

Dynamic Mechanisms without Money

Yingni Guo

Johannes Hörner*

October 19, 2026

Abstract

We analyze the design of dynamic mechanisms without transfers. The designer uses future allocation decisions to elicit private information. Values evolve according to a two-state Markov chain. We solve for the optimal allocation rule, which permits a simple implementation: the agent is entitled to a number of units “no questions asked”; this number is reduced by one when he requests a unit and revised, although not necessarily upward, when he forgoes one. Unlike with transfers, inefficiencies are backloaded, and both immiseration and its polar opposite are possible long-run outcomes. In the continuous-time limit, we show that persistence hurts.

Keywords: Mechanism design; principal-agent; token mechanisms.

JEL codes: C73, D82.

1 Introduction

This paper is concerned with the dynamic allocation of resources when transfers are not allowed and information regarding their optimal use is privately known to an individual. The informed agent is strategic rather than necessarily truthful. We seek a mechanism that maximizes expected discounted surplus subject to incentive compatibility.

Both the societal decision not to allow money—for economic, physical, legal, or ethical reasons—and the sequential nature of the problem are assumed. Services, tasks, and attention must often be allocated before future demands are known.

*We thank Junchi Li for excellent research assistance. Guo: Department of Economics, Duke University; email: yingni.guo@duke.edu. Hörner: CNRS (Toulouse School of Economics); email: joh.horner@gmail.com.

Throughout, we assume that the good to be allocated is perishable.¹ Absent private information, the allocation problem is trivial: the good should be provided if and only if its value exceeds its cost.² However, in the presence of private information, and in the absence of transfers, linking future allocation decisions to current decisions is the only instrument available to society to elicit truthful information. Our goal is to understand this link.

Our main results are a characterization of the optimal mechanism and an intuitive indirect implementation for it. In essence, the agent should be granted an entitlement, corresponding to a certain number of units of the good that he can claim “no questions asked.” This entitlement is updated according to his choice: whenever the agent desires the unit, his entitlement is reduced by one unit; whenever he forgoes it, his entitlement is also revised, although not necessarily upward. Furthermore, we demonstrate that this results in simple dynamics: the efficient allocation is used whenever promise keeping and incentive compatibility permit it. Eventually, provision becomes permanent or ceases permanently. These results contrast with those from static design with multiple units (*e.g.*, Jackson and Sonnenschein, 2007) and from dynamic mechanism design with transfers (*e.g.*, Battaglini, 2005).

Formally, our good can take one of two values during each round. Values are serially correlated. The binary assumption is certainly restrictive, and general characterizations become more difficult with more than two persistent types, even with transfers (see Battaglini and Lamba, 2019).³ We begin with the i.i.d. case, which suffices to illustrate many of the insights of our analysis, before proving the results in full generality. The cost of providing the good is fixed and known. Hence, it is efficient to assign the good during a given round if and only if the value is high. We cast our problem of solving for the efficient mechanism (given the values, cost, and agent’s discount factor) as one faced by a disinterested principal with commitment who determines when to supply the good as a function of the agent’s reports. There are no transfers, no certification technology, and no signals concerning the agent’s value, even *ex post*.

As mentioned above, the optimal policy can be implemented through a “budget” mechanism in which the appropriate unit of account is the number of units that the agent is

¹Many allocation decisions involve goods or services that are perishable, such as how a nurse or a worker should divide their time; which patients should receive scarce medical resources (blood or treatments); or which investments and activities should be approved by a firm.

²With constant provision costs and no intertemporal constraints, the complete-information problem separates across rounds. Optimal ordering of perishable goods raises different issues; an early contribution is Fries (1975).

³In Section 5.2, we consider the case of a continuum of types that are independently distributed over time.

entitled to receive in consecutive rounds with “no questions asked.” While the updating process when the agent forgoes a unit is somewhat subtle, it is independent of the principal’s belief concerning the agent’s type. The only role of the prior belief is to specify the initial budget. Unlike in a token mechanism, the total (discounted) number of units the agent receives is not fixed: depending on the sequence of reports, he might ultimately receive few or many units.⁴ Eventually, the agent is either granted the unit forever (“retirement”) or never again. Hence, the agent may escape immiseration (Thomas and Worrall, 1990). The positive lower bound on values keeps the marginal cost of delivering utility bounded near the highest promise. Retirement therefore occurs with positive probability. Attainability of the highest promise alone does not suffice: with continuous types whose support reaches zero, the marginal cost diverges and immiseration occurs almost surely.

In Section 5.1, we study the continuous-time limit, in which the flow value of the good changes according to a two-state Markov chain. This allows us to demonstrate that persistence hurts: as the Markov chain becomes more persistent, efficiency decreases, although the agent might benefit.

Allocation problems in the absence of transfers are plentiful, and it is not our purpose to survey them here. We believe that our results can inform practices concerning how to implement algorithms to improve allocations. For example, consider nurses who must decide whether to take seriously alarms that are triggered by patients. The opportunity cost of nurses’ time is significant. Patients, however, appreciate quality time with nurses irrespective of whether their condition necessitates it. This discrepancy produces a challenge with which every hospital must contend: ignore alarms and risk that a patient with a serious condition is not attended to, or heed all alarms and overwhelm the nurses. “Alarm fatigue” is a serious problem that health care must confront (see, for instance, Sendelbach, 2012). We suggest the best approach for trading off the risk of neglecting a patient in need against that of attending to one who simply cries wolf.⁵

Related Literature. Our work is closely related to the bodies of literature on mechanism design with transfers and on “linking incentive constraints.” Section 4.5 explains why transfers matter, and Section 3.4 explains why the dynamic nature of the relationship does; hence, we will be brief here.

⁴The budget mechanism is not a “bankruptcy” mechanism in the sense of Radner (1986) because the specific ordering of the reports matters.

⁵Clearly, our mechanism is much simpler than existing electronic nursing workload systems. However, none appears to seriously consider strategic agent behavior as a constraint.

The obvious benchmark work is Battaglini (2005), who considers our general model but allows transfers.⁶ Another important difference is his focus on revenue maximization, a meaningless objective in the absence of transfers. The results of his work are diametrically opposed to ours. In Battaglini’s model, efficiency improves over time (exact efficiency is eventually obtained with probability 1). Here, the allocation converges to permanent provision or permanent nonprovision, independent of the report. Efficiency need not decrease along every history. The agent’s utility can increase or decrease depending on the history that is realized: receiving the good forever is clearly the best possible outcome, while never receiving it again is the worst. Krishna, Lopomo and Taylor (2013) provide an analysis of liquidity constraints (payments flow only from the principal to the agent) with independently distributed costs, suggesting that restricting transfers affects both the optimal contract and the dynamics. Note that there is an important exception to the quasilinearity commonly assumed in the dynamic mechanism design literature, namely, Garrett and Pavan (2015).

“Linking incentive constraints” refers to the notion that as the number of identical decision problems increases, linking them allows the designer to improve on the isolated problem. See Fang and Norman (2006) and Jackson and Sonnenschein (2007) for papers specifically devoted to this idea (which dates back at least to Radner, 1981; Rubinstein and Yaari, 1983). Hortala-Vallve (2010) provides an interesting analysis of the inefficiencies that must be incurred away from the limit, and Cohn (2010) demonstrates the suboptimality of the mechanisms that are commonly used, even regarding the rate of convergence. Our focus is on the exactly optimal mechanism for a fixed degree of patience, not on proving the asymptotic optimality of a specific mechanism (numerous mechanisms yield asymptotic optimality). We derive the rate of convergence and show that it is the same as for token mechanisms. Another important difference from most of these papers is that our problem is truly dynamic in the sense that the agent does not know future values but learns them as they are realized. Section 3.4 elaborates on this distinction.

The notion that virtual budgets could be used as intertemporal instruments to discipline agents with private information has appeared in several papers in economics. Möbius (2001) might be the first to suggest that tracking the difference in the number of favors granted (with two agents) and using it to decide whether to grant new favors is a simple but powerful way of sustaining cooperation in long-run relationships. See also Athey and Bagwell (2001), Kalla (2010), and Abdulkadiroğlu and Bagwell (2012). While these token mechanisms are known to be suboptimal (as is clear from our characterization of the optimal mechanism),

⁶See Zhang (2012) for a detailed analysis of Battaglini’s model, and see also Fu and Krishna (2019).

they have desirable properties nonetheless: properly calibrated, they yield an approximately efficient allocation as the discount factor approaches one. To our knowledge, Hauser and Hopenhayn (2008) come closest to solving for the optimal mechanism (within the class of perfect public equilibria). Their numerical analysis allows them to assess the performance of the simple token rule (according to which each favor is weighted equally, independent of the history), showing that this rule might be too simple (the loss relative to the first best can reach 30% of surplus). Remarkably, their analysis suggests that the optimal (Pareto-efficient) strategy shares many features with the optimal policy that we derive in our one-player world: the incentive constraint always binds, and the efficient policy is followed unless it is inconsistent with promise keeping (that is, when promised utilities are too extreme). Our model can be regarded as a game with one-sided incomplete information in which the production cost of the principal is known to the agent. There are some differences, however. First, our principal has commitment and hence is neither tempted to act opportunistically nor bound by individual rationality. Second, this principal maximizes efficiency rather than a private objective. Third, there is a technical difference: our limiting model in continuous time corresponds to the Markovian case in which flow values switch according to a Poisson process. In Hauser and Hopenhayn’s model, the lump-sum value arrives according to a Poisson process, and the process is memoryless. Li, Matouschek and Powell (2017) study power dynamics under relational contracting in a model with private information and public shocks. Their public shocks change the environment rather than reveal the agent’s private information.

More generally, that allocation rights to other (or future) units can be used as a “currency” to elicit private information has long been recognized and dates to Hylland and Zeckhauser (1979), who first explained the extent to which this can be viewed as a pseudo-market. Casella (2005) develops a similar idea within the context of voting rights. Miralles (2012) solves a two-unit version of our problem with more general value distributions, but his analysis is not dynamic: both values are (privately) known at the outset. A dynamic two-period version of a similar problem is analyzed by Abdulkadiroğlu and Loertscher (2007).

In all variants considered here, observed values would allow efficient provision: the good would be supplied if and only if its value exceeded its cost. Due to private information, it is necessary to distort the allocation: after some histories, the good is provided forever; after others, the good is never provided again. In this sense, the scarcity of provision is endogenously determined to elicit information. There is a large body of literature in operations research treating this scarcity as exogenous—there are only finitely many opportunities to provide the good, and the problem is then when to exercise these opportunities. Impor-

tant early contributions to this literature include Derman, Lieberman and Ross (1972) and Albright (1974). Their analyses suggest a natural mechanism that can be applied in our environment: the agent receives a certain number of “tokens” and uses them whenever he pleases.

Exactly optimal mechanisms have been computed in related environments. In a closely related setting, Frankel (2016) also derives an optimal mechanism. His model allows for more than two types and more than two actions, but it focuses on the case of serially independent types, which is also the starting point of our analysis. A key difference is that in his model the agent’s preferences are independent of the state, which makes the problem more tractable. Gershkov and Moldovanu (2010) consider a dynamic allocation problem related to Derman, Lieberman and Ross (1972) in which agents possess private information regarding the value of obtaining the good. In their model, agents are myopic and the scarcity of the resource is exogenously given. In addition, transfers are allowed. They demonstrate that the efficient policy of Derman, Lieberman and Ross (which is very different from ours) can be implemented via appropriate transfers.

A related literature considers the problem of optimal stopping in the absence of transfers; see, in particular, Kováč, Krähmer and Tatur (2014). The difference reflects the nature of the good, namely, whether it is durable or perishable. When only one unit is desired and waiting is possible, this represents a stopping problem, as in their paper. With a perishable good, a decision must be made during each round. As a result, the two problems give rise to different incentives and optimal contracts. In the stopping case, forgoing the current unit has option value when the value is low and future prospects are good. This is not the case here—incentives to forgo the unit must be endogenously generated via promises. In the stopping case, there is only one history of outcomes that does not terminate the game. Here, policies differ not only in when the good is first provided but also thereafter.

Finally, while the motivations of the papers differ, the techniques we use for the i.i.d. benchmark borrow numerous ideas from Thomas and Worrall (1990), as we explain in Section 3, and our intellectual debt is considerable.

Section 2 introduces the model. Section 3 solves the i.i.d. benchmark, introducing most of the ideas of the paper, while Section 4 solves the general model. Section 5 extends the results to continuous time and to continuous types. Section 6 concludes.

2 The Model

Time is discrete, indexed by $n = 0, 1, \dots$, and the horizon infinite. There are two parties, a disinterested principal and an agent. During each round, the principal can produce an indivisible unit of a good at a cost $c > 0$. The agent's value (or *type*) during round n , v_n , is a random variable that takes one of two values, l or h . We assume that $0 < l < c < h$, so that supplying the good is efficient if and only if the value is high, but the agent's value is always positive.

The value evolves according to a Markov chain:

$$\mathbf{P}[v_{n+1} = h \mid v_n = h] = 1 - \rho_h, \quad \mathbf{P}[v_{n+1} = l \mid v_n = l] = 1 - \rho_l,$$

for all $n \geq 0$, where $\rho_h, \rho_l > 0$, except in the benchmark of perfect persistence discussed below. The (invariant) probability of h is $q := \rho_l / (\rho_h + \rho_l)$.⁷ Our leading case draws the initial value from the invariant distribution, so $\mathbf{P}[v_0 = h] = q$. Section 4 also considers other initial beliefs. The (unconditional) expected value of the good is denoted $\mu := \mathbf{E}[v_n] = qh + (1 - q)l$. We make no assumptions regarding how μ compares to c .

Let $\kappa := 1 - \rho_h - \rho_l$ be a measure of the persistence of the Markov chain. Throughout, we assume that $\kappa \geq 0$ or, equivalently, $1 - \rho_h \geq \rho_l$; that is, the distribution over tomorrow's type conditional on today's type being h first-order stochastically dominates the distribution conditional on today's type being l .⁸ Two interesting special cases occur when $\kappa = 1$ and $\kappa = 0$. The former corresponds to perfect persistence, while the latter corresponds to independent values.

The agent's value is private information. Specifically, at the beginning of each round, the value is drawn and the agent is informed of it.

Players are impatient and share a common discount factor $\delta \in [0, 1)$.⁹ In the binary model, we assume $\delta > l/\mu$ and $\delta > 1/2$. In the i.i.d. case, the first restriction places the lower peripheral interval within the linear part of the value function. The second ensures strict concavity above that part. We also exclude the boundary case $\delta\rho_l(h - l) = (1 - \delta)l$ in (3). Both strict inequalities are allowed. Section 5.2 states its discount-factor restrictions separately.

⁷When $\rho_h = \rho_l = 0$, the value is constant, and $q \in (0, 1)$ denotes the probability that it is h .

⁸The role of this assumption, which is commonly adopted in the literature, and what occurs in its absence, when values are negatively serially correlated, are discussed in Sections 4.3 and 4.4.

⁹The common discount factor is important. We view our principal as a social planner rather than an actual player, trading off the agent's utility against the social cost of providing the good. Since the planner internalizes the agent's utility, it is natural to assume that his discount factor coincides with the agent's.

Let $x_n \in \{0, 1\}$ denote the supply decision during round n ; *e.g.*, $x_n = 1$ means that the good is supplied during round n .

Our focus is on identifying the (constrained) efficient mechanism defined below. Hence, we assume that the principal internalizes both the cost of supplying the good and its value to the agent. We solve for the principal's preferred mechanism.

Thus, given an infinite history $\{x_n, v_n\}_{n=0}^{\infty}$, the principal's realized *payoff* is defined as follows:

$$(1 - \delta) \sum_{n=0}^{\infty} \delta^n x_n (v_n - c).$$

The agent's realized *utility* is defined as follows:¹⁰

$$(1 - \delta) \sum_{n=0}^{\infty} \delta^n x_n v_n.$$

Throughout, payoff and utility refer to the expectation of these values given the relevant player's information. The *ex ante* utility belongs to the interval $[0, \mu]$. The (risk-neutral) agent seeks to maximize his utility.

We now introduce or emphasize several important assumptions maintained throughout our analysis.

There are no transfers. This is our point of departure from most of the previous research on dynamic mechanism design. Note also that our objective is efficiency, not revenue maximization. With transfers, there is a trivial mechanism that achieves efficiency: supply the good if and only if the agent pays the fixed price c .

There is no *ex post* signal regarding the realized value of the agent—not even payoffs are observed. Depending on the context, it might be realistic to assume that a (possibly noisy) signal of the value occurs at the end of each round, independent of the supply decision. In other applications, it might make more sense to assume instead that this signal occurs only if the good is supplied (*e.g.*, a firm discovers the productivity of a worker who is hired). Conversely, statistical evidence might arise only when the good is not supplied, if supplying it averts a risk (a patient calling for care or police calling for backup). Presumably, the optimal mechanism will differ according to the monitoring structure. Understanding what happens without *any* signal is the natural first step.

The principal commits *ex ante* to a possibly randomized mechanism, as in dynamic mechanism design. By contrast, models of trading favors assume no commitment and study perfect public equilibria. The same distinction applies to Li, Matouschek and Powell (2017).

¹⁰Throughout, the term *payoff* describes the principal's objective and *utility* describes the agent's.

The good is perishable. Hence, previous choices neither constrain future allocations nor affect future values. If the good were perfectly durable and only one unit were demanded, the problem would be one of stopping, as in Kováč, Krähmer and Tatur (2014).

Due to commitment by the principal, we focus, without loss of generality, on policies in which the agent truthfully reports his type during each round and the principal commits to a (possibly random) supply decision as a function of the current report and the entire history of past reports.

Formally, a direct mechanism or *policy* is a collection $(\mathbf{x}_n)_{n=0}^\infty$, with $\mathbf{x}_n : \{l, h\}^{n+1} \rightarrow [0, 1]$,¹¹ mapping a sequence of reports by the agent into a probability of supplying the good during a given round. Our definition exploits the fact that, because preferences are time-separable, the policy may be considered independent of past realized supply decisions. A direct mechanism defines a decision problem for the agent. A reporting strategy is a collection $(\mathbf{m}_n)_{n=0}^\infty$, where $\mathbf{m}_n : \{l, h\}^n \times \{l, h\} \rightarrow \Delta(\{l, h\})$ maps previous reports and the value during round n into a report for that round.¹² The policy is *incentive compatible* if truth-telling (that is, reporting the current value faithfully, independent of past reports) is an optimal reporting strategy.

Our first objective is to solve for the *optimal* policy, that is, for the policy that maximizes the principal's payoff subject to incentive compatibility. The resulting payoff is the *value* of the problem. Second, we would like to find a simple indirect implementation of this policy. Finally, we wish to understand the payoff and utility dynamics under this policy.

3 The i.i.d. Benchmark

We begin our investigation with the case in which values are i.i.d. over time; that is, $\kappa = 0$. This environment is related to that of Thomas and Worrall (1990), although delivering utility in units of the good rather than in unbounded consumption leads to markedly different dynamics. See Section 4 for positively persistent types, $0 < \kappa < 1$, and the benchmark of perfect persistence.

With independent values, it is well known that attention can be further restricted to policies that can be represented by a tuple of functions $p_h, p_l : [0, \mu] \rightarrow [0, 1]$, $u_h, u_l : [0, \mu] \rightarrow [0, \mu]$ mapping a promised utility U onto the probabilities $p_h(U), p_l(U)$ of supplying the good during this round and the continuation utilities $u_h(U), u_l(U)$ at the beginning of the next

¹¹For simplicity, we use the same symbols l, h for the possible agent reports as for the values of the good.

¹²Without loss of generality, we assume that this strategy does not depend on past values, given past reports, as the decision problem from round n onward does not depend on these past values.

round, indexed by the agent's current report. (When no confusion may arise, we omit their dependence on U .) These functions must be consistent in the sense that, given U , the probabilities of supplying the good and the continuation utilities yield the agent a utility of U . This is "promise keeping." We stress that U is the *ex ante* utility in a given round; that is, it is computed before the agent's value is realized. The reader is referred to Spear and Srivastava (1987) and Thomas and Worrall (1990) for details.¹³

Because such a policy is Markovian with respect to the promised utility U , the principal's payoff is also a function of U only. Hence, solving for the optimal policy and the (principal's) value function $W : [0, \mu] \rightarrow \mathbf{R}$ amounts to a Markov decision problem. Given discounting, the optimality equation characterizes both the value and the set of optimal policies. For any fixed $U \in [0, \mu]$, the optimality equation states the following:

$$W(U) = \sup_{p_h, p_l, u_h, u_l} \{ (1 - \delta) (qp_h(h - c) + (1 - q)p_l(l - c)) + \delta (qW(u_h) + (1 - q)W(u_l)) \} \quad (\text{OBJ})$$

subject to incentive compatibility and promise keeping, namely,

$$(1 - \delta)p_h h + \delta u_h \geq (1 - \delta)p_l h + \delta u_l, \quad (\text{IC}_H)$$

$$(1 - \delta)p_l l + \delta u_l \geq (1 - \delta)p_h l + \delta u_h, \quad (\text{IC}_L)$$

$$U = (1 - \delta) (qp_h h + (1 - q)p_l l) + \delta (qu_h + (1 - q)u_l), \quad (\text{PK})$$

$$(p_h, p_l, u_h, u_l) \in [0, 1] \times [0, 1] \times [0, \mu] \times [0, \mu].$$

The incentive-compatibility and promise-keeping conditions are denoted IC (IC_H , IC_L) and PK. This optimization program is denoted $\mathcal{P}$.

The first step is to calculate the value function W as well as the optimal policy. Because the principal is free to choose the promise with which the relationship starts, we are also interested in solving for the *initial promise* U^* , the maximizer of the value function W .

3.1 Complete Information

Consider the benchmark case in which there is complete information; that is, consider $\mathcal{P}$ without the IC constraints. As the values are i.i.d., we can assume, without loss of generality,

¹³Not every policy can be represented in this fashion, as the principal need not treat two histories leading to the same promised utility identically. However, because such histories are equivalent from the agent's viewpoint, the principal's payoff must be maximized by some policy that can be so represented.

that p_h, p_l are constant over time. Given U , the principal chooses p_h and p_l to maximize

$$qp_h(h - c) + (1 - q)p_l(l - c),$$

subject to $U = qp_hh + (1 - q)p_ll$. This yields the following result.

Lemma 1 *Under complete information, the optimal constant policy is*

$$\begin{cases} p_h = \frac{U}{qh}, & p_l = 0 & \text{if } U \in [0, qh], \\ p_h = 1, & p_l = \frac{U - qh}{(1 - q)l} & \text{if } U \in [qh, \mu]. \end{cases}$$

The value function, denoted $\bar{W}$, is equal to

$$\bar{W}(U) = \begin{cases} (1 - \frac{c}{h}) U & \text{if } U \in [0, qh], \\ (1 - \frac{c}{l}) U + cq(\frac{h}{l} - 1) & \text{if } U \in [qh, \mu]. \end{cases}$$

Hence, the initial promise (the maximizer of $\bar{W}$) is qh .

That is, unless $U = qh$, the optimal policy (p_h, p_l) differs from the efficient choice $(1, 0)$. To deliver $U < qh$, the principal chooses to scale down the probability of supplying the good when the value is high, maintaining $p_l = 0$. Similarly, for $U > qh$, the principal is forced to supply the good with positive probability even when the value is low to satisfy promise keeping.

While this is the only optimal policy that is constant, there are many other (non-constant) optimal policies. We will encounter some in the sequel.

We call $\bar{W}$ the complete-information value function. It is piecewise linear (see Figure 1). Plainly, it is an upper bound on the value function under incomplete information.

3.2 The Optimal Mechanism

We now solve for the optimal policy under incomplete information in the i.i.d. case. We first provide an informal derivation of the solution. It follows from two observations (formalized in Theorem 1). First,

The efficient supply choice $(p_h, p_l) = (1, 0)$ is made “as long as possible.”

To understand this qualification, note that if $U = 0$ (resp., $U = \mu$), promise keeping allows no latitude: the good cannot (resp., must) be supplied, independent of the report. More generally, if $U \in [0, (1 - \delta)qh]$, promise keeping alone rules out $p_h = 1$ because doing

so would give the agent a utility of at least $(1 - \delta)qh > U$. Incentive compatibility widens this interval further; Theorem 1 gives the exact threshold. For such low promised utilities, “as long as possible” means as efficient as promise keeping and incentive compatibility allow. This implies that a high report leads to a continuation utility of 0, with today’s supply probability adjusting accordingly. An analogous interpretation applies to promised utilities close enough to μ .

These two peripheral intervals vanish as $\delta \rightarrow 1$ and are ignored for the remainder of this discussion. For every other promised utility, we claim that it is optimal to make the statically efficient supply choice. Intuitively, there is never a better time to redeem part of the promised utility than when the value is high. During such rounds, the interests of the principal and the agent are aligned. Conversely, there cannot be a worse opportunity to repay the agent what he is due than when the value is low because tomorrow’s value cannot be lower than today’s.

This simple observation already implies that the dynamics of the inefficiencies are the opposite of those in Battaglini’s model with transfers. Here, inefficiencies are backloaded.

As the supply decision is efficient as long as possible, the high-type agent has no incentive to pretend to be a low type. However,

The incentive constraint of the low-type agent always binds.

The reason that IC_L binds is standard: the agent is risk neutral, so the principal’s payoff must be a concave function of U (otherwise, he could offer a lottery that the agent would accept and that would make the principal better off). Concavity implies that there is no gain in spreading continuation utilities u_h, u_l beyond what is required for IC_L to be satisfied.

Because we are left with two variables (u_h, u_l) and two constraints (IC_L and PK), we can immediately solve for the optimal policy. The agent is always willing to state that his value is high, so his utility can be computed *as if* he followed this reporting strategy, namely,

$$U = (1 - \delta)\mu + \delta u_h, \text{ or } u_h = \frac{U - (1 - \delta)\mu}{\delta}.$$

Because U is a weighted average of u_h and μ , and $\mu \geq U$, it follows that $u_h \leq U$. The promised utility necessarily decreases after a high report. To compute u_l , note that the high-type agent is unwilling to report a low value: obtaining the good gives him $(1 - \delta)(h - l)$ more than what would make him indifferent between the two reports. Hence, defining $\underline{U} := q(h - l)$, it holds that

$$U = (1 - \delta)\underline{U} + \delta u_l, \text{ or } u_l = \frac{U - (1 - \delta)\underline{U}}{\delta}.$$

Because U is a weighted average of $\underline{U}$ and u_l , it follows that $u_l \leq U$ if and only if $U \leq \underline{U}$. In that case, even a low report leads to a decrease in the continuation utility, albeit a smaller decrease than if the report had been high and the good provided.

The following theorem (proved in the appendix) summarizes this discussion with the necessary adjustments on the peripheral intervals.

Theorem 1 *An optimal policy, unique for $U \geq \underline{U}$ and for $U = 0$, is*

$$p_h = \min \left\{ 1, \frac{U}{(1-\delta)\mu} \right\}, \quad p_l = \max \left\{ 0, 1 - \frac{\mu - U}{(1-\delta)l} \right\}.$$

Given these values of (p_h, p_l) , continuation utilities are

$$u_h = \frac{U - (1-\delta)p_h\mu}{\delta}, \quad u_l = \frac{U - (1-\delta)(p_l l + p_h \underline{U})}{\delta}.$$

For reasons that will become clear shortly, this policy is not uniquely optimal for $0 < U < \underline{U}$. We now turn to a discussion of the utility dynamics and of the shape of the value function, which are closely related. This discussion revolves around the following lemma, which requires some notation. Let $(B_j)_{j \geq 0}$ be independent Bernoulli random variables with parameter q , and define

$$X := (1-\delta) \sum_{j=0}^{\infty} \delta^j (\underline{U} + l B_j),$$

so that $X \in [\underline{U}, \mu]$ and $\mathbf{E}[X] = qh$. We write $x_+ := \max\{x, 0\}$.

