

The Dynamics of Verification when Searching for Quality

ZIHAO LI* AND JONATHAN LIBGOBER**

*Columbia University

**University of Southern California

April 18, 2026

ABSTRACT. An agent samples projects over time, observing quality for each, while a principal can select at most one. The principal values quality, whereas the agent only wants a project chosen. Transfers are unavailable, but the principal can verify quality by paying a cost. We fully characterize the dynamics of verification by determining optimal mechanisms for this problem. With a low verification cost and a long horizon, the optimal mechanism involves a deterministic selection rule that initially discriminates on quality but chooses a project irrespective of its quality at a deadline. Verification occurs with an intermediate probability before the deadline, declining over time. We show how these conclusions change if the verification cost is high or the horizon is short, and under certain forms of imperfect commitment. Our analysis provides guidelines on how dynamics interact with the benefits of verification.

KEYWORDS. Costly verification, dynamic contracts, mechanism design without transfers.

CONTACT. ZL3366@COLUMBIA.EDU, LIBGOBER@USC.EDU. We thank Hector Chade, Thomas Chaney, Yeon-Koo Che, Eddie Dekel, Drew Fudenberg, Daniel Gottlieb, Marina Halac, Johannes Hörner, Navin Kartik, Deniz Kattwinkel, Jan Knoepfle, Pablo Kurlat, Yingkai Li, Bart Lipman, Qingmin Liu, Rob Metcalfe, Beatrice Michaeli, Andrea Prat, Simon Quach, Danny Sokol, Rakesh Vohra, Kai Hao Yang and Beixi Zhou for discussions, various seminar audiences for excellent feedback, and the editor and three anonymous referees for many helpful comments. Jonathan Libgober thanks the Cowles Foundation and Yale University for their hospitality during part of this research. This paper was originally written while Zihao Li was at Princeton University, and he gratefully acknowledges the Princeton University Department of Economics for helpful discussions and encouragement.

1. Introduction

This paper presents a dynamic principal-agent model in which a principal seeks to make a selection from a set of possibilities. These possibilities emerge over time, and our main question is how the selection mechanism should evolve as a result. While our model is theoretical, it aims to capture essential trade-offs of broad practical relevance. One setting where the issues we study are particularly salient is the problem firms face when seeking to enter a new market by acquiring a company. In these situations, the task of determining a viable target typically falls to the CEO, who may know whether a given acquisition will end up being profitable, but whose primary incentive is often simply in having an acquisition go through.¹ While the CEO can make proposals over time, the board ultimately decides whether to approve one, along with whether to independently audit quality on their own first. Our focus is on how such activities may be used to counteract the aforementioned incentive conflict between the board and the CEO.

Our model broadly captures the key features of this application, in particular studying interactions between two capabilities highlighted: First, the retention of *decision rights* to specify choices to overcome the subordinate's bias. And second, the ability to *verify*, at a cost, whether the subordinate has misrepresented their information. A distinguishing feature of our exercise stems from the observation that in situations where particular choices are sought, new options can be found if a given prospect is deemed inadequate. The dynamic aspect of this setting raises a fundamental question: How should verification and selection rules evolve over time?

Economists have emphasized the significance of information asymmetries within organizations, particularly in shaping incentives and organizational design (since at least Alchian and Demsetz (1972) and Jensen and Meckling (1976)). Continuing in this tradition, we study the use of costly verification, focusing on the dynamics that arise when prospects emerge over time. We show how the optimal evolution of selection rules and verification depends on, for instance, the ease of verification and degree of agent bias. The literature on organizational governance has studied the interaction between discretion on the one hand and oversight in the form of monitoring or verification on the other—not only in the context of boards managing CEOs (Adams and Ferreira, 2007), but also for within-organization capital budgeting problems (Harris and Raviv, 1996). We contribute to this work by providing prescriptions and predictions for settings featuring dynamic selection. Correspondingly, we strive toward *minimality* to provide clear intuition for our conclusions, modeling the conflict and available mechanisms in a stylized manner.

¹Empirical evidence supports this assumption. Bliss and Rosen (2001); Harford and Li (2007) document that acquiring CEOs generally benefit financially even when the deal harms shareholder value. The preference for acquisition by CEOs can reflect incentives for “empire building” and deal completion; for instance, Grinstein and Hribar (2004) show that CEOs frequently receive large cash bonuses simply for closing a transaction.




Figure 1: Dynamics of verification with a long horizon and low verification cost

That said, two features clarify our precise contribution: First, we posit a single selection is sought (e.g., only one target can be acquired), consistent with past work on delegated project selection surveyed below. Second, we disallow transfers (e.g., if the board cannot effectively tie compensation to long-term profitability), consistent with past work in organizational economics which has observed that within-firm contracting generally does not feature such payments.² Both properties are defining features of Ben-Porath et al. (2014) which inspired a large body of work on costly verification without transfers. But while much of this work has focused on static environments, our interest is in how verification evolves *dynamically* when new options emerge over time.

We summarize our model as follows. In each period, the agent learns the quality of a single new *project*, where quality is binary-valued quality (low or high) drawn i.i.d. according to a known distribution.³ We refer to the total length of time available to sample projects as the *horizon*, which we assume for most of the paper is infinite but allow to be finite in an extension to draw a contrast for our results. After observing the quality of that period’s project, the agent reports to the principal. We allow the principal to fully specify when to perform (costly) verification and whether to select that period’s project; we relax commitment in our extensions. Quality determines the principal’s payoff from the selected project (gross verification costs), whereas the agent obtains one unit of utility whenever selection occurs; outside options (e.g., retaining the capital necessary for the acquisition) for both players are 0. Both parties discount future payoffs with a common discount factor δ , but to avoid trivialization, we assume the principal would prefer to wait for high quality were verification costless.

Our headline result is that when the horizon is long and the verification cost is low, optimal

²A previous version of this paper showed that our main message is unchanged if transfers are allowed when the agent is protected by limited liability, provided the value of selection is sufficiently large relative to the value of transfers.

³Equivalently, a default low-quality option is available, with high-quality arriving with fixed probability each period.

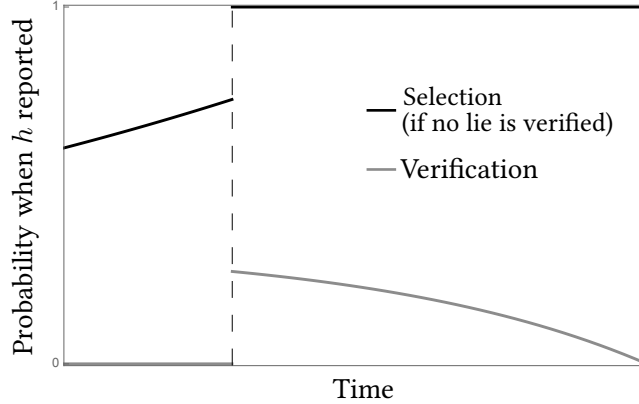


Figure 2: Dynamics of verification and selection with a high verification cost (and long horizon)

verification dynamics satisfy a property we dub *decreasing skepticism*: The principal verifies claims of high quality with an intermediate probability that declines over time. Figure 1 plots a representative solution for the evolution of verification. We interpret the decreasing probability of verification as reflecting the principal’s diminishing concern of misrepresentation, explaining our terminology. The verification probability hits zero at some (endogenously chosen) terminal date, at which point the principal selects a project irrespective of quality. We refer to this time as the *deadline*. Before this deadline, the principal never selects a project following a report of low quality, but always selects one after a report of high quality—provided that no verification occurs or verification confirms the report. If verification confirms a lie, outside options are selected. The agent, in turn, follows a strategy of proposing selection only if the project is of high quality.

For some intuition, suppose project selection exogenously takes the form specified above—i.e., the principal only accepts high-quality projects until a deadline at which anything is selected. While the principal wants to select only once high quality is realized, the agent always prefers immediate selection. Thus, verification can dissuade an agent with a low-quality project from misrepresenting, lest it reveal their lie and end the possibility of any selection at all. The higher the agent’s payoff when reporting low quality, the lower the verification probability needs to be. Now, a project will finally be selected once the deadline arrives. So while selection only occurs before the deadline if high quality arrives, the agent’s payoff absent selection also increases as the deadline approaches. Thus, the agent’s losses following the detection of a lie—namely, the forgone future continuation value—increase over time as well. The principal can therefore save on costs by decreasing the verification probability.

Our message that the verification probability decreases over time has important qualifiers. The above intuition is incorrect if verification costs are high or the interaction horizon is short. In these cases, different verification dynamics emerge.

