

Authority in Relational Contracts with Asymmetric Information

Harvey Upton*

13th January 2026

Abstract

This paper studies the allocation of decision rights in relational contracts where the agent privately observes his costs in each period. When the agent has authority to select the action, the presence of an adverse selection problem typically reduces the attainable surplus from the relationship, relative to the case of symmetric information. In contrast, under principal authority, we show that private information has zero impact on the attainable surplus. Specifically, the optimal action schedule derived under symmetric information remains implementable with private information. Accordingly, when contracts must be self-enforcing and are necessarily second best, private information reduces the attainable surplus if and only if the informed party has authority.

1 Introduction

A significant literature has documented the importance of relational contracts for sustaining cooperation between firms (see e.g. Baker et al. 2008 and Gil & Zananone 2017). Economic theory suggests that efficiently allocating decision rights—for instance, using formal contracts or through changes in asset ownership—increases the range of actions that can be sustained as part of an informal agreement, and is hence an important consideration for firms when building relationships with one another (Baker et al. 2002, 2011). However, the optimal allocation of authority must also account for the fact that, in many interactions, parties have private information regarding the costs and benefits of a particular course of action. It is tempting to think that the presence of asymmetric information provides a rationale for allocating decision rights to the informed party in the relational agreement; in this paper, we argue that the opposite is true.

We study an infinitely-repeated interaction between two firms, hereafter the principal and the agent. In each period, an action must be chosen, regarding which the parties have opposing interests. Specifically, a higher action increases the principal’s benefit from the relationship, but raises the agent’s costs. The agent’s costs are also affected by random shocks, which are i.i.d. in each period and are privately observed by the agent after contracts have been signed, creating an adverse selection problem. In contrast, the principal’s benefit is known to both parties. Since the action is observable to both the principal and the agent, but non-verifiable to third parties, relational contracts can be used to sustain behavior that is infeasible in one-shot games, thereby increasing the surplus from the relationship. In this environment, we wish to understand

*Humboldt-Universität zu Berlin; harvey.upton@hu-berlin.de

Humboldt-Universität zu Berlin, School of Business and Economics, 10178 Berlin, Germany. This project was financially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 403041268 – TRR 266 Accounting for Transparency. I thank Anja Schöttner and participants at the TRR 266 Annual Conference 2025 in Berlin for valuable comments.

how the optimal allocation of decision rights over the action is influenced by presence of the agent’s private information.

The case of an adverse selection problem with relational contracts when the agent (the informed party) has decision rights over the action has been studied by Levin (2003). As in the classical adverse selection problem, the agent has an incentive to overstate his type (i.e. to claim that his costs are higher than they truly are) in order to extract a higher rent. To prevent this, the contract must be designed such that the bonus schedule is incentive compatible, thereby inducing the agent to always select the action associated with his true type. Since incentive compatibility requires increased bonuses, the set of implementable action schedules shrinks relative to an environment with symmetric information, reducing the attainable surplus from the relationship. In other words, when the agent has authority, the fact that his information is private is detrimental to the parties’ ability to sustain cooperation as part of a relational agreement.

Our contribution is to study how the agent’s private information influences the outcome of relational contracting when the principal (the uninformed party) has decision rights over the action, allowing for a comparison with Levin’s (2003) results. We begin our analysis by deriving the optimal relational contract under principal authority for the case of symmetric information. With formal contracts, the parties would implement the efficient action associated with each realization of the agent’s costs. We show that, when contracts are relational, the optimal action schedule under symmetric information can be split into two distinct regions. First, for low cost realizations, where the first-best action is high, the principal’s temptation to deviate by implementing the maximum possible action is low. In this case, the dynamic enforcement constraint is slack, so that the first-best can be implemented and the action schedule is downward sloping. Second, for sufficiently high cost realizations, the first-best action is small, and hence cannot be feasibly implemented due to the principal’s deviation incentives. In this case, the dynamic enforcement constraint binds, so that the action schedule becomes flat, and, to curtail the principal’s incentive to renege, above the first-best; that is, pooling occurs, with the same superoptimal action being implemented for several different realizations of the agent’s costs.

We then move on to the case of adverse selection and show that the surplus-maximizing action schedule under symmetric information remains implementable—and hence optimal—when the agent’s cost realization in each period is private information. The intuition behind this finding is as follows. As in the case of agent authority, incentive compatibility requires paying additional bonuses to induce separation of types. For low cost realizations, these additional bonuses are strictly positive; however, since the dynamic enforcement constraint is slack over this region, as explained above, this is only problematic if the required increase in payments is sufficiently large. We show that, in order to implement the optimal action schedule under symmetric information, the necessary size of these bonuses is exactly equal to the slackness in the constraint. Moreover, for high cost realizations, the agent has no incentive to overstate his true cost realization, since there is pooling over this region. It follows that no additional bonus is required relative to the symmetric information case, leaving the dynamic enforcement constraint binding, as before; the constraint then holds with equality for all possible cost realizations, so that the symmetric information action schedule remains implementable.

The result implies that, in stark contrast to the case of agent authority, the presence of an adverse selection problem has zero impact on the attainable surplus when the principal has decision rights over the action. Accordingly, when relational contracts cannot achieve the first-best, one-sided private information reduces the relationship’s surplus *if and only if* the informed party has authority. As a result, if it is optimal to allocate authority to the principal under symmetric information, then this is also the case when the agent has private

information. Conversely, the result also implies the existence of situations where agent authority is optimal when the cost realization is commonly observed by both parties, but not when it is privately observed. More broadly, our findings suggest that when firms must sustain cooperation using relational agreements, there are advantages to allocating decision rights to the less informed party, as this can reduce incentives to deviate from the agreement.

Our results also shed light on Levin’s (2003) findings regarding the optimal structure of relational contracts under adverse selection, and, in particular, the extent to which they are dependent on a particular allocation of decision rights. For example, we show that when the first-best is not attainable, the optimal action schedule always features pooling, partial or complete, regardless of which party is responsible for choosing the action. Accordingly, our results support Levin’s (2003) prediction that relational agreements are more flexible when the surplus generated by the relationship is greater. However, other results do not extend. Under agent authority, any contract that is not first-best exhibits inefficiencies throughout the whole action schedule, similar to the classical adverse selection problem where inefficiency is a probability-one event. In contrast, when the principal holds decision rights, second-best action schedules typically implement efficient actions for many different types.

