right-hand side does not change, and this implies that x_I must go up. Again, here is a parameter change that strains incentives, but paradoxically leads to an *increase* in the extent of informal contracting.

Finally, note that an increase in the discount factor δ has the intuitive effect of decreasing informal contracting, just as in the case of non-interior solution. The following proposition summarizes the results we just highlighted.[34]

Proposition 4.

(i) At a non-interior optimum ($x_S = 0$ or $x_C = 0$), any small parameter change that improves incentives (an increase in δ, Π, or p, or a decrease in writing costs) increases x_I.

(ii) At an interior optimum ($x_S > 0, x_C > 0$), an increase in Π has no effect on x_I, even though it improves incentives; an increase in writing costs or uncertainty leads to an increase in x_I, even though it strains incentives; and an increase in δ has the intuitive effect of increasing x_I.

Thus far, we have focused on the tradeoff between formal and informal contracting. We now consider the tradeoff between spot and contingent contracting. Recall that in Section 3 we examined this tradeoff in a situation of purely formal contracting; here we ask how this tradeoff changes when informal contracting is available.

We start with a preliminary observation, which follows directly from the discussion above. Whereas in the case of purely formal contracting the only relevant consideration for the choice between spot and contingent rules was the comparison between their costs, in the presence of informal contracting the incentive properties of these rules also become important. And as we pointed out above, contingent rules have better incentive properties than spot rules, because the cost of contingent rules is paid upfront, whereas the cost of spot rules is incurred over time. As a consequence, *the presence of informal contracting tends to favor contingent rules over spot rules*. More concretely, it is easy to show the following: a task may be governed by contingent rule even though a spot rule is less costly ($\frac{\delta}{1-\delta}(1-p)rc^a < c^e$), whereas the reverse case is not possible. Thus, the model predicts that the relative importance of contingent versus spot contracting should be higher in relationships where formal and informal contracting coexist, compared with relationships that are governed exclusively by formal contracting.

Next we consider the comparative-statics effects of parameter changes on the importance of spot and contingent contracting, as captured by x_C and x_S. We focus on the case in which the solution is interior, so that both contingent and spot rules are used. Note that, although x_I can be determined simply by using (FOC), in order to determine x_C and x_S we need to also use the conditions (IC') and $x_I + x_C + x_S = 1$.

Consider first an increase in uncertainty or an equiproportional increase in writing costs. These parameter changes make spot contracting more costly and also strain incentives. Intuitively, then, the effect is a substitution of spot rules with contingent rules. Recall also from the analysis above that these parameter changes increase x_I, and therefore we have $\Delta x_C < |\Delta x_S|$. An increase in Π has the opposite effect: it improves incentives without changing the cost of spot and contingent contracting, and hence it leads to a more intensive use of spot rules, which are less costly; because Π does not appear in (FOC), x_I does not change and $\Delta x_S = -\Delta x_C$. The effect of an increase in the durability of the relationship is ambiguous. The reason is that an increase in δ increases the cost of spot contracting, but at the same time it improves incentives, thus decreasing the attractiveness of contingent rules, so the net effect cannot be signed. The following proposition summarizes the relevant comparative-statics results.

Proposition 5. At an interior optimum: (i) an increase in uncertainty or writing costs leads to a decrease in x_S and an increase in x_C, with $\Delta x_C < |\Delta x_S|$; (ii) an increase in Π leads to an increase in x_S and a decrease in x_C, with $|\Delta x_C| = \Delta x_S$.

[34] In this proposition, when we speak of an increase in x_I we mean it in the weak sense. Also, by "increase in writing costs" we mean an equiproportional increase in c^a and c^e.

It is interesting to compare these results with those we found in the case of purely formal contracting. There, a change in profits or an equiproportional change in writing costs had no effect on the tradeoff between spot and contingent rules, because these parameter changes do not affect the relative cost of the two rules. Here these parameter changes do have an effect, because they have an impact on the incentive constraint (IC'). Changes in uncertainty, on the other hand, have the same qualitative impact as under pure formal contracting, because they affect cost considerations and incentive considerations in the same direction.

5. Conclusion

■ In this concluding section, we discuss briefly some possible extensions of the model.

Our analysis focused on the case in which writing costs are relatively small. In the working paper version (Battigalli and Maggi, 2003), we also consider the case of large writing costs. The main change in results is that it may no longer be optimal to implement the first best for some (or all) of the tasks. In particular, two additional possibilities emerge: it might be optimal to regulate a task by *rigid rule*, that is, by writing a noncontingent clause once and for all,[35] or to leave a task to the agent's *discretion* with no informal agreement to take the efficient action. If we interpret a contract implementing the first-best outcome as a "complete" contract, then the main implication of large writing costs is that they generate contractual incompleteness. In the working paper we also show that, other things equal, tasks characterized by low surplus are left to the agent's discretion, intermediate-surplus tasks are regulated by rigid rules, and high-surplus tasks are regulated in a first-best way (by formal or informal contracting), similarly as in our static model (Battigalli and Maggi, 2002).

We considered only "variable" writing costs, that is, costs that increase with the number of modifications to the contract. One may also consider the impact of "quasi-fixed" writing costs, that is, costs that are incurred every time the contract is modified.[36] Quasi-fixed writing costs would affect the mode of formal contracting in two ways: first, they would favor contingent contracting over spot contracting, because the latter involves contract modifications over time and the former does not; second, if spot contracting were nevertheless efficient, there would be a tendency to postpone some contract modifications so as to "cluster" several modifications in a single period. This would complicate the analysis without adding much insight.