Regarding the former case, deterministic project selection is *suboptimal* if the cost of verification

is large—in particular, we show that (some) randomization of the selection rule is optimal if and only if verification ensures a negative ex-post principal payoff. In that case, optimal mechanisms must be augmented with an initial “screening via randomization” phase. In this phase, the principal does *not* utilize the verification technology. Instead, the principal incentivizes truth-telling by reacting to a report of high quality by selecting the outside option with intermediate probability (and otherwise selecting that period’s project).⁴ This screening benefits the principal by inducing the agent to sample again after a low-quality draw rather than falsely claiming high quality, which could risk forgoing selection entirely. Over time, the probability of selecting outside options declines.

The key take-away from the above discussion is that verification is optimally *backloaded*. Specifically, the verification phase, which resembles the optimal mechanism described with low verification costs, begins only once the randomization phase ends. Figure 2 illustrates a representative two-phase mechanism, with non-degenerate selection probabilities in the first phase and non-degenerate verification probabilities in the second. The dashed line depicts the time at which phases switch. While the principal either randomizes the project selection rule or verifies with positive probability (but never both), verification always occurs after randomization. While the principal verifies reports of high quality with a probability that declines over time *once the verification phase starts*, the verification probability is initially zero while randomizing selection. These dynamics yield non-monotonic verification probabilities.⁵ Briefly, the intuition behind backloading of the verification phase stems from the observation that the *agent’s* expected payoff is higher in any period where the principal utilizes verification. Traditional intuition from the dynamic contracting literature suggests these rewards should be backloaded to relax incentive constraints. In our case, utilizing verification in later periods leads to accumulated gains in earlier periods, implying that backloaded verification can be useful even if the verification cost is high.

The second shortcoming of the intuition for decreasing skepticism is that when the horizon is *short*, the verification probability may increase over time when low quality is worse for the principal than her outside option. This case is precisely when the principal and agent disagree about the optimal decision in the *static* (one-period) problem. Here, the optimal one-period mechanism involves *no selection following a low-quality draw* and selection-with-verification following high-quality (provided verification is not too costly).⁶ If the conflict between principal

⁴Mechanisms involving screening via randomization were shown to be optimal in a similar model without costly verification by Kovac et al. (2013); we discuss the relationship with our exercise extensively.

⁵One could interpret decreasing probabilities of withholding high-quality projects as “decreasing skepticism in selection,” analogous to decreasing skepticism in verification; our paper simply focuses on verification dynamics.

⁶One interpretation of this mechanism is that the agent *requests verification*, after which the efficient action is taken. This implementation coincides with the “escape clause” mechanisms of Halac and Yared (2020); from this perspective,




Figure 3: Dynamics of verification with a short horizon and static bias (and low verification cost).

and agent is large in the one-period problem, then the principal should verify *with probability 1* in the final period with a short horizon. The optimality of verification in the static problem underlies the emergence of increasing skepticism. This difference flips the above intuition and yields opposite dynamics. Figure 3 illustrates a representative solution. Before the deadline, the principal can dispense with verification to some extent, as the possibility of a future selection dissuades the agent from misrepresenting quality and risking punishment. But as the prospect that a selection might never occur grows, the verification probability must increase.

This discussion distinguishes the roles of *static* and *dynamic bias* in determining optimal verification dynamics. Decreasing skepticism necessarily obtains if there is no static bias, in that both principal and agent would do at least as well from selecting rather than not, despite dynamic bias, in that the principal prefers waiting, whereas the agent does not. In our setting, static bias means that verification has value even in the one-shot problem. But given a sufficiently long horizon, *any* static bias can be overcome. As the probability of being stuck with low quality shrinks, environments with static bias resemble those without it, and decreasing skepticism is optimal. Still, whether skepticism is decreasing or increasing is endogenously determined by the dynamic nature of our problem. Our results delineate the amount of static bias necessary to reverse the verification dynamics given any fixed horizon length.

While these results assume commitment, versions of our main message survive even when such commitment is absent. We show this by determining principal-optimal outcomes in versions of the model that replace full commitment with a notion of renegotiation-proofness (Section 6.2) and with within-period commitment only (Section 6.3). The optimal mechanism in the commitment benchmark involves two simple absorbing outcomes: (i) allocating after sufficiently many l reports, and (ii) withholding allocation indefinitely if allocation does not occur after a report of h . Commitment ensures that both outcomes are implementable. But if this is no longer

our analysis describes how such mechanisms should be modified once dynamics are present.

the case once commitment is relaxed, then the verification probabilities must change to preserve incentive compatibility. However, the key insight that l reports induce higher continuation utility promises, which in turn cause verification probabilities to decrease, continues to hold even under the alternative assumptions we study.

In the context of our leading application, our results identify how the appropriate exercise of verification activities should vary over time, findings that are relevant for determining whether a failure to engage in due diligence amounts to gross negligence, for instance.⁷ We obtain empirically relevant predictions on how institutional factors—such as CEO bias, the size of the acquisition pool, and the ease of verification—affect the level and evolution of oversight in dynamic selection settings. However, the intuition behind decreasing skepticism strikes us as broadly relevant to selection problems with costly verification where resampling is possible.⁸ Our findings would be relevant to cases where the board instead seeks to be acquired, provided the relevant dimensions of acquirer quality are similarly not immediately observable to the board. Costly verification can take a variety of forms—e.g., the costs associated with hiring a third-party, or simply the need to undertake extra effort with less expertise to determine quality independently. While our model may apply more to some situations than others, Graham et al. (2020) provides broad empirical support for the general finding that optimal oversight of CEOs by boards decreases over time. We believe our model captures salient elements of many selection problems, with our analysis being relevant to such situations as well.

2. Related Literature

We study a dynamic version of delegated project selection, where a biased agent has information regarding the payoffs associated with each project in some set. Static models studying project selection have a long tradition in economic theory, applicable to organizational economics (Aghion and Tirole, 1997), corporate finance (Berkovitch and Israel, 2004), and merger choice (Armstrong

⁷A prominent example of such a case is the 2011 acquisition of Autonomy by Hewlett-Packard, which was sought in order to pivot HP toward software. While Autonomy was bought for \$11 billion, HP subsequently recorded an \$8.8 billion write-down attributed to accounting improprieties. Shareholders alleged the board had engaged in gross negligence by rushing due diligence, reportedly limiting the financial review to just six hours. See Fox (2021) for a detailed description of this case.

⁸Adams et al. (2010) provide an authoritative discussion of the role of boards in corporate governance, both in terms of monitoring the information available to CEOs and guiding organizational decisions. The interest in the former is driven in part by the requirement of the Sarbanes-Oxley Act of 2002 that a board's audit committee maintain independence from the organization; this point is discussed in Song and Thakor (2006); Goel and Thakor (2008); DeFond et al. (2005). In particular, incentives of auditors should depend solely on the work provided by the audit, rather than other organizational characteristics, a feature of our model.

and Vickers, 2010; Nocke and Whinston, 2013).⁹ The dynamic decision problem partly resembles McCall (1970)'s (single-agent) job search model, where job offers yield i.i.d. payoffs available for one period. Lewis and Ottaviani (2008); Lewis (2012) study principal-agent versions of this model.¹⁰

A defining feature of our model is the ability to verify private information at a cost. While Townsend (1979) introduced costly verification in a principal-agent model with transfers, we follow work where monetary transfers are infeasible—notably Ben-Porath et al. (2014), characterizing optimal costly verification for (static) multiagent object (or project) selection. Other previously studied applications include voting (Erlanson and Kleiner, 2020), delegation (Halac and Yared, 2020), or more general mechanism design problems (Ben-Porath et al., 2019).¹¹ In dynamic settings, Popov (2016) studies repeated risk sharing with costly verification, whereas Epitropou and Vohra (2019) studies a version of Ben-Porath et al. (2014) with short-lived agents arriving sequentially. Piskorski and Westerfield (2016) study a principal-agent project-management problem featuring moral hazard. Outside of project selection, Wang et al. (2016) studies a dynamic mechanism design problem where a regulator, able to use both transfers and costly inspections, incentivizes a firm to disclose an adverse event arriving at a Poisson rate. The optimal policy features inspection at deterministic dates, with rewards varying between inspections and resetting afterward.

Since our principal cannot directly reward the agent, the agent prefers to terminate search earlier. This conflict is common in organizations. Varas (2018); Grenadier et al. (2016); Escobar and Zhang (2021) study different ways organizations can navigate agent preferences for sooner action. Our paper builds on the principal-agent stopping problem in Kovac et al. (2013) most directly. The most significant substantive difference is the availability of costly verification in our setting, absent there. Verification provides an additional tool for the principal, and we show how it influences the tradeoffs that would otherwise give rise to randomized selection rules, as in Kovac et al. (2013). Still, our analysis applies and extends the innovations of Kovac et al. (2013).