Related Literature. As discussed in the foregoing, our results are closely related to those of Levin (2003), which we further discuss in Section 3. Malcomson (2015) extends Levin’s (2003) analysis to a setting where the principal’s benefit from a particular action also varies in each period. However, neither of these papers study the implications of differing allocations of decision rights for the outcome of relational contracting. Baker et al. (2011) study how authority should be optimally allocated to facilitate relational agreements. They consider an abstract symmetric information setting with an arbitrary number of parties and decision rights to be allocated, showing that authority should be distributed such that the parties’ aggregate reneging temptation is minimized. Schöttner & Upton (2025) build on their work by considering an interaction between a firm and a worker, where the choice between employment and self-employment of the worker determines the allocation of authority over the design of the work process. They show that employment relationships (principal authority) are typically more rigid, and better suited to favorable production environments with large gains to trade. While their analysis investigates the impact of a moral hazard problem under self-employment (agent authority), they do not consider adverse selection problems.¹

The paper is also related to studies that analyze relational contracting when one party has private information that is relevant to the interaction in future periods. Malcomson (2016) considers an environment where the agent’s costs are privately known and persistent over time, showing that full separation of types is not possible when post-revelation continuation payoffs are on the Pareto frontier (see also Yang 2013). Fahn & Klein (2019) study situations where, at the time when bonuses are paid, the principal is privately informed about the relationship’s value in the subsequent period. As a result, she is tempted to announce ‘bad news’ in order to reduce the bonus; to prevent this, the optimal contract may distort the agent’s action following such announcements, reducing surplus.² Halac (2012) studies a model where the principal privately knows her outside option, and hence the true value of the relationship. Optimal contracts feature either pooling or a positive probability of separation, so that efficiency is similarly reduced relative to the symmetric information benchmark. More broadly, Malcomson (2012) provides an extensive overview of the theoretical

¹Our analysis also relates to papers that study how the allocation of decision rights affects implicit incentives for cooperation in repeated interactions; see e.g. Baker et al. (1999), Englmaier et al. (2010) and Upton (2024).

²Along similar lines, Li & Matouschek (2013) consider an environment where the principal privately observes her opportunity costs of paying the worker’s bonus. The optimal relational contract then features periods of conflict, where effort and profits temporarily decline.

relational contracting literature.³

There is also a significant literature that studies the optimal allocation of authority in principal-agent relationships where parties can make use of formal contracts. Many of these papers find that asymmetric information favors allocating decision rights to the informed party. For instance, Bester & Krämer (2017) study an environment where the parties can assign both decision rights and exit rights, the latter of which gives one party the ability to prematurely leave the relationship, destroying surplus. They show that, in this case, the first-best outcome can always be achieved by allocating decision rights to the agent (informed party), and exit rights to the principal (uninformed party). In contrast, the current paper shows that when the parties' interaction must be sustained by relational agreements, the presence of asymmetric information favors granting authority over the action to the uninformed party.

Finally, from a political science perspective, Luo et al. (2025) study the allocation of 'power' in relational contracts, which confers private information over the revenue generated by the interaction in each period, as well as the authority to allocate this revenue between parties. They show that when transfers between parties are restricted, alternation of power is necessary to sustain long run cooperation.

2 Model Setup

We consider an infinitely-repeated game in discrete time between a principal and an agent, both of whom are risk-neutral and discount the future at a common rate $\delta \in (0, 1)$. In each period, an action $a \in [a^{min}, a^{max}]$ must be chosen. A higher action benefits the principal, but is more costly for the agent, creating a conflict of interest between the parties. Specifically, the principal's benefit associated with a particular action is $v(a)$, where $v' > 0$ and $v'' \leq 0$ for all $a \in [a^{min}, a^{max}]$. The agent's costs are $c(a, \theta)$, where $\theta \in [\underline{\theta}, \bar{\theta}] \subset (0, \infty)$ is a random variable independently and identically distributed between periods according to the function $F(\cdot)$ with $F'(\cdot) = f(\cdot) > 0$. We assume that $c_a > 0$ and $c_{aa} > 0$ for all $a \in (a^{min}, a^{max}]$, and also that $c_\theta > 0$ and $c_{a\theta} > 0$, so that a higher realization of θ implies higher costs and marginal costs for the agent. We further assume that $v'(a^{max}) = c'(a^{min}) = 0$. The implemented action a is observable by both parties, but non-verifiable to third parties; in contrast, θ is observed in each period solely by the agent, but only after contracts have been signed.

For a particular realization of the agent's cost parameter θ , the first-best action $a^*(\theta)$ solves:

$$\max_a v(a) - c(a, \theta) \quad (1)$$

Our restrictions on v and c imply that $a^*(\theta)$ is the unique solution to:

$$v'(a) - c_a(a, \theta) = 0 \quad (2)$$

and that $a^*(\theta) \in (a^{min}, a^{max})$ for all θ . Applying the implicit function theorem, it is straightforward to verify that $a^*(\theta)$ is strictly decreasing in θ . The first-best surplus from the relationship is:

$$S^* = \int_{\underline{\theta}}^{\bar{\theta}} [v(a^*(\theta)) - c(a^*(\theta), \theta)] f(\theta) d\theta \quad (3)$$

which is assumed to satisfy $S^* > \bar{\pi} + \bar{u}$, where $\bar{\pi} \geq 0$ and $\bar{u} \geq 0$ are the outside options of the principal and

³See also the recent Symposium on the Economics of Relational Contracts in the *Journal of Institutional and Theoretical Economics*, 179(3-4).

agent, respectively. Finally, we assume that the expected surplus from implementing either a^{min} or a^{max} regardless of the realized θ is less than the parties' combined outside options.⁴

The main focus of the paper concerns the allocation of decision rights over the action a . While we briefly discuss existing results for the case of agent authority, in our formal analysis we focus on the case where the principal has the right to select the action in each period. We assume that once a contract has been agreed and an action has been chosen, both parties must incur the ensuing consequences; in particular, this implies that the agent has no means of escaping the costs associated with a particular action, e.g. by prematurely leaving the relationship.

The timing of the game in each period is as follows:

1. The principal decides whether to offer a contract to the agent and, if so, the terms of this contract. If a contract is offered, the agent either accepts or rejects. If a contract is offered and accepted then the game proceeds; otherwise, the game ends and parties receive their outside options in all subsequent periods.
2. The agent's cost parameter θ is realized and observed by the agent.
3. The agent reports the cost parameter to the principal.
4. The principal chooses the action a .
5. Transfers between parties are made.

Since neither the action nor the associated costs and benefits are verifiable, in a single period interaction the parties have no means to create incentives. This implies that either a^{min} or a^{max} will be implemented regardless of the agent's cost realization, depending on which party has decision rights. By assumption, this generates a negative surplus, so that the parties prefer not to interact. However, in an infinitely repeated relationship, the parties can use informal agreements to discipline their behavior, allowing them to sustain a productive relationship. Formally, we define a relational contract as being self-enforcing if it represents a perfect public equilibrium of the repeated game (Levin 2003). As argued by Abreu (1988), the set of self-enforcing contracts is largest when punishments are most severe, which, in the current environment, corresponds to the parties breaking off trade following a deviation from their agreement. From the foregoing, this is the statically optimizing behavior, and implies that both parties receive their outside option in each period for the remainder of the game.