We assumed that the principal makes the payment before the agent acts, and that the payment is exactly the one specified in the contract (i.e., we did not consider the possibility of "bonuses"). It is clear that each of these assumptions is without loss of generality given the other. If bonuses are not allowed, it does not matter when the payment m occurs, because the principal has to pay m independently of whether the agent has shirked on the informally regulated tasks or not. On the other hand, if every payment has to be made before the agent acts, bonuses are redundant, because the best incentive to keep the agent from shirking is still to hold him down to his maxmin from period $t + 1$ if he shirks in period t. We now argue that even the *joint* assumption about timing and bonuses is without loss of generality. Suppose that the principal is allowed to pay an informal bonus immediately after the agent acts. In this case, the agent has a stronger incentive not to shirk, because shirking will prevent him from enjoying an immediate reward. On the other hand, with the modified assumption, the principal has an incentive to renege on the informally promised bonus, whereas with our current assumption he can only omit to offer the formal contract specified by the equilibrium. Due to the quasi-linearity of the utility functions, these two effects cancel out in the "aggregate" incentive constraint of Section 4.

[35] This feature seems broadly consistent with the empirical findings by Kahan and Klausner (1997). In their analysis of bond indentures, they argue that contract terms are modified less often than the first best would call for (see their Section II.3).

[36] Gray (1976, 1978) and Dye (1985b) present macroeconomic models that feature recontracting costs of this kind.

Similarly, it can be shown that there is no benefit from writing formal multiperiod contracts. To gain intuition, suppose there is only one task. Even though committing to wages for the next T periods would remove the principal's incentive to renege on these payments, this would worsen the agent's incentives, because if the agent shirks on an informally regulated task the punishment is delayed. Again, due to the quasi-linearity of the utility functions, these two effects cancel out in the aggregate incentive constraint. And if the task is regulated formally, a multiperiod contract offers no gain relative to a sequence of one-period contracts.

Finally, in this article, we have implicitly assumed away an alternative mode of governance that could avoid the costs of writing detailed contracts, namely giving *authority* to the principal. If the principal could instruct the agent on what actions to take in each period, there would be no need to specify contingencies or actions in a contract. How would our results change if we allowed for authority as a governance mode?

To address this question, let us distinguish between *formal* and *informal* authority. We speak of formal authority when the principal's authority is specified in the formal contract, and the agent can be punished by courts for disobeying the principal. We speak of informal authority when the principal's orders are not enforced by courts, but by credible punishment mechanisms. Let us focus on formal authority first. It is critical to note that formal authority is enforceable only if two conditions are met: (i) the principal can send verifiable messages to the agent; this requires that messages be written, or at least recorded; and (ii) messages must be expressed in a language understood by the courts. In other words, messages must be *formal*. For this reason, even if a system of formal messages is feasible, it is not clear that its costs would be significantly lower than a system of formal contracts. Informal authority is a more common mode of governance in real organizations.[37] In our model, however, there is no role for such an arrangement, due to the assumption of symmetric information. Given that the principal and the agent have the same information, informal authority cannot improve on an informal contract as we defined it, because in the latter arrangement the agent knows what actions to take under any contingency, and hence there is no need for further instructions from the principal. A role for informal authority would probably arise if the principal had private information. An extension of the model in this direction is left for future research.

Appendix

■ *Proof of Proposition 1.* We need to prove that only the rules $\mathcal{C}_k$, $\mathcal{S}_k$, and $\mathcal{E}_k$ can be optimal. Due to additive separability and cross-sectional independence, the optimum can be found task by task. We have to prove that there are only three candidate optimal rules: $\mathcal{C}_k$, $\mathcal{S}_k$, and $\mathcal{E}_k$.

The Markov state for task k at time t has four dimensions:

$$\left(g_{k,t-1}, M^a_{k,t-1}, M^e_{k,t-1}, s_{k,t}\right) \in \{C_k, R_k, \overline{R}_k, D_k\} \times \left\{0^a_k, 1^a_k\right\} \times \left\{0^e_k, 1^e_k\right\} \times \{0_k, 1_k\}.$$

For example, $(g_{k,t-1}, M^a{}_{k,t-1}, M^e{}_{k,t-1}, s_{k,t}) = (R_k, 1^a_k, 0^e_k, 1_k)$ means that the clause from the previous period is R_k, the primitive sentence a_k has been used in the past (indeed, it has been used at time $t-1$ when the clause was R_k), the primitive sentence e_k has not been used in the past, and event e_k occurred at time t.

We call $g_{k,t-1}$ the "default clause," as this clause is binding in period t unless the contract is modified.

Note that not all $4 \times 2^3 = 32$ four-tuples are possible because a nondiscretionary clause ($g_{k,t-1} \neq D_k$) implies that $M^a_{k,t-1} = 1^a_k$ and possibly that $M^e_{k,t-1} = 1^e_k$ (if $g_{k,t-1} = C_k$).