In our model, high quality arrives at an exponentially distributed time in the limit as period

⁹In these papers, firms (as agents) have private information regarding the profitability of mergers and an antitrust authority (as principal) can decide which mergers to approve. Nocke and Whinston (2010) introduce dynamics into this setting by studying a repeated version of this problem.

¹⁰Chade and Kovrijnykh (2016) studies delegated search featuring moral hazard, where the agent exerts effort to acquire information about that period's project. Ulbricht (2016) incorporates ex-ante adverse selection jointly with moral hazard, while Bajoori and Wirtz (2022) considers a version of that problem without transfers or moral hazard.

¹¹Also related is Mylovanov and Zapechelnyuk (2017) on *ex-post* punishments instead of *ex-ante* verification, where the prospect of an exogenous fine deters lying. Another related technology previously studied in dynamic settings is *costly inspections*, where inspections provide information about past actions. By contrast, costly verification in our setting only provides information about the *current* period's project. Ball and Knoepfle (2023); Varas et al. (2020) consider infinite-horizon problems where an underlying state evolves according to a Poisson process; the resulting solution is stationary (unlike ours). Marinovic et al. (2018) study a version of the latter model with disclosure.

length vanishes. Other work uses similar formulations to study verifiable disclosure. Costly verification produces hard information as in disclosure, but with probability endogenously determined by the principal as a way of managing the agency conflict. Knoepfle and Salmi (2024) consider disclosure policy design when information is generated by product adoption by forward-looking agents. Che and Hörner (2018) study a similar problem with myopic agents. Marinovic and Varas (2016) consider disclosure of an underlying state in a model where this state switches between good and bad according to a Poisson process. Gratton et al. (2018); Zhou (2024) consider an agent receiving information privately and deciding when to start an exogenous disclosure process; the latter in particular also studies the incentives to *stop* disclosing. Our use of costly verification as part of a project-selection problem is a distinguishing feature of our exercise. Verification influences project selection by essentially *rewarding the agent with commitment*; our intuitions thus hinge on the principal acquiring the hard information, rather than the agent revealing it.

Beyond disclosure models, Antler et al. (2023) study a two-player game where players acquire information via a Poisson learning technology one after the other, meaning the second mover effectively can verify the first mover’s claims. In a setting absent verifiable information, Green and Taylor (2016) consider incentivizing an agent to reveal privately-observed intermediate breakthroughs arriving at an exponential rate, showing “soft deadlines” with stochastic termination are optimal. Curello and Sinander (2025) study a principal, able to determine a utility profile for herself and the agent, incentivizing an agent to disclose a breakthrough when biased *against* doing so (unlike in our case, where the agent prefers selection to occur *sooner*). Curello and Sinander (2025) identify a deadline structure for mechanisms satisfying an undomination criterion.

The single-selection constraint in our model is also present in Ben-Porath et al. (2014) and much subsequent work. A contrasting assumption is that the principal faces a stream of projects which do not a priori interact. Among these papers, Malenko (2019) is likely closest, studying capital allocation with both costly verification (like us) and transfers (unlike us). While details differ, it is precisely the single-selection constraint in our problem that yields non-trivial verification dynamics *even absent* a conflict between principal and agent in the static problem. By contrast, Malenko (2019) obtains nontrivial dynamics via the static conflict that results from the presence of an intensive margin. Without costly verification, Lipnowski and Ramos (2020); Guo and Hörner (2018) study repeated allocation facing a biased agent. While we study the implications of relaxed commitment, our main model assumes full commitment.

3. Model

A principal (she, e.g., the board of directors) and an agent (he, e.g., the CEO) interact in discrete time, at times $t = 1 \cdot \Delta, 2 \cdot \Delta, \dots$, where we take the horizon to be infinite throughout the paper except in Section 6.1. The principal can make a selection once and only once and must decide when to do so. We refer to selection and allocation interchangeably. At each time t , the agent observes the quality of a newly sampled *project*¹² (e.g., a new potential acquirer); we interchangeably refer to project quality as “project type.” Project type can take one of two numerical values, denoted by $\theta \in \Theta = \{h, l\}$ (where $h > \max\{l, 0\}$), with h referring to “high quality” and l “low quality.” In every period, $\theta = h$ with probability $q = \lambda\Delta \in (0, 1)$ and $\theta = l$ with probability $1 - q$, independently across rounds. We let $r = qh + (1 - q)l$ denote the expected value of quality.

The value of selection at time t depends on project quality at time t only. If the principal selects at time t when the project type is θ , then the principal obtains (undiscounted) payoff equal to θ . We take the agent’s value of selection to be normalized to 1; that is, the agent only cares about the probability of selection and not on the project type. The principal can alternatively withhold the allocation, which delivers a payoff of 0 to both the principal and the agent at time t .

The agent observes project quality perfectly¹³ and privately in every period, and does not incur costs when evaluating projects.¹⁴ The agent also has no private information about *future* projects and only learns the quality of a given project once it is drawn. The principal cannot (costlessly) observe project quality, despite knowing that the agent observes a new project every period (so that there is no moral hazard, as in Malenko (2019); Escobar and Zhang (2021))—for instance, if the firm can costlessly see the prospective acquirer does in fact exist, even without knowing its quality. At each time, the principal can verify that period’s project at a cost $c > 0$; doing so perfectly reveals the project type. The agent’s cost of providing information is zero.

The two parties share a common discount factor $\delta = e^{-\rho\Delta}$. Thus, if the allocation of a project of type θ occurs at time t , and the principal verifies the agent’s type at times $t_1, \dots, t_k \leq t$, then from the perspective of time 1, the payoff of the agent is δ^{t-1} , whereas the principal’s is:

$$\delta^{t-1}\theta - c \sum_{j=1}^k \delta^{t_j-1}.$$

¹²The “project” terminology follows Armstrong and Vickers (2010) on delegation applied to merger choice.

¹³The assumption on observed project quality is common in the literature on verification, as in Ben-Porath et al. (2014); it is relaxed in Khalfan (2023), illustrating different distortions in the second-best solution when verification contains *additional* information about project quality relative to when the agent observes this perfectly.

¹⁴Assuming sampling were costly for the agent would not meaningfully change the analysis, but the costless benchmark is sensible when sampling is not difficult.

We briefly discuss the solution when the type is observable, which we use as our first-best benchmark. In this case, observing a type h always implies allocation should occur. Following type l , the principal can still allocate or wait further. The payoff from waiting for a type h agent is the solution for V in the following equation:

$$V = qh + \delta(1 - q)V \Rightarrow V = \frac{qh}{1 - \delta(1 - q)}.$$

From now on, we assume that the first-best solution does not involve immediate allocation—that is, allocation in the initial period with probability 1 (irrespective of project quality). Thus, we take $qh + (1 - q)l < \frac{qh}{1 - \delta(1 - q)}$; we note that this holds in the $\Delta \rightarrow 0$ limit when $l(\lambda + \rho) < h\lambda$.

We highlight that the principal *may or may not* do worse from allocating a project of quality l than not allocating. If $l < 0$, this presents a second source of conflict. Still, allocations of low-quality projects when $l < 0$ may be useful toward incentivizing the agent to be truthful about project quality. This type of bias is more common in static settings, which often feature agents preferring “higher actions” and principals preferring actions more tailored to the state.

The principal can access two independent public randomization devices in each period, whose outcomes in period t are $a_t \in A_t = [0, 1]$ and $b_t \in B_t = [0, 1]$ drawn according to some continuous distribution.

- A (stochastic) *stage mechanism* at period t is a mapping $m_t : A_t \rightarrow (S_t, x_t, p_t)$ where S_t is the message space for the agent, $x_t : S_t \times B_t \rightarrow \{0, 1\}$ specifies the verification decision, and $p_t : S_t \times \{h, l, \emptyset\} \times B_t \rightarrow \{0, 1\}$ specifies the allocation decision.
- We denote by H_t^p the set of histories on which the principal can condition her mechanism at each period t , with typical element:

$$h_t^p = (m_1, a_1, s_1, b_1, r_1, \dots, m_{t-1}, a_{t-1}, s_{t-1}, b_{t-1}, r_{t-1}),$$

where m_j is the mechanism announced at period j , a_j and b_j denote the realizations of the two public randomization devices, s_j is the agent’s message at period j , and $r_j \in \{h, l, \emptyset\}$ denotes the verification outcome at period j . Since the game ends immediately if the principal allocates, it is without loss to omit the allocation outcomes.