3 Agent Authority

Before we study the case of principal authority, it is helpful to consider the outcome of contracting when the agent is responsible for choosing the action. This environment has been studied by Levin (2003), whose findings we briefly summarize here. Under agent authority, the parties design a stationary relational contract, which consists of a fixed payment w , an action schedule $a(\theta)$ and a corresponding bonus schedule $b(\theta)$. This contract is then offered to the agent in each period, and must be designed such that it satisfies the following

⁴Formally, we impose:

$$\max \left\{ \int_{\underline{\theta}}^{\bar{\theta}} [v(a^{min}) - c(a^{min}, \theta)] f(\theta) d\theta, \int_{\underline{\theta}}^{\bar{\theta}} [v(a^{max}) - c(a^{max}, \theta)] f(\theta) d\theta \right\} < \bar{\pi} + \bar{u}$$

properties. First, the contract must ensure both parties' participation. Second, the bonus schedule must be *incentive compatible*; that is, it must be that following a realization of θ , the agent must prefer to implement the associated action $a(\theta)$ and receive the corresponding bonus $b(\theta)$, rather than misrepresent his true cost realization. Formally, we require that the combination of action and transfer $\{a(\theta), b(\theta)\}$ yields a weakly higher payoff to the agent than any other combination $\{a(\hat{\theta}), b(\hat{\theta})\}$ for all $\theta, \hat{\theta} \in [\underline{\theta}, \bar{\theta}]$. Third, for all realizations of θ , the contract must ensure that the agent prefers to implement $a(\theta)$, rather than go 'off-schedule' and implement the lowest feasible action, a^{min} . Fourth, both parties must be willing to adhere to the agreed-upon transfers $b(\theta)$, rather than renege on the relational agreement.

Solving for the optimal contract yields the following result.

Proposition 1 (Levin, 2003). *If an optimal contract exists under agent authority with adverse selection, it implements an action schedule that takes one of three forms:*

1. *First-best:* $a(\theta) = a^*(\theta)$ for all θ .
2. *Partial Pooling:* $a(\theta)$ is constant on $[\underline{\theta}, \theta^c]$, and strictly decreasing on $(\theta^c, \bar{\theta}]$ for some $\theta^c \in (\underline{\theta}, \bar{\theta})$.
3. *Pooling:* $a(\theta)$ is the same for all θ .

In either second-best scenario, $a(\theta) < a^(\theta)$ for all θ .*

Under symmetric information, discretionary bonuses must be sufficiently high to ensure the agent does not deviate off-schedule and implement a^{min} . With adverse selection, these bonuses must be further increased to make the contract incentive compatible. For sufficiently high δ , the principal can credibly promise large discretionary transfers, so that the first-best can be implemented. However, for lower values of δ , these transfers are no longer feasible, so that the action schedule must be adjusted in order for the contract to be self-enforcing.

Two features of second-best relational contracts under agent authority with adverse selection particularly stand out. First, since the size of discretionary transfers is limited, the principal faces restrictions on the set of implementable actions. For sufficiently low realizations of θ , she offers the largest feasible transfer, subject to the dynamic enforcement constraint; incentive compatibility then implies pooling for all such realizations of the agent's cost parameter. When δ is low, this pooling extends to all cost types, so that the action becomes constant throughout the schedule. Second, regardless of whether a particular cost type is pooled with others, the associated action becomes inefficient for *all* realizations of θ . Intuitively, due to the adverse selection problem, raising the action of a particular type requires increasing the slope of the bonus schedule. Since the size of bonus payments is constrained, incentive compatibility then necessitates a reduction in the action of some other cost type. It then follows that implementing efficient effort for any type has zero benefit at the margin, but a positive shadow cost, and hence cannot be optimal; this leads to inefficiencies throughout the schedule.

While Levin (2003) does not explicitly discuss the impact of asymmetric information on efficiency, it follows from his arguments that when contracts are second-best, private information has a negative impact on the surplus attainable from the relationship. In particular, the need for the contract to be incentive compatible implies a monotonicity constraint on the action schedule, which is binding at the optimum for some cost realizations, resulting in pooling. Indeed, Schöttner & Upton (2025) show in a related model that when information is symmetric, the optimal action schedule under agent authority never exhibits pooling, and often implements efficient actions for a subset of cost realizations. Due to the incentive compatibility requirement, neither of these features extend to the case of adverse selection.

4 Principal Authority

In this section, we analyze the relational contracting problem when the principal has authority over the action to be implemented. We first consider a symmetric information environment, where the agent's cost parameter is observed by both parties in each period. We then use these findings to analyze the environment with asymmetric information.

4.1 Symmetric Information

Suppose there exists some self-enforcing relational contract that generates a positive surplus. As argued by Levin (2003), this surplus can be arbitrarily distributed between the two parties via up-front payments, which are made at the beginning of the relationship and do not affect subsequent incentives. Accordingly, the parties always prefer to implement the relational contract that achieves the maximum attainable surplus. For the case of agent authority, Levin (2003) further shows that we can without loss of generality restrict attention to stationary contracts. Slightly adjusting his arguments in a few places, this result can be extended to the current environment with principal authority and symmetric information.

Proposition 2. *If an optimal contract exists under principal authority with symmetric information, there are stationary contracts that are optimal.*

A stationary relational contract consists of three elements. First, a fixed wage payment w , which can be incorporated into a formal contract and hence can be credibly promised in each period regardless of its size.⁵ Second, an action schedule $a(\theta)$, that specifies the action that the principal should implement following each realization of θ and, third, a corresponding bonus schedule $b(\theta)$. Since the latter two elements cannot be incorporated into a formal contract, they must be designed such that the parties' relational agreement is self-enforcing. In the current environment, an agreement is self-enforcing if, at each stage of the game and for any realization of θ , both parties prefer to adhere to their respective promises rather than renege on the agreement and subsequently earn their outside option in each period thereafter.

Formally, the parties' problem is to choose the relational contract that maximizes the surplus from the relationship:

$$\max_{w, a(\theta), b(\theta)} \int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} - \bar{\pi} \quad (4)$$

subject to:

$$\int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - b(\theta)] f(\theta) d\theta - w \geq \bar{\pi} \quad (5)$$

$$w + \int_{\underline{\theta}}^{\bar{\theta}} [b(\theta) - c(a(\theta), \theta)] f(\theta) d\theta \geq \bar{u} \quad (6)$$

$$v(a(\theta)) - b(\theta) - w + \frac{\delta}{1-\delta} \left[\int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - b(\theta)] f(\theta) d\theta - w \right] \geq v(a^{max}) - w + \frac{\delta}{1-\delta} \bar{\pi} \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \quad (7)$$

$$w + b(\theta) - c(a(\theta), \theta) + \frac{\delta}{1-\delta} \left[w + \int_{\underline{\theta}}^{\bar{\theta}} [b(\theta) - c(a(\theta), \theta)] f(\theta) d\theta \right] \geq w - c(a(\theta), \theta) + \frac{\delta}{1-\delta} \bar{u} \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \quad (8)$$

⁵Alternatively, the fixed payment w could be legally unenforceable, but paid contemporaneously with the agent's decision to accept the contract. This would not affect the paper's results.

The parties maximize the relationship surplus (4), subject to the participation constraints (5) and (6), and the two dynamic enforcement constraints (7) and (8). The first of these, (7), ensures that for each θ , the principal prefers to adhere to the relational agreement. The left-hand side of (7) is her current period payoff from implementing the agreed-upon action $a(\theta)$ and paying the corresponding bonus $b(\theta)$, along with the discounted future benefit from the relationship. The right-hand side of (7) is the payoff from her most profitable deviation—choosing the highest possible action a^{max} and withholding any bonus payments—in which case she receives her outside option in each subsequent period. The constraint (8) can be explained similarly, and ensures that the agent does not withhold discretionary payments in cases where $b(\theta) < 0$.