Lemma 2 *For all $U \in [0, \mu]$,*

$$W(U) = \left(1 - \frac{c}{h}\right) U - \frac{c(h-l)}{hl} \mathbf{E}[(U - X)_+]. \quad (1)$$

In particular, W is continuous and concave on $[0, \mu]$, continuously differentiable on $(0, \mu)$, linear and equal to $\bar{W}$ on $[0, \underline{U}]$, and strictly concave on $[\underline{U}, \mu]$. Furthermore,

$$\lim_{U \downarrow 0} W'(U) = 1 - \frac{c}{h}, \quad \lim_{U \uparrow \mu} W'(U) = 1 - \frac{c}{l}.$$

For intuition, consider the following functional equation for W that we obtain from Theorem 1 (ignoring again the peripheral intervals for the sake of the discussion):

$$W(U) = (1-\delta)q(h-c) + \delta q W\left(\frac{U - (1-\delta)\mu}{\delta}\right) + \delta(1-q)W\left(\frac{U - (1-\delta)\underline{U}}{\delta}\right).$$

Hence, taking for granted the differentiability of W stated in the lemma,

$$W'(U) = qW'(u_h) + (1 - q)W'(u_l).$$

Fix $U_0 \in [0, \mu]$ and let $U_{n+1} = u_h(U_n)$ if $v_n = h$ and $U_{n+1} = u_l(U_n)$ if $v_n = l$. In probabilistic terms, $W'(U_n) = \mathbf{E}[W'(U_{n+1})]$ given the information at round n . That is, $W'(U_n)$ is a bounded martingale and must therefore converge.¹⁴ This martingale was first uncovered by Thomas and Worrall (1990), and hence, we refer to it as the TW-martingale. Strict concavity on $(\underline{U}, \mu)$ and the separation of u_h and u_l rule out an interior limit. The proof of Lemma 3 gives a direct argument for absorption at the endpoints, including the peripheral intervals. Moreover, because $u_h < U$ and $u_l \leq U$ on the interval $(0, \underline{U}]$, this interval is a transient region for the process: once there, U_n decreases and converges to 0. Together, these observations imply that U_n converges, and its limit U_∞ is either 0 or μ . In particular, if the process begins in the interval $[0, \underline{U}]$, the limit must be 0, and the TW-martingale implies that W' must be constant on this interval—hence the linearity of W .¹⁵

While $W'(U_n)$ is a martingale, U_n is not. In the region where $(p_h, p_l) = (1, 0)$, the optimal policy yields

$$U_n = (1 - \delta)qh + \delta\mathbf{E}[U_{n+1}],$$

utility drifts up or down (stochastically) according to whether U_n is above or below qh . Intuitively, delivering the promise U_n amounts to providing a constant flow of U_n in every round. If the flow delivered, qh , is below U_n , the expected continuation utility must exceed U_n to make up the shortfall.

This raises the question of the initial promise U^* : does it lie above or below qh , and with what probabilities does the process $\{U_n\}_{n=0}^\infty$ with $U_0 = U^*$ converge to 0 and to μ ? The answer is provided by the TW-martingale. Indeed, U^* is characterized by $W'(U^*) = 0$ (uniquely so, given strict concavity on $[\underline{U}, \mu]$). Hence,

$$0 = W'(U^*) = \mathbf{P}[U_\infty = 0 \mid U_0 = U^*]W'(0) + \mathbf{P}[U_\infty = \mu \mid U_0 = U^*]W'(\mu),$$

where $W'(0)$ and $W'(\mu)$ are the one-sided derivatives given in Lemma 2. Hence,

$$\frac{\mathbf{P}[U_\infty = 0 \mid U_0 = U^*]}{\mathbf{P}[U_\infty = \mu \mid U_0 = U^*]} = \frac{(c - l)/l}{(h - c)/h}. \quad (2)$$

¹⁴It is bounded because W is concave, and hence, W' is bounded by its limits at 0 and μ , given in the lemma.

¹⁵This yields multiple optimal policies on this interval: IC_L need not hold with equality, and, with $p_l = 0$, any pair of continuation utilities in $[0, \underline{U}]$ satisfying IC_L , IC_H , and PK is optimal.

The initial promise is chosen to yield this ratio of absorption probabilities at 0 and μ . This ratio is independent of the discount factor, even though the random walk is discrete and its step size depends on δ . Hence, both long-run outcomes are possible irrespective of how patient the players are. Depending on the parameters, U^* can be above or below qh , the complete-information initial promise, as is easy to confirm through examples. Lemma 10 shows that U^* is decreasing in the cost: the random walk $\{U_n\}$ depends on c only via the choice of the initial promise, and, by (2), a larger cost makes absorption at 0 more likely, which requires a lower initial promise.

We record this discussion in the following lemma.

Lemma 3 *The process $\{U_n\}_{n=0}^\infty$ with $U_0 = U^*$ converges to 0 or μ , a.s., and the ratio of the absorption probabilities is given by (2).*




Figure 1: Value functions W and $\bar{W}$ for $(\delta, l, h, q, c) = (.85, .40, .60, .60, .50)$.

Our result that both extreme outcomes occur with positive probability contrasts with Thomas and Worrall’s (1990) celebrated “immiseration” result. Both analyses rest on the TW-martingale; the difference lies in the range of W' . In Thomas and Worrall (1990), utility is bounded above, but approaching this bound requires arbitrarily large consumption. Hence, $W'(U)$ diverges to $-\infty$ as U approaches this bound: delivering utility becomes prohibitively costly near the top. In their model, $W'(U_n)$ is a nonpositive martingale, so it converges to a finite limit; this rules out the upper bound, where W' is infinite. Since incentives require continuation utilities to keep being spread, U_n cannot settle at an interior point either, so it must diverge to $-\infty$, where W' tends to 0. Here, the positive lower bound l on values keeps the marginal provision cost bounded by $1/l$. Consequently, $W'(U) \rightarrow 1 - c/l$ as $U \rightarrow \mu$ (Lemma 2), both 0 and μ are possible limits, and the probability of each depends on the initial promise, as captured by (2). Retirement is a real possibility.

3.3 Implementation

As mentioned in the introduction, the optimal policy is not a token mechanism because the total (discounted) number of units the agent receives is not fixed.¹⁶ However, the policy admits a simple indirect implementation in terms of a budget that can be described as follows. Let $f := (1 - \delta)\underline{U}$, and $g := (1 - \delta)\mu - f = (1 - \delta)l$.

Provide the agent with an initial budget of U^* . At the beginning of each round, charge him a fixed fee equal to f . If the agent requests the unit, supply it and charge a price g for it. Increase his budget at the interest rate $(\frac{1}{\delta} - 1)$ each round—provided that all of this is feasible. The mechanism thus operates as a two-part tariff, denominated in utils: a fee paid every round, and a price paid per unit received.

This scheme might become infeasible for two reasons. First, his budget U might no longer cover both the fee and the price of a unit. Then, scale the round down—fee, price, and unit alike—by the fraction $U/(f + g)$, so that a request exhausts his budget. Second, his budget might be so close to μ that it is no longer possible to pay him the interest on his budget. Then, return the excess to him by awarding him the unit even though he did not request it, at a conversion rate that is also given by the price g .¹⁷

For budgets below $\underline{U}$, the agent is “in the red,” and even if he does not request a unit, his budget shrinks over time. If his budget is above $\underline{U}$, he is “in the black,” and forgoing a unit increases the budget. When doing so pushes the budget above $\mu - (1 - \delta)l$, the agent “breaks the bank” and reaches μ —an absorbing state—if he forgoes once more.

This structure is reminiscent of results in research on optimal financial contracting (see, for instance, Biais et al., 2007), a literature that assumes transfers.¹⁸ Under their assumption that the agent is more impatient than the principal, the upper dividend boundary is reflecting: payments keep the agent’s continuation utility from exceeding it. The lower liquidation boundary is absorbing. There are several important differences, however. Most importantly, the agent is not paid in the intermediate region: promises are the only source of incentives. In our environment, the agent receives the good if his value is high, achieving efficiency in this intermediate region.

¹⁶To be clear, this is not an artifact of discounting: the optimal policy in the finite-horizon undiscounted version of our model can be derived along the same lines (using the binding IC_L and PK constraints), and the number of units obtained by the agent is also history-dependent in that case.

¹⁷Formally, he receives the fraction x of the unit solving $U - f - xg = \delta\mu$: his balance after the fee and this award is exactly $\delta\mu$, and the interest payment brings it to μ .

¹⁸The setups also differ in other respects. They allow two instruments: downsizing the firm and payments. Additionally, theirs is a moral hazard problem because the agent can divert cash flows from the project.

3.4 Token Mechanisms

A discussion of the relationship of our results to those in environments with transfers is relegated to Section 4.5 because the environment considered in Section 4 is the counterpart to Battaglini (2005). However, because token mechanisms are typically introduced in i.i.d. environments, we make a few observations here concerning the connection between our results and those of Jackson and Sonnenschein (2007) to explain why our dynamic analysis is substantially different from the static one with many copies.

There are two distinct issues. The first concerns the agent’s knowledge, and two settings must be distinguished. In Jackson and Sonnenschein (2007), the agent is a *prophet* (in the language of stochastic processes): he knows the entire realization of the value process from the outset. In our environment, he is a *forecaster*: he learns each value at the beginning of the corresponding round, so that the process of his reports must be adapted to the realized values. The second issue is whether token mechanisms are optimal. For the purpose of asymptotic analysis (when the discount factor tends to 1 or the number of equally weighted copies T tends to infinity), token mechanisms are asymptotically optimal, whether the agent is a prophet or a forecaster. Because the emphasis in Jackson and Sonnenschein is on asymptotic analysis, they focus on the prophetic case and on a token mechanism, deriving a rate of convergence for this mechanism (the loss relative to the first best is of order $1/\sqrt{T}$); they also extend their results to the case in which decisions are implemented over time.

A token budget K fixes the total discounted quantity: a unit in round n uses $(1 - \delta)\delta^n$ of the budget, with fractional provision when the remaining budget is insufficient. Not only are these mechanisms asymptotically optimal whether the agent is a prophet or a forecaster, but, with binary values, they perform identically in the two settings: foreknowledge of the realized sequence is of no use to an agent endowed with a fixed total (discounted) number of units. Forgoing a low-value unit is costless as long as his endowment does not cover all remaining rounds, as no subsequent unit can be worth less; and claiming a high-value unit is never a mistake, as no subsequent unit can be worth more.

Away from this limit ($\delta \rightarrow 1$ or $T \rightarrow \infty$), however, token mechanisms are optimal in neither setting. In our forecaster environment, this follows from Theorem 1: a report affects not only whether the agent obtains the current unit but also the total (discounted) number of units he obtains.¹⁹ When the agent is a prophet, this can be seen in the finite-horizon undiscounted case: unlike a token mechanism, which lets the agent place a fixed number of tokens wherever he pleases, the optimal mechanism offers a menu trading off the expected

¹⁹To be clear, token mechanisms are not optimal even without discounting.

number of units against how reliably they land on the copies he designates.²⁰ Indeed, the agent’s private information pertains not only to whether a given unit has a high value but also to how many units are high, and token mechanisms elicit no information in this regard.

As the prophet has more information than the forecaster, the optimal mechanisms in the two settings differ. How these two mechanisms compare in terms of efficiency loss is ambiguous *a priori*. Because the prophet has more information regarding the number of high-value units, the mechanism must satisfy more incentive-compatibility constraints (which harms efficiency) but might induce a better fit between the number of units he actually receives and the number he should receive. It is not difficult to construct examples (say, for $T = 3$) in which the comparison could go either way depending on the parameters. Asymptotically, however, the difference disappears at the level of rates, as the next lemma states.

Lemma 4 *Whether the agent is a forecaster or a prophet, the loss relative to $q(h - c)$ converges to zero at rate $\sqrt{1 - \delta}$. In the forecaster case,*

$$|W(U^*) - q(h - c)| = \Theta(\sqrt{1 - \delta}).$$

In both settings, the rate is the same as with token mechanisms. Token mechanisms do not attempt to elicit the number of high-value units; by the central limit theorem, this number fluctuates around its mean by an amount of order $\sqrt{1 - \delta}$, so a token mechanism misallocates a fraction of units of that order—hence the rate $\Theta(\sqrt{1 - \delta})$. In the case of a prophet, the cost of incentive compatibility is such that the optimal mechanism eliminates a fraction of this inefficiency, but not its order of magnitude.²¹

For a forecaster, the rate can be seen from (1). In the language of Section 3.3, X is the present value of the fees and prices charged to an agent who requests the unit exactly when his value is high; it is random only because the number of high-value rounds is. The agent breaks the bank if his initial budget exceeds X and runs out of budget otherwise, so the loss is driven by the mismatch between the initial promise and X , which is of the order of the standard deviation of X , that is, of $\sqrt{1 - \delta}$. In sum, the budget mechanism improves on token mechanisms in the size of the loss, not in the rate at which it vanishes.²²

²⁰The characterization of the optimal mechanism for a prophetic agent is tangential to our analysis and is thus omitted.

²¹See Eilat and Pauzner (2011) for an optimal mechanism in a related setting.

²²For the parameters of Figure 1, the efficiency loss is 0.0092 under the optimal mechanism and 0.0113, or 23% more, under the best token mechanism.

4 The General Markov Model

We now return to the general model in which values are persistent rather than independent.

As an initial exercise, consider the case of perfect persistence, $\rho_h = \rho_l = 0$. If types never change, the principal cannot use future allocations as an instrument to screen the agent. We revert to the static problem, for which the solution is to either always supply the good (if $\mu \geq c$) or never do so.

This case suggests that persistence plays an ambiguous role *a priori*. Because current types differ in the probability they assign to being a high type tomorrow, they value any given promise of future allocations differently; one might therefore hope that persistence sharpens this instrument, making it easier to elicit truthful reports. However, the case of perfect persistence demonstrates that correlation diminishes the scope for using future allocations as “transfers”: although utilities remain separable over time, the value of future allocations now depends on the agent’s current type, so they are no longer the type-neutral currency they are under independence. Definite comparative statics are obtained in the continuous-time limit; see Section 5.1.

The martingale arguments that served us well with independent values are no longer useful. Worse, *ex ante* utility is no longer a valid state variable. To understand why, note that with independent values, the agent can evaluate each report based only on his current type, the supply probability, and the continuation utility, the latter two as functions of his report. With correlated values, he cannot: how much a promised continuation is worth to him depends on how the principal intends to deliver it. This is problematic because the agent can deviate, unbeknownst to the principal, in which case the continuation utility computed by the principal, given his incorrect belief regarding the agent’s type tomorrow, is not the same as the continuation utility under the agent’s belief.

The Markov property suggests the fix. Conditional on tomorrow’s type, today’s type carries no further information about the future. Hence, a pair of promised *ex interim* utilities for tomorrow, one for each type the agent may then report, suffices for him to compute his utility today regardless of whether he deviates: he weights the pair by his own transition probabilities. Of course, tomorrow’s type is not observable. We must therefore use the utility he receives from truthfully reporting his type tomorrow. This creates no difficulty, as truthful reporting remains optimal tomorrow even after a lie today (given today’s report and tomorrow’s type, today’s type is irrelevant to the agent’s continuation problem). This is the one-shot deviation principle: when the agent considers lying today, there is no loss in assuming that he reports truthfully from tomorrow on.

We are not the first to observe that, when types are serially correlated, the state variable must include the vector of *ex interim* utilities, one for each possible report, rather than the *ex ante* utility. See Fernandes and Phelan (2000), Cole and Kocherlakota (2001), Doepke and Townsend (2006), and Zhang and Zenios (2008).

To evaluate the principal's payoff, we must also specify his belief regarding the agent's current type. Let ϕ denote the probability that the principal assigns to the high type. This belief can take only three values on path: q in the initial round, $1 - \rho_h$ after a high report, and ρ_l after a low one. Nonetheless, we treat ϕ as an arbitrary element of the unit interval.

4.1 The Program

The characterizations in this section assume $0 < \kappa < 1$. Independence is treated in Section 3; perfect persistence is the benchmark described above. We continue to use both cases when discussing the limiting geometry of the feasible set. As discussed above, the principal's optimization program, cast as a dynamic programming problem, requires three state variables: his belief, $\phi = \mathbf{P}[v = h] \in [0, 1]$, and the pair of promised *ex interim* utilities that he delivers as a function of the current report, $U = (U(h), U(l))$. The highest utility μ_h (resp., μ_l) that can be given to an agent whose type is high (resp., low), which is obtained by always supplying the good, solves

$$\mu_h = (1 - \delta)h + \delta(1 - \rho_h)\mu_h + \delta\rho_h\mu_l, \quad \mu_l = (1 - \delta)l + \delta(1 - \rho_l)\mu_l + \delta\rho_l\mu_h;$$

that is,

$$\mu_h = h - \frac{\delta\rho_h(h - l)}{1 - \delta + \delta(\rho_h + \rho_l)}, \quad \mu_l = l + \frac{\delta\rho_l(h - l)}{1 - \delta + \delta(\rho_h + \rho_l)}.$$

We note that

$$\mu_h - \mu_l = \frac{1 - \delta}{1 - \delta + \delta(\rho_h + \rho_l)}(h - l).$$

This gap decreases in δ , vanishing as $\delta \rightarrow 1$ (provided $\rho_h + \rho_l > 0$).

A policy is now a tuple of functions $p_h, p_l : [0, \mu_h] \times [0, \mu_l] \times [0, 1] \rightarrow [0, 1]$, $u_h, u_l : [0, \mu_h] \times [0, \mu_l] \times [0, 1] \rightarrow [0, \mu_h] \times [0, \mu_l]$ mapping the state (U, ϕ) onto the probabilities p_h, p_l of supplying the good during this round and the pairs u_h, u_l of continuation *ex interim* utilities at the beginning of the next round, indexed by the agent's current report. (When no confusion may arise, we omit their dependence on (U, ϕ) .) After a report of h , the pair $u_h = (u_h(h), u_h(l))$ specifies the continuation utility for each report the agent may make tomorrow; similarly for u_l . These functions must be consistent in the sense that, given $U = (U(h), U(l))$, the probability of supplying the good and the continuation utilities that

follow each report yield a truthful agent exactly the utility promised for that report. This is again “promise keeping.” We stress that $U(h)$ and $U(l)$ are *ex interim* utilities in a given round; that is, they are computed after the agent’s value is realized.

Define the function $W : [0, \mu_h] \times [0, \mu_l] \times [0, 1] \rightarrow \mathbf{R} \cup \{-\infty\}$ that solves the following program for all $U \in [0, \mu_h] \times [0, \mu_l]$ and $\phi \in [0, 1]$:

$$W(U, \phi) = \sup \left\{ \phi ((1 - \delta)p_h(h - c) + \delta W(u_h, 1 - \rho_h)) \right. \\ \left. + (1 - \phi) ((1 - \delta)p_l(l - c) + \delta W(u_l, \rho_l)) \right\},$$

over $p_h, p_l \in [0, 1]$ and $u_h, u_l \in [0, \mu_h] \times [0, \mu_l]$, subject to promise keeping and incentive compatibility, namely,

$$U(h) = (1 - \delta)p_h h + \delta(1 - \rho_h)u_h(h) + \delta\rho_h u_h(l) \quad (\text{PK}_H)$$

$$\geq (1 - \delta)p_l h + \delta(1 - \rho_h)u_l(h) + \delta\rho_h u_l(l) \quad (\text{IC}_H)$$

and

$$U(l) = (1 - \delta)p_l l + \delta(1 - \rho_l)u_l(l) + \delta\rho_l u_l(h) \quad (\text{PK}_L)$$

$$\geq (1 - \delta)p_h l + \delta(1 - \rho_l)u_h(l) + \delta\rho_l u_h(h) \quad (\text{IC}_L)$$

with the convention that the supremum is $-\infty$ whenever the feasible set is empty. Note that, for each ϕ , $W(\cdot, \phi)$ is concave (by the linearity of the constraints in the promised utilities). An optimal policy is a map from (U, ϕ) into (p_h, p_l, u_h, u_l) that achieves the supremum in the program above.

4.2 Complete Information

Proceeding as with independent values, we briefly derive the solution under complete information, that is, dropping (IC_H) and (IC_L) . Write $\bar{W}$ for the resulting value function. Absent promise-keeping constraints, the efficient policy is to supply the good if and only if the type is high. Let v_h^* (resp., v_l^*) denote the utility that a high (resp., low) type obtains under this policy. The pair (v_h^*, v_l^*) satisfies

$$v_h^* = (1 - \delta)h + \delta(1 - \rho_h)v_h^* + \delta\rho_h v_l^*, \quad v_l^* = \delta(1 - \rho_l)v_l^* + \delta\rho_l v_h^*,$$

which yields

$$v_h^* = \frac{h(1 - \delta(1 - \rho_l))}{1 - \delta(1 - \rho_h - \rho_l)}, \quad v_l^* = \frac{\delta h \rho_l}{1 - \delta(1 - \rho_h - \rho_l)}.$$

When a high type's promised utility $U(h)$ is in $[0, v_h^*]$, the principal supplies the good only if the type is high. Therefore, the payoff is $U(h)(1 - c/h)$. When $U(h) \in (v_h^*, \mu_h]$, the principal always supplies the good if the type is high. To deliver the promised utility, the principal also supplies the good when the agent's type is low. The payoff is $v_h^*(1 - c/h) + (U(h) - v_h^*)(1 - c/l)$. We proceed analogously given $U(l)$ (notice that the problems of delivering $U(h)$ and $U(l)$ are uncoupled). In summary, $\bar{W}(U, \phi)$ is given by

$$\begin{cases} \phi \frac{U(h)(h-c)}{h} + (1 - \phi) \frac{U(l)(h-c)}{h} & \text{if } U \in [0, v_h^*] \times [0, v_l^*], \\ \phi \frac{U(h)(h-c)}{h} + (1 - \phi) \left(\frac{v_l^*(h-c)}{h} + \frac{(U(l) - v_l^*)(l-c)}{l} \right) & \text{if } U \in [0, v_h^*] \times [v_l^*, \mu_l], \\ \phi \left(\frac{v_h^*(h-c)}{h} + \frac{(U(h) - v_h^*)(l-c)}{l} \right) + (1 - \phi) \frac{U(l)(h-c)}{h} & \text{if } U \in [v_h^*, \mu_h] \times [0, v_l^*], \\ \phi \left(\frac{v_h^*(h-c)}{h} + \frac{(U(h) - v_h^*)(l-c)}{l} \right) + (1 - \phi) \left(\frac{v_l^*(h-c)}{h} + \frac{(U(l) - v_l^*)(l-c)}{l} \right) & \text{if } U \in [v_h^*, \mu_h] \times [v_l^*, \mu_l]. \end{cases}$$

For each ϕ , $\bar{W}(\cdot, \phi)$ is piecewise linear (a “tilted pyramid”) with a maximum at $v^* = (v_h^*, v_l^*)$. The partial derivatives of $\bar{W}$ with respect to $U(h)$ and $U(l)$ (defined except where $U(h) = v_h^*$ and $U(l) = v_l^*$, respectively) lie in the interval $[1 - c/l, 1 - c/h]$. The upper bound corresponds to the most efficient allocation, whereas the lower bound corresponds to the least efficient allocation.

4.3 Feasible and Incentive-Feasible Utilities

One difficulty in using *ex interim* utilities as state variables rather than *ex ante* utility is that the dimensionality of the problem increases with the cardinality of the type set. Another related difficulty is that it is not obvious which vectors of utilities are feasible given the incentive constraints. Promising to allocate all future units to the agent in the event that his current report is high while allocating none if this report is low is not incentive compatible.

The set of *feasible* utility pairs (that is, the largest bounded set of pairs $U = (U(h), U(l))$ such that (PK_H) and (PK_L) can be satisfied with continuation pairs in the set itself) is easy to describe. Because the two promise-keeping equations are uncoupled, it is the set $[0, \mu_h] \times [0, \mu_l]$ itself (as was already implicit in Section 4.2).

What is challenging is to solve for the incentive-compatible, feasible (in short, incentive-feasible) utility pairs: these are *ex interim* utility pairs for which one can find supply probabilities and continuation pairs that make truthful reporting optimal, with the continuation pairs satisfying the same property in turn.

Definition 1 *The incentive-feasible set, $V \subseteq \mathbf{R}^2$, is the set of ex interim utility pairs in round 0 that are obtained under some incentive-compatible policy.*

It is standard to show that V is the largest bounded set such that for each $U \in V$ there exist $p_h, p_l \in [0, 1]$ and two pairs $u_h, u_l \in V$ solving (PK_H) – (IC_L) .²³

Our first step toward the optimal mechanism is to solve for V . To obtain intuition regarding its structure, let us review some of its elements. Both $(0, 0)$ and (μ_h, μ_l) belong to V . It suffices to never or always supply the unit, independently of the reports, which is incentive compatible. More generally, for any integer $m \geq 0$, the principal can supply the unit for the first m rounds, independently of the reports, and never supply the unit afterward. We refer to such policies as pure *frontloaded* policies because they deliver a given number of units as quickly as possible. A mixed frontloaded policy randomizes over two pure frontloaded policies with consecutive integers $m, m + 1$. Similarly, we define a pure *backloaded* policy as one that does not supply the unit for the first m rounds but does afterward, independently of the reports. (Mixed backloaded policies are defined analogously.)

Fix $U(h) \in (0, \mu_h)$ and consider the backloaded and frontloaded policies that each deliver $U(h)$ to the high type. The low type prefers the backloaded policy, which affords him “more time” to switch from his initial low value to a high value.

The utility pairs delivered by backloading and frontloading obey simple recursions, whose solutions are as follows. First, for every integer $m \geq 0$, let

$$\begin{aligned}\bar{u}_h^m &= \delta^m \mu_h - \delta^m (1 - q)(\mu_h - \mu_l)(1 - \kappa^m), \\ \bar{u}_l^m &= \delta^m \mu_l + \delta^m q(\mu_h - \mu_l)(1 - \kappa^m),\end{aligned}$$

and set $\bar{u}^m := (\bar{u}_h^m, \bar{u}_l^m)$, the pair under the backloaded policy with switching time m . Second, for every integer $m \geq 0$, let

$$\begin{aligned}\underline{u}_h^m &= (1 - \delta^m)\mu_h + \delta^m (1 - q)(\mu_h - \mu_l)(1 - \kappa^m), \\ \underline{u}_l^m &= (1 - \delta^m)\mu_l - \delta^m q(\mu_h - \mu_l)(1 - \kappa^m),\end{aligned}$$

and set $\underline{u}^m := (\underline{u}_h^m, \underline{u}_l^m)$, the pair under the frontloaded policy with switching time m . The sequence $\bar{u}^m$ is decreasing (in both its components) as m increases, with $\bar{u}^0 = (\mu_h, \mu_l)$ and $\lim_{m \rightarrow \infty} \bar{u}^m = 0$. Similarly, $\underline{u}^m$ is increasing, with $\underline{u}^0 = 0$ and $\lim_{m \rightarrow \infty} \underline{u}^m = (\mu_h, \mu_l)$.

Not only is backloading better than frontloading for the low-type agent, holding the high-type agent’s utility fixed: the two are extremal, yielding the low type’s best and worst utilities among all incentive-compatible policies. Formally,

²³Incentive-feasibility is closely related to self-generation (see Abreu, Pearce and Stacchetti, 1990), though it pertains to the different types of a single agent rather than to different players in a repeated game. The distinction is not merely a matter of interpretation because a high type can become a low type and *vice versa*, which represents a situation with no analogue in repeated games. Nonetheless, the proof of this characterization is identical.

Lemma 5 *It holds that*

$$V = \text{co}\{\bar{u}^m, \underline{u}^m : m = 0, 1, 2, \dots\}.$$

That is, V is a polygon with a countable infinity of vertices (and two accumulation points): the $\bar{u}^m$ trace its upper boundary, and the $\underline{u}^m$ its lower boundary. See Figure 2 for an illustration.




Figure 2: The set V for $(\delta, \rho_h, \rho_l, l, h) = (9/10, 1/3, 1/4, 1/4, 1)$.