- A *stage mechanism strategy* M_t is a mapping from H_t^p to the set of all stage mechanisms at period t . A (*dynamic*) *mechanism* \mathcal{M} is a sequence of stage mechanism strategies, namely:

$$\mathcal{M} = (M_1, M_2, \dots).$$

- An *agent strategy* at period t is a mapping $\sigma_t^a : H_t^a \rightarrow \Delta(S_t)$, where H_t^a is the history the agent has access to at period t . A typical element is:

$$h_t^a = (m_1, a_1, \theta_1, s_1, b_1, r_1, \dots, m_t, a_t, \theta_t),$$

where $\theta_j \in \Theta$ is the type at period j , and all other components are as described above.

In principle one could allow the agent to send messages after verification, or allow multiple stages of verification within period. However, these capabilities would not change our analysis so we will omit them from our model. Our baseline model assumes that the principal can commit to an arbitrary dynamic mechanism, although we will subsequently relax this commitment power.

We conclude with a discussion of our model. The literature on corporate governance has studied project selection recognizing the importance of CEO private information (e.g., Baldenius et al. (2014); Gregor and Michaeli (2024), as well as those discussed in the introduction), which may arise from a variety of channels—for instance, experience with the industry or access to network connections (see El-Khatib et al. (2015) for a discussion of this idea). However, the literature has also recognized the potential for CEO bias; Décaire and Sosyura (2024) empirically show how a CEO’s private asset holdings influence company investment decisions. Allowing these possibilities to emerge over time incorporates a realistic feature into these settings (Feng et al., 2024).

While our leading example involves boards tasking a CEO with making a choice, our framework also extends to capital budgeting within organizations more broadly. For instance, headquarters may need to allocate capital to a division whose manager is biased toward more investment. In these applications, the assumption of a single selection (investment) may seem more limiting, but in practice, organizations have many reasons to limit consideration to a single investment decision at a time—e.g., if organizational focus is limited (Rotemberg and Saloner, 1994, 2000; Dessein et al., 2016). In our model, if an investment could be made every period, the problem would be trivialized when $l \geq 0$: the principal would simply allocate without verifying in every period. The capacity constraint generates nontrivial dynamics, motivating our focus.

4. Optimal Mechanisms

We now turn to a description of the optimal mechanism for the model in Section 3. In the process, we describe the intuition for why the mechanism takes the form that it does and highlight the practical implications. We defer our more formal derivation to Section 5.

We start by invoking a revelation principle argument for our setting. Concretely, it turns out to be without loss to restrict to mechanisms where:

- At every time t , the agent reports a type $\hat{\theta} \in \{l, h\}$,
- The stage mechanism at time t specifies a verification probability which only depends on the current report, and
- The stage mechanism then decides whether to allocate in that period, where this decision depends on the current report and verification outcome within the period.

Online Appendix E presents a formal proof.¹⁵

Whether the optimal allocation rule is deterministic as a function of type reports and verification outcomes hinges on whether $c \leq h$. In general, the principal may benefit from randomized allocation as a means of facilitating screening.¹⁶ Indeed, if the principal were to allocate deterministically without verification, the agent would simply choose the message sequence yielding the highest discounted probability of allocation, no matter what the actual type sequence were. Screening is thus impossible under such mechanisms. However, suppose the principal were to use a stage mechanism which, after a report of h , allocates with intermediate probability and otherwise *never* allocates. Promises to increase this probability over time can induce an agent with a type- l project to report truthfully, exploiting the agent's indifference between receiving more favorable terms tomorrow and the possibility of allocation today. While the principal may fail to allocate after some h draws under such a mechanism, such losses can be worthwhile if they allow the principal to discriminate on quality.

However, when $c \leq h$, mechanisms that rely on randomized selection in order to screen the agent can be improved by instead using the verification technology. To see why, suppose the principal uses a mechanism that screens in the way described above—specifically, allocating in period t following a report of h with probability p_t and withholding forever with probability $1 - p_t$. Denote the agent's continuation payoff following a report of l under this mechanism as δu_{t+1} . While incentive compatibility requires $p_t \leq \delta u_{t+1}$, optimality requires this hold with equality (since the principal would maximize the probability of allocating following type- h). If the principal instead uses verification, she can promise to withhold allocation only when a lie is verified. Suppose she were to do so in period t , keeping the continuation mechanism otherwise unchanged. If the principal verifies with probability x_t , then incentive compatibility requires

¹⁵The version of the revelation principle we prove considers various cases of principal commitment power, which we study in extensions. While the version for commitment follows from known results (e.g., Sugaya and Wolitzky (2021)), the elaborations we show are new to the best of our knowledge.

¹⁶As far as we know, Kovac et al. (2013) first identified such screening via randomization mechanisms as optimal in stopping problems. One of our contributions is to show how verification enables the principal to dispense with this randomization, and in particular that verification is optimally backloaded.

$1 - x_t \leq \delta u_{t+1}$. Since verification is costly, x_t should be minimized. Thus, this inequality should bind, so that $x_t = 1 - \delta u_{t+1}$. So, while the cost of verification following a report of h is $c(1 - \delta u_{t+1})$, the probability of allocating at this history increases by $1 - \delta u_{t+1}$. This modification maintains incentive compatibility, while yielding a net gain for the principal proportional to $h - c$.

While this particular modification no longer benefits the principal when $c > h$, there may still be gains to verification even in this case. The conclusion that the principal may optimally utilize verification even when doing so ensures negative ex-post payoffs is a consequence of dynamics. Indeed, when $\delta = 0$ the principal would never verify when $c > h$; similar observations hold in other static settings (e.g., Ben-Porath et al. (2014), where verification is not useful if more costly than the maximum benefit from allocation). When δ is sufficiently large, verifying with even a small probability can incentivize truthful reporting, since the cost of delay is small relative to the risk of forgoing allocation entirely. Importantly, however, this logic alone does not explain why verification can dominate stochastic allocation. Identifying gains to verification over randomized selection when $c > h$ is more subtle, and hinges on how verification in one period influences incentives in *other* periods. We develop this argument in detail in Section 4.2. Overall, our model delivers an empirically relevant consequence of the presence of dynamics: verification may yield benefits even when its costs are large relative to the gains from quality.

4.1. Decreasing Skepticism with $c \leq h$

We start with the headline setting, when verification costs are sufficiently low (that is, $c \leq h$):

Theorem 1. *Suppose $c \leq h$, and set $T^* := \lceil \frac{(1-q)(\delta r - l)}{qc(1-\delta)} \rceil$. The optimal mechanism lasts for a finite number of periods, and can be implemented as follows:*

1. *The agent reports whether the current project type is h or l in the first $T^* - 1$ periods; when the agent reports that the project type is l , the interaction continues to the next period.*
2. *If the agent reports h in the i th period, the principal verifies with probability*

$$\frac{(1 - \delta) [1 - [(1 - q)\delta]^{T^* - i}]}{1 - [(1 - q)\delta]}, \quad (1)$$

where this verification probability is such that when the project is type l , the agent is indifferent between truth-telling and misreporting. Allocation occurs if (i) the principal does not verify or (ii) the agent is verified to be truthful.

3. *No allocation occurs if the agent is ever verified to have lied. On path, the agent is truthful.*

4. Any project is selected—irrespective of quality—at time T^* if the agent does not report that the project type is h in the first $T^* - 1$ periods.

The first notable feature of the optimal mechanism is the existence of an *endogenous deadline*, T^* , at which point allocation occurs with probability 1. This feature arises since the agent is optimally given successively larger utility promises until the only sufficiently large reward involves allocation with probability 1. This property of backloading rewards to induce truth-telling is common in dynamic incentive problems and often leads to an endogenous deadline of the kind stated. The deadline itself balances the value of waiting for additional draws against the costs of overcoming the agency conflict:

- The term $\delta r - l$ reflects the relative value of another quality draw compared to selecting low quality. When an additional sample is more valuable by either increasing h or decreasing l , the interaction is (weakly) longer.
- A lower verification cost makes it easier for the principal to overcome the agency conflict, making her more willing to wait longer.
- Fixing r and l , lowering λ (where we recall that $q = \lambda\Delta$) decreases the arrival rate of high-quality projects, leading the principal to set a (weakly) larger deadline to achieve sufficient confidence one will arrive.

Theorem 1 also characterizes the optimal dynamics of verification: only reports of h are ever verified, with a verification probability (1) that is *strictly decreasing* over time. This declining verification probability represents the *decreasing skepticism* feature highlighted above. To see why it emerges, note that due to the deadline structure, the agent obtains a payoff of 1 in the final period. Before this deadline, allocation occurs only once an h project is drawn. Under such an allocation rule, the agent's payoff increases over time as the deadline approaches. While the agent would be tempted to misreport in hopes of being allocated, the cost of doing so is the chance the lie is found out, in which case no allocation is ever made. Since the agent's expected utility increases over time, the loss associated with losing out on allocations correspondingly decreases. Thus, less verification is needed to induce truth-telling as the mechanism progresses.