By standard arguments from the relational contracting literature (e.g. MacLeod & Malcomson 1989), the parties' problem (4)-(8) can be reduced to choosing the surplus-maximizing action schedule, subject to an aggregate dynamic enforcement constraint:

$$\frac{\delta}{1-\delta} \left[\int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} - \bar{\pi} \right] \geq v(a^{max}) - v(a(\theta)) \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \quad (9)$$

attained by summing (7) and (8). The formal arguments are outlined in the appendix. The left-hand side of (9) is the future discounted surplus from the relationship, while the right-hand side is the principal's gain from a deviation that implements the maximum possible action. The following proposition solves for the action schedule that maximizes (4), subject to this constraint.

Proposition 3. *The optimal relational contract under symmetric information implements the action schedule:*

$$a^S(\theta) = \begin{cases} a^*(\theta) & \underline{\theta} \leq \theta < \theta^c \\ a^c & \theta^c \leq \theta \leq \bar{\theta} \end{cases} \quad (10)$$

where a^c is independent of θ , $\theta^c \in [\underline{\theta}, \bar{\theta}]$ and $a^c > a^*(\theta)$ for all $\theta \in (\theta^c, \bar{\theta}]$.

Proposition 3 shows that the optimal action schedule under principal authority with symmetric information can be split into two sections. We distinguish between three different cases. First, for δ sufficiently high, (9) remains slack for all θ when the first-best action schedule is implemented, in which case $\theta^c = \bar{\theta}$ and $a^S(\theta) = a^*(\theta)$ everywhere. Second, for moderate values of δ , (9) begins to bind for some, but not all, cost realizations; Figure 1 illustrates an example of this case.⁶ For low realizations of $\theta < \theta^c$, the first-best action is relatively high, in which case the principal's benefit from reneging on the agreement by implementing the maximum action (i.e. the right-hand side of eqn. 9) is small. Accordingly, the first-best can be implemented here, so that $a^S(\theta) = a^*(\theta)$ for $\theta < \theta^c$. In contrast, for high realizations of $\theta > \theta^c$, the first-best action is low, so that the right-hand side of (9) becomes prohibitively large. Accordingly, the constraint binds. To reduce the principal's temptation to deviate, the action schedule must then be adjusted upward over this region, and becomes constant with $a^S(\theta) = a^c$, above the first-best. The critical value θ^c separates these two regions, and satisfies $a^S(\theta^c) = a^*(\theta^c) = a^c$.

As we reduce δ further, the flat section of the action schedule becomes both higher (a^c increases) and wider (θ^c decreases), until eventually δ is such that the first-best action cannot be implemented for any θ . In this third case, we have $\theta^c = \underline{\theta}$, and the schedule becomes flat throughout with $a^S(\theta) = a^c > a^*(\theta)$ for all

⁶Specifically, Figure 1 illustrates an example where $a \in [0, 1]$, $v(a) = a$ and $c(a, \theta) = \frac{1}{2}\theta a^2$, with θ uniformly distributed over $[1.5, 3]$ and $\delta = 0.718$.




Figure 1: The first-best action schedule, $a^*(\theta)$ (dashed curve) and the optimal action schedule under symmetric information, $a^S(\theta)$ (solid curve).

cost realizations. Eventually, as $\delta \rightarrow 0$, a relational agreement can only be sustained for actions sufficiently close to a^{max} , in which case a profitable relationship may fail to exist.

4.2 Adverse Selection

We now turn our attention to the case where the agent privately observes the cost parameter θ in each period. As in the previous subsection, our goal is to derive the relational contract that maximizes the surplus generated by the parties' interaction. Note that since Proposition 2 applies only in the case of symmetric information between parties, it is possible that this contract is non-stationary. Our analysis proceeds as follows. First, we initially restrict attention to stationary contracts, and, in this environment, solve for the contract that maximizes the parties' surplus. Second, at the end of the section, we show that no other non-stationary relational agreement can generate a higher surplus.

Compared to the case of symmetric information, the presence of adverse selection introduces additional incentive compatibility requirements that the contract must satisfy. Specifically, in each period, once the contract has been agreed upon and the agent has privately observed his cost realization θ , we have the following sequence of events. First, the agent is required to report this realization to the principal. However, since θ is privately observed, he is free to report some $\hat{\theta} \neq \theta$ that yields a more favorable outcome. Moreover, since the principal cannot observe such a deviation, she cannot punish the agent by breaking off trade.⁷ It follows that the combination of action schedule $a(\theta)$ and bonus schedule $b(\theta)$ must be designed such that the relational agreement is incentive compatible. Formally, we require the following set of constraints:

$$b(\theta) - c(a(\theta), \theta) \geq b(\hat{\theta}) - c(a(\hat{\theta}), \theta) \quad \forall \theta, \hat{\theta} \in [\underline{\theta}, \bar{\theta}] \quad (11)$$

which guarantee that the agent reports his type honestly.⁸

Second, after receiving the agent's report, the principal decides which action to implement. After a report

⁷In principle, one could allow the agent to refuse to make a report, or report $\hat{\theta} \notin [\underline{\theta}, \bar{\theta}]$. However, since this is an observable deviation, the principal would optimally respond to such behavior by implementing $a = a^{max}$, refusing to pay any bonus transfers and breaking off trade in subsequent periods; clearly, the agent cannot benefit from this.

⁸Note that (11) is identical to the set of incentive compatibility constraints when the agent has decision rights.

of $\hat{\theta}$, the principal can either implement $a(\hat{\theta})$, or deviate from the relational agreement by implementing some $a \neq a(\hat{\theta})$. Since any deviation from $a(\hat{\theta})$ is observed by the agent and punished by refusing to interact in future periods, her most advantageous deviation is to set $a = a^{max}$, as in the case of symmetric information. Third, the parties decide whether to make their promised discretionary transfers. The constraints (7) and (8) guarantee that neither party can benefit from these deviations, for any realization of θ . Overall, the optimal stationary relational contract under adverse selection solves (4)-(8), along with the additional incentive compatibility requirements (11). The following lemma simplifies this problem.

Lemma 1. *An action schedule $a(\theta)$ is implementable under principal authority with adverse selection by a stationary relational contract if and only if (i) $a(\theta)$ is non-increasing and (ii) we have:*

$$\begin{aligned} \frac{\delta}{1-\delta} \left[\int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} - \bar{\pi} \right] \\ \geq v(a^{max}) - v(a(\theta)) + c(a(\theta), \theta) - c(a(\bar{\theta}), \bar{\theta}) + \int_{\theta}^{\bar{\theta}} c_{\theta}(a(x), x) dx \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \end{aligned} \quad (12)$$

Lemma 1 shows that the additional incentive compatibility requirements (11) under adverse selection result in two key differences, relative to the case of symmetric information. In particular, by standard arguments from the contract theory literature, incentive compatibility implies a monotonicity constraint on the action schedule, along with a restriction on the set of bonus schedules that can be incorporated into the contract. The proof of Lemma 1 shows that the schedule

$$\hat{b}(\theta) \equiv c(a(\theta), \theta) - c(a(\bar{\theta}), \bar{\theta}) + \int_{\theta}^{\bar{\theta}} c_{\theta}(a(x), x) dx \quad (13)$$

is incentive compatible for any non-increasing action schedule; adding $\hat{b}(\theta)$ to the right-hand side of (9)—the aggregate dynamic enforcement constraint under symmetric information—yields the corresponding restriction under adverse selection, (12). Intuitively, relative to symmetric information, under adverse selection the agent must be offered (weakly) increased bonuses to induce truthful reporting of his cost parameter; this tightens (12), thereby reducing the set of implementable action schedules for sufficiently low δ .