A direct computation shows that, for $\kappa < 1$,

$$\lim_{m \rightarrow \infty} \frac{\bar{u}_l^{m+1} - \bar{u}_l^m}{\bar{u}_h^{m+1} - \bar{u}_h^m} = \lim_{m \rightarrow \infty} \frac{\underline{u}_l^{m+1} - \underline{u}_l^m}{\underline{u}_h^{m+1} - \underline{u}_h^m} = 1.$$

When the switching time m is large, the impact on the agent's utility of increasing it further is essentially independent of his initial type. The slope of either boundary segment corresponding to round m is

$$\frac{(1 - \delta)\mu - q(\mu_h - \mu_l)\kappa^m(1 - \delta\kappa)}{(1 - \delta)\mu + (1 - q)(\mu_h - \mu_l)\kappa^m(1 - \delta\kappa)}.$$

Thus the slopes are at most 1 and approach 1 as $m \rightarrow \infty$.

Because $(\mu_l - v_l^*)/(\mu_h - v_h^*) > 1$, the vector v^* is outside V . This should not be surprising: with private information, the low-type agent derives information rents, so that if the high-type agent's utility were first-best, the low-type agent's utility would be too high. For any belief ϕ , the complete-information value function $\bar{W}(U, \phi)$ increases as U moves within V toward v^* .

It is instructive to study how the shape of V varies with persistence. When $\kappa = 0$, so that values are independent over time, the low-type agent prefers to receive the good tomorrow rather than today (adjusting for discounting), but he is indifferent among all later dates: from tomorrow on, his expected value is μ whatever the date. The case is analogous for the high-type agent. As a result, all the vertices $\{\bar{u}^m\}_{m=1}^\infty$ (resp., $\{\underline{u}^m\}_{m=1}^\infty$) are aligned and V is a parallelogram with vertices 0 , (μ_h, μ_l) , $\bar{u}^1$, and $\underline{u}^1$. As κ increases, the imbalance between type utilities increases, tilting V toward the $U(h)$ -axis; in the limit of perfect persistence, V collapses onto the segment joining 0 and (μ_h, μ_l) . The low-type agent is then indifferent between frontloading and backloading because no amount of time allows his type to change. See Figure 3.




Figure 3: Impact of persistence, as measured by $\kappa \geq 0$, holding $\rho_h/\rho_l = 4/3$ fixed, for $(\delta, l, h) = (9/10, 1/4, 1)$. The three sets have $\kappa = 0$, $5/12$, and $593/600$ (strong persistence).

The structure of V relies on the assumption $\kappa \geq 0$. If values were negatively correlated over time, then frontloading and backloading would not be policies spanning the boundary of V . To see this, consider perfect negative serial correlation. Then, supplying the unit if and only if the round is odd (resp., even) favors (resp., hurts) the low-type agent relative to the high-type agent. These two policies, together with never and always supplying the unit, achieve the extreme points of V . A negative correlation thus requires a separate (but analogous) treatment, motivating our focus on $\kappa \geq 0$.

Importantly, frontloading and backloading are not the only ways to achieve boundary utilities. The lower boundary corresponds to policies that (starting from this boundary) assign as high a probability as possible to the good being supplied for every high report while promising continuation utilities that make (IC_L) always bind. Similarly, the upper boundary corresponds to those policies that assign as low a probability as possible to the

good being supplied for low reports while promising continuation utilities that make (IC_H) always bind. Frontloading and backloading are examples of the two classes, respectively.

4.4 The Optimal Mechanism

Not every incentive-feasible utility pair is on path given the optimal initial promise and the resulting dynamics. Irrespective of the sequence of reports, some pairs simply never arise. While it is necessary to solve for the value function and optimal policy on the entire domain V , we first focus on the subset of V that is on path. We relegate discussion of the optimal policy for other utility pairs to Section 4.6.




Figure 4: Dynamics of utility on the lower boundary for $(\delta, \rho_h, \rho_l, l, h) = (9/10, 1/3, 1/4, 1/4, 1)$.

This subset is the lower boundary—the polygonal chain spanned by pure frontloading. Furthermore, two observations from the i.i.d. case remain valid. First, the statically efficient supply choice is made as long as feasibility allows, and second, the continuation promises are chosen so that the agent is indifferent between the two reports when his type is low.

To understand why such a policy yields utilities on the “frontloading” boundary (as mentioned at the end of Section 4.3), note that, because the agent is indifferent between the two reports whenever his type is low, he is willing to always report high, irrespective of his type. Because the principal then supplies the good, from the agent’s perspective, the pair of utilities can be computed as if frontloading were the policy being implemented.

From the principal's perspective, however, it matters that this is not the actual policy. As in the i.i.d. benchmark, the payoff is higher under the optimal policy than under frontloading: as long as feasibility allows, the delivery of the promised utility is delayed so that units are supplied only when the value is high, each yielding $h - c > 0$ rather than $l - c < 0$.

Hence, after a high report, as in the i.i.d. case, continuation utility declines.²⁴ Specifically, u_h is computed, as under frontloading, as the solution to the following system:

$$U(v) = (1 - \delta)v + \delta \mathbf{E}_v[u_h], \quad v = l, h,$$

where $(U(h), U(l))$ is given. Here, $\mathbf{E}_v[u_h]$ is the expectation of the utility pair u_h given that the current type is v (e.g., for $v = h$, $\mathbf{E}_h[u_h] = (1 - \rho_h)u_h(h) + \rho_h u_h(l)$).

The promised u_l does not admit such an explicit formula because it is specified by (PK_L) and the requirement that it lies on the lower boundary. In fact, u_l might be lower or higher than U (see Figure 4). Under a certain condition, u_l is lower than U when U is low enough.²⁵ The condition has a simple geometric interpretation: if the half-open line segment $(0, v^*]$ intersects the lower boundary,²⁶ then u_l is lower than U if and only if U lies below the point of intersection, which we denote by $\underline{U}$.²⁷ If instead there is no such intersection, then u_l is higher than U for all U , and we set $\underline{U} = 0$. The intersection exists if and only if

$$\frac{h - l}{l} > \frac{1 - \delta}{\delta \rho_l}, \quad (3)$$

that is, if and only if the low-type persistence is sufficiently low.²⁸ In this case, utility declines even after a low report when U lies below $\underline{U}$: the low-type agent expects a high value sufficiently soon and often that the efficient policy would yield too high a utility. When the low-type persistence is high, this does not occur. For utility pairs on the lower boundary, the principal achieves the complete-information payoff when $U \leq \underline{U}$ or $U = (\mu_h, \mu_l)$. These are the only such pairs if $v_h^* \leq \bar{u}_h^1$, as established in Lemma 6.

We summarize this discussion with the following theorem, a special case of Theorem 3 below.

²⁴Because the lower boundary is upward sloping, the *ex interim* utilities of both types vary in the same way. Accordingly, we use terms such as “higher” or “lower” utility and write $U < U'$ for the componentwise order.

²⁵As in the i.i.d. case, u_l is always higher than u_h .

²⁶This line has the equation $U(l) = \frac{\delta \rho_l}{1 - \delta(1 - \rho_l)} U(h)$.

²⁷With some abuse of notation: this point is the natural extension of $\underline{U} \in \mathbf{R}$ as introduced in Section 3.

²⁸This condition is satisfied in the i.i.d. case due to our assumption that $\delta > l/\mu$.

Theorem 2 *For any $U = (U(h), U(l))$ on the lower boundary of V , an optimal policy specifies the constrained-efficient supply probabilities*

$$p_h = \min \left\{ 1, \frac{U(h)}{(1-\delta)h} \right\}, \quad p_l = \max \left\{ 0, 1 - \frac{\mu_l - U(l)}{(1-\delta)l} \right\},$$

and continuation pairs u_h, u_l on the lower boundary of V such that (IC_L) always binds. For $\phi \in [\rho_l, 1 - \rho_h]$, an optimal initial promise $U^* > \underline{U}$ can also be chosen on the lower boundary.

While the implementation in the i.i.d. case is described in terms of a “utility budget,” inspired by the use of *ex ante* utility as a state variable, the analysis of the Markov case reveals the appropriate unit of account—the number of units that the agent is entitled to claim in a row with “no questions asked.” Because the utility pairs on the lower boundary correspond to frontloading for a possibly non-integer number of rounds, we denote such a policy by a pair $(m, \lambda) \in (\mathbf{N}_0 \cup \{\infty\}) \times [0, 1)$, with the interpretation that the good is supplied during the first m rounds and, with probability λ , during the following round as well, and write $\underline{u}^{m, \lambda} = (\underline{u}_h^{m, \lambda}, \underline{u}_l^{m, \lambda}) := (1 - \lambda) \underline{u}^m + \lambda \underline{u}^{m+1}$ for the corresponding utility pair. If $m = \infty$, the good is always supplied, yielding utility $\underline{u}^\infty := (\mu_h, \mu_l)$.

We may think of the optimal policy as follows. During a given round, the agent holds the promise (m, λ) . If the agent requests the unit, he receives it whenever $m \geq 1$, and the promise becomes $(m - 1, \lambda)$; when $m = 0$, he receives it with probability λ , and the promise becomes $(0, 0)$. If instead the agent forgoes the unit, the promise is revised to the pair (m', λ') that solves

$$\frac{\underline{u}_l^{m, \lambda}}{\delta} = \mathbf{E}_l [\underline{u}^{m', \lambda'}] = (1 - \rho_l) \underline{u}_l^{m', \lambda'} + \rho_l \underline{u}_h^{m', \lambda'},$$

provided that such a pair with finite m' exists.²⁹ In the i.i.d. case, these dynamics reduce to those described in Section 3.3.

It is perhaps surprising that the optimal policy can be derived but less surprising that comparative statics are difficult to obtain except numerically. We hold the ratio ρ_h/ρ_l fixed, so that the invariant probability q , and hence the expected value μ , are unaffected, and we vary the persistence κ . Numerically, a higher κ leads to a lower payoff; see Figure 5. When $\kappa = 1$, types never change, and we are left with the static problem (for the parameters of the figure, it is best not to provide the good). As κ decreases, types change more rapidly, and promises approach the type-neutral currency that they are under independence.

²⁹No such pair exists when the promise (m, λ) is already too large, namely, when $\underline{u}_l^{m, \lambda} \geq \delta \mathbf{E}_l [\underline{u}^\infty]$. In that case, the unit is awarded with the probability x that solves $\frac{\underline{u}_l^{m, \lambda} - x(1-\delta)l}{\delta} = \mathbf{E}_l [\underline{u}^\infty]$, and the promise becomes ∞ .




Figure 5: The principal's optimal payoff as a function of the persistence κ , holding $q = 3/7$ fixed ($\rho_h/\rho_l = 4/3$), for $(\delta, l, h, c) = (9/10, 1/4, 1, 3/5)$. The dashed line is the complete-information payoff $q(h - c)$.

As mentioned above, these comparative statics are numerical rather than formal. The promised utilities follow a random walk with unequal, state-dependent step sizes along the lower boundary of V , and there is little hope of establishing this monotonicity more formally here; see Section 5.1, where a result along these lines is derived. Nonetheless, it is not serial dependence but positive correlation that is detrimental. It is tempting to believe that any form of serial dependence is bad because it endows the agent with private information that pertains not only to today's value but also to all future ones, and eliciting private information is usually costly. However, conditional on today's type, the agent holds no residual private information regarding future types: his beliefs are pinned down by the Markov chain. Indeed, with perfectly negatively correlated values, the complete-information payoff would be easy to achieve: offer the agent a choice between receiving the good in all odd or all even rounds. As $\delta > l/h$ (in fact, we assumed that $\delta > l/\mu$), the low-type agent would select the odd rounds. Just as with a higher discount factor, a more negative correlation (or a less positive one) makes future promises more effective incentives because preference misalignment is shorter-lived.

Finally, consider the long-run dynamics. Given any initial promise above $\underline{U}$, other than (μ_h, μ_l) itself, finitely many consecutive reports of l (or h) suffice for the promised utilities to reach (μ_h, μ_l) (or $(0, 0)$). Hence, because $U^* > \underline{U}$, both long-run outcomes have strictly positive probability, and absorption occurs almost surely.³⁰ However, we are unable to derive

³⁰At every history, absorption occurs within a uniformly bounded number of rounds with probability

the absorption probabilities because we know of no analogue of the TW-martingale.

4.5 A Comparison with Transfers as in Battaglini (2005)

Our model can be regarded as the no-transfer counterpart of Battaglini (2005).

At first glance, the difference in results is striking. A main finding of Battaglini, “no distortion at the top,” has no counterpart in our model. With transfers, efficient provision occurs *forever* once the agent is revealed to be a high type. Moreover, even along the history in which efficiency is not achieved in finite time, namely, an uninterrupted string of low reports, efficiency is asymptotically approached. Here, in contrast, the long-run outcome is inefficient with probability one: the good is supplied forever or never again—either can arise—and no further communication is needed. In summary, with transfers, inefficiencies are frontloaded; here, they are backloaded.

The difference can be understood as follows. To begin with, Battaglini’s results rely on revenue maximization as the objective. With transfers, efficiency is trivial: supply the good if and only if the agent pays the price c .

Once revenue maximization is the objective, transfers reverse the incentive constraints: it is no longer the low type who would like to mimic the high type but the high type who would like to avoid paying his entire value by claiming that his value is low. To prevent this, the high type must be granted information rents, and his incentive constraint binds. Ideally, the principal would charge for these rents *before* the agent is privately informed, when their expected value is still common knowledge. When values are i.i.d., this poses no difficulty: the rents can be expropriated one round in advance. With correlation, the two types value these rents differently, as the likelihood of a high value in the future depends on the current type. For rents sufficiently far in the future, however, the current type has little effect on their conditional expectation (provided that $\kappa < 1$), so that their value can “almost” be extracted. It is then optimal for the principal to maximize the surplus, offering a nearly efficient contract at all sufficiently distant dates.

Money thus plays two roles. First, it allows promises to “clear” on the spot without allocative distortions, and thereby prevents backloaded inefficiencies; promises of future units are a poor substitute for money in this regard. Even if payments could not be made in advance, spot payments would suffice to restore efficiency if efficiency were the objective. Second, money allows value to be transferred from the agent to the principal before information becomes private. Information rents then no longer impede efficiency in the remote

bounded away from zero.

future: with transfers, inefficiencies arise only over the short run.

A plausible intermediate case arises when money is available but the agent is protected by limited liability, so that payments can flow in only one direction, from the principal to the agent. The principal maximizes the social surplus net of payments.³¹ In this case, we demonstrate in the appendix (Lemma 11) that, when values are i.i.d., an optimal policy without transfers exists if and only if $c \leq 2l$. At $c = 2l$, optimal policies with transfers also exist. The interpretation is as follows: $c - l$ is the cost to the principal of supplying a unit when the value is low, whereas l is the cost to the agent of forgoing such a unit. If buying off the agent is at least as costly as supplying the good when the value is low, the principal can follow the optimal policy without money.

4.6 The General Solution

Theorem 2 follows from the analysis of the optimal policy on the entire domain V , to which this subsection turns: it solves the program of Section 4.1. Because only the utility pairs on the lower boundary are relevant on path, the reader may elect to skip this subsection.

First, we further divide V into subsets and introduce two sequences of utility pairs for this purpose. Write $\underline{U} = (\underline{U}_h, \underline{U}_l)$ and, for every integer $m \geq 0$, define $v^m = (v_h^m, v_l^m)$ by

$$\begin{aligned} v_h^m &= \delta^m ((1 - q)\underline{U}_l + q\underline{U}_h + (1 - q)\kappa^m(\underline{U}_h - \underline{U}_l)), \\ v_l^m &= \delta^m ((1 - q)\underline{U}_l + q\underline{U}_h - q\kappa^m(\underline{U}_h - \underline{U}_l)), \end{aligned}$$

so that $v^0 = \underline{U}$, and define

$$\underline{V} := \text{co}(\{v^m : m = 0, 1, \dots\} \cup \{\underline{u}^m : m = 0, 1, \dots, \underline{u}^m \leq \underline{U}\}). \quad (4)$$

See the left panel of Figure 6. Note that $\underline{V}$ has a non-empty interior if and only if the low-type persistence is sufficiently low (see (3)). The complete-information payoff can be achieved on this set. The following lemma also gives a converse under an additional restriction on δ .

Lemma 6 *For all $U \in \underline{V} \cup \{(\mu_h, \mu_l)\}$ and all ϕ ,*

$$W(U, \phi) = \bar{W}(U, \phi).$$

Conversely, suppose that $v_h^* \leq \bar{u}_h^1$, which holds when δ is high enough. If $U \in V \setminus (\underline{V} \cup \{(\mu_h, \mu_l)\})$, then $W(U, \phi) < \bar{W}(U, \phi)$ for all $\phi \in (0, 1)$.

³¹If payments were costless to the principal, efficiency would be easily achieved: pay the agent c if the report is low, and nothing otherwise.




Figure 6: The set V for $(\delta, \rho_h, \rho_l, l, h) = (9/10, 1/3, 1/4, 1/4, 1)$. Left panel: the boundaries P_b and P_t , the chain P , and the subsets V_t , V_b , and $\underline{V}$. Right panel: the optimal policy.

Second, for every integer $m \geq 0$, define $\hat{u}^m := (\hat{u}_h^m, \hat{u}_l^m)$ by

$$\begin{aligned}\hat{u}_h^m &= \mu_h - (1 - \delta)h - \delta^{m+1} \left((1 - q)l + qh + (1 - q)\kappa^{m+1}(\mu_h - \mu_l) \right), \\ \hat{u}_l^m &= \mu_l - (1 - \delta)l - \delta^{m+1} \left((1 - q)l + qh - q\kappa^{m+1}(\mu_h - \mu_l) \right).\end{aligned}\tag{5}$$

Let P_b and P_t be the closures of the chains generated by $\{\underline{u}^m\}_{m \geq 0}$ and $\{\bar{u}^m\}_{m \geq 0}$, respectively. They are the lower and upper boundaries of V . The identity

$$\hat{u}^m = \underline{u}^{m+1} - (1 - \delta)(h, l) = \delta(\mathbf{E}_h[\underline{u}^m], \mathbf{E}_l[\underline{u}^m])$$

shows that the chain P generated by $\{\hat{u}^m\}_{m \geq 0}$ is the image of P_b under the discounted transition map. In particular, $\hat{u}^0 = 0$ and $\hat{u}^m \rightarrow \bar{u}^1$. Let V_t be the points of V weakly above P , and let V_b be the remaining points. For $p_l = 0$, PK_L and a binding IC_H place the continuation pair u_l on P_b precisely when $U \in P$.

We now define a policy (which, as we will see below, is optimal), leaving aside for now the choice of the initial promise.

Definition 2 For all $U \in V$, set

$$p_h = \min \left\{ 1, \frac{U(h)}{(1 - \delta)h} \right\}, \quad p_l = \max \left\{ 0, 1 - \frac{\mu_l - U(l)}{(1 - \delta)l} \right\},\tag{6}$$

and set $u_h \in P_b$. If $U \in V_b$, set $u_l \in P_b$; if $U \in V_t$, choose u_l so that (IC_H) binds.

The definition imposes one constraint on each continuation pair: u_h must lie on P_b , whereas u_l must lie on P_b or make (IC_H) bind, depending on the case. Together with the two promise-keeping equations, these are four constraints, which uniquely determine (u_h, u_l) . One may verify that the supply probabilities and the choices of u_h, u_l also imply that (IC_L) binds for all $U \in P_b$. A surprising property of this policy is that it is independent of the principal's belief: given U , the belief is irrelevant to the optimal choices. (The belief does determine the optimal initial promise on the lower boundary, as well as the principal's payoff.) The reason is the monotonicity property recorded in the theorem below: promise keeping pins down the expectation of each continuation pair under the corresponding continuation belief, and, along each line on which this expectation is constant, the value function is weakly increasing in the pair's h -component. Whatever the current belief, each continuation pair is therefore pushed toward a higher h -component until a constraint binds, so that the constraints alone pin down the policy.

The right panel of Figure 6 illustrates the dynamics of the optimal policy. Unless U is too close to 0 or (μ_h, μ_l) , the efficient supply choice $(p_h, p_l) = (1, 0)$ is made. A report of h shifts the promise to the left, onto the lower boundary. Away from $\underline{V}$, a report of l shifts it to the right, toward higher utilities: below the polygonal chain P , the promise again moves onto the lower boundary; above the chain, (IC_H) binds; and if the promise is on the upper boundary, it stays there. The latter follows from the backloading recursion in Lemma 5: a low report shortens the remaining delay by one round.

For completeness, we also define the subsets of V over which promise keeping alone prevents the efficient choice $(p_h, p_l) = (1, 0)$: let $V_h := \{U \in V : U(l) > \bar{u}_l^1\}$ and $V_l := \{U \in V : U(h) < \underline{u}_h^1\}$. The choice $(p_h, p_l) = (1, 0)$ is feasible at U given promise keeping if and only if $U \in V \setminus (V_h \cup V_l)$.

Theorem 3 *For any promise $U \in V$ and belief $\phi \in [0, 1]$, the policy of Definition 2 is optimal. For $\phi \in [\rho_l, 1 - \rho_h]$, an optimal initial promise U^* can be chosen in $P_b \cap (V \setminus \underline{V})$ and weakly increasing in ϕ .*

Furthermore, for $\phi \in \{\rho_l, 1 - \rho_h\}$, the value function $W(U, \phi)$ is weakly increasing in $U(h)$ along each line on which $\phi U(h) + (1 - \phi)U(l)$ is constant.

Since $U^* \in P_b$, the structure of the optimal policy implies that the promised utilities never leave P_b . As in the i.i.d. case, and by the same arguments, the one-sided derivative of W along the lower boundary approaches the corresponding derivative of $\bar{W}$ as U approaches either (μ_h, μ_l) or the set $\underline{V}$. As a result, the initial promise U^* is bounded away from both.

5 Extensions

Two modeling choices deserve discussion. First, we have opted for a discrete-time framework because it embeds the case of independent values—a natural starting point for which there is no counterpart in continuous time. However, this choice comes at a price. With independent values, shortening the round makes the players more patient but also makes types change more often, so the two effects cannot be disentangled. With Markovian types they can, since transition probabilities can be scaled with the round length so as to hold persistence fixed while patience varies; even so, the discrete-time analysis delivers comparative statics in persistence only numerically (Section 4.4). We remedy this deficiency below by resorting to continuous time, where patience and persistence are separate primitives and payoffs can be solved for explicitly.

Second, we have assumed that the agent’s value is binary. As is well known (see Battaglini and Lamba, 2019, for instance), it is difficult to make progress with more types, even with transfers, unless strong assumptions are imposed. In the i.i.d. case, this is nonetheless possible. Below, we consider the case of a continuum of types, which allows us to evaluate the robustness of our various findings.

5.1 Continuous Time

To make further progress, we examine the limiting stochastic process of utility and payoff under the usual Poisson scaling: the round length $\Delta > 0$ is taken to 0, while the discount factor and the transition probabilities scale as $\delta = e^{-r\Delta}$, $\rho_h = \lambda_h\Delta$, and $\rho_l = \lambda_l\Delta$, where $r > 0$ is the discount rate and $\lambda_h, \lambda_l > 0$ are the transition rates. That is, we let $(v_t)_{t \geq 0}$ be a continuous-time Markov chain (by definition, a right-continuous process) with values in $\{l, h\}$, initial probability $q = \lambda_l/(\lambda_h + \lambda_l)$ of h , and transition rates λ_h, λ_l . Let $T_0, T_1, T_2, \dots$ be the corresponding random times at which the value switches (setting $T_0 = 0$ if the initial value is l , so that, by convention, $v_t = l$ on any interval $[T_{2k}, T_{2k+1})$).

The limit of the optimal discrete-time policies defines a tuple of continuous-time processes that follow deterministic trajectories between consecutive switching times. First, the belief $(\phi_t)_{t \geq 0}$ of the principal takes values in $\{0, 1\}$ after the initial instant. Namely, $\phi_t = 0$ over any interval $[T_{2k}, T_{2k+1})$, and $\phi_t = 1$ otherwise. Second, the utilities of the agent $(U_{h,t}, U_{l,t})_{t \geq 0}$ are functions of his type. Finally, the expected payoff of the principal, $(W_t)_{t \geq 0}$, is computed according to his belief ϕ_t .

The pair of processes $(U_{h,t}, U_{l,t})_{t \geq 0}$ takes values in V , which is now obtained by taking

the limit (as $\Delta \rightarrow 0$) of the formulas for $\{\underline{u}^m, \bar{u}^m\}_{m \in \mathbf{N}_0}$. In particular, one obtains that the lower boundary is given in parametric form by

$$\begin{aligned}\underline{u}_h(\tau) &= (1 - e^{-r\tau})\mu_h + e^{-r\tau}(1 - e^{-(\lambda_h + \lambda_l)\tau})(1 - q)(\mu_h - \mu_l), \\ \underline{u}_l(\tau) &= (1 - e^{-r\tau})\mu_l - e^{-r\tau}(1 - e^{-(\lambda_h + \lambda_l)\tau})q(\mu_h - \mu_l),\end{aligned}$$

where $\underline{u}(\tau) := (\underline{u}_h(\tau), \underline{u}_l(\tau))$ is the utility pair delivered by the frontloaded policy that supplies the good on $[0, \tau)$, independently of the reports. Here, as before,

$$(\mu_h, \mu_l) = \left(h - \frac{\lambda_h}{\lambda_h + \lambda_l + r}(h - l), l + \frac{\lambda_l}{\lambda_h + \lambda_l + r}(h - l) \right)$$

is the utility pair achieved by providing the good forever. The upper boundary is now given by

$$\begin{aligned}\bar{u}_h(\tau) &= e^{-r\tau}\mu_h - e^{-r\tau}(1 - e^{-(\lambda_h + \lambda_l)\tau})(1 - q)(\mu_h - \mu_l), \\ \bar{u}_l(\tau) &= e^{-r\tau}\mu_l + e^{-r\tau}(1 - e^{-(\lambda_h + \lambda_l)\tau})q(\mu_h - \mu_l),\end{aligned}$$

where $\bar{u}(\tau) := (\bar{u}_h(\tau), \bar{u}_l(\tau))$ is the utility pair delivered by the backloaded policy that supplies the good on $[\tau, \infty)$. Finally, the set $\underline{V}$ is either degenerate or defined by those utility vectors in V lying below the graph of the curve defined by

$$\begin{aligned}v_h(\tau) &= e^{-r\tau}((1 - q)\underline{U}_l + q\underline{U}_h) + e^{-r\tau}e^{-(\lambda_h + \lambda_l)\tau}(1 - q)(\underline{U}_h - \underline{U}_l), \\ v_l(\tau) &= e^{-r\tau}((1 - q)\underline{U}_l + q\underline{U}_h) - e^{-r\tau}e^{-(\lambda_h + \lambda_l)\tau}q(\underline{U}_h - \underline{U}_l),\end{aligned}$$

where $\underline{U} = (\underline{U}_h, \underline{U}_l)$ is the intersection, other than 0, of the lower boundary with the line $U_l = \frac{\lambda_l}{\lambda_l + r}U_h$ (with $\underline{U} = 0$ if there is no such intersection). It is immediate that $\underline{V}$ has a non-empty interior if and only if (*cf.* (3))

$$\frac{h - l}{l} > \frac{r}{\lambda_l}.$$

The condition holds when the agent is sufficiently patient or the low type is sufficiently transient; when it fails, the complete-information payoff cannot be achieved for any utility pair aside from 0 and (μ_h, μ_l) . This is the limit of Lemma 6; its condition $v_h^* \leq \bar{u}_h^1$ holds for all sufficiently small Δ .

Figure 7 illustrates this construction. Note that the boundary of V is smooth except at 0 and (μ_h, μ_l) . It is also easy to verify that the limit of the chain P defined by (5) coincides with the lower boundary: V_b is asymptotically empty.