We comment on some other features of the optimal mechanism. The verification path the principal chooses is determined by the deadline, and set to make an agent with a l project indifferent between lying (and risking losing out on allocation if the principal verifies) and continuing to the next period. Fixing the deadline, these incentives depend only on the discount factor and the probability of drawing a high quality project. But in addition to these parameters, the principal's

payoff and the cost of verification influence the principal’s optimal choice of deadline. As c decreases, the deadline increases as well; in the $c \rightarrow 0$ limit, the deadline approaches ∞ , and the principal’s payoff converges to the payoff achieved in the first-best.

4.2. Backloaded Verification when $c > h$

We now turn to optimal mechanisms when verification costs are high—i.e., when verification ensures negative ex-post principal payoff. A version of the argument presented before Section 4.1 shows that any optimal dynamic mechanism in this case should be stochastic. If $c > h$, a mechanism from Theorem 1 could be improved by instead imposing incentive compatibility in the first period by randomizing the allocation decision (i.e., following an h report, either allocating in that period or never allocating at all) without utilizing verification in that period. Indeed, $c > h$ is a necessary and sufficient condition for this modification to be beneficial.

One simple observation is that the principal would never optimally *simultaneously* utilize verification and randomize allocation in the same period (except perhaps in knife-edge cases)—within any period, there are no interactions between these capabilities. Thus, in any period, the principal induces incentive compatibility using one tool or the other.

Define \bar{c} to be the maximum cost such that verification is used in the principal’s optimal mechanism whenever $c < \bar{c}$. In Online Appendix G, we calculate \bar{c} explicitly,¹⁷ but for our main observations it suffices to note that it is not hard to find parameters such that $\bar{c} > h$ (for instance, taking ρ small so that $\delta = e^{-\rho\Delta}$ is close to 1). We have the following:

Theorem 2. *Suppose $h < c < \bar{c}$. Let $T_R = \lfloor \frac{\log(h/c)}{\log(1-q)} \rfloor + 1$ and $T^* = T_R + \lfloor \frac{(h-c-\frac{l-\delta r}{1-\delta})(1-q)}{qc} \rfloor$. The optimal mechanism involves two phases:*

- A randomization phase when $t \leq T_R$, where:
 - If the agent reports l in period t , the interaction continues to the next period, and
 - If the agent reports h in period t , the mechanism randomizes between allocating at time t or never allocating, with the randomization probability set so an agent with a type- l project is indifferent between misreporting and truth-telling.¹⁸

¹⁷Our explicit calculation of \bar{c} in Online Appendix G shows that it is “increasing with plateaus” as a function of δ (i.e., increasing and flat on alternating intervals).

¹⁸Note that if the allocation probability is set so an agent with a type- l project is indifferent between misreporting and truth-telling, an agent with a type- h project will also be indifferent between misreporting and truth-telling. This feature contrasts with the verification phase, where incentives are strict for the high type.

- A verification phase at all times after T_R and before the deadline. The mechanism in the verification phase coincides with the description from Theorem 1 (conditional on the deadline).

At the deadline, allocation occurs irrespective of quality; the deadline is either at $t = T^*$ or $t = T^* + 1$.¹⁹

The main insight of Theorem 2 is that when c is so high that verification yields a negative ex-post payoff for the principal, it is optimal for randomization to occur *before* utilizing the verification device—In other words, *verification is backloaded*. The intuition driving backloaded verification stems from the observation that utilizing it acts as a reward for the *agent*: it guarantees that allocation occurs *whenever* h is realized, albeit at a significant cost to the principal when $c > h$. Under randomization, since the agent obtains the same payoff whether h or l is realized, the prospect of future verification therefore makes the agent more willing to be truthful for earlier l realizations. This change, in turn, allows the principal to increase allocation probabilities when h is realized in the randomization phase. The longer the randomization phase, the more periods in which these gains accumulate. Since verifying earlier yields the same static costs but does not facilitate the aforementioned accumulation of gains in the randomization phase, we conclude that backloading verification is a feature of the optimal mechanism.²⁰

We illustrate this accumulation of gains more precisely. Consider a mechanism \mathcal{M} that (i) allocates with probability 1 at some arbitrarily chosen time T^* (provided allocation has not occurred before T^*), and (ii) does not utilize verification up to and including time s . Assume that \mathcal{M} involves randomization probabilities following a report of h at or before time s determined in the way described by Theorem 2. For mechanisms of this form, we can straightforwardly determine how the principal's payoff changes when switching from \mathcal{M} to a mechanism otherwise identical to \mathcal{M} except that the verification technology is utilized at time s .

If the agent's payoff under \mathcal{M} starting at time $s+1$ is u_{s+1} , then incentive compatibility requires that allocation following a report of h in period $s' \leq s$ occurs with probability $\delta^{s+1-s'} u_{s+1}$. Now, if instead the mechanism utilizes verification in period s , the verification probability should optimally be set as $1 - \delta u_{s+1}$, since this makes an agent with type l indifferent between misreporting and truth-telling—increasing the verification probability would simply introduce unnecessary added costs. But since the added verification only occurs if h is first realized at time s , and since this event occurs with probability $q(1 - q)^{s-1}$, the added *cost* of verification in period s is:

¹⁹Our proof explicitly delineates when the deadline is T^* as opposed to $T^* + 1$. The ambiguity arises since the optimal continuation path for the principal may involve *replacing* one period of the verification phase with randomization. Fixing the length of the randomization phase, the length of the verification phase thus depends on whether this replacement delivers a higher payoff to the principal. See the proof for the precise condition.

²⁰Several papers have identified backloaded rewards as a property of optimal contracts that arises to relax incentive constraints; see, for instance, Lazear (1981); Harris and Holmstrom (1982); Thomas and Worrall (1988); Ray (2002). Our contribution is to show that this phenomenon also arises when verification is costly but not prohibitive.

$$\overbrace{\delta^{s-1}(1-q)^{s-1}q}^{\text{Discounted probability } h \text{ is first drawn at } s} \cdot \overbrace{(1-\delta u_{s+1})}^{\text{Probability of verification if } h \text{ is first drawn at } s} \cdot c. \quad (2)$$

To identify the benefits, notice that verification ensures that the principal obtains h with probability 1 if first drawn in period s , implying a gain of h multiplied by the same factor on c in (2). Crucially, however, since verification increases the probability of allocation, the agent's utility under the modified mechanism starting at period s increases by $q(1-\delta u_{s+1})$ with verification at time s , from δu_{s+1} to $q + (1-q)\delta u_{s+1}$. So, the principal can increase the probability of allocation following h with probability $\delta^{s-s'}q(1-\delta u_{s+1})$ in period $s' < s$. We have that the principal's *stage* payoff in period s' therefore increases by:

$$\overbrace{\delta^{s'-1}(1-q)^{s'-1}q}^{\text{Discounted probability } h \text{ is first drawn at } s'} \cdot \overbrace{\delta^{s-s'}q(1-\delta u_{s+1})}^{\text{Increased allocation probability if } h \text{ is first drawn at } s'} \cdot h \quad (3)$$

Comparing (2) and (3), we see that all terms involving δ cancel. We conclude that changing the mechanism in the way described leads to a loss in period s proportional to $(c-h) \cdot (1-q)^{s-1}q$, but gains proportional to $h \cdot (\sum_{s'=1}^{s-1} (1-q)^{s'-1}) (q)^2$. In particular, while the costs in (2) are incurred only at time s , the gains in (3) are obtained in *every period before* s .

The accumulation of gains under verification is concretely reflected in this calculation by the factor $\sum_{s'=1}^{s-1} (1-q)^{s'-1}$ scaling the benefits from the verification phase. This delivers backloading. While the value of introducing verification at time s depends on discounting, this argument shows the benefits of verification scale proportionally to its cost times a function that increases in the number of periods before verification. Therefore, *whether* doing so is beneficial depends only on the length of the initial randomization phase and q ; the discount factor matters only in terms of its influence on the deadline. If it is optimal to switch from randomization to verification at time s , then it is also optimal to do so at any time after s .

In fact, using the identity $\sum_{s'=1}^{s-1} (1-q)^{s'-1} = \frac{1-(1-q)^s}{q}$, algebra reveals that this argument implies verification at time s is optimal whenever:

$$c(1-q)^{s-1} < h.$$

This observation yields the form of T_R in Theorem 2. However, we emphasize that this argument does not constitute a proof of Theorem 2. First, it assumes that the mechanism takes the form identified in the Theorem. By contrast, Theorem 2 claims that such mechanisms are optimal among *all* those described in Section 3. Second, Theorem 2 *also* shows that verification costs will influence the deadline, which our proof takes into account but this argument does not. Still, the

simplicity of this intuition strikes us as suggestive that the backloading phenomenon emerges due to fundamental economic forces, underscoring its potential empirical relevance.