By Lemma 1, the optimal stationary contract implements the non-increasing action schedule that maximizes the surplus from the relationship, subject to the constraint (12). Clearly, the parties cannot do better than implementing the optimal action schedule under symmetric information, $a^S(\theta)$, which is non-increasing. The question is whether the tightened constraint under adverse selection, (12), allows for this. Perhaps surprisingly, the answer to that question is yes.

Proposition 4. *The optimal action schedule under symmetric information, $a^S(\theta)$, is also implementable under adverse selection.*

By the foregoing arguments, to prove Proposition 4 it is sufficient to show that $a^S(\theta)$ satisfies (12) for all realizations of θ . We begin by defining:

$$\gamma(\theta) \equiv v(a^{max}) - v(a^S(\theta)) + c(a^S(\theta), \theta) - c(a^S(\bar{\theta}), \bar{\theta}) + \int_{\theta}^{\bar{\theta}} c_{\theta}(a^S(x), x) dx \quad (14)$$

as the right-hand side of (12) given action schedule $a^S(\theta)$; this is the current-period benefit to the principal

from a deviation when implementing $a^S(\theta)$ using the incentive compatible bonus schedule $\hat{b}(\theta)$, as a function of θ . The key observation of Proposition 4 is that $\gamma(\theta)$ is constant over $[\underline{\theta}, \bar{\theta}]$, which we show in the appendix. A sketch of the proof is as follows. Taking the total derivative of (14) yields:

$$\gamma'(\theta) = -[v'(a^S(\theta)) - c_a(a^S(\theta), \theta)] \frac{da^S(\theta)}{d\theta} \quad (15)$$

for any θ where the derivative exists. Recall that the optimal action schedule under symmetric information $a^S(\theta)$ can be split into two sections. First, for high cost realizations $\theta \in [\theta^c, \bar{\theta}]$, we have $a^S(\theta) = a^c$ and the action is constant; this implies that $da^S(\theta)/d\theta = 0$ in (15), so that $\gamma'(\theta) = 0$ over this section. Intuitively, since these high cost types are pooled together, the agent has no incentive to overstate his costs, so there is no need to use discretionary payments to induce separation; accordingly, $\hat{b}(\theta) = 0$ over this entire region. As there is no difference in either the implemented action or bonus payments over $[\theta^c, \bar{\theta}]$, deviation incentives $\gamma(\theta)$ are also constant.

Second, for low cost realizations $\theta \in [\underline{\theta}, \theta^c)$, we have $a^S(\theta) = a^*(\theta)$. By the definition of $a^*(\theta)$, (2), this implies that the square bracket in (15) is equal to zero, so that $\gamma'(\theta) = 0$ also over this section. Intuitively, for $\theta \in [\underline{\theta}, \theta^c)$, the principal implements an action that is strictly greater than a^c . On the one hand, this reduces the principal's deviation incentives, since the benefit from deviating by implementing the highest possible action, $v(a^{max}) - v(a(\theta))$, is reduced. On the other, to ensure incentive compatibility she must pay strictly positive discretionary bonuses, which increase her incentives to renege. When $a(\theta) = a^*(\theta)$, these two effects exactly offset one another, so that the principal's deviation incentives are the same for any $\theta \in [\underline{\theta}, \theta^c)$, as they are for $\theta \in [\theta^c, \bar{\theta}]$. Hence, $\gamma(\theta)$ is constant over the whole schedule.

Accordingly, to show that (12) is satisfied for all $\theta \in [\underline{\theta}, \bar{\theta}]$, it is sufficient to show that it is satisfied for one such θ . However, this is trivial: since $\hat{b}(\theta) = 0$ for any $\theta \in [\theta^c, \bar{\theta}]$, (12) reduces to (9), the dynamic enforcement constraint under symmetric information, which is satisfied for all such θ . It then follows from Lemma 1 that $a^S(\theta)$ is implementable, and therefore the optimal action schedule under adverse selection.

Finally, the fact that $a^S(\theta)$ can be implemented using a stationary contract under adverse selection implies that there is no other non-stationary relational agreement that can improve on this outcome. To see this, suppose to the contrary that there exists a non-stationary contract that implements an action schedule $a^\dagger(\theta)$ that generates a strictly higher surplus than $a^S(\theta)$. It then follows that the same contract could also implement $a^\dagger(\theta)$ in a symmetric information environment, and, by Proposition 2, the associated surplus could be attained by a stationary contract. However, this clearly contradicts the notion that the optimal contract under symmetric information implements $a^S(\theta)$, as shown by Proposition 3.

4.3 Discussion

Proposition 4 implies that the presence of private information regarding the agent's costs has no impact on what can be achieved by relational agreements when the principal has authority, in contrast to the case where the agent chooses the action. While this result appears surprising, the underlying intuition is relatively straightforward. Loosely speaking, relational contracts become more difficult to sustain as the promised action becomes further away from the 'default action' implemented in the absence of contracts, which, under principal authority, is a^{max} . Accordingly, with symmetric information, deviation incentives are greatest for high cost realizations, where the efficient action is necessarily low. As a result, the optimal contract curtails incentives to deviate by increasing the specified action for those values of θ for which the dynamic enforcement constraint (9) binds. In contrast, for lower cost realizations, (9) remains slack and the

first-best is implemented.

The presence of an adverse selection problem requires offering additional incentives, via the discretionary transfers, to induce separation. But, crucially, since these transfers must be designed to stop the agent from overstating his costs, rather than understating them, they are highest for those low cost realizations where the dynamic enforcement constraint is slack. Moreover, since pooling already occurs under symmetric information for those with high costs, an agent with a realization $\theta > \theta^c$ does not benefit from overstating his costs, and hence requires no bonus, leaving the dynamic enforcement constraint unchanged. Put simply, the problems caused by asymmetric information occur at the opposite end of the schedule to those caused by the need for the contract to be self-enforcing.

This same line of thinking illustrates why private information *does* reduce the attainable surplus when the agent has authority, as discussed in Section 3. In this case, the default action is at the lower end of the action space, a^{min} , so that relational contracts are most difficult to sustain for the high actions that are efficient with low cost realizations. But since the allocation of authority does not affect the agent's incentives to overstate his costs under adverse selection, these are exactly those realizations for which higher discretionary bonuses are required for incentive compatibility. Here, both problems occur at the same end of the schedule and therefore compound one another.

Proposition 4 also allows for a comparison of how the action schedule varies under different allocations of authority.