Figure 7: Incentive-feasible set for $(r, \lambda_h, \lambda_l, l, h) = (1, 3, 9/4, 1/4, 1)$. The upper and lower boundaries are traced by $(\bar{u}_h(\tau), \bar{u}_l(\tau))$ and $(\underline{u}_h(\tau), \underline{u}_l(\tau))$ as τ ranges over $[0, \infty]$.

The Poisson limit allows an explicit solution for payoffs. We sketch the derivation.

How does τ —the denomination of utility on the lower boundary—evolve over time? Along the lower boundary, it evolves continuously. On any interval over which h is continuously reported, it evolves deterministically, with increments

$$d\tau_h := -dt.$$

However, when l is reported, the evolution is more complicated. Algebra yields

$$d\tau_l := \frac{g(\tau)}{\mu - q(h-l)e^{-(\lambda_h+\lambda_l)\tau}} dt,$$

where

$$g(\tau) := q(h-l)e^{-(\lambda_h+\lambda_l)\tau} + le^{r\tau} - \mu,$$

and $\mu = qh + (1-q)l$, as before.

The increment $d\tau_l$ is negative if τ maps onto a utility vector in $\underline{V}$ and positive otherwise. If $\underline{V}$ has a non-empty interior, let $\hat{\tau}$ denote the largest value of τ for which $\underline{u}(\tau)$ lies in $\underline{V}$, that is, $\underline{u}(\hat{\tau}) = \underline{U}$; it is simply the positive root of g . Otherwise, set $\hat{\tau} = 0$.

The evolution of utility is not continuous for utility pairs that are not on the lower boundary. A high report leads to a vertical jump down to the lower boundary: the utility of the low type falls discontinuously, while the utility of the high type does not move. See Figure 7. The latter is intuitive: by promise keeping, the utility of the high-type agent cannot

jump, as a report affects the allocation only over an instant, which contributes negligibly to his utility. A low report, in contrast, involves no jump: the agent's utility drifts.

Our goal is to derive the principal's value function. Because his belief is degenerate, except at the initial instant, we write $W_h(\tau)$ (resp., $W_l(\tau)$) for the payoff when (he assigns probability one to the event that) the agent's type is currently high (resp., low). By definition of the policy that is followed, the value functions solve the following paired system of equations:

$$W_h(\tau) = rdt(h - c) + \lambda_h dt W_l(\tau) + (1 - rdt - \lambda_h dt) W_h(\tau + d\tau_h) + \mathcal{O}(dt^2),$$

and

$$W_l(\tau) = \lambda_l dt W_h(\tau) + (1 - rdt - \lambda_l dt) W_l(\tau + d\tau_l) + \mathcal{O}(dt^2).$$

On $(\hat{\tau}, \infty)$, the functions W_h, W_l are twice differentiable, as the appendix verifies. We then obtain the differential equations

$$(r + \lambda_h)W_h(\tau) = r(h - c) + \lambda_h W_l(\tau) - W_h'(\tau) \quad (7)$$

and

$$(r + \lambda_l)W_l(\tau) = \lambda_l W_h(\tau) + \frac{g(\tau)}{\mu - q(h - l)e^{-(\lambda_h + \lambda_l)\tau}} W_l'(\tau) \quad (8)$$

subject to the following boundary conditions.³² First, at $\tau = \hat{\tau}$, the value must coincide with the complete-information payoff $\bar{W}$. That is, $W_h(\hat{\tau}) = \bar{W}_h(\hat{\tau})$ and $W_l(\hat{\tau}) = \bar{W}_l(\hat{\tau})$, where $\bar{W}_h(\tau)$ and $\bar{W}_l(\tau)$ denote the complete-information payoffs at $\underline{u}(\tau)$ when the type is high and low, respectively; we also write $\bar{W}(\tau) := q\bar{W}_h(\tau) + (1 - q)\bar{W}_l(\tau)$. At $\hat{\tau}$, equation (8) reduces to $(r + \lambda_l)W_l(\hat{\tau}) = \lambda_l W_h(\hat{\tau})$. The complete-information boundary values satisfy this identity because $\underline{U}_l = \lambda_l \underline{U}_h / (r + \lambda_l)$. Second, as $\tau \rightarrow \infty$, it must hold that the always-supply payoff is approached. Hence,

$$\lim_{\tau \rightarrow \infty} W_h(\tau) = \mu_h - c, \quad \lim_{\tau \rightarrow \infty} W_l(\tau) = \mu_l - c.$$

Despite having variable coefficients, this system can be solved in closed form. See Appendix C for the solution upon which the following comparative statics are based.

Lemma 7 *Fix $\bar{\lambda}_h, \bar{\lambda}_l$ and let $\lambda_h = \bar{\lambda}_h/\pi$, $\lambda_l = \bar{\lambda}_l/\pi$. For every fixed frontloading duration τ , the limiting payoff $W(\tau) := qW_h(\tau) + (1 - q)W_l(\tau)$ is nonincreasing in the persistence π . Moreover,*

$$\lim_{\pi \rightarrow 0} (\bar{W}(\tau) - W(\tau)) = 0, \quad \lim_{\pi \rightarrow \infty} \max_{\tau} W(\tau) = \max\{\mu - c, 0\}.$$

³²These equations evaluate the limit of the optimal discrete-time policies. The appendix proves convergence of their payoffs. We use this limiting formulation throughout this subsection.

The proof is in Appendix C. The comparison holds at a fixed duration τ , whose utility pair depends on persistence. Taking the maximum over τ gives the same ranking of the optimal initial payoffs. With independent types, the agent's preferences are quasilinear in promised utility, so that the only source of inefficiency derives from the bounds on this currency. When types are correlated, the promised utility is no longer independent of today's type in the agent's preferences, reducing the degree to which this currency can be used to efficiently provide incentives. With perfectly persistent types, there is no longer any leeway, and we are back to the inefficient static outcome. Figure 8 illustrates the value function and the complete-information payoff for two levels of persistence.




Figure 8: Value function W (solid) and complete-information payoff $\bar{W}$ (dashed) plotted against τ , for $(\lambda_h, \lambda_l, r, l, h, c) = (10/(4\pi), 1/(4\pi), 1, 1/4, 1, 2/5)$ and $\pi = 1, 4$. The dashed curves coincide outside approximately $[0.158, 1.589]$.

What about the agent's utility? We note that the utility of both types is increasing in τ . Indeed, because a low type is always willing to claim that his value is high, we may compute his utility from the expected discounted value of provision over the duration τ during which he would receive the good by always reporting high. However, persistence plays an ambiguous role in determining the agent's utility. When $\mu > c$, the static solution is to always supply the good—the best possible outcome for the agent—so that perfect persistence is his favorite outcome; when $\mu < c$, the static solution is to never supply it, his worst possible outcome. Hence, persistence tends to improve the agent's situation precisely when $\mu > c$.³³

As $r \rightarrow 0$, the principal's value converges to the complete-information payoff $q(h - c)$. We conclude with the rate of convergence; the result parallels Lemma 4, with the discount

³³For $(r, \bar{\lambda}_h, \bar{\lambda}_l, l, h, c) = (1, 3, 9/4, 1/4, 1, 0.45)$, numerical evaluation of Proposition 2 gives optimal *ex ante* utilities approximately 0.4132, 0.4042, and 0.4206 at $\pi = 1, 4$, and 8, respectively.

rate r playing the role of $1 - \delta$.

Lemma 8 *As $r \downarrow 0$, with λ_h, λ_l fixed,*

$$|\max_{\tau} W(\tau) - q(h - c)| = \Theta(\sqrt{r}).$$

5.2 Continuous Types

Suppose values are drawn independently from an atomless distribution F on $[\underline{v}, 1]$, where $0 \leq \underline{v} < 1$. Its density is positive and continuous on $(\underline{v}, 1)$, and $\underline{v} < c < 1$. Write $\mu = \mathbf{E}[v]$. The results for general distributions hold for every $0 < \delta < 1$. The power-distribution envelope identity also holds for every $0 < \delta < 1$. The three-regime characterization below maintains $\delta > 1/2$.

Lemma 9 *The complete-information payoff $\bar{W}$ is strictly concave. At each promise, the optimal policy is unique up to null sets and provides the good above a threshold. This threshold decreases continuously from 1 to $\underline{v}$ as the promise increases from zero to μ . Continuation utility equals the current promise after every realized value, almost surely.*

At a fixed promise U , write $p(v)$ for provision and $u(v)$ for continuation utility. Endpoint values of these rules denote limits from the interior of the type support.

Theorem 4 *The value function W is concave on $[0, \mu]$ and continuously differentiable on $(0, \mu)$. For $0 < U < \mu$, $W(U) < \bar{W}(U)$, and every optimal policy satisfies*

$$\mathbf{P}\{0 < p(v) < 1\} > 0.$$

Thus optimal provision randomizes over a set of types of positive probability. Incentive compatibility makes provision nondecreasing and continuation utility nonincreasing in the report. A deterministic threshold would require a downward jump in continuation utility. Spreading this jump over nearby types lowers continuation cost unless the cost is affine between the two promises. The appendix rules out the remaining case using differentiability of the value function and the optimality conditions. This argument does not require strict concavity or policy uniqueness.³⁴

Let u_n be an optimal continuation rule selected after the history at date n , and let $U_{n+1} = u_n(v_n)$. The selection may depend on the entire history.

³⁴The maintained assumptions do not guarantee uniqueness of the optimal policy up to null sets, even at a fixed initial promise. A counterexample is available upon request. Stronger conditions ensure uniqueness; in particular, it holds for the power distributions considered in Proposition 1.

Theorem 5 *For every $U_0 \in (0, \mu)$ and every selection of optimal policies, U_n converges almost surely to a limit in $\{0, \mu\}$. If $\underline{v} = 0$, the limit is zero almost surely. If $\underline{v} > 0$, each endpoint has positive probability. The limiting probabilities do not depend on the selection of optimal policies. They are also the same across optimal initial promises.*

Power distributions $F(v) = v^a$, $a \geq 1$, admit a simpler marginal-value argument. The construction in the appendix gives expected marginal continuation cost equal to current marginal cost, even when a continuation bound binds. Hence

$$W'(U) = \int_0^1 W'(u(v)) dF(v), \quad 0 < U < \mu. \quad (9)$$

Derivatives at endpoints denote one-sided limits. Thus $1 - c - W'(U_n)$ is a nonnegative martingale. Its convergence excludes the upper endpoint, where W' tends to $-\infty$; a compactness argument excludes interior limits. For other distributions, binding continuation bounds contribute terms that cannot be discarded on the basis of differentiability alone. The general proof instead obtains one-sided inequalities from feasible perturbations and rules out an invariant distribution supported in the interior. It applies to every optimal selection, without strict concavity or continuity of the policy.

We next characterize provision for this family. The *lower transition promise* U_- is the infimum of interior promises at which an optimal policy has $p(1) = 1$. For $a > 1$, the *upper transition promise* U_+ is the infimum of interior promises at which an optimal policy has $p(0+) > 0$. For $a = 1$, set $U_+ = \mu$.

Proposition 1 *For $F(v) = v^a$, $a \geq 1$, and $\delta > 1/2$, the optimal allocation and continuation rules at each interior promise are unique up to null sets. Moreover, $0 < U_- < U_+ \leq \mu$, with $U_+ < \mu$ if $a > 1$.*

If $0 < U < U_-$, there is $s \in (0, 1)$ such that $p(v) = 0$ for $v < s$, and p is strictly increasing on $(s, 1)$. The lower continuation bound binds: $u(1) = 0$. Moreover, $p(1) < 1$ and $u(0) < \mu$.

If $U_- < U < U_+$, there are $0 < s < t < 1$ such that $p(v) = 0$ for $v < s$, p is strictly increasing on (s, t) , and $p(v) = 1$ for $v > t$. Both continuation bounds are slack.

If $a > 1$ and $U_+ < U < \mu$, there is $t \in (0, 1)$ such that $p(0+) > 0$, p is strictly increasing on $(0, t)$, and $p(v) = 1$ for $v > t$. The upper continuation bound binds: $u(0) = \mu$. The lower continuation bound is slack.

At U_- , $p(1) = 1$ and $u(1) = 0$. For $a > 1$, at U_+ , $p(0+) = 0$ and $u(0) = \mu$. At $U = 0$, provision and continuation utility are zero. At $U = \mu$, provision is one and continuation utility is μ .

The descriptions allow an upward jump at a cutoff separating zero, interior, or full provision. For $a > 1$, high promises require positive provision even to the lowest types. This third regime is absent under the uniform distribution. An indirect implementation now requires a nonlinear menu: the agent chooses a provision probability and the corresponding continuation promise. Incentive compatibility makes the prescribed choice optimal. Such an implementation only requires selecting an optimal rule; it does not require uniqueness for general distributions.

Markovian Types. With three or more persistent types, we do not have a characterization of the incentive-feasible set or the optimal policy. The binary analysis suggests studying which incentive constraints bind on its boundary. Any extension must specify the transition matrix: the ordering of the types alone does not establish that frontloading and backloading are extremal. We leave these questions open.

6 Concluding Comments

Renegotiation-Proofness. Our optimal policy is not renegotiation-proof. After a history that leads to zero continuation utility, both parties would prefer to start afresh. With transfers, an optimal contract can be renegotiation-proof (Battaglini, 2005). Strong renegotiation-proofness (Farrell and Maskin, 1989) imposes further restrictions on continuation payoffs. It might lead to a lower reflecting boundary, but this is a conjecture. Establishing it would require deriving the admissible continuation set and checking the incentive constraints again.

Public Signals. While assuming no statistical evidence allows us to clarify how the principal can exploit the repeated allocation decision to mitigate the inefficiency entailed by private information, there are many applications for which such evidence is available. Such a signal would depend on the current type and possibly on the action chosen by the principal. For instance, if we interpret the decision as filling a position (as in labor market applications), the principal receives feedback on the quality of the applicant only if he is hired. If the good insures the agent against a risk whose cost may be high or low, a principal who fails to provide it may discover that the agent's claim was warranted.

Incomplete Information Regarding the Process. Thus far, we have assumed that the agent's type is drawn from a distribution that is common knowledge. In practice, the agent might have superior information regarding the frequency with which high values arrive. If

the agent knows the distribution from the beginning, this is persistent private information. A recursive formulation would require continuation promises indexed by the possible distributions. Incentive constraints must compare reports of the distribution using the true law of future values. Adding only an initial incentive constraint to the scalar-promise formulation of Section 3 does not suffice.

Alternatively, the agent may learn the distribution from successive realizations. His posterior can then be private, including after deviations. A recursive formulation would have to track the principal’s belief over these posteriors and continuation utilities for each possible private posterior. A scalar continuation promise generally does not suffice. Characterizing the optimal mechanism in either case remains open. Token and budget mechanisms may nevertheless provide useful benchmarks.

References

- Abdulkadiroğlu, Atila, and Kyle Bagwell.** 2012. “The Optimal Chips Mechanism in a Model of Favors.” Working paper, Duke University and Stanford University.
- Abdulkadiroğlu, Atila, and Simon Loertscher.** 2007. “Dynamic House Allocations.” Working paper, Duke University and University of Melbourne.
- Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1990. “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring.” *Econometrica*, 58(5): 1041–1063.
- Albright, S. Christian.** 1974. “Optimal Sequential Assignments with Random Arrival Times.” *Management Science*, 21(1): 60–67.
- Athey, Susan, and Kyle Bagwell.** 2001. “Optimal Collusion with Private Information.” *RAND Journal of Economics*, 32(3): 428–465.
- Battaglini, Marco.** 2005. “Long-Term Contracting with Markovian Consumers.” *American Economic Review*, 95(3): 637–658.
- Battaglini, Marco, and Rohit Lamba.** 2019. “Optimal Dynamic Contracting: The First-Order Approach and Beyond.” *Theoretical Economics*, 14(4): 1435–1482.
- Biais, Bruno, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet.** 2007. “Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications.” *The Review of Economic Studies*, 74(2): 345–390.

- Casella, Alessandra.** 2005. “Storable Votes.” *Games and Economic Behavior*, 51(2): 391–419.
- Cohn, Zachary.** 2010. “A Note on Linked Bargaining.” *Journal of Mathematical Economics*, 46(2): 238–247.
- Cole, Harold L., and Narayana R. Kocherlakota.** 2001. “Efficient Allocations with Hidden Income and Hidden Storage.” *The Review of Economic Studies*, 68(3): 523–542.
- Derman, Cyrus, Gerald J. Lieberman, and Sheldon M. Ross.** 1972. “A Sequential Stochastic Assignment Problem.” *Management Science*, 18(7): 349–355.
- Doepke, Matthias, and Robert M. Townsend.** 2006. “Dynamic Mechanism Design with Hidden Income and Hidden Actions.” *Journal of Economic Theory*, 126(1): 235–285.
- Eilat, Ran, and Ady Pauzner.** 2011. “Optimal Bilateral Trade of Multiple Objects.” *Games and Economic Behavior*, 71(2): 503–512.
- Fang, Hanming, and Peter Norman.** 2006. “Overcoming Participation Constraints.” Cowles Foundation for Research in Economics, Yale University Cowles Foundation Discussion Paper 1511R. Revised April 2006.
- Farrell, Joseph, and Eric Maskin.** 1989. “Renegotiation in Repeated Games.” *Games and Economic Behavior*, 1(4): 327–360.
- Fernandes, Ana, and Christopher Phelan.** 2000. “A Recursive Formulation for Repeated Agency with History Dependence.” *Journal of Economic Theory*, 91(2): 223–247.
- Frankel, Alexander.** 2016. “Discounted Quotas.” *Journal of Economic Theory*, 166: 396–444.
- Fries, Brant E.** 1975. “Optimal Ordering Policy for a Perishable Commodity with Fixed Lifetime.” *Operations Research*, 23(1): 46–61.
- Fu, Shiming, and R. Vijay Krishna.** 2019. “Dynamic Financial Contracting with Persistent Private Information.” *RAND Journal of Economics*, 50(2): 418–452.
- Garrett, Daniel F., and Alessandro Pavan.** 2015. “Dynamic Managerial Compensation: A Variational Approach.” *Journal of Economic Theory*, 159: 775–818.

- Gershkov, Alex, and Benny Moldovanu.** 2010. “Efficient sequential assignment with incomplete information.” *Games and Economic Behavior*, 68(1): 144–154.
- Hauser, Christine, and Hugo Hopenhayn.** 2008. “Trading Favors: Optimal Exchange and Forgiveness.” Working paper, UCLA.
- Hortala-Vallve, Rafael.** 2010. “Inefficiencies on linking decisions.” *Social Choice and Welfare*, 34(3): 471–486.
- Hylland, Aanund, and Richard Zeckhauser.** 1979. “The Efficient Allocation of Individuals to Positions.” *Journal of Political Economy*, 87(2): 293–314.
- Jackson, Matthew O., and Hugo F. Sonnenschein.** 2007. “Overcoming Incentive Constraints by Linking Decisions.” *Econometrica*, 75(1): 241–257.
- Kalla, Simo J.** 2010. “Essays in Favor-Trading.” PhD diss. University of Pennsylvania.
- Kováč, Eugen, Daniel Krähmer, and Tymon Tatur.** 2014. “Optimal Stopping in a Principal-Agent Model with Hidden Information and No Monetary Transfers.” Working paper, University of Bonn.
- Krishna, R. Vijay, Giuseppe Lopomo, and Curtis R. Taylor.** 2013. “Stairway to Heaven or Highway to Hell: Liquidity, Sweat Equity, and the Uncertain Path to Ownership.” *RAND Journal of Economics*, 44(1): 104–127.
- Li, Jin, Niko Matouschek, and Michael Powell.** 2017. “Power Dynamics in Organizations.” *American Economic Journal: Microeconomics*, 9(1): 217–241.
- Miralles, Antonio.** 2012. “Cardinal Bayesian Allocation Mechanisms without Transfers.” *Journal of Economic Theory*, 147(1): 179–206.
- Möbius, Markus.** 2001. “Trading Favors.” Working paper, Harvard University.
- Radner, Roy.** 1981. “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship.” *Econometrica*, 49(5): 1127–1148.
- Radner, Roy.** 1986. “Repeated Moral Hazard with Low Discount Rates.” In *Uncertainty, Information, and Communication: Essays in Honor of Kenneth J. Arrow, Vol. III.*, ed. Walter P. Heller, Ross M. Starr and David A. Starrett, 25–64. Cambridge University Press.

- Rubinstein, Ariel, and Menahem E. Yaari.** 1983. “Repeated Insurance Contracts and Moral Hazard.” *Journal of Economic Theory*, 30(1): 74–97.
- Sendelbach, Sue.** 2012. “Alarm Fatigue.” *Nursing Clinics of North America*, 47(3): 375–382.
- Spear, Stephen E., and Sanjay Srivastava.** 1987. “On Repeated Moral Hazard with Discounting.” *The Review of Economic Studies*, 54(4): 599–617.
- Thomas, Jonathan, and Tim Worrall.** 1990. “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem.” *Journal of Economic Theory*, 51(2): 367–390.
- Zhang, Hao.** 2012. “Analysis of a Dynamic Adverse Selection Model with Asymptotic Efficiency.” *Mathematics of Operations Research*, 37(3): 450–474.
- Zhang, Hao, and Stefanos Zenios.** 2008. “A Dynamic Principal-Agent Model with Hidden Information: Sequential Optimality Through Truthful State Revelation.” *Operations Research*, 56(3): 681–696.

A Proofs for Section 3

Proof of Theorem 1 and Lemma 2. Consider a concave continuation value $\widehat{W}$ satisfying $\widehat{W}(0) = 0$, $\widehat{W}(\mu) = \mu - c$, and $\widehat{W} \leq \bar{W}$. The two incentive constraints imply $p_h \geq p_l$. Holding these probabilities fixed, move u_h, u_l toward their weighted mean until IC_L binds. This preserves PK, feasibility, and IC_H , and weakly raises the continuation payoff by concavity. Thus we may impose equality in IC_L . Based on PK and the binding IC_L , we solve for u_h, u_l as functions of p_h, p_l and U :

$$u_h = \frac{U - (1 - \delta)p_h(qh + (1 - q)l)}{\delta}, \quad (10)$$

$$u_l = \frac{U - (1 - \delta)(p_h q(h - l) + p_l l)}{\delta}. \quad (11)$$

We want to show that an optimal policy is such that (i) either u_h as defined in (10) equals 0 or $p_h = 1$; and (ii) either u_l as defined in (11) equals μ or $p_l = 0$. Write $\widehat{W}(U; p_h, p_l)$ for the payoff from using p_h, p_l as probabilities of assigning the good, and using promised utilities as given by (10)–(11). Because $\widehat{W}$ is concave, with $\widehat{W}(0) = 0$, $\widehat{W}(\mu) = \mu - c = \bar{W}(\mu)$, and

$\widehat{W} \leq \bar{W}$, the slope of any chord of $\widehat{W}$ lies in $[1 - c/l, 1 - c/h]$. Substituting u_h and u_l into the optimality equation, we get that, for any fixed $p_h < p'_h$, writing $\Delta := p'_h - p_h$ and u'_h, u'_l for the corresponding utilities, so that $u_h - u'_h = (1 - \delta)\mu\Delta/\delta$ and $u_l - u'_l = (1 - \delta)\underline{U}\Delta/\delta$,

$$\begin{aligned} & \widehat{W}(U; p'_h, p_l) - \widehat{W}(U; p_h, p_l) \\ &= (1 - \delta)q(h - c)\Delta - \delta q \left[\widehat{W}(u_h) - \widehat{W}(u'_h) \right] - \delta(1 - q) \left[\widehat{W}(u_l) - \widehat{W}(u'_l) \right] \\ &\geq (1 - \delta)q(h - c)\Delta - \left(1 - \frac{c}{h}\right) [\delta q(u_h - u'_h) + \delta(1 - q)(u_l - u'_l)] \\ &= (1 - \delta)q\Delta \left\{ h - c - \left(1 - \frac{c}{h}\right) (\mu + (1 - q)(h - l)) \right\} = 0, \end{aligned}$$

where the inequality plugs in the upper bound $1 - c/h$ for the two chord slopes, and the last equality uses $\mu + (1 - q)(h - l) = h$. It follows that there is no loss (and possibly a gain) in increasing p_h , unless feasibility prevents this. An entirely analogous argument, plugging in the lower bound $1 - c/l$ for the chord slope of $\widehat{W}$ between u'_l and u_l , implies that $\widehat{W}(U; p_h, p_l)$ is nonincreasing in p_l . Both continuation utilities are nonincreasing in each supply probability. Therefore, either $u_h \geq 0$ binds or p_h equals 1. Similarly, either $u_l \leq \mu$ binds or p_l equals 0. The resulting probabilities satisfy $p_h \geq p_l$, so IC_H holds.

We apply this argument to the function in (1). Since $X \in [\underline{U}, \mu]$ and $\mathbf{E}[X] = qh$, this function is continuous and concave, has the required endpoint values, and lies below $\bar{W}$ by Jensen's inequality.

We now verify that, with the proposed W as the continuation value, the maximization in (OBJ), subject to (IC_H) , (IC_L) , and (PK), gives $W(U)$ for every $U \in [0, \mu]$. We consider three cases.

By definition,

$$X \stackrel{d}{=} (1 - \delta)(\underline{U} + lB) + \delta X', \quad (12)$$

where B is Bernoulli with parameter q , and X' has the same distribution as X and is independent of B .

First, suppose $U \in [(1 - \delta)\mu, \mu - (1 - \delta)l]$. The maximizing choice is $p_h = 1$, $p_l = 0$, with

$$u_h = \frac{U - (1 - \delta)\mu}{\delta}, \quad u_l = \frac{U - (1 - \delta)\underline{U}}{\delta}.$$

Conditioning on B in (12) gives

$$\mathbf{E}(U - X)_+ = \delta[q\mathbf{E}(u_h - X)_+ + (1 - q)\mathbf{E}(u_l - X)_+].$$

Since $\delta[qu_h + (1 - q)u_l] = U - (1 - \delta)qh$, substituting (1) for the continuation payoffs in

(OBJ) gives $W(U)$. Thus

$$W(U) = (1 - \delta)q(h - c) + \delta qW\left(\frac{U - (1 - \delta)\mu}{\delta}\right) + \delta(1 - q)W\left(\frac{U - (1 - \delta)\underline{U}}{\delta}\right).$$

Next, suppose $U \leq (1 - \delta)\mu$. The maximizing choice is $p_h = \frac{U}{(1 - \delta)\mu}$, $p_l = 0$, $u_h = 0$, $u_l = \frac{U}{\delta\mu}$. Because $\delta > l/\mu$, $U \leq (1 - \delta)\mu < \underline{U}$, $u_l \leq \frac{(1 - \delta)l}{\delta} < \underline{U}$. Thus the proposed W is linear at the current and continuation utilities. By (PK), the maximized payoff is

$$\left(1 - \frac{c}{h}\right) [(1 - \delta)qp_h h + \delta(1 - q)u_l] = \left(1 - \frac{c}{h}\right) U = W(U).$$

Finally, suppose $U \geq \mu - (1 - \delta)l$. The maximizing choice is $p_h = 1$, $p_l = 1 - (\mu - U)/((1 - \delta)l)$, $u_l = \mu$, and $u_h = (U - (1 - \delta)\mu)/\delta$. Since $U - (1 - \delta)\underline{U} \geq \delta\mu \geq \delta X$, conditioning on B in (12) gives

$$\mathbf{E}(U - X)_+ = \delta q \mathbf{E}(u_h - X)_+ + (1 - q)[U - (1 - \delta)\underline{U} - \delta qh].$$

Substituting (1) for $W(u_h)$ and using $W(\mu) = \mu - c$ and the preceding identity, we obtain $W(U)$ as the maximized payoff in (OBJ). Equivalently,

$$W(U) = (1 - q)\left(U - c - (U - \mu)\frac{c}{l}\right) + (1 - \delta)q(\mu - c) + \delta qW\left(\frac{U - (1 - \delta)\mu}{\delta}\right).$$

The proposed W is therefore a fixed point of the Bellman operator. Since this operator is a contraction with modulus δ , its bounded fixed point is unique. This proves that (1) is the unique solution of (OBJ), subject to (IC_H), (IC_L), and (PK).