5. Deriving the Optimal Mechanism

We describe the recursive formulation underlying our formal analysis. We characterize the principal's utility by a value function, $V(\cdot)$, using the agent's reservation utility, u , as the state variable. This characterization allows us to determine the maximal value attainable by the principal for each promised utility u , and hence to identify the utility level delivered to the agent in an optimal dynamic mechanism. Crucially, for any reservation utility u , our analysis also characterizes how the principal can optimally deliver this utility—via a stage mechanism and a continuation utility promise. This establishes a direct mapping from the solution of $\max_u V(u)$ to an optimal mechanism.

As mentioned in Section 4, the revelation principle implies we can assume that at every time, the principal decides:

- whether or not to verify the agent's type given a report of $\hat{\theta}$, which we denote $x_{\hat{\theta}}$
- whether to allocate as a function of the report and verification outcome; we let $\hat{p}_{\hat{\theta}}$ denote the allocation decision if the agent reports $\hat{\theta}$ and there is no verification, and $p_{\hat{\theta}\theta}^*$ denote the allocation decision if the agent reports $\hat{\theta}$ and the principal verifies the agent's type is θ .

We let $\hat{u}_{\hat{\theta}}$ denote the promised continuation payoff following a $\hat{\theta}$ report and no verification, and $u_{\hat{\theta}\theta}^*$ denote the promised continuation payoff if the principal verifies the type is θ after a $\hat{\theta}$ report.

5.1. Preliminary Simplifications

Appendix A.1 presents our complete Bellman equation, without any simplifications. To start our analysis, we provide some simple observations that will help us reduce the number of variables we must optimize over. The most immediate observations are the following:

Lemma 1. *The value function $V(\cdot)$ is a concave function supported on $[0, 1]$. The following are properties of the optimal solution in the Bellman equation:*

1. $u_{lh}^* = u_{hl}^* = 0$ and $p_{hl}^* = p_{lh}^* = 0$.
2. $p_{hh}^* = 1$.
3. $x_l = 0$.

These results are relatively straightforward and follow familiar intuition. Concavity of $V(\cdot)$ follows from that the principal has access to public randomization devices. Intuition for the enumerated properties follow from noting that a high report is one which the principal should be skeptical of and that verified lies should not be rewarded. Notice that since the low type is never verified in the optimal mechanism (since $x_l = 0$), we do not have to consider allocation probabilities after the verification of the low type.

We next determine how the principal should deliver utility to the agent over time on-path:

Lemma 2. *In the solution to the Bellman equation, the following properties hold:*

1. $\hat{u}_h = u_{hh}^* = 0$.
2. Let $\hat{a} \equiv \hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l$ denote the total value delivered to the low type. The optimal allocation probability \hat{p}_l and continuation utility \hat{u}_l are given by:

$$(\hat{p}_l, \hat{u}_l) = \begin{cases} (0, \frac{\hat{a}}{\delta}) & \text{if } \hat{a} \leq \delta, \\ (\frac{\hat{a}-\delta}{1-\delta}, 1) & \text{if } \hat{a} > \delta. \end{cases}$$

The first part follows from the fact that it is better to promise as much utility to the agent as possible *immediately* once h is realized, and not defer any promises to the future. Indeed, future utility promises only matter to the agent if allocation occurs with probability less than 1. But in this case, setting a positive utility promise to an h -type agent leaves scope to increase the allocation to this type. Since the best the principal can do is obtain a payoff of h , increasing the allocation probability—while simultaneously lowering the continuation utility to keep the agent indifferent—can only make the principal better off. The second part inverts this logic for the low type, implying that \hat{p}_l should not be any larger than necessary to satisfy the constraints on the promised utility of the agent. Here, this means that the principal should set the allocation probability to be 0 if and only if the latter is strictly less than 1—that is, to only start allocating once it is no longer possible to further increase the agent’s continuation utility.

Lastly, we have the following result:

Lemma 3. *The low type’s incentive compatibility constraint holds with equality; the high type’s incentive constraint holds even if dropped.*

The intuition for this result follows from the fact that the information asymmetry is less significant for the high type than the low type. While the high type may *also* be indifferent between reporting truthfully and lying, we need not consider their incentives in our optimization problem.

Using these lemmas, our problem now involves four free variables— $x_h, \hat{p}_h, \hat{p}_l, \hat{u}_l$ —and, in addition to the restrictions on \hat{p}_l identified in Lemma 2, two other constraints: First, the low type’s incentive compatibility constraint:

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = (1 - x_h)\hat{p}_h, \quad (4)$$

and second, a promise-keeping constraint:

$$u = q[x_h + (1 - x_h)\hat{p}_h] + (1 - q)[\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l].$$

Note that in particular we can write (4) as an equality, rather than an inequality, due to the observation in Lemma 3 that this constraint binds.

5.2. Three-Region Form

Though simpler, the Bellman equation still involves choice variables that interact with each other non-linearly—specifically, the allocation and verification probabilities for the high type. We address this nonlinearity by performing a simple change of variables. The idea is to replace this non-linear term with a single variable representing the probability a high-type report is not verified but allocated. Specifically, we set $y = (1 - x_h)\hat{p}_h$, and rewrite the problem as:

$$V(u) = \max_{x_h, y, \hat{p}_l, \hat{u}_l} q[(h - c)x_h + hy] + (1 - q)[l\hat{p}_l + (1 - \hat{p}_l)\delta V(\hat{u}_l)] \quad (5)$$

subject to the low type’s incentive compatibility constraint:

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = y,$$

where, by Lemma 2, $\hat{p}_l = 0$ and $\hat{u}_l = y/\delta$ if $y \leq \delta$, while $\hat{p}_l = (y - \delta)/(1 - \delta)$ and $\hat{u}_l = 1$ if $y > \delta$. Furthermore, since the incentive-compatibility constraint binds, promise-keeping becomes:

$$u = y + qx_h. \quad (6)$$

Our general analysis involves calculating the derivative of the value function to determine how the value function changes as u decreases from 1. Let $x_h(u)$ and $y(u)$ denote the solutions for x_h and y in the principal’s maximization problem for some promised utility u . The following result identifies a three-region form of the value function, characterizing how the optimal stage mechanism changes as the promised utility changes and identifying how to set $x_h(u)$ as u changes:

Lemma 4. *Suppose $c < \bar{c}$. There exist $u_{low} < u_{high}$ such that:*

1. $x_h(u) = 0$ when $u \in [0, u_{low}]$.
2. $\frac{dV(u)}{du} = h - c$ when $u \in [u_{low}, u_{high}]$, and
3. $x_h(u) + y(u) = 1$ when $u \in [u_{high}, 1]$.

An important observation is that the term $h - c$ determines how much a change in the verification probability x_h can affect the value function. To see this more formally, suppose that if y is fixed, say at y^* , and consider an increase in x_h to $x_h + \varepsilon$. Using the identity (6), since y is fixed at y^* , this modification leads to an increase in u of $q \cdot \varepsilon$, while the objective in (5) increases by $q(h - c) \cdot \varepsilon$. We thus conclude that, on any region where $y(u)$ is fixed, the slope of the value function must be $h - c$ —the second case of Lemma 4. Now, since concavity of the value function implies $\frac{dV(u)}{du}$ is weakly decreasing, the equality $\frac{dV(u)}{du} = h - c$ can only hold in an intermediate interval, which inspection reveals is $[y^*, y^* + q(1 - y^*)]$ for some value of y^* . Outside of this interval, the feasibility constraints on $x_h(u)$ must be binding. Which feasibility constraint binds hinges on how the slope of the value function compares to $h - c$. In the third case, the slope of the value function is less than $h - c$, so $x_h(u)$ is maximized and $x_h(u) + y(u) = 1$. In the first case the slope is greater than $h - c$, so $x_h(u)$ is minimized, implying that $x_h(u) = 0$. Concavity implies that the region where $x_h(u) = 0$ is to the left of the region where $\frac{dV}{du} = h - c$, while the region where $x_h(u) + y(u) = 1$ is to the right.

Using Lemma 4, we therefore set $x_h = 1 - y$ for high values of u and $x_h = 0$ for low values of u . In either of these cases, we see that y is a linear function of u . Lemma 4 also implies that in the region where x_h adjusts as u changes, the slope of the value function is exactly $h - c$.

5.3. Tracing Procedure

Our previous observations allow us to derive the solution to this Bellman equation by “tracing out” the value function. Briefly, the idea is to start at $u = 1$, and to use the solution for $V(u)$ at higher values of u to infer the solution at lower values of u , using the structure of the solution described so far. We now describe how this procedure works.