Also note that by assumption (SC), optimality implies that the efficient action is always implemented, hence the empty clause D_k is never used, otherwise the principal would lose the opportunity to obtain a positive net surplus in the current period without any future advantage. Therefore, except for the initial state $(D_k, 0^a_k, 0^e_k, 1_k)$, we can disregard states where $M^a_k = 0^a_k$, as well as states where $g_k = D$ and either $M^a_k = 1^a_k$ or $M^e_k = 1^e_k$, because these states are unreachable under any optimal policy. Finally, note that optimality implies that if clause C has been used in the past, it must be kept forever. Once this clause is in place, the external state becomes irrelevant. The following table summarizes the relevant states, checked by $\times$ and labeled (from now on we drop the k index).

[37] But see MacLeod and Chakravorty (2004) for a case of formal authority in construction contracts.

	$(0^a, 0^e, 0)$	$(1^a, 0^e, 0)$	$(0^a, 1^e, 0)$	$(0^a, 0^e, 1)$	$(1^a, 1^e, 0)$	$(1^a, 0^e, 1)$	$(0^a, 1^e, 1)$	$(1^a, 1^e, 1)$
D	-	-	-	$\times\ (D.1)$	-	-	-	-
R	-	$\times\ (R.0)$	-	-	-	$\times\ (R.1)$	-	-
$\overline{R}$	-	$\times(\overline{R}.0)$	-	-	-	$\times(\overline{R}.1)$	-	-
C	-	-	-	-	$\times\ (C)$	-	-	$\times\ (C)$

Coalescing the states where clause C is in place, we obtain six relevant states.

- $D.1$: the initial state (empty clause, $s(e) = 1$),
- $R.0$: the default clause is R and $s(e) = 0$,
- $R.1$: the default clause is R and $s(e) = 1$,
- $\overline{R}.0$: the default clause is $\overline{R}$ and $s(e) = 0$,
- $\overline{R}.1$: the default clause is $\overline{R}$ and $s(e) = 1$,
- C: the default clause is C.

There are four possible actions, corresponding to the four possible clauses (including the empty one) that can be implemented at time t. Taking into account that the efficient action is always implemented, in each state there are only two candidate optimal clauses. Let $\psi^*(X)$ and $c^*(X)$, respectively, denote the optimal policy and the PDV of writing costs under this policy. Clearly, $\psi^*(C) = C$ and $c^*(C) = 0$ (the contingent clause is maintained forever at zero cost).

For the other five relevant states, there are only two candidate optimal actions: $\psi^*(D.1) \in \{C, R\}$, $\psi^*(R.0) \in \{C, \overline{R}\}$, $\psi^*(R.1) \in \{C, R\}$, $\psi^*(\overline{R}.0) \in \{C, \overline{R}\}$, and $\psi^*(\overline{R}.1) \in \{C, R\}$.

Next note that replacing the default clause with C in Markov states $R.1$ or $\overline{R}.0$ cannot be optimal, because the same flow of benefits can be obtained at lower cost by keeping the default clause and postponing its replacement with C until the external state changes. Thus $\psi^*(R.1) = R$, $\psi^*(\overline{R}.0) = \overline{R}$ and the values are as follows:

$$\begin{aligned}
c^*(D.1) &= \min\{c(a) + c(e), c(a) + \delta[pc^*(R.1) + (1 - p)c^*(R.0)]\},\\
c^*(R.0) &= \min\{r \cdot c(a) + c(e), r \cdot c(a) + \delta[(1 - q)c^*(\overline{R}.1) + qc^*(\overline{R}.0)]\},\\
c^*(R.1) &= \delta[pc^*(R.1) + (1 - p)c^*(R.0)],\\
c^*(\overline{R}.0) &= \delta[(1 - q)c^*(\overline{R}.1) + qc^*(\overline{R}.0)],\\
c^*(\overline{R}.1) &= \min\{r \cdot c(a) + c(e), r \cdot c(a) + \delta[pc^*(R.1) + (1 - p)c^*(R.0)]\}.
\end{aligned}$$

For example, consider state $D.1$. If it is optimal to choose C, then $c^*(C) = c(a) + c(e)$; if it is optimal to choose R, then with probability p there is a transition to $R.1$ in the next period (discounted by δ) and with probability $(1 - p)$ there is a transition to state $R.0$. Thus, in this second case, $c^*(D.1) = c(a) + \delta[pc^*(R.1) + (1 - p)c^*(R.0)] : c(a)$ is paid immediately and $\delta[pc^*(R.1) + (1 - p)c^*(R.0)]$ is the PDV of future writing costs. Similar considerations apply to the other relevant Markov states (recall that when $s = 0$, the transitions to $s = 0$ and $s = 1$ occur with probability q and $(1 - q)$, respectively).