We next show that X has no atoms. Otherwise, its maximal atom is attained at finitely many points. Choose the largest one U' . Conditioning on B in (12) gives

$$\mathbf{P}[X = U'] = q\mathbf{P}\left[X = \frac{U' - (1 - \delta)\mu}{\delta}\right] + (1 - q)\mathbf{P}\left[X = \frac{U' - (1 - \delta)\underline{U}}{\delta}\right].$$

Both probabilities must equal $\mathbf{P}[X = U'] > 0$. Thus $\frac{U' - (1 - \delta)\mu}{\delta}$ is at least $\underline{U}$, so $U' > \underline{U}$ and $\frac{U' - (1 - \delta)\underline{U}}{\delta} > U'$, contradicting the choice of U' as the largest point with maximal atomic mass.

We next show that every nonempty open interval contained in $[\underline{U}, \mu]$ has strictly positive probability under the distribution of X . Since $\delta > 1/2$, the set

$$\left\{\underline{U} + (1 - \delta)l \sum_{j=0}^{n-1} \delta^j b_j : n \geq 1, b_0, \dots, b_{n-1} \in \{0, 1\}\right\}$$

is a dense subset of $[\underline{U}, \mu]$. Given any nonempty open subinterval, choose a value from this set strictly inside it. Matching its finite Bernoulli sequence and then requiring sufficiently

many additional zeros makes the remaining tail small enough that X lies in the interval. Hence every nonempty open subinterval of $[\underline{U}, \mu]$ has positive probability.

Since X has no atoms, (1) gives a continuous derivative

$$W'(U) = 1 - \frac{c}{h} - \frac{c(h-l)}{hl} \mathbf{P}[X \leq U]. \quad (13)$$

Since every open subinterval of $[\underline{U}, \mu]$ has positive probability for X , W' is strictly decreasing there, so W is strictly concave on $[\underline{U}, \mu]$. Also, $X \geq \underline{U}$ implies $W(U) = (1 - c/h)U = \bar{W}(U)$ on $[0, \underline{U}]$. Finally, since $\mathbf{P}[X > 0] = 1$ and $\mathbf{P}[X < \mu] = 1$, (13) gives $\lim_{U \downarrow 0} W'(U) = 1 - c/h$ and $\lim_{U \uparrow \mu} W'(U) = 1 - c/l$.

It remains to establish policy uniqueness. For $U > \underline{U}$, the rule just verified has $u_l > \underline{U}$. Averaging this rule with any other optimal rule and using strict concavity above $\underline{U}$ forces the same low-report continuation. If IC_L were slack, contracting the continuation utilities toward their mean would strictly raise the payoff. Thus IC_L binds. The preceding chord inequalities are then strict when p_h can be increased: the low-report chord intersects $(\underline{U}, \mu)$. They are also strict when p_l can be decreased, since every nontrivial chord has slope greater than $1 - c/l$. This gives the stated rule uniquely. At $U = \underline{U}$, equality with the complete-information bound requires $p_l = 0$ and $u_h, u_l \leq \underline{U}$. IC_L and PK imply

$$\underline{U} \leq (1 - \delta)\underline{U}p_h + \delta u_l \leq \underline{U}.$$

Hence $p_h = 1$, $u_l = \underline{U}$, and IC_L binds, which determines u_h . At $U = 0$, promise keeping determines the rule. ■

Lemma 10 *W has a unique maximizer, $U^* = \arg \max_U W(U)$. U^* is decreasing in c .*

Proof. By Lemma 2, W is increasing on $[0, \underline{U}]$ and strictly concave on $[\underline{U}, \mu]$, so it has a unique maximizer U^* ; since $W'(\underline{U}) = 1 - c/h > 0$ and $\lim_{U \uparrow \mu} W'(U) = 1 - c/l < 0$, U^* lies in $(\underline{U}, \mu)$ and solves $W'(U^*) = 0$. By (13), this condition is equivalent to

$$\mathbf{P}[X \leq U^*] = \frac{l(h-c)}{c(h-l)}. \quad (14)$$

The right-hand side of (14) is strictly decreasing in c . The law of X does not depend on c , and $\mathbf{P}[X \leq U]$ is nondecreasing in U . Therefore, if c increases, $\mathbf{P}[X \leq U^*]$ must strictly decrease, which requires U^* to strictly decrease. ■

Proof of Lemma 3. Use the Bernoulli sequence defining X to generate reports, and set $Y_0 = U^*$ and

$$Y_{n+1} = \frac{Y_n - (1 - \delta)(\underline{U} + lB_n)}{\delta}.$$

Then

$$\delta^n Y_n = U^* - (1 - \delta) \sum_{j < n} \delta^j (\underline{U} + l B_j) \longrightarrow U^* - X.$$

The processes U_n and Y_n agree until a peripheral interval is entered. In the lower interval, $U_n < \underline{U}$ remains below $\underline{U}$ and reaches zero at the next high report. Moreover, $\underline{U} - Y_{n+1} \geq (\underline{U} - Y_n)/\delta$, so $Y_n \rightarrow -\infty$. In the upper interval, the processes continue to agree after a high report; after a low report, $U_{n+1} = \mu$ and $Y_{n+1} \geq \mu$. The latter inequality is preserved thereafter. If neither absorption occurs, Y_n remains bounded, implying $X = U^*$, an event of probability zero by Lemma 2. These observations give

$$\mathbf{P}[U_\infty = \mu] = \mathbf{P}[X < U^*], \quad \mathbf{P}[U_\infty = 0] = \mathbf{P}[X > U^*].$$

Equation (14) gives (2). ■

Proof of Lemma 4. We first consider the forecaster. Using $\mathbf{E}[X] = qh$ in (1), we obtain

$$q(h - c) - W(U) = \left(1 - \frac{c}{h}\right) \mathbf{E}(X - U)_+ + \left(\frac{c}{l} - 1\right) \mathbf{E}(U - X)_+.$$

For every $U \in [0, \mu]$, the triangle inequality and Jensen's inequality give

$$\mathbf{E}|X - qh| \leq \mathbf{E}|X - U| + \mathbf{E}|U - qh| \leq 2\mathbf{E}|X - U|.$$

Since both coefficients in the loss expression are positive, this provides a lower bound uniform in U . For the upper bound, evaluate at $U = qh$, where both expected positive parts equal $\mathbf{E}|X - qh|/2$, and use $W(U^*) \geq W(qh)$. Thus

$$\frac{1}{2} \min \left\{ 1 - \frac{c}{h}, \frac{c}{l} - 1 \right\} \mathbf{E}|X - qh| \leq q(h - c) - W(U^*) \leq \frac{c(h - l)}{2hl} \mathbf{E}|X - qh|. \quad (15)$$

It remains to bound $\mathbf{E}|X - qh|$. By the definition of X , since $\mathbf{E}[X] = qh$ and the B_j are independent with $\text{Var}(B_j) = q(1 - q)$,

$$\sigma^2 := \mathbf{E}(X - qh)^2 = \text{Var}(X) = l^2(1 - \delta)^2 \sum_{j=0}^{\infty} \delta^{2j} \text{Var}(B_j) = l^2 q(1 - q) \frac{1 - \delta}{1 + \delta}.$$

We estimate the fourth moment of $X - qh$ as follows,

$$\begin{aligned}
\mathbf{E}(X - qh)^4 &= l^4(1 - \delta)^4 \sum_{j=0}^{\infty} \delta^{4j} \mathbf{E}(B_j - q)^4 + 6l^4(1 - \delta)^4 q^2(1 - q)^2 \sum_{0 \leq i < j} \delta^{2i+2j} \\
&\leq l^4(1 - \delta)^4 \sum_{j=0}^{\infty} \delta^{4j} \mathbf{E}(B_j - q)^4 + 3l^4(1 - \delta)^4 q^2(1 - q)^2 \left(\sum_{j=0}^{\infty} \delta^{2j} \right)^2 \\
&\leq l^4 q(1 - q) \frac{(1 - \delta)^4}{1 - \delta^4} + 3\sigma^4 \\
&= \left(3 + \frac{1 - \delta^2}{q(1 - q)(1 + \delta^2)} \right) \sigma^4 \\
&\leq 4\sigma^4 \quad (\text{for } \delta \text{ sufficiently close to } 1).
\end{aligned}$$

Combining the fourth-moment bound with Hölder's inequality yields

$$\frac{\sigma}{2} \leq \frac{\sigma^3}{\sqrt{\mathbf{E}(X - qh)^4}} \leq \mathbf{E}|X - qh| \leq \sigma.$$

Since $\sigma = \Theta(\sqrt{1 - \delta})$ and the coefficients in (15) are strictly positive and independent of δ , we conclude that

$$q(h - c) - W(U^*) = \Theta(\sqrt{1 - \delta}).$$

We now consider the prophet. Let $H = (1 - \delta) \sum_{j \geq 0} \delta^j B_j$ be the discounted quantity of high values. A token budget $K = q$ gives loss at most $\max\{h - c, c - l\} \mathbf{E}|H - q| = O(\sqrt{1 - \delta})$. It remains to prove a lower bound for every mechanism. We first consider finitely many equally weighted copies. Let us start with the simplest case: a fixed number of copies $2N + 1$, and $q = 1/2$.³⁵ We relax the problem by publicly revealing a partition of the copies into three sets. Conditional on the smaller of the numbers of high and low copies being $m \leq N$, select uniformly a set of m high copies and a set of m low copies. Reveal these sets and their values, but do not reveal the values of the remaining $2N + 1 - 2m$ copies. The remaining copies are either all high or all low. These two possibilities have equal conditional probability: the two complementary counts have the same probability, and the partition rule assigns the same likelihood to each compatible partition in the two cases.

Within each disclosed set the copies can be treated symmetrically. Let p_0, p_1, p_2 be their supply probabilities, indexed by the report on the undisclosed set. The conditional loss is

³⁵We take the number of copies to be odd for simplicity; the argument for an even number of copies is identical.

bounded below by the minimum of

$$\begin{aligned}\mathcal{L} = & \mathbf{E}[(1 - p_0)m(h - c) + p_2m(c - l)] \\ & + \frac{2N + 1 - 2m}{2} \left[(1 - p_1((2N + 1 - m, m)))(h - c) + p_1((m, 2N + 1 - m))(c - l) \right],\end{aligned}$$

subject to incentive compatibility. Here p_0 applies to the disclosed high copies, p_2 to the disclosed low copies, and p_1 to the remaining copies. The IC constraint of the agent with m high copies is

$$\begin{aligned}hm p_0((m, 2N + 1 - m)) + l(2N + 1 - 2m) p_1((m, 2N + 1 - m)) + lm p_2((m, 2N + 1 - m)) \\ \geq hm p_0((2N + 1 - m, m)) + l(2N + 1 - 2m) p_1((2N + 1 - m, m)) \\ + lm p_2((2N + 1 - m, m)),\end{aligned}$$

which implies

$$\begin{aligned}l(2N + 1 - 2m) & \leq hm[1 - p_0((2N + 1 - m, m))] + lmp_2((m, 2N + 1 - m)) \\ & \quad + l(2N + 1 - 2m)[1 - p_1((2N + 1 - m, m)) + p_1((m, 2N + 1 - m))] \\ & \leq \frac{2h}{\min\{h - c, c - l\}} \mathcal{L}.\end{aligned}$$

For each information set, $\mathcal{L} \geq (2N + 1 - 2m)\eta$, where $\eta = \frac{l}{2h} \min\{h - c, c - l\}$. Hence, the inefficiency is bounded below by (adding up over all information sets)

$$\sum_{m=0}^N 2 \binom{2N + 1}{m} \left(\frac{1}{2}\right)^{2N+1} (2N + 1 - 2m)\eta = \frac{2\Gamma(N + \frac{3}{2})}{\sqrt{\pi}\Gamma(N + 1)}\eta \sim \frac{2\sqrt{2N + 1}}{\sqrt{2\pi}}\eta.$$

This bound continues to hold in the presence of additional disclosed copies of arbitrary weights. In the low state's incentive constraint, a disclosed high copy can help only by receiving less provision in the high state, and a disclosed low copy only by receiving more in the low state. Each unit of such a distortion contributes at most h to incentives and costs at least $\min\{h - c, c - l\}/2$ in conditional expected surplus. Thus the same inequality applies, with the loss on the disclosed copies included in $\mathcal{L}$.

If $q > 1/2$, independently reveal each high copy with probability $\lambda = 2 - 1/q$. Every unrevealed copy is then high with probability $1/2$. If $q < 1/2$, reveal each low copy with probability $2 - 1/(1 - q)$ instead. In either case, the number N' of unrevealed copies satisfies $N'/(2N + 1) \rightarrow 2 \min\{q, 1 - q\}$ in probability. Apply the preceding partition argument to these copies. Conditional on N' , their expected absolute high-low count difference is

asymptotic to $\sqrt{2N'/\pi}$, and is bounded by $\sqrt{N'}$. Bounded convergence therefore gives the lower bound

$$(1 + o(1)) \frac{2\sqrt{\min\{q, 1-q\}}}{\sqrt{\pi}} \eta \sqrt{2N+1}.$$

For discounting, we first justify replacing a weighted copy by independent smaller copies. Suppose a copy of weight w is replaced by independent copies of weights $w_1, \dots, w_k$, with $\sum_i w_i = w$. Given any incentive-compatible mechanism for the original problem, run it once for each new copy. In run i , use that copy's report in place of the original copy. Allocate the new copy according to this run, and allocate every other good with the weighted average of its probabilities across the runs, using weights w_i/w . For every realized type vector, utility is the same weighted sum of utilities in the original mechanism. Truthful reporting maximizes every summand, so it maximizes their sum. Each run has the original distribution of types, and expected payoff is unchanged. The optimal payoff therefore cannot decrease under this replacement.

Fix $\alpha \in (0, 1)$ and set $N = \lfloor \log \alpha / \log \delta \rfloor$. Then $\delta^N \rightarrow \alpha$ as $\delta \rightarrow 1$. Reveal all values from round N onward. For each $k = 0, \dots, N-1$, split the weight δ^k into $\lfloor \delta^k / \delta^N \rfloor$ independent copies of weight δ^N and, if needed, one independent remainder. Reveal each remainder's value. The disclosed remainders and tail may still be used for incentives; the preceding bound includes their losses. These disclosures further relax the problem. The undisclosed copies have equal weights, so their loss is asymptotically at least

$$\frac{2\sqrt{\min\{q, 1-q\}}}{\sqrt{\pi}} \eta (1-\delta) \delta^N \sqrt{\sum_{k=0}^{N-1} \left\lfloor \frac{\delta^k}{\delta^N} \right\rfloor}.$$

For $\alpha = 1/2$, the sum is at least N , and $(1-\delta)N \rightarrow \log 2$. The lower bound is therefore of order $\sqrt{1-\delta}$, as required. ■

B Proofs for Section 4

Proof of Lemma 5. We focus on the case with $1 - \rho_h > \rho_l$. We first make a few observations about V :

Because $(0, 0)$ is delivered by never providing the good and (μ_h, μ_l) is delivered by always providing the good, $(0, 0)$ and (μ_h, μ_l) are the lowest and highest utility pairs that the principal can deliver. It follows that $V \subset [0, \mu_h] \times [0, \mu_l]$. Moreover, for any $(U(h), U(l)) \in V$, $0 < U(h) < \mu_h$ if and only if $0 < U(l) < \mu_l$.

Utility pairs in $\text{co}\{\bar{u}^m, \underline{u}^m : m = 0, 1, 2, \dots\}$ can be implemented without eliciting the agent's information. It follows that $\text{co}\{\bar{u}^m, \underline{u}^m : m = 0, 1, 2, \dots\} \subset V$.

The set V is compact. Allocation probabilities belong to a compact countable product of copies of $[0, 1]$. Discounted utilities are continuous in this product topology, and incentive compatibility is closed. The image of the incentive-compatible policies is therefore compact. The set is convex because mixtures of policies preserve incentives and average their utilities.

Let $SE(x) := \min_{(x, U(l)) \in V} U(l)$ and $NW(x) := \max_{(x, U(l)) \in V} U(l)$. Then $SE(x)$ and $NW(x)$ are the lowest and the highest utilities to a current low type, given that a current high type gets x . The convexity of V implies that $SE(\cdot)$ is convex and $NW(\cdot)$ is concave.

Write $r_0 = l/h$, $a_0 = 1 - \rho_l - r_0\rho_h \in (0, 1)$, and $b_0 = (1 - r_0)(\rho_l + r_0\rho_h) > 0$. Suppose $\min_{U \in V} (U(l) - r_0U(h)) = -M < 0$. At a minimizing pair, IC_L and PK_H give

$$-M \geq \delta a_0(u_h(l) - r_0u_h(h)) + \delta b_0u_h(h) \geq -\delta a_0M > -M,$$

a contradiction. Thus $SE(x) \geq r_0x$. Current provision with probability $x/((1 - \delta)h)$ attains equality for $x \leq (1 - \delta)h$. Similarly, let $M = \max_{U \in V} (U(l) - r_0U(h))$. IC_H and PK_L give $M \leq \delta a_0M + \delta b_0\mu_h$. The always-supply pair satisfies $\mu_l - r_0\mu_h = \delta a_0(\mu_l - r_0\mu_h) + \delta b_0\mu_h$. It follows that $NW(x) \leq \mu_l + r_0(x - \mu_h)$. Partial current provision followed by permanent provision attains this bound for $x \geq \mu_h - (1 - \delta)h$.

Convexity of SE and concavity of NW , together with these endpoint segments, imply that both one-sided boundary slopes are at least l/h . Below, SE' denotes the right derivative and NW' the left derivative.

For any $x \in (0, \mu_h)$, $SE(x)$ is given by the following program:

$$\begin{aligned} \min_{p_h, p_l \in [0, 1], u_h, u_l \in V} & U(l) \\ \text{s.t., } x &= (1 - \delta)p_h h + \delta((1 - \rho_h)u_h(h) + \rho_h u_h(l)) \\ &\geq (1 - \delta)p_l h + \delta((1 - \rho_h)u_l(h) + \rho_h u_l(l)) \\ U(l) &= (1 - \delta)p_l l + \delta(\rho_l u_l(h) + (1 - \rho_l)u_l(l)) \\ &\geq (1 - \delta)p_h l + \delta(\rho_l u_h(h) + (1 - \rho_l)u_h(l)). \end{aligned}$$

The four constraints are (PK_H) , (IC_H) , (PK_L) , (IC_L) , respectively. Because we assume that $x \in (0, \mu_h)$, the minimal $U(l)$ must be strictly positive. It follows that in the solution either $p_l > 0$ or $u_l > (0, 0)$. We make three observations regarding the solution of this program:

(IC_L) must bind. If not, we can reduce p_l or u_l . This has no impact on (PK_H) and makes (IC_H) easier to satisfy. This reduction will lower the objective $U(l)$.

u_h must be on the lower boundary $SE(\cdot)$. If not, we can replace u_h with

$$\left(u_h(h) + \varepsilon, u_h(l) - \varepsilon \frac{1 - \rho_h}{\rho_h}\right)$$

for small $\varepsilon > 0$. This has no impact on the promised utility to the current high type, and thus no impact on (PK_H) , (IC_H) . This lowers the right-hand side of (IC_L) , so allows us to lower the objective $U(l)$.

Either $p_h = 1$ or $u_h = (0, 0)$. If not, we can replace p_h, u_h with

$$p_h = p_h + \frac{\delta(h(1 - \rho_h) + l\rho_h)}{(1 - \delta)h^2}\varepsilon, \quad u_h(h) = u_h(h) - \varepsilon, \quad u_h(l) = u_h(l) - \frac{l}{h}\varepsilon.$$

This change has no impact on (PK_H) , (IC_H) , but lowers the right-hand side of (IC_L) . Hence, we can further lower the objective $U(l)$. We can replace u_h with the new utility pair because we have shown that $SE'(\cdot)$ is at least $\frac{l}{h}$.

These observations show that the utility pair $(x, SE(x))$ is delivered by a policy in which (i) the current low type is indifferent between reporting low and high, (ii) the continuation utility pair after a high report is still on the lower boundary $SE(\cdot)$, (iii) p_h is as high as possible. Combining (i) and (ii), we conclude that the current low type's utility can be calculated as if he always reports high. According to (iii), the good is provided as long as possible after high reports. This shows that the lower boundary $SE(\cdot)$ is characterized by the frontloading policies.

Using a similar argument, we next show that $NW(\cdot)$ is given by the backloading policies. For any $x \in (0, \mu_l)$, $NW^{-1}(x)$ is given by the following program:

$$\begin{aligned} \min_{p_h, p_l \in [0, 1], u_h, u_l \in V} \quad & U(h) \\ \text{s.t.}, \quad & U(h) = (1 - \delta)p_h h + \delta((1 - \rho_h)u_h(h) + \rho_h u_h(l)) \\ & \geq (1 - \delta)p_l h + \delta((1 - \rho_h)u_l(h) + \rho_h u_l(l)) \\ & x = (1 - \delta)p_l l + \delta(\rho_l u_l(h) + (1 - \rho_l)u_l(l)) \\ & \geq (1 - \delta)p_h l + \delta(\rho_l u_h(h) + (1 - \rho_l)u_h(l)). \end{aligned}$$

Because we assume that $x \in (0, \mu_l)$, the minimal $U(h)$ must be strictly positive. It follows that in the solution either $p_h > 0$ or $u_h > (0, 0)$. We again make three observations:

(IC_H) must bind. If not, we can reduce p_h or u_h . This has no impact on (PK_L) and makes (IC_L) easier to satisfy. This reduction will lower the objective $U(h)$.

u_l must be on the upper boundary $NW(\cdot)$. If not, we can replace u_l with

$$\left(u_l(h) - \varepsilon, u_l(l) + \varepsilon \frac{\rho_l}{1 - \rho_l}\right)$$

for small $\varepsilon > 0$. This has no impact on the promised utility to the current low type, and thus no impact on (PK_L) , (IC_L) . This lowers the right-hand side of (IC_H) , so allows us to lower the objective $U(h)$.

Either $p_l = 0$ or $u_l = (\mu_h, \mu_l)$. If not, we can replace p_l, u_l with

$$p_l = p_l - \frac{\delta(h\rho_l + l(1 - \rho_l))}{(1 - \delta)hl}\varepsilon, \quad u_l(h) = u_l(h) + \varepsilon, \quad u_l(l) = u_l(l) + \frac{l}{h}\varepsilon.$$

This change has no impact on (PK_L) , (IC_L) , but lowers the right-hand side of (IC_H) . Hence, we can further lower the objective $U(h)$. We can replace u_l with the new utility pair because we have shown that $NW'(\cdot)$ is at least $\frac{l}{h}$.

These observations show that the utility pair $(NW^{-1}(x), x)$ is delivered by a policy in which (i) the current high type is indifferent between reporting low and high, (ii) the continuation utility pair after a low report is still on the upper boundary $NW(\cdot)$, (iii) p_l is as low as possible. This shows that the upper boundary $NW(\cdot)$ is characterized by the backloading policies. ■

Proof of Lemma 6. For $A \subset V$, let $\mathcal{B}_0(A)$ be the promises delivered by $p_l = 0$, $p_h \in [0, 1]$, and $u_h, u_l \in A$ satisfying (PK_H) – (IC_L) . Let V_0 be the compact set of promises implemented by policies that never supply the low type. Equivalently, it is the largest set with $V_0 \subset \mathcal{B}_0(V_0)$.

If $\underline{U} = 0$, the excluded equality and (3) imply $\delta(1 - \rho_l + \rho_l h/l) < 1$. Let $M = \sup_{U \in V_0} U(l)$. The lower boundary of V gives $U(h) \leq (h/l)U(l)$. Since $p_l = 0$, promise keeping implies

$$M \leq \delta(1 - \rho_l + \rho_l h/l)M.$$

Hence $M = 0$, so $V_0 = \{0\} = \underline{V}$. The zero promise is implemented by never supplying the good. We can therefore assume $\underline{U} > 0$ in the construction below.

We first show that $\underline{V} \subset V_0$. By (4), the sequence $\{v^m\}_{m \geq 0}$ solves the system of equations, for all $m \geq 0$:

$$v_h^{m+1} = \delta(1 - \rho_h)v_h^m + \delta\rho_h v_l^m, \quad v_l^{m+1} = \delta(1 - \rho_l)v_l^m + \delta\rho_l v_h^m,$$

and $v_l^1 = v_l^0$ (which is easily verified given the conditions that $v^0 = \underline{U}$ and that $\underline{U}$ lies on the line $U(l) = \frac{\delta\rho_l}{1 - \delta(1 - \rho_l)}U(h)$). First, for any $m \geq 1$, v^m can be supported by setting $p_h = p_l = 0$ and $u_h = u_l = v^{m-1}$. Therefore, v^m can be delivered with p_l being 0 and continuation utilities in $\underline{V}$. Second, we show that v^0 can itself be obtained with continuation utilities in $\underline{V}$. This one is obtained by setting $(p_h, p_l) = (1, 0)$, setting (IC_L) as a binding constraint, and $u_l = v^0$ (again one can check that u_h is in $\underline{V}$ and that (IC_H) holds). Third, for any $U \in P_b \cap \underline{V}$,

U can be delivered with p_l being 0 and continuation utilities on $P_b \cap \underline{V}$. This shows that $\underline{V} \subseteq V_0$, because the extreme points of $\underline{V}$ can be supported with p_l being 0 and continuation utilities in $\underline{V}$, and all other utility vectors in $\underline{V}$ can be written as a convex combination of these extreme points.

For the reverse inclusion, let $y = \max_{U \in V_0} U(l)$. PK_L and the lower boundary of V imply

$$y \leq \delta[\rho_l S E^{-1}(y) + (1 - \rho_l)y].$$

The definition of $\underline{U}$ makes $\underline{U}_l$ the largest y satisfying this inequality. Hence $U(l) \leq v_l^0 = v_l^1$ on V_0 . For $m \geq 1$, set

$$s_m = \frac{v_l^{m-1} - v_l^m}{v_h^{m-1} - v_h^m} = \frac{q(1 - \kappa^{m-1})}{q + (1 - q)\kappa^{m-1}}, \quad \lambda_m = (-s_m, 1).$$

In particular, $s_1 = 0$ and $0 \leq s_m < 1$. Let A be the discounted transition matrix, so $v^m = A v^{m-1}$. Direct multiplication gives, for $m \geq 2$,

$$\lambda_m A = \delta(1 - \rho_l - s_m \rho_h) \lambda_{m-1},$$

with a positive coefficient. When $p_l = 0$, PK_L and IC_H imply $U(l) = (A u_l)_l$ and $U(h) \geq (A u_l)_h$. Induction therefore gives, for every $U \in V_0$,

$$\lambda_m U \leq \lambda_m A u_l \leq \delta(1 - \rho_l - s_m \rho_h) \lambda_{m-1} v^{m-1} = \lambda_m v^m = \lambda_m v^{m-1}.$$

These halfspaces, together with V , delimit $\underline{V}$. Thus $V_0 \subset \underline{V}$.

Next, we argue that the complete-information payoff is achieved in $\underline{V}$. It is clear that any policy that never gives the unit to the low type must be optimal. This is a feature of the policies that we have described to obtain the boundary of $\underline{V}$.

For the converse, assume $v_h^* \leq \bar{u}_h^1$ as in the lemma. We show that the complete-information payoff cannot be achieved for $U \notin \underline{V} \cup \{(\mu_h, \mu_l)\}$.

Since $\underline{V}$ is the largest fixed point of $\mathcal{B}_0$, starting from any utility $U \notin \underline{V}$, there is a positive probability that the unit is given (after some history that has positive probability) when the value is low. This implies that the complete-information payoff cannot be achieved for $U \leq v^*$, $U \notin \underline{V}$.

For $U \geq v^*$, achieving the complete-information payoff requires that the agent gets the unit whenever he reports h . This means that the agent can report h and get the unit all the time. The agent's interim utility is effectively (μ_h, μ_l) . Therefore, the complete-information payoff is achieved only at $U = (\mu_h, \mu_l)$.