Corollary 1. *The action schedule implemented by the optimal relational contract under adverse selection with principal authority takes one of three forms:*

1. *First-best:* $a(\theta) = a^*(\theta)$ for all θ .
2. *Partial Pooling:* $a(\theta)$ is strictly decreasing on $[\underline{\theta}, \theta^c)$ with $a(\theta) = a^*(\theta)$, and constant on $[\theta^c, \bar{\theta}]$ for some $\theta^c \in (\underline{\theta}, \bar{\theta})$.
3. *Pooling:* $a(\theta)$ is the same for all θ .

Compared with Proposition 1, the optimal action schedules under principal and agent authority share several similarities. For sufficiently high δ the first-best is implemented everywhere, while for sufficiently low δ there is complete pooling. Between these two extremes, for middling values of δ , both allocations of authority feature regions where there is pooling of actions despite differing cost realizations. This supports the notion that, with relational contracts, relationships between firms become more flexible as the attainable surplus from the relationship increases (Levin 2003).

However, there are also a number of notable differences between the two different allocations of authority. First, partial pooling occurs at different ends of the action schedule; i.e. for different realizations of θ . Second, under principal authority, inefficiencies take the form of an overprovision of the action relative to the efficient level, compared to an underprovision with agent authority. Both of these differences follow from the impact of authority over parties' incentives to renege on the relationship, as discussed in the foregoing. Third, under principal authority it is possible to have inefficiencies for a strict subset of cost realizations, with the first-best implemented elsewhere. This is never the case under agent authority, where the action schedule is either efficient everywhere or nowhere. Intuitively, the incentive compatible bonus required to induce separation from low cost types is dependent on the actions implemented throughout the schedule. Accordingly, when a binding dynamic enforcement constraint prevents the first-best from being implemented for low cost realizations under agent authority, it is optimal to adjust the action schedule everywhere: reducing

the actions of high cost agents decreases the extent to which the actions of low cost agents must be lowered. The inefficiency thus cascades upward throughout the action schedule as a result of agents' incentives to overstate their costs. In contrast, under principal authority, low cost realizations already implement the first best, which cannot be improved upon, so that this rationale is absent.

Our results also contrast to findings from the literature that studies standard adverse selection problems (i.e. in single period games and with formally contractable actions).⁹ A persistent result from this literature is that distortions from the first-best occur throughout the schedule, except for the most efficient type (see e.g. Bolton & Dewatripont 2004); in other words, inefficiency is a probability one event. This is similar to the case of agent authority under relational contracting, as discussed above, where the action schedule is often inefficient everywhere. In contrast, in the current setting where the principal has authority, the optimal action schedule often implements the first-best action for an interval of the most efficient types, with inefficiencies only occurring for realizations of sufficiently high costs, i.e. with probability strictly less than one. Moreover, in the standard setting, under certain regularity conditions on the distribution function, optimal contracts perfectly separate the actions of agents with different cost realizations. The results of Corollary 1 and Levin (2003) emphasize that pooling of different types is a common feature of relational contracts.

To conclude the section, we discuss the implications for our results for the optimal allocation of decision rights in relational contracts.

Corollary 2. *If principal authority is optimal under symmetric information, it is also optimal under adverse selection. In contrast, there exist cases where agent authority is preferable under symmetric information, but not in the presence of adverse selection.*

Our findings suggest that the existence of private information regarding the agent's costs favors allocation of decision rights to the principal. This follows straightforwardly from the fact that adverse selection can reduce the attainable surplus when the agent has authority, but does not affect the surplus under principal authority. It is important to note, however, that allocating authority to the principal is not always optimal under adverse selection; for some specifications of the model, principal authority creates large inefficiencies solely due to problems stemming from the need for contracts to be self-enforcing, so that agent authority can remain optimal even with informational asymmetries. This is especially the case if a^{max} is large, relative to the first-best action schedule $a^*(\theta)$, so that the principal's benefit from deviating is significant.¹⁰

More generally, our results suggest that when contracts must be self-enforcing, *ceteris paribus*, decision rights should be allocated to the party that has an informational disadvantage regarding the costs and benefits of different courses of action. To illustrate this idea, consider a variation of the model in which the agent's cost parameter θ is fixed, but the principal's benefit can vary between periods according to a parameter $\omega \sim G(\omega)$: e.g. $v(a, \omega)$, where v is strictly increasing in both arguments, with $v_{a\omega} > 0$. In this case, it is straightforward to verify that the efficient action $a^*(\omega)$ is increasing in ω . If the principal has private information regarding ω , he is always tempted to overstate his benefit, so that incentive compatibility implies a tighter dynamic enforcement constraint, in particular, following low realizations of ω . This is especially

⁹In this setting, the allocation of decision rights plays no role. Since actions are verifiable, a formal contract specifies a set of permissible actions $\{a(\theta) \mid \theta \in [\underline{\theta}, \bar{\theta}]\}$. The agent chooses an action from this set: directly, under agent authority, or indirectly, by reporting the associated θ , under principal authority. In contrast, with relational contracts, the party with authority can deviate 'off-schedule', i.e., by implementing some $\tilde{a} \notin \{a(\theta) \mid \theta \in [\underline{\theta}, \bar{\theta}]\}$, with the optimal deviation varying depending on the allocation of decision rights.

¹⁰Schöttner & Upton (2025) argue that under symmetric information, principal authority dominates agent authority in favorable production environments with large gains to trade; i.e. where low cost realizations are common or the value of the action to the principal is high. Corollary 2 implies that this result extends to the current setting with adverse selection, since the attainable surplus under principal authority is unchanged.

problematic under principal authority, since efficiency calls for relatively low actions in response to such realizations, creating large incentives for the principal to deviate by setting $a = a^{max}$. In contrast, the agent can credibly commit to implementing low actions when he has authority, since his incentive to deviate by setting $a = a^{min}$ in this case is small; the need to increase transfers due to incentive compatibility is then of lesser consequence. Indeed, one can show that in such a setup, private information on the side of the principal does not reduce the attainable surplus from the relationship under agent authority, similar to the results derived in the foregoing.¹¹

5 Conclusion

The allocation of authority plays an important role in determining the surplus that can be generated when firms must use relational contracts to sustain relationships. This paper investigates the optimal allocation decision rights when one party possesses private information regarding the efficiency of different actions. We show that when the uninformed party has authority, the presence of private information has zero impact on the maximum surplus that can be generated by the parties' interaction, relative to the case of symmetric information. In contrast, when the informed party is responsible for choosing the action, the need for the contract to be incentive compatible typically reduces the attainable surplus. As a result, the presence of asymmetric information tends to favor the allocation of decision rights to the uninformed party.

6 Appendix

Proof of Proposition 2. Two minor adjustments to Levin's (2003) proof are required; to ease exposition, we adopt here his notation. First, in his proof of Lemma 1, instead of increasing $\pi(\varphi)$, we consider the implications of increasing $u(\varphi)$. Since there is no action constraint for the agent, this does not affect his behavior, and relaxes the constraint that ensures he will adhere to discretionary payments. Accordingly, following the remaining arguments in the proof, any optimal contract must be sequentially optimal.