It is intuitively clear that $c^*(R.1) \le c^*(\overline{R}.0)$, with equality when $p = q$: in both Markov states, the external environment matches the default clause and it is not necessary to change this clause, but the likelihood of a future mismatch is higher in Markov state $\overline{R}.0$, hence expected future costs are higher. Therefore,

$$\delta[pc^*(R.1) + (1 - p)c^*(R.0)] \le \delta[(1 - q)c^*(\overline{R}.1) + qc^*(\overline{R}.0)].$$

If $p > q$, we can consider three generic cases:

(i) $c(e) < \delta[pc^*(R.1) + (1 - p)c^*(R.0)]$: then $\psi^*(D.1) = C$ and the optimal rule is $\mathcal{C}$.
(ii) $c(e) > \delta[(1 - q)c^*(\overline{R}.1) + qc^*(\overline{R}.0)]$: then $\psi^*(D.1) = \psi^*(\overline{R}.1) = R$, $\psi^*(R.0) = \overline{R}$, and the optimal rule is $\mathcal{S}$.
(iii) $\delta[pc^*(R.1) + (1 - p)c^*(R.0)] < c(e) < \delta[(1 - q)c^*(\overline{R}.1) + qc^*(\overline{R}.0)]$: then $\psi^*(D.1) = \psi^*(\overline{R}.1) = R$, $\psi^*(R.0) = C$, and the optimal rule is $\mathcal{E}$ (note that state $\overline{R}.1$ is not reachable under this rule).

This shows that only the rules $\mathcal{C}$, $\mathcal{S}$, and $\mathcal{E}$ can be optimal. The rest of the proof is in the main text (preceding Proposition 1) and need not be replicated here.

Proof of Proposition 2.

(i) In the symmetric-persistence case, the relevant thresholds for the parameter $\gamma = \frac{c(e)}{r \cdot c(a)}$ collapse to the same value $\overline{\gamma}_{\mathcal{C}/\mathcal{S}} = \overline{\gamma}_{\mathcal{E}/\mathcal{S}} = \overline{\gamma}_{\mathcal{C}/\mathcal{E}} = \frac{\delta(1-p)}{(1-\delta)}$, which is decreasing in p. By Proposition 1, the number of tasks regulated by a spot rule is therefore increasing in p. (Using the general expressions for $\overline{\gamma}_{\mathcal{C}/\mathcal{S}}$ and $\overline{\gamma}_{\mathcal{E}/\mathcal{S}}$ it is also easy to check that these thresholds are decreasing in p keeping q fixed. This proves the claim made in footnote 20 that the number of tasks regulated by a spot rule is increasing if, for all tasks k, p_k increases with q_k fixed, which is another way to decrease uncertainty.)

(ii) We verify that the threshold $\overline{\gamma}_{\mathcal{E}/\mathcal{S}}$ is increasing in δ,

$$\frac{\partial \overline{\gamma}_{\mathcal{E}/\mathcal{S}}}{\partial \delta} = (1-q)\frac{[1-\delta(2p-1)]\{(1-\delta)+\delta(1-q)[1-\delta(p+q-1)]\}}{(1-\delta)^2[1-\delta(p+q-1)]^2} > 0.$$

By Proposition 1, it follows that the number of tasks regulated by a spot rule is decreasing in δ.

Proof of Lemma 1. (**i**) Suppose that task partition $\mathbf{\Psi} = (\mathbf{I}, \mathbf{C}, \mathbf{E}, \mathbf{S})$ is used from a certain period, that we normalize as period 1, with initial Markov state X. For every infinite sequence of shock vectors $(s_1, s_2, \ldots)$, partition $\mathbf{\Psi}$ determines the writing costs in each period. We let $\hat{c}(\mathbf{\Psi} \mid X, s_1, \ldots, s_\ell)$ denote the expected PDV of writing costs from period ℓ on, discounted back to period ℓ, conditional on $(X, s_1, \ldots, s_\ell)$, and we let $\hat{c}_{+1}(\mathbf{\Psi} \mid X, s_1, \ldots, s_\ell)$ denote the expected PDV of writing costs from period $\ell + 1$ on, discounted back to period $\ell + 1$, conditional on $(X, s_1, \ldots, s_\ell)$.

Consider period ℓ; suppose that the agent is expected to be the recipient of the *net* surpluses generated by $\mathbf{\Psi}$ in all *future* periods if he does not shirk, and to get a PDV of zero in period $\ell + 1$ if he shirks. Then he has no incentive to shirk if

$$\delta[\Pi(\mathbf{N}) - D(\mathbf{N}) - \widehat{c}_{+1}(\mathbf{\Psi} \mid X, s_1, \ldots, s_\ell)] \geq d(\mathbf{I}). \tag{IC*}$$

For each beginning-of-period state X, select a task partition $\mathbf{\Psi}(X)$ that solves the following problem:

$$\begin{gathered} \min_{\mathbf{\Psi}} \widehat{c}(\mathbf{\Psi} \mid X) \\ \text{s.t. } \forall \ell \geq 1, \forall (s_1, \ldots, s_\ell) \text{ (IC*) hold.} \end{gathered} \tag{P*}$$

Problem (P*) has a solution because the set of task partitions is finite and, by the small costs assumption (SC), a fully formal task partition satisfies (IC*). We consider strategy profiles that, starting with a Markov state X, govern tasks according to $\mathbf{\Psi}(X)$ on the path, and according to $\mathbf{\Psi}(X')$ in the beginning-of-period subgame after a deviation that leads to state X'. If such a strategy profile were subgame perfect, by definition it would be a self-enforcing contracting plan. Furthermore, by construction, such a contracting plan would be optimal. We show that such optimal contracting plans exist if δ is large enough.