The case $U(h) > v_h^*$ and $U(l) < v_l^*$ is impossible because v^* lies below the increasing lower boundary of V . We are left with $U \in V$ such that $U(l) > v_l^*$ and $U(h) < v_h^*$. Given that $U(l) > v_l^*$, the complete-information payoff requires that if the agent reports low today, then he gets the unit after the high report from tomorrow on. This implies that the agent will get at least $\bar{u}^1$ as his interim utilities, contradicting the presumption that $U(h) < v_h^*$, since $v_h^* \leq \bar{u}_h^1$. This condition is equivalent to $h(1 - \delta)(1 - \delta\kappa) \leq \delta\rho_h l$, which holds when δ is high enough.

■

Proof of Theorems 2 and 3.

We first examine a relaxed problem in which (IC_H) is dropped. We define the function $\tilde{W} : V \times \{\rho_l, 1 - \rho_h\} \rightarrow \mathbf{R} \cup \{-\infty\}$, that solves the following program: for any $U \in V, \phi \in \{\rho_l, 1 - \rho_h\}$,

$$\begin{aligned} \tilde{W}(U, \phi) = & \sup_{p_h, p_l, u_h, u_l} \phi \left((1 - \delta)p_h(h - c) + \delta\tilde{W}(u_h, 1 - \rho_h) \right) \\ & + (1 - \phi) \left((1 - \delta)p_l(l - c) + \delta\tilde{W}(u_l, \rho_l) \right), \end{aligned}$$

over $p_h, p_l \in [0, 1]$ and $u_h, u_l \in V$, subject to (PK_H) , (PK_L) , (IC_L) . Given that (IC_H) is dropped, this is a relaxed version of the original problem. It holds that $\tilde{W}(U, \phi) \geq W(U, \phi)$ for any $U \in V, \phi \in \{\rho_l, 1 - \rho_h\}$.

We want to show that the value function $\tilde{W}$ satisfies two properties (I) and (II):

(I) $\tilde{W}(U(h), U(l), 1 - \rho_h)$ is weakly increasing in $U(h)$ along the lines $x = (1 - \rho_h)U(h) + \rho_h U(l)$, and $\tilde{W}(U(h), U(l), \rho_l)$ is weakly increasing in $U(h)$ along the lines $y = \rho_l U(h) + (1 - \rho_l)U(l)$; (II) Given $\tilde{W}$, we define two functions $\tilde{w}_h(x)$ and $\tilde{w}_l(y)$. Let $(U_{h1}(x), U_{l1}(x))$ be the intersection of P_b and the line $x = (1 - \rho_h)U(h) + \rho_h U(l)$. We define $\tilde{w}_h(x) := \tilde{W}(U_{h1}(x), U_{l1}(x), 1 - \rho_h)$ on the domain $[0, (1 - \rho_h)\mu_h + \rho_h\mu_l]$. Similarly, $(U_{h2}(y), U_{l2}(y))$ denotes the intersection of P_b and the line $y = \rho_l U(h) + (1 - \rho_l)U(l)$. We define $\tilde{w}_l(y) := \tilde{W}(U_{h2}(y), U_{l2}(y), \rho_l)$ on the domain $[0, \rho_l\mu_h + (1 - \rho_l)\mu_l]$. For any U , let $X(U) = (1 - \rho_h)U(h) + \rho_h U(l)$ and $Y(U) = \rho_l U(h) + (1 - \rho_l)U(l)$. We want to show that (II.a) $\tilde{w}_h(\cdot)$ and $\tilde{w}_l(\cdot)$ are concave; (II.b) $\tilde{w}'_h, \tilde{w}'_l$ are in $[1 - c/l, 1 - c/h]$ (both one-sided derivatives satisfy these bounds); and (II.c) for any U on P_b

$$\tilde{w}'_{h,-}(X(U)) \geq \tilde{w}'_{l,+}(Y(U)). \quad (16)$$

If $\tilde{W}$ satisfies (I) and (II), then it follows that for the relaxed problem, for any U :

It is optimal to choose u_h and u_l that lie on P_b . First, for any $U \in V$, it is feasible to choose a u_h on P_b since the intersection of (IC_L) and (PK_H) is above P_b . It is also feasible to

choose a u_l on P_b since (IC_H) is dropped. Given property (I), the principal prefers to have u_h, u_l as close to P_b as possible.

It is optimal to choose p_h, p_l as in (6), *i.e.*, to choose p_h and p_l as efficiently as possible. This is due to property (II.b) that $\tilde{w}'_h, \tilde{w}'_l$ are in $[1 - c/l, 1 - c/h]$.

Next, we show that $\tilde{W}$ satisfies (I) and (II) by an argument of value function iterations. Let $\mathcal{W}$ be the set of all functions $\hat{W}(U, 1 - \rho_h)$ and $\hat{W}(U, \rho_l)$ that satisfy (I) and (II). More specifically,

$\hat{W}(U(h), U(l), 1 - \rho_h)$ is weakly increasing in $U(h)$ along the lines $x = (1 - \rho_h)U(h) + \rho_h U(l)$, and $\hat{W}(U(h), U(l), \rho_l)$ is weakly increasing in $U(h)$ along the lines $y = \rho_l U(h) + (1 - \rho_l)U(l)$;

We define $\hat{w}_h$ and $\hat{w}_l$ given $\hat{W}(U, 1 - \rho_h)$ and $\hat{W}(U, \rho_l)$ in the same way as in (II). The functions $\hat{w}_h$ and $\hat{w}_l$ satisfy (II.a)–(II.c).

Start value iteration from the complete-information upper bound $\bar{W}$. In addition to (I) and (II), retain agreement with $\bar{W}$ on $\underline{V}$, the endpoint values, and endpoint slopes $1 - c/h$ at zero and $1 - c/l$ at the maximal promise. These properties hold initially. Agreement on $\underline{V}$ is preserved because the policies used in Lemma 6 achieve $\bar{W}$ and remain in $\underline{V}$; the Bellman operator cannot exceed the complete-information bound.

Let $\hat{W}$ be a continuation value with these properties and $\hat{W}$ its Bellman update. On a segment of P_b , write $U(l) = sU(h) + b$, where $s > 0$. For $U \notin \underline{V}$, define

$$A(U) = \begin{cases} 1 - c/h, & U \in V_l, \\ \hat{w}'_h((U(h) - (1 - \delta)h)/\delta), & U \notin V_l, \end{cases}$$

and

$$B(U) = \begin{cases} 1 - c/l, & U \in V_h, \\ \hat{w}'_l(U(l)/\delta), & U \notin V_h. \end{cases}$$

Derivatives in these formulas are taken where they exist; the one-sided limits give the assertions at the remaining points. These expressions include the fractional-allocation regions. In V_l , $u_h = 0$ and the high-type contribution to the current payoff is $(1 - c/h)U(h)$. In V_h , $u_l = (\mu_h, \mu_l)$ and the marginal low-type contribution is $1 - c/l$. Elsewhere, the high and low continuation pairs lie on P_b , with $u_h < u_l$. Choose a point of P_b strictly between u_h and u_l . Concavity and (16) at that point order every one-sided high-continuation slope above every low-continuation slope. Thus

$$1 - c/l \leq B(U) \leq A(U) \leq 1 - c/h.$$

Both A and B are nonincreasing along P_b on each allocation region. They agree across its endpoints because the continuation-value slopes at zero and at the maximal promise equal the adjacent current-allocation slopes.

For $\phi \in \{\rho_l, 1 - \rho_h\}$, parameterize P_b by $z = \phi U(h) + (1 - \phi)U(l)$. The derivative of the updated value along this segment is

$$\frac{\phi}{\phi + (1 - \phi)s} A(U) + \frac{(1 - \phi)s}{\phi + (1 - \phi)s} B(U).$$

It lies in $[1 - c/l, 1 - c/h]$. The weight on A is increasing in ϕ , which proves the updated marginal-value ordering. Within a segment, both terms decrease with U . At a vertex, the slope s of the convex lower boundary increases, so the weight on A decreases. Since $A \geq B$, the derivative cannot jump upward. This proves concavity across the vertices as well as within each segment, using one-sided derivatives only. On $P_b \cap \underline{V}$, the value equals the linear complete-information payoff with slope $1 - c/h$; the continuation slopes at the exit point have the same value. Continuity handles the accumulation point at (μ_h, μ_l) . The endpoint slopes are preserved by the same formulas. Thus (II) is preserved on all of P_b .

To verify (I), move U along a line of constant $\phi U(h) + (1 - \phi)U(l)$ toward larger $U(h)$. In the interior allocation region the directional derivative of $\hat{W}$ has the sign of

$$\hat{w}'_h((U(h) - (1 - \delta)h)/\delta) - \hat{w}'_l(U(l)/\delta).$$

For fixed $U(l)$, its smallest value occurs at the largest feasible $U(h)$, which lies on P_b . There it is nonnegative by the ordering just proved. In V_l , replace the first term by $1 - c/h$; in V_h , replace the second by $1 - c/l$. The slope bounds again give the required sign. On $\underline{V}$, the complete-information payoff is constant along the line. Continuity joins these regions and proves (I).

The Bellman operator is a contraction on bounded functions on the compact feasible set. Its iterates converge uniformly to $\tilde{W}$. Concavity, the chord-slope bounds, and the monotonicity inequalities pass to the limit. The ordered one-sided slopes in (II.c) also pass to the limit, by first comparing difference quotients on disjoint neighboring intervals and then shrinking those intervals. Hence $\tilde{W}$ satisfies (I) and (II). Its endpoint slopes can also be checked directly, rather than inferred from uniform convergence. At zero, the low-belief right slope D_l satisfies $D_l = \alpha(1 - c/h) + (1 - \alpha)D_l$, with $\alpha = \rho_l/(\rho_l + (1 - \rho_l)l/h) > 0$. Thus both endpoint slopes there equal $1 - c/h$. At the maximal promise, the boundary slope tends to one and the high-belief left slope satisfies $D_h = (1 - \rho_h)D_h + \rho_h(1 - c/l)$. Thus both endpoint slopes there equal $1 - c/l$. This shows that the value function for the relaxed

problem $\tilde{W}$ is in $\mathcal{W}$. If the initial promise is on P_b , then the continuation utility never leaves P_b . Moreover, for $U \in P_b$, (IC_H) is satisfied under the optimal policy of the relaxed problem. This implies that $W = \tilde{W}$ for any $U \in P_b$.

We turn to the original optimization problem by adding back the constraint (IC_H) . The first observation is that we can decompose the optimization problem into two subproblems: (i) choose p_h, u_h to maximize $(1 - \delta)p_h(h - c) + \delta W(u_h, 1 - \rho_h)$ subject to (PK_H) and (IC_L) ; (ii) choose p_l, u_l to maximize $(1 - \delta)p_l(l - c) + \delta W(u_l, \rho_l)$ subject to (PK_L) and (IC_H) .

We want to show that the policy in Definition 2 with respect to p_h, u_h is the optimal solution to the first subproblem. First, we have shown at the beginning of this proof that it is feasible to choose u_h on P_b so that (PK_H) and (IC_L) are satisfied. If we take $\tilde{W}$ as the continuation value function, the policy in Definition 2 is optimal. Given that (i) $\tilde{W}$ is weakly higher than W pointwise, and (ii) $\tilde{W}$ coincides with W on P_b , the policy in Definition 2 with respect to p_h, u_h solves the first subproblem.

For the second subproblem, choose p_l as low as promise keeping allows and $u_l \in P_b$. If $p_l = 0$, PK_L and the identity following (5) show that $(\delta \mathbf{E}_h[u_l], U(l)) \in P$. Thus IC_H holds below P , binds on P , and fails above it. When $p_l > 0$, the choice $u_l = (\mu_h, \mu_l)$ gives the same conclusion by the upper endpoint segment of V . The relaxed policy is therefore feasible throughout V_b and stays on P_b thereafter, proving $W = \tilde{W}$ on V_b .

We can now establish (I) for W without presuming it. On a line of constant $\phi U(h) + (1 - \phi)U(l)$, with $\phi \in \{\rho_l, 1 - \rho_h\}$, W is concave. Near the endpoint on P_b , it equals $\tilde{W}$ and is nondecreasing toward that endpoint. Concavity extends this monotonicity to the whole line. Consequently, for $U \in V_t$, moving u_l toward P_b along PK_L weakly improves the objective until IC_H binds. We may therefore impose equality in IC_H . It remains to show that p_l should be as low as possible.

For a fixed $U \in V_t$, (PK_L) and a binding (IC_H) determine $u_l(h), u_l(l)$ as a function of p_l . Let γ_h, γ_l denote the derivatives of $u_l(h), u_l(l)$ with respect to p_l

$$\gamma_h = \frac{(1 - \delta)(l\rho_h - h(1 - \rho_l))}{\delta(1 - \rho_h - \rho_l)}, \quad \gamma_l = \frac{(1 - \delta)(h\rho_l - l(1 - \rho_h))}{\delta(1 - \rho_h - \rho_l)}.$$

Given that $\rho_l + \rho_h < 1$, it follows that $\gamma_h < 0$. We want to show that it is optimal to set p_l as low as possible. That is, the principal's payoff from the low type, $(1 - \delta)p_l(l - c) + \delta W(u_l, \rho_l)$, decreases in p_l . The principal's payoff decreases in p_l if and only if:

$$\gamma_h \frac{\partial W(U, \rho_l)}{\partial U(h)} + \gamma_l \frac{\partial W(U, \rho_l)}{\partial U(l)} \leq \frac{(1 - \delta)(c - l)}{\delta}, \quad (17)$$

where the left-hand side is the directional derivative of $W(U, \rho_l)$ along the vector (γ_h, γ_l) . We first show that (17) holds for all $U \in P_b \setminus (V_h \cup V_l)$. For any fixed $U \in P_b \setminus (V_h \cup V_l)$, we

have

$$W(U, \rho_l) = \rho_l \left((1 - \delta)(h - c) + \delta \tilde{w}_h \left(\frac{U(h) - (1 - \delta)h}{\delta} \right) \right) + (1 - \rho_l) \delta \tilde{w}_l \left(\frac{U(l)}{\delta} \right).$$

It follows that $\partial W / \partial U(h) = \rho_l \tilde{w}'_h$ and $\partial W / \partial U(l) = (1 - \rho_l) \tilde{w}'_l$. We have shown that $\tilde{w}'_h \geq \tilde{w}'_l$ and $\tilde{w}'_h, \tilde{w}'_l \in [1 - c/l, 1 - c/h]$. Given the linearity of (17) in $\tilde{w}'_h$ and $\tilde{w}'_l$, it is sufficient to show that (17) holds when either (i) $\tilde{w}'_l = \tilde{w}'_h$, or (ii) $\tilde{w}'_l = 1 - c/l$. In case (i), the left-hand side of (17) equals $-(1 - \delta)l\tilde{w}'_h/\delta$, so (17) holds for any $\tilde{w}'_h \in [1 - c/l, 1 - c/h]$. In case (ii), the left-hand side of (17) decreases in $\tilde{w}'_h$ and is maximized when $\tilde{w}'_h = 1 - c/l$. Substituting $\tilde{w}'_l = \tilde{w}'_h = 1 - c/l$, we obtain that (17) holds. In V_l , replace the high-continuation slope by $A = 1 - c/h$; in V_h , replace the low-continuation slope by $B = 1 - c/l$. The same linear inequality applies to every $1 - c/l \leq B \leq A \leq 1 - c/h$, so it also covers both fractional-provision regions and their one-sided limits.

The formulas for γ_h, γ_l give $\gamma_h < 0$ and $\gamma_l/\gamma_h < l/h$. Since both boundary slopes are at least l/h , the ray from $U \in V_t$ in direction $-(\gamma_h, \gamma_l)$ meets P_b before it can meet P_t . Restrict $W(\cdot, \rho_l)$ to the line from this intersection toward U . Concavity makes its one-sided directional derivative nonincreasing. The boundary bound (17) therefore holds throughout the line inside V . This proves optimality of the policy in Definition 2. The two separate maximizations do not involve the current belief, so the policy is optimal for every $\phi \in [0, 1]$.

An optimal initial promise on P_b must exceed $\underline{U}$. If $\underline{U} = 0$, this follows from the positive endpoint slopes. Otherwise, the high-report continuation at $\underline{U}$ lies below $\underline{U}$, and the low-report continuation is $\underline{U}$. Let s be the right slope of P_b there and $D = \tilde{w}'_{l,+}(Y(\underline{U}))$. Taking right derivatives of the Bellman equation gives

$$D = \alpha(1 - c/h) + (1 - \alpha)D, \quad \alpha = \frac{\rho_l}{\rho_l + (1 - \rho_l)s} > 0.$$

Thus $D = 1 - c/h > 0$. The high-belief right derivative has the same value. Every prior in $[\rho_l, 1 - \rho_h]$ therefore gives a payoff that increases just above $\underline{U}$; below it the payoff is the increasing complete-information payoff.

We next show that an optimal initial promise can be chosen on P_b and weakly increasing in the prior belief of the high type. For any $x \in [0, (1 - \rho_h)\mu_h + \rho_h\mu_l]$, let $z(x)$ be $\rho_l U_{h1}(x) + (1 - \rho_l)U_{l1}(x)$. The function $z(x)$ is piecewise linear with z' being positive and increasing in x . Because the two subproblems above are separate and their constraints do not involve the belief, $W(U, \phi)$ is affine in ϕ . Hence, for a prior belief $\phi = \phi_0(1 - \rho_h) + (1 - \phi_0)\rho_l$ with $\phi_0 \in [0, 1]$, the principal's payoff is $\phi_0 W(U, 1 - \rho_h) + (1 - \phi_0)W(U, \rho_l)$; these beliefs include the invariant probability q . We want to show that the maximum of this payoff is achieved on P_b for any ϕ_0 .

Suppose no point of P_b achieves the maximum, and let $\tilde{U} \in V \setminus P_b$ be a maximizer. Let U^0 (or U^1) denote the intersection of P_b and $(1 - \rho_h)U(h) + \rho_h U(l) = (1 - \rho_h)\tilde{U}(h) + \rho_h \tilde{U}(l)$ (or $\rho_l U(h) + (1 - \rho_l)U(l) = \rho_l \tilde{U}(h) + (1 - \rho_l)\tilde{U}(l)$). Given that $\rho_l < 1 - \rho_h$, it follows that $U^0 < U^1$. Given that $\tilde{U}$ achieves the maximum, it must be true that

$$\begin{aligned} W(U^1, 1 - \rho_h) - W(U^0, 1 - \rho_h) &< 0, \\ W(U^1, \rho_l) - W(U^0, \rho_l) &> 0. \end{aligned}$$

We show that this is impossible by arguing that for any $U^0, U^1 \in P_b$ and $U^0 < U^1$, $W(U^1, 1 - \rho_h) - W(U^0, 1 - \rho_h) < 0$ implies that $W(U^1, \rho_l) - W(U^0, \rho_l) < 0$. We have

$$\begin{aligned} W(U^1, 1 - \rho_h) - W(U^0, 1 - \rho_h) &= \int_{s^0}^{s^1} \tilde{w}'_h(s) ds, \\ W(U^1, \rho_l) - W(U^0, \rho_l) &= \int_{s^0}^{s^1} \tilde{w}'_l(z(s)) z'(s) ds, \end{aligned}$$

where $s^0 = (1 - \rho_h)U^0(h) + \rho_h U^0(l)$, $s^1 = (1 - \rho_h)U^1(h) + \rho_h U^1(l)$. Let s^* maximize $\tilde{w}_h$ on $[s^0, s^1]$. Because $\tilde{w}_h$ is concave, $\tilde{w}'_h \geq 0$ on $[s^0, s^*)$ and $\tilde{w}'_h \leq 0$ on $(s^*, s^1]$. Given that $\tilde{w}'_h(s) \geq \tilde{w}'_l(z(s))$ and that z' is positive and increasing,

$$\int_{s^0}^{s^1} \tilde{w}'_l(z(s)) z'(s) ds \leq \int_{s^0}^{s^1} \tilde{w}'_h(s) z'(s) ds \leq z'(s^*) \int_{s^0}^{s^1} \tilde{w}'_h(s) ds < 0.$$

This shows the optimal initial promise is on P_b .

The optimal U^* is chosen such that $X(U^*)$ maximizes $\phi_0 \tilde{w}_h(x) + (1 - \phi_0) \tilde{w}_l(z(x))$. Let x_l denote the largest maximizer of $\tilde{w}_l(z(x))$. For $x < x_l$, (16) implies $\tilde{w}'_h(x) \geq \tilde{w}'_l(z(x)) \geq 0$, so that the objective is weakly increasing. For $x > x_l$, $\tilde{w}'_l(z(x))$ is negative and decreasing, and z' is positive and increasing, so that $\tilde{w}'_l(z(x)) z'(x)$ is decreasing; hence, the objective is concave for $x \geq x_l$. We can therefore take $X(U^*) \geq x_l$, and, using left derivatives, at $x = X(U^*)$,

$$\phi_0 \tilde{w}'_h(X(U^*)) + (1 - \phi_0) \tilde{w}'_l(z(X(U^*))) z'(X(U^*)) \geq 0.$$

If $X(U^*) > x_l$, then $\tilde{w}'_l(z(X(U^*))) \leq 0$, so that $\tilde{w}'_h(X(U^*)) \geq 0$; if $X(U^*) = x_l$, both terms are weakly positive. In either case, the derivative above is weakly positive for any $\phi'_0 > \phi_0$. Choosing the largest maximizer at each belief therefore gives an optimal initial promise that is weakly increasing in ϕ_0 . ■

B.1 Transfers with Limited Liability

Here, we consider the case in which values are i.i.d. and only the principal can pay the agent. The principal maximizes his payoff net of payments. The following lemma compares transfers with inefficient provision.

Lemma 11 *An optimal policy without transfers exists if and only if $c \leq 2l$. At $c = 2l$, optimal policies with transfers also exist.*

Proof. We first show that the principal makes transfers if $c - l > l$. Suppose not. The optimal mechanism is the same as the one characterized in Theorem 1. When U is sufficiently close to μ , we want to show that it is “cheaper” to provide incentives using transfers. Given the optimal allocation (p_h, u_h) and (p_l, u_l) , if we reduce u_l by ε and make a transfer of $\delta\varepsilon/(1 - \delta)$ to the low type, the IC_L/PK constraints are satisfied. When u_l is sufficiently close to μ , the principal’s payoff increment is close to $\delta(c/l - 1)\varepsilon - \delta\varepsilon = \delta(c/l - 2)\varepsilon$, which is strictly positive if $c - l > l$. This contradicts the fact that the allocation (p_h, u_h) and (p_l, u_l) is optimal. Therefore, the principal makes transfers if $c - l > l$.

Now suppose $c \leq 2l$. Extend the no-transfer value to $U \geq \mu$ by

$$\widehat{W}(U) = \begin{cases} W(U), & 0 \leq U \leq \mu, \\ \mu - c - (U - \mu), & U \geq \mu. \end{cases}$$

This is concave, and all its chord slopes lie in $[-1, 1 - c/h]$. For $U \geq \mu$, it is the complete-information upper bound and is attained by always supplying the good and paying $U - \mu$ per round.

We verify the Bellman inequality for $U \leq \mu$. A current transfer t_v can be replaced by increasing u_v by $(1 - \delta)t_v/\delta$. This preserves every report’s utility and weakly raises the objective evaluated at $\widehat{W}$, since its slopes are at least -1 . Incentive compatibility implies $p_h \geq p_l$. Retain this restriction and relax IC_H . Concavity allows IC_L to bind by moving the continuation utilities toward each other while preserving their expectation. The formulas (10)–(11) then apply. The upper chord-slope bound permits increasing p_h until $p_h = 1$ or $u_h = 0$, by the argument of Theorem 1. In either case $u_h \leq \mu$. If $u_l > \mu$, increase p_l until $u_l = \mu$. This is feasible because $u_h \leq \mu$ and $u_l - u_h = (1 - \delta)l(p_h - p_l)/\delta$. On this portion the objective’s derivative with respect to p_l is

$$(1 - \delta)(1 - q)(2l - c) \geq 0.$$

Both continuations now lie in $[0, \mu]$. The maximization is the relaxed problem verified in Theorem 1, and its optimum satisfies IC_H . Thus $\widehat{W}$ is an upper bound for the problem with transfers, attained without transfers when $U \leq \mu$. The verification extends to the infinite horizon because finite promised utility and nonnegative flow utility imply that discounted expected continuation utility tends to zero.

At $c = 2l$, replace permanent provision at the promise μ by supplying after a high report and paying l after a low report. Both reports retain the same utility, and the payoff change is $(1 - q)(c - 2l) = 0$. Since μ is reached with positive probability under the optimal no-transfer policy, this gives an optimal policy with transfers on path. ■

C Proofs for Section 5

We first justify the discrete-time limit on the lower boundary. For a round length Δ , parameterize the interpolated frontloading promises by $\tau = (m + \lambda)\Delta$. Their utility pairs converge uniformly on compact intervals to $\underline{u}(\tau)$. Away from zero and infinity, a high report changes τ by $-\Delta$, and a low report changes it by

$$\Delta \frac{g(\tau)}{\mu - q(h - l)e^{-(\lambda_h + \lambda_l)\tau}} + O(\Delta^2),$$

uniformly on compact intervals. The latter follows by expanding PK_L along the interpolated frontloading curve; its limiting low-type derivative is $re^{-r\tau}[\mu - q(h - l)e^{-(\lambda_h + \lambda_l)\tau}] > 0$. Couple the discrete type chains so their geometric holding times, multiplied by Δ , converge to the exponential holding times of the continuous chain. There are almost surely finitely many switches on each finite time interval. Between switches, the preceding updates converge to the stated differential equations by the Euler approximation for locally Lipschitz drifts. The discounted provision payoffs therefore converge until exit from $[\varepsilon, R]$ or a fixed time T .

The stopping restrictions can be removed uniformly. At $\tau \leq \varepsilon$, the remaining expected discounted quantity is at most $U_v/l = O(\varepsilon + \Delta)$ for either current type. At $\tau \geq R$, its shortfall from permanent provision is at most $(\mu_v - U_v)/l = O(e^{-rR} + \Delta)$. The corresponding payoff errors obey the same bounds, since $|v - c|$ is bounded. The payoff beyond time T is bounded by a constant times e^{-rT} . Letting $\Delta \rightarrow 0$, then $\varepsilon \downarrow 0$ and $R, T \rightarrow \infty$, proves convergence of the optimal discrete-time payoffs on the lower boundary, uniformly on compact initial durations. The endpoint bounds also imply convergence of their maxima. This establishes the limiting interpretation used in Section 5.1.

Based on (7) and (8), we directly work with the expected payoff $W(\tau) = qW_h(\tau) + (1 -$

$q)W_l(\tau)$. Using the system of differential equations, we get

$$\begin{aligned} & g(\tau)((r + \lambda_h)W'(\tau) + W''(\tau)) \\ & = (h - l)q\lambda_h e^{-(\lambda_h + \lambda_l)\tau} W'(\tau) + \mu(r(\lambda_h + \lambda_l)W(\tau) + \lambda_l W'(\tau) - r\lambda_l(h - c)). \end{aligned} \quad (18)$$

Write $\Lambda = \lambda_h + \lambda_l$ and define

$$\bar{W}_1(\tau) = (1 - e^{-r\tau})(1 - c/h)\mu, \quad \chi(\tau) = \mu e^{-r\tau} - (1 - q)l, \quad k(\tau) = g(\tau)/\Lambda.$$

The unique zero τ_0 of χ exceeds $\hat{\tau}$: $\chi(\hat{\tau}) = ql + q(h - l)e^{-(r+\Lambda)\hat{\tau}} > 0$. For $\tau > \hat{\tau}$, set

$$\psi(\tau) = r - \frac{e^{r\tau}\chi(\tau)}{k(\tau)}, \quad E(\tau) = \exp\left(-\int_{\tau_0}^{\tau} \psi(s)ds\right), \quad D = \int_{\hat{\tau}}^{\infty} \frac{\Lambda e^{2rt}E(t)}{g(t)}dt.$$

Proposition 2 *The limiting payoff equals $\bar{W}_1(\tau)$ for $0 \leq \tau \leq \hat{\tau}$. For $\tau > \hat{\tau}$, it is*

$$W(\tau) = \bar{W}_1(\tau) + \frac{c(h - l)}{hlD} \left[\chi(\tau) \int_{\hat{\tau}}^{\tau} \frac{\Lambda e^{2rt}E(t)}{g(t)}dt - e^{r\tau}E(\tau) \right].$$

Proof. Substituting $w = W - \bar{W}_1$ into (18) gives

$$k(\tau)(w''(\tau) + rw'(\tau)) = r\mu w(\tau) + e^{r\tau}\chi(\tau)w'(\tau), \quad (19)$$

with $w(\hat{\tau}) = 0$ and $w(\infty) = -(1 - l/h)(1 - q)c$. Since $\chi' = -r\mu e^{-r\tau}$, the quantity $\Phi = r\mu w + e^{r\tau}\chi w'$ satisfies $\Phi' = (r - \psi)\Phi$. Thus $\Phi = Ce^{r\tau}E(\tau)$.