Second, in Levin's (2003) proof of Theorem 2, when constructing a stationary contract we choose discretionary bonus payments that leave the *principal's* incentives unchanged. Specifically, we set the stationary discretionary payments:

$$b^*(\varphi) \equiv b(\varphi) - \frac{\delta}{1-\delta}\pi(\varphi) + \frac{\delta}{1-\delta}\pi^* \quad (16)$$

and choose the stationary fixed payment such that the principal's per-period expected payoff is π^* . Following the remaining arguments in the proof, it is straightforward to verify that the necessary constraints are satisfied, so that any optimal contract can be replicated by a stationary contract that generates the same surplus. $\square$

Proof of Proposition 3. The parties' problem is to maximize (4) subject to (5)-(8). The aggregate dynamic enforcement constraint (9) is attained by summing (7) and (8), and is therefore a necessary condition. Suppose the parties set $w = \bar{u} + \int_{\underline{\theta}}^{\bar{\theta}} [c(a(\theta), \theta)] f(\theta) d\theta$ and $b(\theta) = 0$ for all θ , in which case the principal extracts the entire surplus, while the agent earns his outside option; (5) and (6) are thus satisfied. Moreover, (8) holds with equality, while (7) reduces to (9). Accordingly, (9) is also a sufficient condition.

¹¹In an extension to their model, Fahn & Klein (2019) consider an alternative timing where the principal's private information about period t is revealed at the beginning of the same period, rather than during period $t - 1$. As a result, the principal's incentive to lie in order to reduce the bonus (in period $t - 1$) is absent, so that the outcome becomes equivalent to the case of symmetric information. However, they consider binary types, and do not analyze the allocation of authority.

The parties therefore maximize (4) subject to (9). Let the solution to this problem be denoted by $a^{opt}(\theta)$ and define $a^c = \inf_{\theta} a^{opt}(\theta)$. Using (9), we must have:

$$\frac{\delta}{1-\delta} \left[\int_{\underline{\theta}}^{\bar{\theta}} [v(a^{opt}(\theta)) - c(a^{opt}(\theta), \theta)] f(\theta) d\theta - \bar{u} - \bar{\pi} \right] \geq v(a^{max}) - v(a^c) \quad (17)$$

Claim 1. $a^{opt}(\theta)$ also solves the alternative problem of maximizing (4) subject to the constraint:

$$a(\theta) \geq a^c \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \quad (18)$$

Proof of Claim 1. Suppose to the contrary that $a^{opt}(\theta)$ does not solve the problem, and let us denote the actual solution by $\xi(\theta)$. Since $a^{opt}(\theta)$ satisfies (18), it must be that $\xi(\theta)$ generates a strictly higher surplus (4) than $a^{opt}(\theta)$. However, (17) then implies that $\xi(\theta)$ satisfies (9), and therefore maximizes (4) subject to (9), contradicting the fact that $a^{opt}(\theta)$ is the solution to our original problem. $\square$

Accordingly, we identify $a^{opt}(\theta)$ by maximizing (4) subject to (18), using pointwise optimization. The associated Lagrangian is:

$$\mathcal{L} = \int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} - \bar{\pi} + \rho(\theta) [a(\theta) - a^c] \quad (19)$$

with first-order condition:

$$[v'(a(\theta)) - c_a(a(\theta), \theta)] f(\theta) + \rho(\theta) = 0 \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \quad (20)$$

We have two possible cases for each realization of θ . First, if $\rho(\theta) = 0$ then:

$$v'(a(\theta)) - c_a(a(\theta), \theta) = 0 \quad (21)$$

so that $a(\theta) = a^*(\theta)$. Second, if $\rho(\theta) > 0$ then we must have $a(\theta) = a^c$. Moreover, from (20):

$$[v'(a(\theta)) - c_a(a(\theta), \theta)] f(\theta) = -\rho(\theta) < 0 \quad (22)$$

$$\Rightarrow v'(a(\theta)) < c_a(a(\theta), \theta) \quad (23)$$

so that $a(\theta) = a^c > a^*(\theta)$. Accordingly, for all θ such that $a^*(\theta) \geq a^c$, we cannot have $a(\theta) = a^c$; hence, $a(\theta) = a^*(\theta)$ for these θ . Moreover, for all θ such that $a^*(\theta) < a^c$, we clearly cannot have $a(\theta) = a^*(\theta)$ without violating (18), so that $a(\theta) = a^c$. Altogether, we conclude that $a(\theta) = a^*(\theta)$ if and only if $a^*(\theta) \geq a^c$, with $a(\theta) = a^c$ otherwise. The optimal action schedule can then be expressed as (10). $\square$

Proof of Lemma 1. Define $U(\theta) \equiv w + b(\theta) - c(a(\theta), \theta)$. By standard arguments from the contract theory literature (see e.g. Stole 2001), the set of incentive compatibility constraints (11) are equivalent to a requirement that $a(\theta)$ is non-increasing and:

$$U(\theta_1) - U(\theta_0) = - \int_{\theta_0}^{\theta_1} c_{\theta}(a(x), x) dx \quad (24)$$

for all $\theta_0, \theta_1 \in [\underline{\theta}, \bar{\theta}]$.¹² This then implies that:

$$U(\theta) = U(\bar{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(a(x), x) dx \quad (25)$$

for any θ , or, inserting the definition of $U(\theta)$ and adding the agent's future discounted net benefit from the interaction to each side of (25):

$$\begin{aligned} w + b(\theta) - c(a(\theta), \theta) + \frac{\delta}{1-\delta} \left[w + \int_{\underline{\theta}}^{\bar{\theta}} [b(\theta) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} \right] \\ = w + b(\bar{\theta}) - c(a(\bar{\theta}), \bar{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(a(x), x) dx + \frac{\delta}{1-\delta} \left[w + \int_{\underline{\theta}}^{\bar{\theta}} [b(\theta) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} \right] \end{aligned} \quad (26)$$

Any self-enforcing contract must satisfy both (7) and (6). Summing (7) and (26), rearranging and simplifying yields:

$$\begin{aligned} \frac{\delta}{1-\delta} \left[\int_{\underline{\theta}}^{\bar{\theta}} [v(a(\theta)) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} - \bar{\pi} \right] \\ \geq v(a^{max}) - v(a(\theta)) + b(\bar{\theta}) + c(a(\theta), \theta) - c(a(\bar{\theta}), \bar{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(a(x), x) dx \\ + \frac{\delta}{1-\delta} \left[w + \int_{\underline{\theta}}^{\bar{\theta}} [b(\theta) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} \right] \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \end{aligned} \quad (27)$$

Finally, rearranging (6) for $\theta = \bar{\theta}$ implies:

$$b(\bar{\theta}) \geq -\frac{\delta}{1-\delta} \left[w + \int_{\underline{\theta}}^{\bar{\theta}} [b(\theta) - c(a(\theta), \theta)] f(\theta) d\theta - \bar{u} \right] \quad (28)$$