Let $\hat{c}(X; \delta)$ denote the value of (P*) and define

$$\underline{v}(\delta) = \frac{\Pi(\mathbf{N}) - D(\mathbf{N})}{1-\delta} - \max_X \hat{c}(X; \delta).$$

Because $\max_X \hat{c}(X; \delta) \leq \sum_k [c(e_k) + c(a_k)]$, $\delta \underline{v}(\delta) \to \infty$ as $\delta \to 1$. Therefore, the equation $\Pi(\mathbf{N}) - D(\mathbf{N}) = \delta \underline{v}(\delta)$ has a largest solution δ^*, and $\Pi(\mathbf{N}) - D(\mathbf{N}) \leq \delta \underline{v}(\delta)$ for each $\delta \geq \delta^*$. Note that δ^* depends on the writing costs (and on other parameters). Because

$$\Pi(\mathbf{N}) - D(\mathbf{N}) = \delta^* \underline{v}(\delta^*) \geq \frac{\delta^*}{1-\delta^*}[\Pi(\mathbf{N}) - D(\mathbf{N})] - \sum_{k \in \mathbf{N}} [c(e_k) + c(a_k)],$$

it follows that $\delta^* \searrow \frac{1}{2}$ as $\sum_{k \in \mathbf{N}} [c(e_k) + c(a_k)] \longrightarrow 0$.

Now assume that $\delta \geq \delta^*$. We suppress in our notation the dependence of values on δ.

For each player i and Markov state X, we define a strategy profile that "punishes" i, keeping him at his maxmin.

If no deviations occur, the principal makes amendments as prescribed by task partition $\mathbf{\Psi}(X)$. After a deviation (by himself or the agent), the principal makes amendments according to task partition $\mathbf{\Psi}(X')$, where X' is the new Markov state. Wages are determined according to two punishment phases $\mathcal{P}^p$ and $\mathcal{P}^a$. The system starts in phase $\mathcal{P}^i$, where $i \in \{p, a\}$ is the player to be kept at his maxmin. As soon as player j deviates from his strategy (to be specified below), the system switches *immediately* to phase $\mathcal{P}^j$. *If there are no further deviations, the system switches to (or stays in) phase* $\mathcal{P}^p$ *at the beginning of the following period.*

The *strategy of the principal* in terms of *wage offers* is as follows. When the system is in phase $\mathcal{P}^p$ the offered wage is the gross profit $\Pi(\mathbf{N})$ minus the writing costs, so that the principal breaks even and all the net surplus goes to the agent. If at the beginning of a period the system is in phase $\mathcal{P}^a$ (meaning that either it is the starting period and $i = a$, or the agent has just deviated), the new Markov state is X', and s realizes, then the principal offers a low wage $m(\mathcal{P}^a, X', s)$ that makes the agent indifferent between getting zero payoff in every period on the one side, and accepting the offered contract and becoming the recipient of all *future* net surpluses on the other side. Recall that in such a situation the principal chooses clauses according to $\mathbf{\Psi}(X')$. Thus, the indifference condition is satisfied by

$$\begin{aligned} m(\mathcal{P}^a, X', s) &= D(\mathbf{N}) - \delta\left[\frac{\Pi(\mathbf{N}) - D(\mathbf{N})}{1-\delta} - \widehat{c}_{+1}(\mathbf{\Psi}(X') \mid X', s)\right] \\ &= -\left[\frac{\Pi(\mathbf{N}) - D(\mathbf{N})}{1-\delta} - \widehat{c}(\mathbf{\Psi}(X') \mid X', s)\right] + [\Pi(\mathbf{N}) - c(\mathbf{\Psi}(X') \mid X', s)], \end{aligned}$$

where $c(\mathbf{\Psi}(X') \mid X', s)$ are the writing costs in the given period. The first equality follows almost immediately from the verbal definition; the second—which is used below—is obtained with simple algebra. Note that $m(\mathcal{P}^a, X, s)$ may be negative.

We now describe the *strategy of the agent*. If the principal has not just deviated, the agent accepts the offered contract and behaves as prescribed by the current task partition, that is, the one determined after the last deviation. If the principal has just deviated offering contract $(\mathbf{g}', m')$ and moving the Markov state to $X' = (s, \mathbf{g}', \mathbf{M}')$, then the agent accepts, and shirks on nonformal tasks, if and only if

$$m' - D(\mathbf{g}', s) > \delta v(X'), \tag{A1}$$

where $D(\mathbf{g}', s)$ is the total disutility implied by the set of clauses $\mathbf{g}'$ given s.

We show that these strategies are immune to one-shot deviations and hence they form an optimal contracting plan.[38] By construction, the *agent has no incentive to deviate*. In phase $\mathcal{P}^a$ the agent is indifferent between accepting and rejecting. By (IC*) he has no incentives to shirk when he is supposed to behave as prescribed by (the informal rules in) the current task partition. Furthermore, (IC*) also implies that the agent has no incentive to reject the contractual offers prescribed by the principal's strategy in phase $\mathcal{P}^p$.

Next suppose that the shock vector is s and the principal deviates offering $(\mathbf{g}', m')$ and thus moving the Markov state to $X' = (s, \mathbf{g}', \mathbf{M}')$. The system enters (or stays in) phase $\mathcal{P}^p$. If $m' - D(\mathbf{g}', s) \le \delta v(X')$, the agent is supposed to reject. The expected payoff if the agent conforms is $\delta v(X')$ because he will be the recepient of all future net surpluses. The maximum gain from a one-shot deviation is $m' - D(\mathbf{g}', s)$, because after the deviation the system enters phase $\mathcal{P}^a$ where the agent gets his maxmin (zero). Therefore rejection is indeed a best response.