Let $I(\tau)$ denote the integral from $\hat{\tau}$ to τ in the statement. Since $(e^{r\tau}E)' = \chi I'$, differentiating the proposed expression gives $w' = c(h - l)\chi' I/(hlD)$ and

$$\Phi(\tau) = -\frac{c(h - l)r\mu}{hlD}e^{r\tau}E(\tau).$$

It therefore solves (19) throughout $(\hat{\tau}, \infty)$, including at τ_0 . At the left endpoint, $E \rightarrow 0$ and $I \rightarrow 0$. Indeed, g has a simple zero with positive derivative at $\hat{\tau}$, except when $\hat{\tau} = 0$ and $g'(0) = 0$; in that case it has a double zero. The equation $E'/E = -\psi$ gives a positive power of $\tau - \hat{\tau}$ in the first case, and exponential decay in $1/\tau$ in the second. In both cases E/g is integrable. At infinity, $E(\tau)$ is asymptotic to a positive constant times $e^{-(r+\lambda_h)\tau}$, so $0 < D < \infty$ and $e^{r\tau}E \rightarrow 0$. The two boundary values follow.

There is at most one bounded continuous solution with these boundary values. A nonzero difference of two solutions, vanishing at both endpoints, has a positive interior maximum or a negative interior minimum. At either extremum, (19) contradicts $k > 0$ and $r\mu > 0$. Conditioning the limiting payoff on the next type switch gives the paired integral equations.

On compact subintervals with nonzero drifts, differentiation gives (7)–(8) and hence (19). The limiting payoff is therefore the stated solution. On compact subintervals of $(\hat{\tau}, \infty)$, the paired equations have smooth coefficients and nonzero drift coefficients. They give the asserted twice differentiability of W_h, W_l there. ■

Proof of Lemma 7. Fix $\pi > \tilde{\pi}$ and use tildes for quantities at the lower persistence. Since $\Lambda = (\bar{\lambda}_h + \bar{\lambda}_l)/\pi$,

$$\frac{\partial k(\tau)}{\partial \pi} = \frac{k(\tau) + q(h-l)\tau e^{-\Lambda\tau}}{\pi} > 0 \quad \text{when } \tau > 0 \text{ and } k(\tau) \geq 0.$$

Thus $\hat{\tau}(\pi) \leq \hat{\tau}(\tilde{\pi})$ and $k \geq \tilde{k} > 0$ on $(\hat{\tau}(\tilde{\pi}), \infty)$. The proof of Proposition 2 gives $\Phi < 0$, hence $w'' + rw' < 0$. Because $\bar{W}_1$ and χ do not depend on π , the difference $d = w - \tilde{w}$ satisfies

$$kd'' + (rk - e^{r\tau}\chi)d' - r\mu d = -(k - \tilde{k})(\tilde{w}'' + r\tilde{w}') \geq 0.$$

At $\hat{\tau}(\tilde{\pi})$, $d \leq 0$ since $w \leq 0 = \tilde{w}$, and $d \rightarrow 0$ at infinity. A positive interior maximum would make the left-hand side negative. Thus $d \leq 0$ on this interval. Below $\hat{\tau}(\tilde{\pi})$, $\tilde{w} = 0 \geq w$ as well. This proves the payoff comparison, including when $\hat{\tau} > 0$.

It remains to prove the two limits. Recall that $W \leq \bar{W} \leq \min\{\bar{W}_1, \bar{W}_2\}$, where $\bar{W}_2$ is defined in the proof of Lemma 8 below, and that $\bar{W}_2 - \bar{W}_1 = \frac{c(h-l)}{hl}\chi$.

For the first limit, the function w_- defined in that proof satisfies $w_- \leq w$ on $[\hat{\tau}, \infty)$ for all parameter values. Because $\sqrt{\chi^2 + \varepsilon} - |\chi| \leq \sqrt{\varepsilon}$ for all $\varepsilon > 0$, it follows that, for all $\tau \geq \hat{\tau}$,

$$0 \leq \min\{\bar{W}_1(\tau), \bar{W}_2(\tau)\} - W(\tau) \leq \min\left\{0, \frac{c(h-l)}{hl}\chi(\tau)\right\} - w_-(\tau) \leq \frac{c(h-l)}{4hl} \sqrt{\frac{r\mu l}{\lambda_h + \lambda_l}}.$$

The same bound holds trivially for $\tau < \hat{\tau}$, where $W = \bar{W}_1 \leq \bar{W}_2$. Hence, $\bar{W}(\tau) - W(\tau)$ converges to 0 as $\pi \rightarrow 0$, uniformly in τ .

For the second limit, $W(0) = 0$ and $\lim_{\tau \rightarrow \infty} W(\tau) = \mu - c$ imply $\max_{\tau} W(\tau) \geq \max\{\mu - c, 0\}$. Conversely, $\bar{W}_h(\tau) \leq (1 - c/h)\underline{u}_h(\tau)$ and $\bar{W}_l(\tau) \leq (1 - c/h)v_l^* + (1 - c/l)(\underline{u}_l(\tau) - v_l^*)$, where $v_l^* = h\lambda_l/(r + \lambda_h + \lambda_l)$. Substituting the formulas for $\underline{u}_h$ and $\underline{u}_l$, and using $e^{-r\tau}(1 - e^{-(\lambda_h + \lambda_l)\tau}) \leq (\lambda_h + \lambda_l)/r$, we obtain

$$W(\tau) \leq \bar{W}(\tau) \leq (1 - e^{-r\tau})(\mu - c) + \frac{2cq(1-q)(h-l)}{lr}(\lambda_h + \lambda_l).$$

The first term is at most $\max\{\mu - c, 0\}$, and the last term vanishes as $\pi \rightarrow \infty$, so that $\max_{\tau} W(\tau) \rightarrow \max\{\mu - c, 0\}$. ■

Proof of Lemma 8. We first prove an upper bound on the loss and then a lower bound, both of order $\sqrt{r}$. The complete-information bound gives $W(\tau) \leq \min\{\bar{W}_1(\tau), \bar{W}_2(\tau)\}$, where

$\bar{W}_1$ is defined above and

$$\bar{W}_2(\tau) = (1 - e^{-r\tau})(1 - c/l)\mu + q(h/l - 1)c.$$

These expressions follow by evaluating two of the affine maps whose lower envelope defines $\bar{W}$ at $\underline{u}(\tau)$; see Section 4.2.

We use the notation from the proof of Lemma 7. In particular, $w = W - \bar{W}_1$ satisfies (19) on $(\hat{\tau}, \infty)$, where $k > 0$, with $w(\hat{\tau}) = 0$ and $\lim_{\tau \rightarrow \infty} w(\tau) = -(1 - l/h)(1 - q)c$. The complete-information bounds and the formulas for $\bar{W}_1, \bar{W}_2$ give

$$w(\tau) \leq \min \left\{ 0, \frac{c(h-l)}{hl} \chi(\tau) \right\}, \quad q(h-c) - \bar{W}_1(\tau) = \left(1 - \frac{c}{h}\right) \chi(\tau).$$

Recall that $\chi(\tau_0) = 0$ and $\tau_0 > \hat{\tau}$.

We first establish the upper bound. Define

$$w_-(\tau) = -\frac{c(h-l)}{2hl} \left(\sqrt{\chi(\tau)^2 + \frac{r\mu l}{4(\lambda_h + \lambda_l)}} - \chi(\tau) \right).$$

Since $k(\tau) \leq l(e^{r\tau} - 1)/(\lambda_h + \lambda_l)$,

$$r\mu e^{-2r\tau} k(\tau) \leq \frac{r\mu l}{\lambda_h + \lambda_l} e^{-r\tau} (1 - e^{-r\tau}) \leq \frac{r\mu l}{4(\lambda_h + \lambda_l)}.$$

Using $\chi' = -r\mu e^{-r\tau}$ and $\chi'' + r\chi' = 0$, differentiation gives

$$k(w''_- + rw'_-) - r\mu w_- - e^{r\tau} \chi w'_- \geq 0.$$

Indeed, the left-hand side has the sign of $\chi^2 + \frac{r\mu l}{4(\lambda_h + \lambda_l)} - r\mu e^{-2r\tau} k$, which is nonnegative by the preceding bound. Moreover,

$$w_-(\hat{\tau}) < w(\hat{\tau}) = 0, \quad \lim_{\tau \rightarrow \infty} w_-(\tau) < \lim_{\tau \rightarrow \infty} w(\tau).$$

If $w_- - w$ were positive somewhere on $(\hat{\tau}, \infty)$, these boundary comparisons would imply that it attains a positive maximum at a finite interior point. Subtracting (19) from the inequality for w_- would then give, at that point, $0 \leq k(w_- - w)'' - r\mu(w_- - w) < 0$,

a contradiction. Thus $w_- \leq w$ on $[\hat{\tau}, \infty)$. Evaluating at τ_0 gives

$$0 \leq q(h-c) - \sup_{\tau \geq 0} W(\tau) \leq -w(\tau_0) \leq -w_-(\tau_0) = \frac{c(h-l)}{4hl} \sqrt{\frac{\mu l}{\lambda_h + \lambda_l}} \sqrt{r}.$$

We next establish the lower bound. Define $s = l\sqrt{\frac{q(1-q)}{\lambda_h + \lambda_l}}\sqrt{r}$, and take r small enough that $s \leq \frac{1}{4} \min\{ql, (1-q)l\}$. Since χ is strictly decreasing and

$$\chi(\hat{\tau}) = ql + q(h-l)e^{-(r+\lambda_h+\lambda_l)\hat{\tau}} > ql, \quad \lim_{\tau \rightarrow \infty} \chi(\tau) = -(1-q)l,$$

the interval $|\chi| \leq s$ has two finite endpoints strictly inside $(\hat{\tau}, \infty)$. On this interval, the formula for k gives

$$r\mu e^{-2r\tau}k(\tau) \geq \frac{r}{\lambda_h + \lambda_l}(\chi + (1-q)l)(ql - \chi) \geq \frac{9rq(1-q)l^2}{16(\lambda_h + \lambda_l)} = \frac{9}{16}s^2 \geq \frac{1}{2}s^2.$$

The first inequality discards a nonnegative term, and the second uses $s \leq ql/4$ and $s \leq (1-q)l/4$. On this interval, define $w_+(\tau) = -\frac{c(h-l)}{4hls}(\chi(\tau) - s)^2$. Differentiation yields

$$k(w_+'' + rw_+') - r\mu w_+ - e^{r\tau}\chi w_+' = \frac{c(h-l)r\mu}{4hls} [s^2 - \chi^2 - 2r\mu e^{-2r\tau}k] \leq 0.$$

At $\chi = s$, the affine bounds imply $w \leq 0 = w_+$; at $\chi = -s$, they imply $w \leq -c(h-l)s/(hl) = w_+$. If $w - w_+$ were positive somewhere on this interval, it would therefore attain a positive maximum at an interior point. Subtracting the inequality for w_+ from (19) would give, at that point,

$$0 \leq k(w - w_+)'' - r\mu(w - w_+) < 0,$$

a contradiction. Hence $w \leq w_+$ throughout this interval. Consequently,

$$q(h-c) - W(\tau) \geq \left(1 - \frac{c}{h}\right)\chi + \frac{c(h-l)}{4hls}(\chi - s)^2 \geq \frac{s(h-c)(c-l)}{c(h-l)}.$$

The last inequality follows by minimizing the quadratic, whose minimizer is

$$\chi = s \left(1 - \frac{2l(h-c)}{c(h-l)}\right) \in (-s, s).$$

For $\chi \geq s$, the bound $W \leq \bar{W}_1$ gives a loss of at least $(1-c/h)s$. For $\chi \leq -s$, the bound $W \leq \bar{W}_2$ gives a loss of at least $(c/l-1)s$. Both exceed $s(h-c)(c-l)/(c(h-l))$. The lower bound therefore holds for every $\tau \geq 0$, and hence

$$q(h-c) - \sup_{\tau \geq 0} W(\tau) \geq \frac{l(h-c)(c-l)}{c(h-l)} \sqrt{\frac{q(1-q)}{\lambda_h + \lambda_l}} \sqrt{r}.$$

For sufficiently small r , the upper bound at τ_0 also gives $W(\tau_0) > \mu - c = \lim_{\tau \rightarrow \infty} W(\tau)$, since $q(h-c) > \mu - c$. Continuity therefore implies that the supremum is attained. Combining the two bounds proves that $q(h-c) - \max_{\tau \geq 0} W(\tau) = \Theta(\sqrt{r})$. ■

C.1 Continuous Types

The *minimal expected discounted quantity* is $C(U) = (U - W(U))/c$. It is convex and continuous on $[0, \mu]$, with $C(0) = 0$ and $C(\mu) = 1$. Convexity follows by mixing feasible mechanisms. Write $\bar{C} = (U - \bar{W}(U))/c$ for the *complete-information cost*.

Proof of Lemma 9. For a current expected utility g , the least costly provision rule is $p(v) = \mathbf{1}_{\{v > t\}}$, where $g = \int_t^1 v dF(v)$. An exchange of provision from a lower type to a higher type proves optimality and uniqueness up to null sets. Its cost is $\bar{C}(g) = 1 - F(t)$, with derivative $1/t$ in the interior. Since t decreases continuously with g , $\bar{C}$ is strictly convex. For any dynamic policy, current quantity is at least $\bar{C}(g)$, and Jensen's inequality gives

$$(1 - \delta)\bar{C}(g) + \delta \mathbf{E}\bar{C}(u) \geq \bar{C}((1 - \delta)g + \delta \mathbf{E}u) = \bar{C}(U).$$

Iterating this inequality and letting the bounded discounted terminal cost vanish gives the lower bound for the infinite-horizon problem. Repeating the threshold for $g = U$ attains it. Equality requires $g = u(v) = U$ almost surely and the unique current threshold rule. The conclusions follow from $\bar{W}(U) = U - c\bar{C}(U)$. ■

For private values, the Bellman equation is

$$C(U) = \min_{p,u} \left\{ (1 - \delta) \int p(v) dF(v) + \delta \int C(u(v)) dF(v) \right\}. \quad (20)$$

Write $J(p, u)$ for the *cost of a feasible rule*, the expression on the right-hand side. The constraints are $0 \leq p \leq 1$, $0 \leq u \leq \mu$, promise keeping, and incentive compatibility. For right-continuous representatives, incentive compatibility is equivalent to p nondecreasing and

$$\delta du(v) = -(1 - \delta)v dp(v). \quad (21)$$

All integrals against monotone rules below are Stieltjes integrals. At the upper endpoint we use left limits. Promise keeping is

$$U = \delta u(\underline{v}) + (1 - \delta) \left[\underline{v} p(\underline{v}) + \int_{\underline{v}}^1 (1 - F(v)) p(v) dv \right]. \quad (22)$$

Bounded monotone rules are sequentially compact in $L^1(F)$. Incentive compatibility passes to the limit, and promise keeping and cost are continuous. Hence the minimum is attained.

Define the *dual values* and the *maximizing promises* by

$$\begin{aligned} D(z) &= \max_{0 \leq x \leq \mu} \{zx - C(x)\}, & \bar{D}(z) &= \max_{0 \leq x \leq \mu} \{zx - \bar{C}(x)\}, \\ \mathcal{Q}(z) &= \arg \max_{0 \leq x \leq \mu} \{zx - C(x)\}, & Q(z) &= \min \mathcal{Q}(z). \end{aligned}$$

Each $\mathcal{Q}(z)$ is an interval. Write $d = \delta/(1 - \delta)$.

Lemma 12 *The endpoint derivatives satisfy*

$$W'_+(0) = 1 - c, \quad W'_-(\mu) = \begin{cases} 1 - c/\underline{v}, & \underline{v} > 0, \\ -\infty, & \underline{v} = 0. \end{cases}$$

Proof. The right derivative of C at zero is finite because randomizing between permanent provision and no provision gives $C(U) \leq U/\mu$.

Fix $t < 1$. Supply with probability ε after reports above t and zero otherwise. Set continuation utility to zero above t and to $(1 - \delta)\varepsilon t/\delta$ below it. This rule is incentive compatible. Letting $\varepsilon \downarrow 0$ in its promise and cost gives

$$C'_+(0) \int_t^1 v dF(v) \leq 1 - F(t).$$

Since $C(U) \geq U$, letting $t \uparrow 1$ yields $C'_+(0) = 1$.

For the other endpoint, supply with probability $1 - \varepsilon$ below $t > \underline{v}$ and one above it. Set continuation utility to μ below t and to $\mu - (1 - \delta)\varepsilon t/\delta$ above it. Incentive compatibility again holds. If $C'_-(\mu)$ is finite, the same limiting argument gives

$$C'_-(\mu) \int_{\underline{v}}^t v dF(v) \geq F(t).$$

Thus $C'_-(\mu) \geq 1/\underline{v}$, with an infinite right-hand side when $\underline{v} = 0$. For $\underline{v} > 0$, comparison with the complete-information cost, with equality at μ , gives the reverse inequality. Since $W' = 1 - cC'$, the claimed endpoint derivatives follow. ■

Lemma 13 *For $0 < U < \mu$,*

$$\overline{C}(U) < C(U) < \frac{U}{\mu}, \quad 1 < C'_-(U) \leq C'_+(U) < 1/\underline{v}.$$

The upper derivative bound is interpreted as infinity when $\underline{v} = 0$. For $1 < z < 1/\underline{v}$,

$$D(z) < \overline{D}(z) = \int (zv - 1)_+ dF(v).$$

Proof. The complete-information cost is strictly convex. Equality in

$$(1 - \delta)\mathbf{E}p + \delta\mathbf{E}C(u) \geq (1 - \delta)\overline{C}(\mathbf{E}[vp]) + \delta\mathbf{E}\overline{C}(u) \geq \overline{C}(U)$$

would require constant continuation U and the complete-information threshold rule. Incentive compatibility excludes this combination. Compactness of bounded monotone rules gives

existence of an optimum, so $C(U) > \bar{C}(U)$. Also, $C(U) \leq U/\mu$ by randomization between the two endpoints. Since $C'_+(0) = 1 < 1/\mu$, this inequality is strict near zero. Convexity and $C(\mu) = 1$ extend strictness to every interior promise.

Since $\bar{C}(U) \geq U$, the lower derivative bound follows from $C'_-(U) \geq C(U)/U > 1$. When $\underline{v} > 0$, strict convexity of $\bar{C}$ gives

$$C(U) > 1 - \frac{\mu - U}{\underline{v}}, \quad C'_+(U) \leq \frac{1 - C(U)}{\mu - U} < \frac{1}{\underline{v}}.$$

For a slope between the endpoint derivatives, the maximizer defining $\bar{D}(z)$ is unique and interior. The strict cost comparison therefore gives $D(z) < \bar{D}(z)$. ■

Lemma 14 *At an interior promise, no optimal rule has a constant left marginal cost equal to $C'_-(U)$, or a constant right marginal cost equal to $C'_+(U)$. Every optimal rule satisfies*

$$\mathbf{P}\{u(v) < U\} > 0, \quad \mathbf{P}\{u(v) > U\} > 0.$$

At endpoints, marginal costs denote the one-sided derivatives.

Proof. Suppose the left marginal costs are constant and equal to $z = C'_-(U)$, or the right marginal costs are constant and equal to $z = C'_+(U)$. Lemma 13 gives $1 < z < 1/\underline{v}$. All continuation promises belong almost surely to $\mathcal{Q}(z)$, a compact interval contained in $(0, \mu)$. Let $p^*(v) = \mathbf{1}_{\{zv \geq 1\}}$. For small $h > 0$, set

$$p_h = (1 - h)p + hp^*, \quad u_h(v) = u(v) + h \frac{1 - \delta}{\delta} \left[\pm 1 - \int_{\underline{v}}^v w \, d(p^* - p)(w) \right].$$

Use the minus sign for left derivatives and the plus sign for right derivatives. The integral lies in $[-1, 1]$. Thus continuation promises move only in the direction of the specified derivative. For small h , the rules satisfy all bounds and incentive compatibility.

Since z supports C at U , the original rule maximizes $zU - J$ over all feasible rules, where J is their cost. Its directional derivative along this perturbation is

$$(1 - \delta) \int (zv - 1)(p^*(v) - p(v)) \, dF(v).$$

The continuation terms cancel. The integrand is nonnegative, so optimality requires $p = p^*$ almost everywhere. Since $u(v) \in \mathcal{Q}(z)$ almost surely, this gives

$$D(z) = (1 - \delta)\bar{D}(z) + \delta D(z),$$

contrary to Lemma 13.

If $u \geq U$ almost surely, replace u by $u - h$ for small $h > 0$. Optimality and monotone convergence of the difference quotients give

$$C'_-(U) \geq \int C'_-(u(v)) dF(v) \geq C'_-(U).$$

Both inequalities are equalities. Hence the left marginal costs are constant, which has just been excluded. If $u \leq U$ almost surely, the feasible perturbation $u + h$ instead gives

$$C'_+(U) \leq \int C'_+(u(v)) dF(v) \leq C'_+(U).$$

The derivatives are bounded on a neighborhood of $[0, U]$, so dominated convergence applies. Constant right marginal costs give the same contradiction. ■

Lemma 15 *Suppose f is positive and continuous on $(\underline{v}, 1)$. If an optimal rule is a non-constant deterministic threshold, C is affine between its two continuation promises. Both continuation promises are interior.*

Proof. Let the threshold be $t \in (\underline{v}, 1)$, with continuation b below t and a above it. Incentive compatibility gives $\delta(b - a) = (1 - \delta)t$. Replace the jump in provision by a linear increase on $[t - \varepsilon, t + \varepsilon]$ and obtain continuation from (21). Outside this interval both controls are unchanged. Inside it, continuation lies in $[a, b]$ and, after rescaling the interval to $[0, 1]$, converges uniformly to $b - (b - a)x$.

The change in provision integrates to zero with respect to Lebesgue measure. Equation (22) therefore gives a change in promised utility of order $O(\varepsilon^2)$. Continuity of f gives a change in current cost of $o(\varepsilon)$. The change in expected continuation cost is

$$2\delta\varepsilon f(t) \left[\int_0^1 C(b - (b - a)x) dx - \frac{C(a) + C(b)}{2} \right] + o(\varepsilon).$$

If C is not affine on $[a, b]$, the expression in brackets is negative. This contradicts optimality and local Lipschitz continuity of C at the original promise. Thus C is affine on $[a, b]$. The endpoint derivatives and Lemma 13 exclude an affine interval adjoining either endpoint. Hence $0 < a < b < \mu$. ■

Proof of Theorem 4. Convexity of C gives concavity of W . Lemma 13 gives $W < \bar{W}$ in the interior. It remains to establish differentiability and randomization. Fix an interior

promise and an optimal rule. Lemma 14 implies that continuation is nonconstant. Write $d = \delta/(1 - \delta)$, $P = p(1)$, $a = u(1)$, and

$$H(v) = \int_{\underline{v}}^v w \, dF(w), \quad T = d \int_{\underline{v}}^1 \frac{-du(v)}{v} = p(1) - p(\underline{v}).$$

Endpoint values denote the limits from the interior.

Let τ be an increasing, continuously differentiable change of reports, equal to the identity outside a compact subset of $(\underline{v}, 1)$. Keep the continuation levels fixed by setting $u_\tau(\tau(v)) = u(v)$. For a choice of upper provision probability P_τ , incentive compatibility determines

$$p_\tau(\tau(v)) = P_\tau - d \int_{(\underline{v}, 1]} \frac{-du(w)}{\tau(w)}.$$

Continuation bounds are unchanged. Provision bounds hold whenever

$$T_\tau := d \int_{\underline{v}}^1 \frac{-du(v)}{\tau(v)} \leq P_\tau \leq 1.$$

Integration by parts gives the resulting promise and cost:

$$U_\tau = (1 - \delta)\mu P_\tau + \delta a + \delta \int_{\underline{v}}^1 \left[F(\tau(v)) - \frac{H(\tau(v))}{\tau(v)} \right] (-du(v)), \quad (23)$$

$$J_\tau = (1 - \delta)P_\tau + \delta C(a) + \delta \int_{\underline{v}}^1 F(\tau(v)) [-dC(u(v))] - \delta \int_{\underline{v}}^1 \frac{F(\tau(v))}{\tau(v)} (-du(v)). \quad (24)$$

Here $-dC(u(v))$ is the Stieltjes measure of the nonincreasing function $C \circ u$. Since f is continuous, both expressions are continuously differentiable along smooth finite-dimensional changes of τ and P_τ . This does not require differentiability of C .

Suppose first that the measure $-du$ has two distinct support points in $(\underline{v}, 1)$. Choose nonnegative, continuously differentiable functions ξ_1, ξ_2 with compact, disjoint supports in that order, and normalize them so that

$$\int_{\underline{v}}^1 \frac{\xi_i(v)}{v^2} (-du(v)) = 1, \quad i = 1, 2.$$

Consider $\tau(v) = v + \varepsilon \xi_1(v) + \eta \xi_2(v)$. Both partial derivatives of T_τ at zero equal $-d$. The implicit function theorem therefore gives a differentiable $\eta(\varepsilon)$ with $\eta(0) = 0$, $\eta'(0) = -1$, and $T_\tau = T$. Set $P_\tau = P$. Both signs of small ε give feasible policies. Since

$$\frac{d}{dv} \left[F(v) - \frac{H(v)}{v} \right] = \frac{H(v)}{v^2},$$

their promise has derivative

$$\left. \frac{dU_\tau}{d\varepsilon} \right|_{\varepsilon=0} = \delta \int_{\underline{v}}^1 \frac{H(v)}{v^2} (\xi_1(v) - \xi_2(v)) (-du(v)) < 0.$$

The inequality follows from the ordering of the supports and the increase of H . Thus (24) gives a differentiable upper bound for C on a neighborhood of U , with equality at U . Convexity implies that C is differentiable at U .

Otherwise, $-du$ is a single atom at some $t \in (\underline{v}, 1)$. Continuation equals $b > a$ below t and a above it, with $T = d(b-a)/t$. Suppose $T < 1$. Move t to $t + \varepsilon$, keeping a and b fixed. If $p(\underline{v}) > 0$, keep $P_\tau = P$. Both signs are feasible, and the promise derivative is

$$\frac{\delta(b-a)H(t)}{t^2} > 0.$$

If $p(\underline{v}) = 0$, keep that probability zero by setting $P_\tau = d(b-a)/(t + \varepsilon)$. Since $P = T < 1$, both signs are again feasible. The promise derivative is now

$$\frac{\delta(b-a)(H(t) - \mu)}{t^2} < 0.$$

The cost is differentiable in each case, so the same upper-bound argument establishes differentiability of C at U .

If $T = 1$, provision is a deterministic threshold. Lemma 15 makes C affine on $[a, b]$, and Lemma 14 gives $a < U < b$. Hence C is differentiable at U in this case as well. This proves differentiability at every interior promise. A convex function that is differentiable on an open interval has a continuous derivative there.