Finding the value of $V(1)$ is immediate, since the only way to deliver reservation utility 1 to the agent is by setting $y = 1$ and $x_h = 0$ —which also implies $p_l = 1$ by Lemma 2. But since $u = y + qx_h$, the constraints on x_h and y also imply that we must have $y > \delta$ whenever u is sufficiently close to 1, so that $\hat{u}_l = 1$ for any such u (pinning down the continuation mechanism to be allocation with probability 1). Inspecting (5) and the constraints, and noting that $x_h = 1 - y$,

we conclude that $V(u)$ is linear in u for an interval of values whose right endpoint equal to 1. Specifically, we have $u = y + q(1 - y)$ and $y = \hat{p}_l + (1 - \hat{p}_l)\delta$ in this interval. Adjusting \hat{p}_l from 1 to 0 traces out the rightmost linear segment of the value function.

Since \hat{p}_l cannot be negative, the situation changes once $\hat{p}_l = 0$, i.e., when $y = \delta$. By Lemma 2, while \hat{p}_l and \hat{u}_l interact multiplicatively (and hence nonlinearly), only one of the two variables changes at a time—in particular, once \hat{p}_l cannot be lowered further, \hat{u}_l adjusts while \hat{p}_l is fixed at 0. At the value of u such that $y = \delta$ —i.e., when switching from the second case to the first case in Part 2 of Lemma 2—the added constraint on the solution implies the slope of the value function changes, yielding a kink.

Continuing further in the region corresponding to Lemma 4’s Case 3, the (binding) low-type IC constraint, $\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = y$, simplifies to $\delta\hat{u}_l = y$ in this next region since $\hat{p}_l = 0$. Since $x_h = 1 - y$ for high values of u by Lemma 4, promise keeping immediately implies $u = q + (1 - q)y$. So, further decreases in u are achieved by decreases in y (which correspond to increases in x_h since $x_h = 1 - y$). Having showed that V is a linear function of u in the region where $\hat{u}_l = 1$ —say, for $u \in [u_1, 1]$ —we have that the value function is *also* linear if $y/\delta \in [u_1, 1]$, which is the case for some interval $[u_2, u_1]$.

In Appendix A.4, we show that the derivative at any u in Lemma 4’s Case 3 is equal to a constant, $\frac{q}{1-q}c$, plus the derivative of the value function evaluated at the continuation utility. In particular, not only is there a kink at u_1 , but *any* kink at u_l generates a further kink at u when $\hat{u}_l(u) = u_l$. The same reasoning holds as we continue to trace out V . For each u in case (3) of Lemma 4, the slope of V increases by a fixed amount each time a successive (lower) kink is reached.

Once the implied slope from these iterations would be greater than $h - c$, we switch to Lemma 4’s Case 2, where $\frac{dV}{du} = h - c$. In this region, changes in u are generated by changes in x_h . Here y is held constant; instead, we trace the value function in this region further by adjusting x_h from $1 - y$ to 0. For lower values of u —i.e., those lower than the value achieved when $x_h = 0$ in this region—we enter Lemma 4’s Case 1. Here we have $x_h = 0$ while y adjusts, with the same observations implying that the value function is still (piecewise) linear. The same reasoning as above implies that the kinks continue to propagate as we lower u further.

Putting all of this together, Figure 4 illustrates the payoff frontier that emerges from this analysis for representative values of the parameters. While our main case of interest is when $c > 0$, observe that when $c = 0$, verification is costless and the principal achieves the first-best. When $c > 0$, verification costs lower the principal’s payoff and introduces additional kinks as described above. Figure 4 also illustrates how continuation values move along the frontier under




Figure 4: Payoff frontier and optimal mechanism dynamics. The dashed line illustrates the first-best frontier. The solid line shows the value function for $h = 1, l = 0, q = 1/2, \delta = 4/5$ and $c = 1/2$. Arrows illustrate the payoff evolution under the optimal dynamic mechanism. Here, the mechanism lasts for up to four periods, and allocation occurs with probability 1 in the final period.

the optimal mechanism. Crucially, since the optimal initial u places the value function at a kink, each successive movement brings the continuation value to a new kink.

5.4. Concluding Observations

Translating these steps into the optimal mechanisms described in Theorems 1 and 2 simply requires (a) determining the value of u such that the *sign* of the derivative of the value function switches, and (b) unravelling the resulting movements in the agent's promised utility to describe the implied optimal mechanism. The above reasoning reduces part (a) to an algebra exercise, since this value of u corresponds to the kink where the slope of the value function switches from positive to negative, which our proofs calculate. Part (b) simply involves associating the optimal u with a deadline, and interpreting each realized \hat{u}_l as the appropriate continuation mechanism following an l report. Specifically, since $V(\cdot)$ is piecewise linear, $\max_u V(u)$ can be solved by setting u to be at a kink of the value function. Lemma 4 allows us to infer the stage mechanism at any u , and the tracing procedure allows us to infer the corresponding \hat{u}_l . This value is either at another kink (corresponding to another interior continuation promise), or $\hat{u}_l = 1$, which corresponds to allocating with probability 1 irrespective of type. The deadline arrives once the latter occurs.

To conclude, we emphasize that Lemma 4 also clarifies why backloaded verification emerges. Recall that the optimal initial u is where the (left) derivative of $V(u)$ switches from negative to positive (when tracing *from* $u = 1$). If $h - c \geq 0$, then the optimum is achieved at a value of u satisfying case (3). Thus, tracing from $u = 1$ down to this value of u does not leave this case, so that the mechanism involves only a verification phase. But if $h - c < 0$, the derivative is still negative




Figure 5: Illustration of the tracing procedure for values of u where $\frac{dV(u)}{du} = h - c$. Arrows point from the value function at some u to $\hat{u}_l(u)$; the dashed part line illustrates the remaining part of the interval before the next kink. In Figure 5a, case (2) of Lemma 4 occurs when the derivative is positive, so this region is never reached in the optimal mechanism on-path. By contrast, the derivative is negative in this region in Figure 5b, so that it is reached on-path in this case.

when case (2) occurs, and hence the optimum is achieved at value of u such that $x_h(u) = 0$. Backloading of verification is a consequence of the observation that the agent's promised utility increases over time, together with the observation that the kinks where $x_h(u) = 0$ are consecutive by Lemma 4. Figure 5 illustrates this contrast by showing the tracing of the value function when u is in the intermediate interval in both the $h \geq c$ and $c > h$ cases.

6. Extensions and Modifications

Our analysis provides a sharp characterization of how entities making a selection should perform costly verification when a biased agent with expertise assesses quality. In acquisition settings of the type studied in this paper, the optimal response to the tradeoff between verification costs and the duty to safeguard shareholder interests involves decreasing the verification probability over time, provided such activities are not too difficult. If such oversight is instead quite costly, then the board should initially rely more heavily on its approval authority to discipline the conflict with the CEO. These results have clear empirical relevance, and we view our work as a roadmap for how costly verification should be used if resampling following a rejection is possible.

We now discuss how these conclusions change under some alternative modeling assumptions.

6.1. Exogenous Deadlines and Increasing Skepticism

As noted in the introduction, one subtlety underlying the emergence of decreasing skepticism is illustrated by comparing to cases where the interaction horizon is *finite*, rather than infinite. In

this section, we let T denote the number of periods in which the principal and agent interact (so that our main model corresponds to $T = \infty$). This modification makes the mechanism's deadline exogenously fixed and potentially very short. This difference can flip the verification dynamics:

Proposition 1. *Let $h - c \geq 0$ and $T < \infty$. Set $\bar{T} := \lfloor \frac{-(1-q)l}{qc} \rfloor$ and let T^* be as in Theorem 1. When $0 < T \leq \bar{T}$, the probability the principal verifies following a report of h increases over time. In this case, the optimal mechanism:*

- *Allocates in a given period only following a report of h that is not a verified lie, and never allocates if any report of h is verified to be a lie.*
- *Verifies reports of h with probability 1 at time T . The verification probability at time $t < T$ is set so that an agent with type $\theta = l$ is indifferent between truth-telling and misreporting. On path, the agent is truthful under this mechanism.*

In Proposition 1, *increasing skepticism* (i.e., verification probabilities increasing over time) emerges as part of the optimal mechanism when $T \leq \lfloor \frac{-(1-q)l}{qc} \rfloor$. Note that this can only arise if the condition $-(1 - q)l \geq qc$ holds, which in turn requires that $l < 0$ —otherwise, $\bar{T} \leq 0$. When $\lfloor \frac{-(1-q)l}{qc} \rfloor < T < T^*$, decreasing skepticism is part of the optimal mechanism, although the horizon is shorter than the optimal deadline in the infinite-horizon solution.