If $m' - D(\mathbf{g}', s) > \delta\, v(X')$, the agent is supposed to accept and this is obviously a best response.

Next we check that the *principal has no incentive to deviate* in phase $\mathcal{P}^p$. The only way for the principal to make a profitable one-shot deviation is to offer a contract $(\mathbf{g}', m')$ (by way of appropriate amendments) satisfying the acceptance condition (A1). The new Markov state is $X' = (s, \mathbf{g}', \mathbf{M}')$. By definition, $v(X') \ge \underline{v}$. Therefore, the net payoff the principal can get by "tempting" the agent is bounded above by

$$\max_{\mathbf{g}', \mathbf{M}'} [\sigma(\mathbf{g}', s) - \delta v(s, \mathbf{g}', \mathbf{M}')] \le \Pi(\mathbf{N}) - D(\mathbf{N}) - \delta \underline{v},$$

where $\boldsymbol{\sigma}(\mathbf{g}', s)$ is the surplus induced by $\mathbf{g}'$ given s. Recall that $\Pi(\mathbf{N}) - D(\mathbf{N}) - \delta\underline{v} \le 0$ because we assumed $\delta \ge \delta^*$. Therefore, the principal has no profitable one-shot deviation in phase $\mathcal{P}^p$.

Now consider phase $\mathcal{P}^a$; let X' be the beginning-of-period Markov state and s the current shock vector. The principal is supposed to offer wage $m(\mathcal{P}^a, X', s)$ and the contractual clauses implied by $\mathbf{\Psi}(X')$ (either because this is the starting period and $X = X'$ or because the agent has just deviated and this is the new task partition). If the principal deviates the system enters phase $\mathcal{P}^p$ immediately and the argument above implies that he cannot get a positive payoff from such deviation. If he does not deviate, he gets

$$\begin{aligned} &\Pi(\mathbf{N}) - c(\mathbf{\Psi}(X') \mid X', s) - m(\mathcal{P}^a, X', s) \\ &= \frac{\Pi(\mathbf{N}) - D(\mathbf{N})}{1-\delta} - \widehat{c}(\mathbf{\Psi}(X') \mid X', s) \ge 0 \end{aligned}$$

and zero afterward, where the equality follows from the definition of $m(\mathcal{P}^a, X', s)$ and the inequality from (IC*). This concludes the proof of (i).

(ii) We first prove that (IC) is necessary for incentive compatibility. Suppose that the tuple $(\mathbf{I}, \mathbf{C}, \mathbf{E}, \mathbf{S}, (m_t)_{t\ge1})$ is part of an SPE. Then it must be the case that the present value of the principal's expected profits is always (weakly) positive and that the agent's expected utility from following $(\mathbf{I}, \mathbf{C}, \mathbf{E}, \mathbf{S}, (m_t)_{t\ge1})$ is (weakly) larger than what the agent can get by accepting the formal contract offered by the principal, "shirking" on the tasks in $\mathbf{I}$ and rejecting all future offers. To write these incentive constraints in a relatively simple form, define

$$\widehat{m}_{t+1}(h_t) \equiv \sum_{k=1}^{\infty} \delta^{k-1} \mathrm{E}(m_{t+k} \mid h_t).$$

For all $t \ge 0$, all $h_{t+1} = (h_t, s_{t+1})$

$$\frac{\Pi(\mathbf{N})}{1-\delta} - \widehat{c}_{t+1}(\mathbf{C}, \mathbf{E}, \mathbf{S} \mid h_t, s_{t+1}) \ge m_{t+1}(h_t, s_{t+1}) + \delta \widehat{m}_{t+2}(h_t, s_{t+1}), \tag{IC$_P^{t+1}$}$$

and for all $t \ge 1$, all h_t

$$m_t(h_t) + \delta \widehat{m}_{t+1}(h_t) - \frac{D(\mathbf{N})}{1-\delta} \ge m_t(h_t) - D(\mathbf{C} \cup \mathbf{E} \cup \mathbf{S}). \tag{IC$_A^t$}$$

[38] It can be assumed without loss of generality that the set of possible transfers is $[-L, L]$ with $L \ge \frac{\Pi(\mathbf{N})}{1-\delta}$. Discounting implies that the compactified game satisfies "continuity at infinity" and hence the one-shot-deviation principle applies (see Fudenberg and Tirole, 1991).

Incentive constraint (IC_A^t) can be written as

$$\delta\widehat{m}_{t+1}(h_t) \geq \frac{\delta}{1-\delta}D(\mathbf{N}) + D(\mathbf{I}).$$

(Note that $D(\mathbf{C}\cup\mathbf{E}\cup\mathbf{S}) = D(\mathbf{N}) - D(\mathbf{I})$). Taking the expected value of both sides of (IC_P^{t+1}) w.r.t. s_{t+1} (conditional on h_t) and combining with the above inequality, we obtain

$$\frac{\delta}{1-\delta}\Pi(\mathbf{N}) - \delta\widehat{c}_{t+1}(\mathbf{C},\mathbf{E},\mathbf{S}\mid h_t) \geq \delta\widehat{m}_{t+1}(h_t) \geq \frac{\delta}{1-\delta}D(\mathbf{N}) + D(\mathbf{I}),$$

which yields (IC).