Finally, suppose a deterministic threshold were optimal. Its two continuation promises are interior, and C is affine between them. The differentiability just proved makes their marginal costs equal to $C'(U)$, since $a < U < b$. This contradicts Lemma 14. Constant provision is also excluded by that lemma and incentive compatibility. Since provision is nondecreasing, every remaining rule assigns a probability in $(0, 1)$ to a set of types of positive F -measure. ■

C.2 Convergence

For $K > 0$, define

$$L_K(U) = \begin{cases} -K, & U = 0, \\ C'_-(U), & 0 < U \leq \mu. \end{cases} \quad (25)$$

Here $L_K(\mu) = \infty$ if $\underline{v} = 0$. When $\underline{v} > 0$, write $C'_+(\mu) = 1/\underline{v}$ in formulas involving the right derivative.

Lemma 16 *There is a finite $K > 0$, depending only on F and δ , such that every optimal rule at every interior promise satisfies*

$$\int L_K(u(v)) dF(v) \leq L_K(U). \quad (26)$$

If $\underline{v} = 0$, every such rule assigns probability zero to $u = \mu$. If $\underline{v} > 0$, it also satisfies

$$C'_+(U) \leq \int C'_+(u(v)) dF(v). \quad (27)$$

Proof. Put $d = \delta/(1 - \delta)$ and fix an optimal rule (p, u) . Neither $u \equiv 0$ nor $u \equiv \mu$ is optimal at an interior promise. Either would make p constant and give cost U/μ , contrary to Lemma 13.

For small $h > 0$, set

$$u_h(v) = (u(v) - h)_+, \quad p_h(v) = p(\underline{v}) + d \int_{(\underline{v}, v]} \frac{-du_h(w)}{w}.$$

Truncation reduces the measure $-du$. Hence $0 \leq p_h \leq p \leq 1$, and these rules are feasible. If $u(1) > 0$, then $p_h = p$ for small h . The Bellman inequality at the new promise gives

$$\int C'_-(u(v)) dF(v) \leq C'_-(U). \quad (28)$$

The integral is finite: the backward difference quotients of C increase to its left derivative, so monotone convergence applies even if the derivative at μ is infinite.

Now suppose $u(1) = 0$ and let $t = \sup\{v : u(v) > 0\}$. Then $t > 0$. Write U_h, J_h for the promise and cost of (p_h, u_h) , and put $z = C'_-(U)$. The mass removed from $-du$, divided by h , converges to a unit mass at t . Integration by parts gives

$$\begin{aligned} \lim_{h \downarrow 0} \frac{U - U_h}{h} &= \delta \left[1 + \frac{1}{t} \int_t^1 (1 - F(v)) dv \right], \\ \lim_{h \downarrow 0} \frac{C(U) - J_h}{h} &= \delta \left[\int_{\{u>0\}} C'_-(u(v)) dF(v) + \frac{1 - F(t)}{t} \right]. \end{aligned}$$

The second limit is initially allowed to be infinite. Since $C(U_h) \leq J_h$, it is finite, and

$$\int_{\{u>0\}} C'_-(u(v)) dF(v) \leq z + \frac{1}{t} \left[z \int_t^1 (1 - F(v)) dv - (1 - F(t)) \right]. \quad (29)$$

Together with (28), this also rules out positive probability of $u = \mu$ when $\underline{v} = 0$.

The cutoffs t of optimal rules with $u(1) = 0$ have a positive lower bound ε , uniform over interior promises. This is immediate if $\underline{v} > 0$. Otherwise, suppose a sequence has $t_j \rightarrow 0$, and put $P_j = p_j(1)$. Incentive compatibility gives

$$p_j(v) = P_j \quad (v > t_j), \quad 0 \leq u_j(v) \leq \frac{t_j P_j}{d}.$$

Along a subsequence, $P_j \rightarrow P$. If $P > 0$, continuity and optimality imply

$$U_j \rightarrow (1 - \delta)\mu P, \quad C(U_j) \rightarrow (1 - \delta)P,$$

contradicting $C(U) < U/\mu$ at the limiting interior promise. If $P = 0$, then $U_j \rightarrow 0$, whereas

$$\frac{C(U_j)}{U_j} \geq \frac{1 - F(t_j)}{\mu + t_j} \rightarrow \frac{1}{\mu} > 1.$$

Here the numerator uses current provision on $(t_j, 1)$, and the denominator uses $U_j \leq (1 - \delta)P_j(\mu + t_j)$. This contradicts $C'_+(0) = 1$ and proves the uniform bound.

The same rules satisfy

$$\delta u(v) = (1 - \delta) \int_{(v, 1]} w \, dp(w) \leq (1 - \delta)(1 - p(v)).$$

Promise keeping therefore gives

$$U \leq (1 - \delta)\mu + \delta \mathbf{E}[(1 - v)u(v)] \leq \mu - \delta\mu^2.$$

Choose $K \geq C'_+(\mu - \delta\mu^2)/\varepsilon$. Since $\int_t^1 (1 - F(v)) \, dv \leq 1 - F(t)$, (29) gives

$$\int_{\{u > 0\}} C'_-(u(v)) \, dF(v) - K \mathbf{P}(u = 0) \leq z.$$

This and (28) prove (26).

For (27), suppose $\underline{v} > 0$ and instead use

$$u_h(v) = \min\{u(v) + h, \mu\}, \quad p_h(v) = p(1) - d \int_{(v, 1]} \frac{-du_h(w)}{w}.$$

Now $p \leq p_h \leq 1$. If $u(\underline{v}) < \mu$, uniform upward shifts give (27) directly. Otherwise, put $s = \inf\{v : u(v) < \mu\}$ and $z = C'_+(U)$. The corresponding forward difference quotients give

$$z \left[1 - F(s) + \frac{1}{s} \int_{\underline{v}}^s v \, dF(v) \right] \leq \int_{\{u < \mu\}} C'_+(u(v)) \, dF(v) + \frac{F(s)}{s}.$$

All derivatives here are bounded by $1/\underline{v}$. Using $\int_{\underline{v}}^s v \, dF(v) \geq \underline{v}F(s)$ yields

$$\int C'_+(u(v)) \, dF(v) - z \geq F(s) \left(\frac{1}{\underline{v}} - z \right) \left(1 - \frac{\underline{v}}{s} \right) \geq 0.$$

This proves the claim, including $s = \underline{v}$. ■

Lemma 17 *No optimal rule at an interior promise satisfies $L_K(u(v)) = L_K(U)$ almost surely, for the K in Lemma 16.*

Proof. At an interior promise, $L_K(U) = C'_-(U) > 1$. The proposed equality excludes continuation zero and implies $C'_-(u(v)) = C'_-(U)$ almost surely. This contradicts Lemma 14. ■

Let $\mathcal{O}$ be the set of optimal triples (U, p, u) , including the two endpoints. Equip the rules with the $L^1(F)$ topology. This set is compact. Indeed, bounded monotone rules have convergent subsequences. Incentive compatibility passes to their limits, and promise keeping and cost are continuous in this topology.

Lemma 18 *There is no probability measure ν on $\mathcal{O}$ whose promise marginal is supported on a compact subset of $(0, \mu)$ and which satisfies*

$$\int_{\mathcal{O}} \left[\int \varphi(u(v)) dF(v) - \varphi(U) \right] d\nu(U, p, u) = 0 \quad (30)$$

for every continuous $\varphi : [0, \mu] \rightarrow \mathbf{R}$.

Proof. Let $\mathcal{K}$ be the support of the promise marginal. Taking $\varphi(U)$ to be the distance from U to $\mathcal{K}$ in (30) shows that $u(v) \in \mathcal{K}$ almost surely, for ν -almost every triple. The integral of this distance is continuous in the rules. It therefore vanishes at every triple in the support of ν .

Put $a = \min \mathcal{K} > 0$. Compactness gives a triple (a, p, u) in that support. Its continuations satisfy $u(v) \geq a$ almost surely. Monotonicity of the left derivative and Lemma 16 give

$$L_K(a) \geq \int L_K(u(v)) dF(v) \geq L_K(a).$$

Thus $L_K(u(v)) = L_K(a)$ almost surely, contradicting Lemma 17. ■

Proof of Theorem 5. By Lemma 16, $L_K(U_n) + K$ is a nonnegative supermartingale. It converges almost surely to a finite limit. When $\underline{v} = 0$, an interior promise assigns zero probability to μ , so this process is finite at every date. The endpoints are absorbing.

By the endpoint derivatives and Lemma 13, $L_K(U_n) \rightarrow -K$ means absorption at zero, and $L_K(U_n) \rightarrow 1$ implies $U_n \rightarrow 0$. When $\underline{v} > 0$, a limit of $1/\underline{v}$ implies $U_n \rightarrow \mu$. Any other possible limit would place all sufficiently late promises in a compact subset of $(0, \mu)$.

We exclude the last possibility without assuming that the optimal rules are stationary or continuous in the promise. Let (p_n, u_n) be the rule selected at date n , and define the

empirical measures

$$\nu_N = \frac{1}{N} \sum_{n=0}^{N-1} \delta_{(U_n, p_n, u_n)}$$

on $\mathcal{O}$, where $\delta_{(U,p,u)}$ denotes a unit mass. For each continuous φ , bounded martingale differences and telescoping give, almost surely,

$$\frac{1}{N} \sum_{n=0}^{N-1} \left[\int \varphi(u_n(v)) dF(v) - \varphi(U_n) \right] \rightarrow 0.$$

Apply this first to a countable uniformly dense family of continuous functions. Uniform approximation then gives the same conclusion for all continuous φ on a common event of probability one. Compactness of $\mathcal{O}$ gives a weakly convergent subsequence of ν_N . The displayed integrands are continuous on $\mathcal{O}$, so its limit satisfies (30). If the promises eventually stay in an interior compact set, the promise marginal of this limit is supported there. Lemma 18 rules this out. This proves convergence to the asserted endpoints.

Suppose now that $\underline{v} > 0$. By (27), $C'_+(U_n)$ is a bounded submartingale, with limiting values 1 and $1/\underline{v}$. The supermartingale $L_K(U_n)$ has limit $1/\underline{v}$ on $\{U_\infty = \mu\}$ and is bounded below by $-K$ elsewhere. Taking expectations gives

$$0 < \frac{C'_+(U_0) - 1}{1/\underline{v} - 1} \leq \mathbf{P}(U_\infty = \mu) \leq \frac{C'_-(U_0) + K}{1/\underline{v} + K} < 1.$$

The strict inequalities follow from Lemma 13.

Lemma 19 below shows that the limiting probabilities do not depend on the selection of optimal rules. Every maximizer of $W(U) = U - cC(U)$ is interior, by the endpoint derivatives and $\underline{v} < c < 1$, and satisfies $C'(U) = 1/c$. The same lemma therefore gives equal limiting probabilities across optimal initial promises. ■

Let (U_n) and $(\tilde{U}_n)$ be two optimal promise processes driven by the same sequence of independent valuations, with deterministic interior initial promises. Optimal rules may depend on their respective histories.

Lemma 19 *Suppose f is positive and continuous on the interior of its support and $C'(U_0) = C'(\tilde{U}_0)$. Then, almost surely,*

$$C'(U_n) = C'(\tilde{U}_n) \quad \text{for every } n \geq 0, \quad \lim_n U_n = \lim_n \tilde{U}_n.$$

Proof. For any two promises with the same derivative, C is affine between them. Mix optimal rules at these promises, pointwise in the report. The mixed rule is feasible at the

mixed promise. Convexity bounds its cost above by the same mixture of optimal costs, which equals the value at that promise. Optimality gives the reverse inequality. Thus the nonnegative Jensen difference in continuation cost has integral zero. For almost every report, C is affine between the two continuation promises. Differentiability in the interior and Lemma 13 imply that their derivatives agree, including when an endpoint is reached. No interior promise shares an endpoint's derivative, by Lemma 13. Conditioning on the histories and applying the preceding argument at each date proves the first assertion by induction. The countable intersection has probability one. The convergence argument above gives endpoint limits for both processes. If $\underline{v} = 0$, both limits are zero. If $\underline{v} > 0$, the derivatives at the two endpoints are different, so the first assertion makes the limits agree. ■

C.3 Power Distributions

In this subsection, $F(v) = v^a$, $a \geq 1$. The construction and envelope identity hold for $0 < \delta < 1$. Write $\mu = a/(a+1)$ and $d = \delta/(1-\delta)$. Theorem 4 gives continuous differentiability of C . Here $\mathcal{Q}(1) = \{0\}$. For $r > 1$, the set is a closed interval contained in $(0, \mu)$. Its endpoints tend to μ as $r \rightarrow \infty$. An interval of positive length corresponds to an affine portion of C .

For a supporting slope $z > 1$, define

$$k_0(z) = \max\{0, 1 - \mu z\}, \quad k_1(z) = \frac{1}{(a+1)z^a}.$$

The construction uses a scalar $k \in [k_0(z), k_1(z))$. When $a = 1$, require $k > 0$. For $k > 0$, set $s = ((a+1)k)^{1/a}$. For $k = 0$ and $a > 1$, set $s = 0$. Let $t \in (s, 1]$ solve

$$z = \frac{a+1}{a^2} \left[\frac{a-1}{t} + t^{a-1} + k \left(\frac{1}{t^{a+1}} - \frac{a+1}{t} \right) \right]. \quad (31)$$

This solution is unique. The expression on the right is decreasing in $t < 1$, equals $(1-k)/\mu$ at $t = 1$, and equals $1/s$ at $t = s > 0$. Define

$$r(v) = \frac{z}{a+1} + \frac{a-1}{av} + \frac{k}{av^{a+1}}. \quad (32)$$

This function is decreasing. Its values satisfy $r(t) \geq 1$ and, when $s > 0$, $r(s) > z > r(t)$.

Choose $q_s \in \mathcal{Q}(r(s))$ when $s > 0$ and $q_t \in \mathcal{Q}(r(t))$. For $s = 0$, set $q_s = \mu$. The

continuation rule is

$$u(v) = \begin{cases} q_s, & 0 \leq v < s, \\ Q(r(v)), & s \leq v < t, \\ q_t, & t \leq v \leq 1. \end{cases} \quad (33)$$

When $s = 0$, the middle expression is used only for $v > 0$, and $u(0) = \mu$. Define the total increase in provision and the allocation rule by

$$T = d \int_{(0,1]} \frac{-du(v)}{v}, \quad p(v) = P - d \int_{(v,1]} \frac{-du(w)}{w}. \quad (34)$$

Thus $P = p(1)$ and $p(0+) = P - T$. The choices of k, q_s, q_t, P must satisfy

$$\begin{aligned} k > k_0(z) : & \quad T = P = 1, \\ k = k_0(z) > 0 : & \quad P = T \leq 1, \\ k = k_0(z) = 0 : & \quad a > 1, \quad P = 1, \quad T \leq 1. \end{aligned} \quad (35)$$

The last case does not occur at $z = 1/\mu$ when $\delta > 1/2$. The promise generated by these choices is

$$U = (1 - \delta)\mu P + \delta\mu q_t + \frac{\delta}{a+1} \int u(v) dF(v). \quad (36)$$

The generalized inverse Q allows jumps. Derivatives at zero denote right derivatives.

For power distributions,

$$\overline{C}(U) = 1 - (1 - U/\mu)^{a/(a+1)}, \quad \overline{D}(z) = \mu z - 1 + \frac{1}{(a+1)z^a}, \quad z > 1.$$

Lemma 13 and the endpoint derivatives imply

$$\max\{0, \mu z - 1\} < D(z) < \overline{D}(z), \quad z > 1. \quad (37)$$

If $q \in \mathcal{Q}(r)$, comparison with μ and with $\overline{C}$ gives

$$r(\mu - q) \leq 1 - C(q) \leq \left(\frac{\mu - q}{\mu} \right)^{a/(a+1)}, \quad \mu - q \leq \frac{1}{\mu^a r^{a+1}}.$$

The last bound also holds when C is replaced by any convex function above $\overline{C}$ that equals one at μ . We refer to it below as the bound on $\mu - q$.

Lemma 20 *For each $z > 1$, an admissible choice in (31)–(36) exists, and k is unique. Every such choice is optimal at its generated promise. Conversely, every optimum at a promise in $\mathcal{Q}(z)$ has this form, up to null sets. At a fixed promise, the optimal rules are unique up to null sets and satisfy (9).*

Proof. Fix an admissible construction. Put

$$m(v) = \begin{cases} r(s), & v < s, \\ r(v), & s \leq v < t, \\ r(t), & v \geq t. \end{cases}$$

When $s = 0$, the first case is absent. The choices in (33) give $m(v) \in \partial C(u(v))$ almost surely. Integration using (31) and $k = s^a/(a+1)$ gives

$$\int m(v) dF(v) = z. \quad (38)$$

For $s = 0$ the calculation uses $a > 1$, which also ensures integrability at zero.

Define two verification functions by

$$\begin{aligned} \alpha(v) &= (1 - \delta) \left[v^a - zv + \frac{zv^{a+1}}{a+1} - k \right], \\ \beta(v) &= -\delta z + \delta \int_0^v m(w) dF(w), \quad H(v) = \alpha(v) - \frac{1 - \delta}{\delta} v \beta(v). \end{aligned}$$

They satisfy

$$\alpha(0) = -(1 - \delta)k, \quad \alpha(1) = (1 - \delta)(1 - \mu z - k), \quad \beta(0) = -\delta z, \quad \beta(1) = 0.$$

For $v \in [s, t]$ with $v > 0$, $\beta(v) = \delta \alpha(v)/((1 - \delta)v)$ and $H = 0$. Off this interval,

$$\left(\frac{H(v)}{v} \right)' = (1 - \delta) f(v) (r(v) - m(v)).$$

Since r is decreasing, $H < 0$ on $(0, s)$ and $(t, 1)$ whenever these intervals are nonempty.

For another feasible policy $(\tilde{p}, \tilde{u})$, write Δ for its difference from the proposed policy. The subgradient inequality and integration by parts give the following bound on its increase in $zU - J$:

$$\Delta(zU - J) \leq -\alpha(1)\Delta p(1) + \alpha(0)\Delta p(0+) + \int H(v) d\Delta p(v) \leq 0. \quad (39)$$

The continuation boundary terms vanish because $\beta(0) + \delta z = \beta(1) = 0$. The remaining boundary signs follow from (35). Also, $H = 0$ on the support of dp , while $H \leq 0$ and $d\tilde{p} \geq 0$. This proves optimality. Equality requires provision to be constant off $[s, t]$, and requires continuation to maximize $m(v)x - C(x)$ at almost every type. Thus every optimizer has the stated form, with the possible endpoint selections q_s, q_t .

We next establish existence of an admissible construction for each $z > 1$. For $0 < \varepsilon < 1$, replace the continuation cost by $C_\varepsilon = (1 - \varepsilon)C + \varepsilon\bar{C}$. This function is strictly convex, so its maximizing promise at each slope is single valued and continuous. It has the same endpoint derivatives as C . For $k \in (k_0, k_1)$, the corresponding $T_\varepsilon(k)$ is continuous. As $k \uparrow k_1$, $s, t \rightarrow 1/z$, so $T_\varepsilon(k) \rightarrow 0$. At $k = k_0$, use the boundary construction if its T is at most one; otherwise the intermediate value theorem gives a k with $T = 1$. For $a = 1$ and $k_0 = 0$, we have $z \geq 2$, $s = 2k$, and $r(s) = z/2 + 1/(4k)$. As $k \downarrow 0$, equation (31) gives $t = 1/2$ if $z = 2$, and $t \rightarrow 0$ with $k/t \rightarrow 0$ if $z > 2$. Thus $m(v) \rightarrow z - 1$ for every fixed $v > 0$, whereas $u(0) \rightarrow \mu$. For fixed ε , the limiting continuation at every positive type is the maximizer of $(z - 1)x - C_\varepsilon(x)$, which is below μ . The gap is therefore independent of v . Since $u(0) - u(v) \leq vT/d$, choosing a sufficiently small fixed v rules out bounded T . Hence $T_\varepsilon(k) \rightarrow \infty$. For $a > 1$, the limit at $k = 0$ is the finite upper construction; the bound on $\mu - q$ controls the integral at zero. The verification above applies to C_ε as a continuation cost and proves that the constructed policy maximizes its one-period dual objective.

Let $\varepsilon \downarrow 0$. Compactness of the allocation and continuation rules, and uniform convergence of C_ε , give an optimizer for the dual problem with continuation cost C . The endpoint choices in (33) retain any limiting jumps. If $k \rightarrow 0$ and $a > 1$, then $q_s \rightarrow \mu$. On the active interval, $r(v) \geq (a - 1)/(av)$. The bound on $\mu - q$ therefore gives $0 \leq u(0) - u(v) \leq Bv^{a+1}$ for $0 < v < t$, where $B = \mu^{-a}(a/(a - 1))^{a+1}$ is independent of ε and k . For $v < s$, the difference is zero. For $0 < \eta < t$, integration by parts yields

$$\begin{aligned} d \int_{(0, \eta]} \frac{-du(v)}{v} &= d \frac{u(0) - u(\eta)}{\eta} + d \int_0^\eta \frac{u(0) - u(v)}{v^2} dv \\ &\leq dB \left(1 + \frac{1}{a}\right) \eta^a. \end{aligned}$$

As $k \rightarrow 0$, equation (31) keeps t bounded away from zero. This estimate therefore passes the integral defining T to the limit at zero. The case $T < 1$ gives $p(0+) > 0$, as permitted in (35). For $a = 1$, the preceding argument also excludes $k \rightarrow 0$ as $\varepsilon \downarrow 0$: every limit of $u(v)$ at a fixed $v > 0$ belongs to $\mathcal{Q}(z - 1)$, whose maximum is below μ . Thus the common positive gap again contradicts $T \leq 1$. The limit $k = k_1$ is also impossible. In that case, $s = t = 1/z$, provision is the threshold rule, and both continuation values belong to $\mathcal{Q}(z)$. The Bellman equality would give

$$D(z) = (1 - \delta)\bar{D}(z) + \delta D(z),$$

contradicting (37). Thus $k < k_1$, and the limiting construction is admissible. Every interior promise has a supporting slope $z > 1$, so the equality conditions in (39) cover every interior promise.

We record a consequence of those equality conditions. Suppose the same nonconstant policy is verified by two parameter pairs $(z_1, k^{(1)})$ and $(z_2, k^{(2)})$. Allocation can vary only in the intersection of their closed active intervals. In that intersection the two marginal functions are

$$r_i(v) = \frac{z_i}{a+1} + \frac{a-1}{av} + \frac{k^{(i)}}{av^{a+1}}.$$

Where $r_1 \neq r_2$, the common continuation rule is locally constant. To see this, choose a neighborhood on which every value of the smaller function is below every value of the larger one. Two different continuation promises on that neighborhood would have ordered subgradient sets, contradicting the fact that each promise supports both marginal values. A jump also requires $r_1 = r_2$ at its location: continuity of each marginal function makes both sides of the jump maximize at that slope. These arguments include the boundaries of the active intervals by taking one-sided limits.

For fixed z , every dual optimizer must attain equality in both verifications. Two different k 's give marginal functions that never coincide. A common optimizer would therefore have constant continuation, contrary to Lemma 14. Hence k is unique. For fixed z, k , continuation on (s, t) is unique almost everywhere: a convex function has only countably many slopes with a nondegenerate maximizing interval, and r is decreasing. In the middle case, $T = 1$ and promise keeping determine q_s, q_t . In fact, with the interior continuation fixed,

$$T = d \left[\frac{q_s}{s} - \frac{q_t}{t} - \int_s^t \frac{u(v)}{v^2} dv \right],$$

and (36) has positive coefficients on both endpoint promises. The two equations are therefore independent. In the lower and upper cases there is only one free endpoint promise, with a positive coefficient in promise keeping. Thus the optimal policy at a fixed U is unique up to null sets.

Since C is differentiable, $m(v) = C'(u(v))$ almost surely. The derivative at zero is one. The value μ is reached only at the zero type in the upper regime, which has probability zero. The marginal function in that regime is integrable because $a > 1$. Equation (38) therefore proves the claimed pointwise envelope identity. ■

For the three-regime characterization, now assume $\delta > 1/2$. The two transitions can be found from one-dimensional equations. For the lower transition, let $s \in (0, 1)$, put

$$z_-(s) = \frac{a+1-s^a}{a}, \quad k = \frac{s^a}{a+1}, \quad t = 1, \quad (40)$$

and use (33). There is a unique s_- at which the interval of possible values of T contains one.

Select q_s so that $T = 1$, and set

$$z_- = z_-(s_-), \quad U_- = \max \mathcal{Q}(z_-).$$

When $a > 1$, the upper transition is obtained by setting $k = s = 0$ in (31). There is a unique $t_+ \in (0, 1)$ at which the interval of possible values of T contains one. Set

$$z_+ = \frac{a+1}{a^2} \left(\frac{a-1}{t_+} + t_+^{a-1} \right), \quad U_+ = \min \mathcal{Q}(z_+).$$

The endpoint selections in these transition equations are necessary because the maximizing sets may contain more than one promise.

Proof of Proposition 1. Lemma 20 gives uniqueness and the form of the optimal rules. In the lower boundary construction, $t = 1$ and

$$r(v) = 1 + \frac{a-1}{a}(v^{-1} - 1) + \frac{s^a}{a(a+1)}(v^{-(a+1)} - 1).$$

As s decreases, the upper slope $r(s)$ increases. For a fixed $r > 1$ in its range, the corresponding value of v decreases. Consequently,

$$T = d \int_0^{q_s} \frac{1}{v(C'(x); s)} dx$$

increases as s decreases. Endpoint selections fill any jump. The ordering is strict away from $s = 1$, since $Q(r) > 0$ for $r > 1$. The limit at $s = 1$ is zero. For $a > 1$, its limit at zero is

$$\frac{d}{(a-1)/a} \int_0^\mu \left(C'(x) - \frac{1}{a} \right) dx = d \frac{a^2}{a^2 - 1} > 1.$$

For $a = 1$, the limit is infinite: on any fixed interval of slopes above one with positive dQ mass, the factor $1/v(r; s)$ diverges. Thus there is a unique lower transition $s_- > 0$.

For $a > 1$, the upper boundary construction has

$$T = \frac{d}{(a-1)/a} \int_{q_t}^\mu \left(C'(x) - \frac{z}{a+1} \right) dx.$$

As t increases, both z and $r(t)$ decrease. The domain of integration expands, and its integrand increases. Endpoint selections again fill any jump. The limit is zero at $t = 0$ and $da^2/(a^2 - 1) > 1$ at $t = 1$. There is therefore a unique upper transition $t_+ < 1$. The marginal thresholds satisfy $z_- < 1/\mu < z_+$.

If the relevant boundary construction has $T \leq 1$, it verifies the optimum, so uniqueness of k excludes another regime. If it has $T > 1$, (35) requires $k > k_0$, giving the middle

regime. At a transition, $T = 1$ pins down the endpoint selection and the promised utility. Monotonicity of supporting slopes gives the stated ordering in U . At the lower transition all promises supported by z_- are on its lower side, with the largest one giving $T = 1$; hence $U_- = \max \mathcal{Q}(z_-)$. At the upper transition the smallest promise supported by z_+ gives $T = 1$; hence $U_+ = \min \mathcal{Q}(z_+)$. For $a = 1$, $k = 0$ is impossible, so there is no third regime. On the active interval, $r(v)$ decreases. Differentiability of C implies that two distinct slopes have distinct maximizing promises. Thus continuation decreases and provision increases on every nonempty subinterval of (s, t) . The increase in provision may include jumps. Off this interval both rules are constant, by the verification equality conditions. When $t = 1$, $r(t) = 1$ gives $u(1) = 0$. When $t < 1$, $r(t) > 1$ gives $u(1) > 0$. When $s > 0$, $r(s) < \infty$ gives $u(0) < \mu$; when $s = 0$, $u(0) = \mu$. Complementarity supplies the stated provision bounds. The endpoint promises admit only perpetual nonprovision and perpetual provision, respectively, up to null sets. ■