Combining this with (27) yields (12). Hence, conditions (i) and (ii) in Lemma 1 are necessarily satisfied by a self-enforcing contract. For sufficiency, consider a contract with a non-increasing action schedule $a(\theta)$, the bonus schedule $\hat{b}(\theta)$ defined by (13) and the fixed payment:

$$\hat{w} = \bar{u} - \int_{\underline{\theta}}^{\bar{\theta}} [\hat{b}(\theta) - c(a(\theta), \theta)] f(\theta) d\theta \quad (29)$$

This implies that the principal extracts the entire surplus from the relationship, while the agent's expected payoff is equal to his outside option; accordingly, (5) and (6) are both satisfied. The bonus schedule $\hat{b}(\theta)$ is incentive compatible, so that (11) also holds. Inserting $\hat{w}$ and $\hat{b}(\theta)$ into (7) yields (12), while (6) reduces to:

$$c(a(\theta), \theta) + \int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(a(x), x) dx - c(a(\bar{\theta}), \bar{\theta}) \geq 0 \quad \forall \theta \in [\underline{\theta}, \bar{\theta}] \quad (30)$$

which holds since $\int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(a(x), x) dx \geq \int_{\underline{\theta}}^{\bar{\theta}} c_{\theta}(a(\bar{\theta}), x) dx$, as $a(\theta)$ is non-increasing. Accordingly, all necessary constraints are satisfied. $\square$

¹²The representation (24) follows from the envelope theorem and is valid under fairly general conditions; see the discussion in Milgrom & Segal (2002).

Proof of Proposition 4. We first consider the case where $\theta^c \in (\underline{\theta}, \bar{\theta})$. Over any interval $[\alpha, \beta] \subset [\underline{\theta}, \bar{\theta}]$ where $a^S(\theta) = a^*(\theta)$ for all $\theta \in [\alpha, \beta]$, we have

$$v(a^{max}) - v(a^S(\theta)) + c(a^S(\theta), \theta) - c(a^S(\bar{\theta}), \bar{\theta}) + \int_{\theta}^{\bar{\theta}} c_{\theta}(a^S(x), x) dx \quad (31)$$

constant in θ by (15). Since $a^c = a^*(\theta^c)$, we let $\alpha = \underline{\theta}$ and $\beta = \theta^c$, which implies that $\gamma(\theta) = \gamma(\theta^c)$ for all $\theta \in [\underline{\theta}, \theta^c]$. Next, note that for any $\theta \in [\theta^c, \bar{\theta}]$ we have $a^S(\theta) = a^c$ and therefore:

$$\hat{b}(\theta) = c(a^S(\theta), \theta) - c(a^S(\bar{\theta}), \bar{\theta}) + \int_{\theta}^{\bar{\theta}} c_{\theta}(a^S(x), x) dx \quad (32)$$

$$= c(a^c, \theta) - c(a^c, \bar{\theta}) + \int_{\theta}^{\bar{\theta}} c_{\theta}(a^c, x) dx = 0 \quad (33)$$

Accordingly, for all $\theta \in [\theta^c, \bar{\theta}]$ we have $\gamma(\theta) = v(a^{max}) - v(a^c) = \gamma(\theta^c)$. We conclude that $\gamma(\theta) = \gamma(\theta^c)$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$.

In the case where $\theta^c = \bar{\theta}$, we have $a^S(\theta) = a^*(\theta)$; setting $\alpha = \underline{\theta}$ and $\beta = \bar{\theta}$ in the foregoing then yields an identical result. When $\theta^c = \underline{\theta}$, $a(\theta) = a^c$ everywhere, and the result similarly follows from the above arguments.

Finally, by the arguments in the main text, (12) must be satisfied for some θ , which in turn implies that (12) holds for all θ when $a(\theta) = a^S(\theta)$. $\square$

References

- Abreu, D. (1988), ‘On the theory of infinitely repeated games with discounting’, *Econometrica* **56**(2), 383–396.
- Baker, G., Gibbons, R. & Murphy, K. J. (1999), ‘Informal authority in organizations’, *Journal of Law, Economics, and Organization* **15**(1), 56–73.
- Baker, G., Gibbons, R. & Murphy, K. J. (2002), ‘Relational contracts and the theory of the firm’, *The Quarterly Journal of Economics* **117**(1), 39–84.
- Baker, G., Gibbons, R. S. & Murphy, K. J. (2011), Relational adaptation. Working paper, Harvard Business School.
- Baker, G. P., Gibbons, R. & Murphy, K. J. (2008), ‘Strategic alliances: Bridges between "islands of conscious power"', *Journal of the Japanese and International Economies* **22**(2), 146–163.
- Bester, H. & Krähmer, D. (2017), ‘The optimal allocation of decision and exit rights in organizations’, *The RAND Journal of Economics* **48**(2), 309–334.
- Bolton, P. & Dewatripont, M. (2004), *Contract Theory*, MIT Press.
- Englmaier, F., Filipi, A. & Singh, R. (2010), ‘Incentives, reputation and the allocation of authority’, *Journal of Economic Behavior & Organization* **76**(2), 413–427.
- Fahn, M. & Klein, N. (2019), ‘Relational contracts with private information on the future value of the relationship: The upside of implicit downsizing costs’, *American Economic Journal: Microeconomics* **11**(4), 33–58.

- Gil, R. & Zananone, G. (2017), ‘Formal and informal contracting: Theory and evidence’, *Annual Review of Law and Social Science* **13**(1), 141–159.
- Halac, M. (2012), ‘Relational contracts and the value of relationships’, *American Economic Review* **102**(2), 750–779.
- Levin, J. (2003), ‘Relational incentive contracts’, *American Economic Review* **93**(3), 835–857.
- Li, J. & Matouschek, N. (2013), ‘Managing conflicts in relational contracts’, *American Economic Review* **103**(6), 2328–2351.
- Luo, Z., Liu, Z., Qiu, Y. & Yu, S. (2025), ‘A relational theory of power alternation’, *Working paper*.
- MacLeod, W. B. & Malcomson, J. M. (1989), ‘Implicit contracts, incentive compatibility, and involuntary unemployment’, *Econometrica* **57**(2), 447–480.
- Malcomson, J. M. (2012), Relational incentive contracts, in R. Gibbons & J. Roberts, eds, ‘The Handbook of Organizational Economics’, Princeton University Press, Princeton, pp. 1014–1065.
- Malcomson, J. M. (2015), ‘Relational incentive contracts with productivity shocks’, *Games and Economic Behavior* **92**, 122–137.
- Malcomson, J. M. (2016), ‘Relational incentive contracts with persistent private information’, *Econometrica* **84**(1), 317–346.
- Milgrom, P. & Segal, I. (2002), ‘Envelope theorems for arbitrary choice sets’, *Econometrica* **70**(2), 583–601.
- Schöttner, A. & Upton, H. (2025), ‘When should a firm employ its workers? A relational contracting approach’, *TRR 266 Accounting for Transparency Working Paper Series No. 105*.
- Stole, L. A. (2001), Lectures on the theory of contracts and organizations. Unpublished manuscript.
- Upton, H. (2024), ‘Implicit incentives and delegation in teams’, *Management Science* **70**(7), 4722–4741.
- Yang, H. (2013), ‘Nonstationary relational contracts with adverse selection’, *International Economic Review* **54**(2), 525–547.