Now suppose that $\delta \geq \delta^*$ and $(\mathbf{I},\mathbf{C},\mathbf{E},\mathbf{S})$ satisfies (IC). We exhibit a self-enforcing contracting plan implementing $(\mathbf{I},\mathbf{C},\mathbf{E},\mathbf{S})$ based on four phases, $\mathcal{N}_1$, $\mathcal{N}_2$, $\mathcal{P}^a$, and $\mathcal{P}^p$. In each "normal" phase $\mathcal{N}_j$ each task is regulated according to task partition $(\mathbf{I},\mathbf{C},\mathbf{E},\mathbf{S})$. The system starts in phase $\mathcal{N}_1$ and then from the following period moves to phase $\mathcal{N}_2$ if no deviation occurs. In phase $\mathcal{N}_1$ (period 1) the transfer m_1 is such that the participation constraints of principal and agent are satisfied (such m_1 must exist because (IC) holds). In phase $\mathcal{N}_2$ and period $t(t\geq 2)$ the transfer is such that all the net surplus goes to the agent,

$$m_t(h_t) = \Pi(\mathbf{N}) - D(\mathbf{N}) - c_t(\mathbf{E},\mathbf{S},h_t).$$

In normal phases, the agent accepts the proposed contract and chooses the efficient action for all the tasks $k\in\mathbf{I}$. As soon as player i deviates, the system switches *immediately* from the current phase to the punishment phase $\mathcal{P}^i$ and continues as described in the proof of (i). Hence, there are no incentives to deviate in the punishment phases.

Because players are punished at their maxmin, the proof that players have no incentive to deviate in a normal phase is almost the same as in (i). Note that first-period participation constraints are satisfied by construction. After the first period the principal is indifferent, and (IC) implies that the agent has no incentive to reject, because he is the recipient of the net surplus, which has positive present value. To see that the agent has no incentive to shirk in a normal phase, simply note that the left-hand side of (IC) is the expected present value of future net benefits to the agent (which he forgoes if he shirks), whereas the right-hand side is the temptation to shirk, that is, the disutility the agent avoids by shirking on informal tasks.

Proof of Proposition 3. Parts (i) and (iii) are obvious. To see that (ii) holds, first note that $D(\mathbf{N}) > \frac{\delta}{1-\delta}[\Pi(\mathbf{N}) - D(\mathbf{N})]$ implies that a fully informal contract violates (IC). Next suppose without loss of generality that $d^{\min} = d_1$, and consider the task partition corresponding to the MPE modified by replacing the formal rule for task 1 with the informal rule. The expected present value of writing costs conditional on any history is bounded above by $\frac{1}{1-\delta}\sum_{k>1}c(a_k)$. If $d_1 \leq \frac{\delta}{1-\delta}[\Pi(\mathbf{N}) - D(\mathbf{N}) - \sum_{k>1}c(a_k)]$, then the modified partition satisfies (IC), which means that a fully formal partition cannot be a solution to problem (P). The latter inequality is satisfied if $d_1 < \frac{\delta}{1-\delta}[\Pi(\mathbf{N}) - D(\mathbf{N})]$ (that is, $[\Pi(\mathbf{N}) - D(\mathbf{N})] - \frac{1-\delta}{\delta}d_1 > 0$) and writing costs are sufficiently small, that is, $\sum_{k>1}c(a_k) \leq [\Pi(\mathbf{N}) - D(\mathbf{N})] - \frac{1-\delta}{\delta}d_1$.

Proof of Proposition 5. (**i**) Write $c^a = c, c^e = \theta c$, where c is a positive scale parameter. We only consider a change in uncertainty (p) and an equiproportional change in writing costs (c). The remaining cases are similar. Differentiating (FOC) with respect to x_I, x_C, x_S, and p, we obtain

$$\Delta x_I = -\frac{\left(c+\frac{d(x_I)}{1+\theta}\right)d(x_I)}{d'(x_I)c}\cdot\frac{\Delta p}{1-p},$$

$$\Delta x_S = \left(\frac{(1-\delta)\left(c+\frac{d(x_I)}{1+\theta}\right)[d(x_I)]^2}{r\delta d'(x_I)(1-p)c^2} + x_S\right)\frac{\Delta p}{1-p}.$$

(FOC) implies $(1-\delta)\,d(x_I) > \delta\,(1-p)rc$. Therefore,

$$\frac{(1-\delta)\left(c+\frac{d(x_I)}{1+\theta}\right)[d(x_I)]^2}{r\delta d'(x_I)(1-p)c^2} > \frac{\left(c+\frac{d(x_I)}{1+\theta}\right)d(x_I)}{d'(x_I)c}.$$

Suppose that $\Delta p < 0$ (increase in uncertainty). Then $\Delta x_S < 0$, $\Delta x_I > 0$, and

$$\begin{aligned}\Delta x_C &= -\Delta x_S - \Delta x_I\\ &= \left(\frac{(1-\delta)\left(c+\frac{d(x_I)}{1+\theta}\right)[d(x_I)]^2}{r\delta d'(x_I)(1-p)c^2} + x_S - \frac{\left(c+\frac{d(x_I)}{1+\theta}\right)d(x_I)}{d'(x_I)c}\right)\frac{|\Delta p|}{1-p} > 0.\end{aligned}$$

Differentiating (FOC) with respect to x_I, x_C, x_S, and c, we obtain

$$\Delta x_I = \frac{d(x_I)}{d'(x_I)}\frac{\Delta c}{c},\quad \Delta x_S = -\left(\frac{(1-\delta)[d(x_I)]^2}{\delta(1-p)rd'(x_I)c} + x_S\right)\frac{\Delta c}{c}.$$

Note that $(1-\delta)\,d(x_I) > \delta\,(1-p)rc$ implies $\frac{(1-\delta)[d(x_I)]^2}{\delta(1-p)rd'(x_I)c} > \frac{d(x_I)}{d'(x_I)}$.

Suppose $\Delta c > 0$. Then $\Delta x_S < 0$, $\Delta x_I > 0$, and

$$\Delta x_C = -\Delta x_S - \Delta x_I = \left(\frac{(1-\delta)[d(x_I)]^2}{\delta(1-p)rd'(x_I)c} + x_S - \frac{d(x_I)}{d'(x_I)} \right) \frac{\Delta c}{c} > 0.$$

(ii) Proposition 4 implies that $|\Delta x_C| = \Delta x_S$. These variations cannot be null, because an increase in Π improves incentives allowing a decrease of the objective function. This is achieved by reducing the fraction of tasks governed by the (*ex ante*) more costly contingent rule, and thus eliminting the slack created by an increase in Π.

References

AL NAJJAR, N., ANDERLINI, L., AND FELLI, L. "Undescribable Events." *Review of Economic Studies*, Vol. 73 (2006), pp. 849–868.

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

——— AND ———. "Incomplete Contracts and Complexity Costs." *Theory and Decision*, Vol. 46 (1999), pp. 23–50.

BAKER, G., GIBBONS, R., AND MURPHY, K. "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics*, Vol. 109 (1994), pp. 1125–1156.

BATTIGALLI, P. AND MAGGI, G. "Rigidity, Discretion, and the Cost of Writing Contracts." *American Economic Review*, Vol. 92 (2002), pp. 798–817.

——— AND ———. "Costly Contracting in a Long-Term Relationship." Working Paper no. 249, IGIER, Bocconi University, 2003.

BOLTON, P. AND FAURE-GRIMAUD, A. "Satisficing Contracts." Columbia University and London School of Economics, Mimeo, 2007.

BULL, C. "The Existence of Self-Enforcing Implicit Contracts." *Quarterly Journal of Economics*, Vol. 102 (1987), pp. 147–159.

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

———. "Optimal Length of Labour Contracts." *International Economic Review*, Vol. 26 (1985b), pp. 251–270.

ECCLES, R. "The Quasifirm in the Construction Industry." *Journal of Economic Behavior and Organization*, Vol. 2 (1981), pp. 335–358.

FUDENBERG, D. AND TIROLE, J., eds. *Game Theory*. Cambridge, Mass.: MIT Press, 1991.

GRAY, J. "Wage Indexation: A Macroeconomic Approach." *Journal of Monetary Economics*, Vol. 2 (1976), pp. 221–236.

———. "On Indexation and Contract Length." *Journal of Political Economy*, Vol. 86 (1978), pp. 1–18.

HART, O. AND HOLMSTROM, B. "The Theory of Contracts." In T. BEWLEY, ed., *Advances in Economic Theory, Fifth World Congress*. Cambridge, UK: Cambridge University Press, 1987.

KAHAN, M. AND KLAUSNER, M. "Standardization and Innovation in Corporate Contracting (or "The Economics of Boilerplate")." *Virginia Law Review*, Vol. 83 (1997), pp. 713–770.

KRASA, S. AND WILLIAMS, S.R. "Limited Observability as a Constraint in Contract Design." *Journal of Economic Theory*, Vol. 134 (2007), pp. 379–404.

LEVIN, J. "Relational Incentive Contracts." *American Economic Review*, Vol. 93 (2003), pp. 835–857.

MACAULAY, S. "Non-Contractual Relations in Business: A Preliminary Study." *American Sociological Review*, Vol. 28 (1963), pp. 55–67.

MACLEOD, W.B. "Complexity and Contract." *Revue d'Economie Industrielle*, Vol. 92 (2000), pp. 149–178.

——— AND CHAKRAVORTY, S. "On the Efficiency of Standard Form Contracts: The Case of Constructions." University of Southern California, Mimeo, 2004.

——— AND MALCOMSON, J.M. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica*, Vol. 57 (1989), pp. 447–480.

MASKIN, E. AND TIROLE, J. "Unforeseen Contingencies and Incomplete Contracts." *Review of Economic Studies*, Vol. 66 (1999), pp. 83–114.

MEIHUIZEN, H.E. AND WIGGINS, S.N. "Information Cascades and Contractual Incompleteness in Natural Gas Contracting." Texas A&M University, Mimeo, 2000.

PEARCE, D. AND STACCHETTI, E. "The Interaction of Implicit and Explicit Contracts in Repeated Agency." *Games and Economic Behavior*, Vol. 23 (1998), pp. 75–96.

SCHMIDT, K. AND SCHNITZER, M. "The Interaction of Explicit and Implicit Contracts." *Economic Letters*, Vol. 48 (1995), pp. 193–199.

TIROLE, J. "Bounded Rationality and Incomplete Contracts." University of Toulouse, Mimeo, 2007.

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