Increasing skepticism arises because, with a short horizon, the interaction *need not* end in allocation, as formalized in the theorem. Intuitively, the verification probability must be 1 in the final period if the principal wishes to discriminate on quality, as is optimal in the cases identified. But since the continuation payoff decreases as the opportunities to draw high-quality dwindle, the agent's continuation utility *decreases* over time. Thus, the agent is *more* tempted to misreport as the interaction progresses, since the punishment after a verified lie—withholding allocation entirely—is less significant. The probability verification occurs must increase to compensate. In other words, with a short horizon the agency problem may increase over time—as the probability allocation never occurs increases—and hence the verification probability increases accordingly. Our observation is that decreasing skepticism necessarily emerges only given a horizon long enough to overcome static bias, in which case the agency conflict relaxes over time as the time at which allocation definitely occur approaches. For a fixed horizon T , Proposition 1 characterizes how large the static bias (parameterized by l) should be before the verification dynamics flip.

We obtain a similar result when $h < c$, although our earlier intuition for the necessity of a randomization phase applies here as well and influences the form of the mechanism. In discussing this case, we focus only on parameter restrictions which can generate increasing skepticism

dynamics. In particular, one necessary (but not sufficient) condition for increasing skepticism is $r < h - c$.²¹ Note that this condition is implicitly assumed in Proposition 1, since if $\bar{T} \geq 1$, then $q(h - c) \geq r$, so $h - c > r$. We therefore maintain this assumption in the following Proposition.

Proposition 2. *Let $r < h - c < 0$ and $T < \infty$. Set $T_1 = \lfloor \frac{\log(h/c)}{\log(1-q)} \rfloor + 1$. Assume $c < \bar{c}$. Then there exists T_2 with $T_1 \leq T_2$ such that*

- *For $T_1 < T \leq T_2$, the optimal mechanism only allocates following an h report that is not verified as a misreport, and never allocates if any misreport is verified. There are two phases:*
 - *The first phase is a randomization phase, where the principal allocates following a report of h with probability equal to the agent’s expected continuation utility, and otherwise never allocates. This phase lasts until T_1 .*
 - *The second phase is a verification phase, where the principal verifies reports of h at time T with probability 1. For earlier times, the verification probability is set so that the low type is indifferent between truth-telling and misreporting.*

In particular, the probability the principal verifies following a report of h increases over time in the verification phase.

- *For $T \leq T_1$, the optimal mechanism never selects a project.*

The proof of Proposition 2 in the Appendix computes the values for T_2 explicitly. This value is less than the deadline described in Theorem 2, as of course otherwise the optimal mechanism under an infinite horizon could be implemented and the verification probability would optimally be decreasing over time. Note the randomization phase length has a fixed length as long as $T > T_1$. When $T \leq T_1$, the horizon is too short for the principal to effectively incentivize the agent enough to obtain a positive payoff; thus, allocation should be withheld entirely. But when $T_1 < T \leq T_2$, the horizon is long enough to make incentivization possible by combining randomization and verification—but not long enough for decreasing skepticism to emerge.

Methodologically, both Proposition 1 and 2 follow similar steps as outlined in Section 5, but using finite horizon Bellman operators instead of finding the fixed point of the infinite horizon

²¹We present a brief sketch of why this is the case. If $r \geq h - c$, then the principal’s value function in the single period problem, say $V^1(u)$, is equal to ru . Piecewise linearity of the value function for finite horizons implies that \hat{u}_t will always be either at an endpoint or a kink in the optimal mechanism. Since there are no intermediate kinks if $V^1(u) = ru$, the mechanism either ends with allocation irrespective of type or never allocating. In the latter case, the solution is the same as the solution with one fewer periods; induction thus shows the optimal mechanism involves never allocating. In the former case, decreasing skepticism emerges for the reasons previously outlined.

operator. Here, the optimization problem becomes:

$$V^T(u) = \max q[(h - c)x_h + hy] + (1 - q)[l\hat{p}_l + (1 - \hat{p}_l)\delta V^{T-1}(\hat{u}_l)],$$

subject to the low type's incentive compatibility constraint and promise-keeping, as well as joint restrictions on \hat{p}_l and \hat{u}_l . All of these constraints are identical to the infinite horizon case—in particular, the restrictions on \hat{p}_l and \hat{u}_l implied by Lemma 2 hold in this case as well. This optimization problem can be solved by backward induction, and the solution will converge to that of Theorem 1 and Theorem 2 as $T \rightarrow \infty$.²²

6.2. Renegotiation

We now illustrate how our methodology applies when relaxing the assumption that the principal can implement arbitrary mechanisms. While the optimal mechanisms identified in Theorems 1 and 2 are simpler than some options the principal could have chosen, they do require withholding allocation indefinitely at some (possibly off-path) histories. If the principal could offer a new mechanism whenever such histories were reached, the agent would be willing to accept any mechanism that led to eventual allocation with positive probability. How would this possibility influence the principal's mechanism choice?

We address these concerns by incorporating a renegotiation-proofness constraint that rules out the possibility that some continuation mechanisms are Pareto dominated by others. This notion of renegotiation is most similar to that proposed by Ray (1994) in the context of repeated games. The difference in our setting is that we do not impose incentive constraints on the principal and thus allow any mechanism satisfying the refinement to be offered, following Zhao (2006)—at any history, the principal can offer a new continuation mechanism that yields a Pareto improvement.²³ Ruling out Pareto-dominated continuations is straightforward in finite horizon interactions using backward inductive reasoning. While this is not possible in infinite-horizon interactions without a terminal history, our notion does so using a recursive formulation. Formally:

Definition 6.1. Fix some $V : [u_{min}, 1] \rightarrow \mathbb{R}$. We say a function $\Gamma V : [\delta u_{min}, 1] \rightarrow \mathbb{R}$ is generated by $V(\cdot)$ if, for every $u \in [\delta u_{min}, 1]$, $\Gamma V(u)$ equals the principal's maximum payoff achievable subject to incentive compatibility and promise-keeping constraints, given that the agent receives utility u and continuation values are determined by $V(\cdot)$. We refer to $\Gamma(\cdot)$ as a generating operator.

²²One subtlety is that, since the finite horizon case does not give us a fixed point of V , we cannot follow the same reasoning as outlined in Section 5 to trace out the value function; indeed, the slope of V^T at a given u need not be the same as the slope of V^{T-1} at the same u .

²³See Bergin and MacLeod (1993) or Watson (2021) for discussions of renegotiation-proofness.

A set of mechanisms \mathcal{M} is renegotiation-proof if the set of payoff pairs induced by mechanisms in \mathcal{M} is exactly $\{(u, V(u)) : u \in [u_{min}, 1]\}$ for some function $V(\cdot)$, where u is the agent's payoff and $V(u)$ is the principal's payoff, and $V(\cdot)$: (1) Is nonincreasing, (2) Satisfies $\Gamma V(u) = V(u)$ for all $u \in [u_{min}, 1]$, and (3) Satisfies $V(u_{min}) \geq \Gamma V(u)$, for all $u \in [\delta u_{min}, 1]$.

We briefly describe why this technical definition captures a natural notion of renegotiation-proofness. In our definition, the generating operator $\Gamma V(u)$ constructs the current value by optimizing over stage outcomes and continuation promises constrained to lie on the set generated by V . The fixed-point condition $\Gamma V(u) = V(u)$ mathematically guarantees that the set of continuation values required to support the current payoffs is identical to the set of payoffs itself. Thus, every continuation mechanism belongs to the set defined by V , so that *recursion* is satisfied. Second, the operator Γ is explicitly defined as a maximization problem (finding the principal's highest feasible payoff for a given agent utility). By requiring V to be the fixed point of this maximization operator, we restrict the set to the Pareto frontier of achievable payoffs, so that *Pareto Optimality* is satisfied. Any mechanism that was strictly Pareto-dominated would lie in the interior of the generated set (yielding a principal payoff strictly less than $\Gamma V(u)$) and thus would not lie on the value function V .

To summarize, this definition captures the following pair of desiderata on any mechanisms \mathcal{M} within a candidate set of renegotiation-proof mechanisms:

1. All continuation mechanisms of \mathcal{M} , conditional on allocation not having yet occurred, is also in the set of renegotiation-proof mechanisms.
2. The expected payoff under \mathcal{M} is not (strictly) Pareto-dominated by that of another mechanism in the set.

Definition 6.1 emerges given an added requirement that any payoff profile that *could* be achieved subject to these requirements is in fact part of the set of renegotiation-proof payoffs.

We emphasize that our goal is to capture the case where the principal would replace some continuation play with a renegotiation-proof mechanism whenever doing so would yield an improvement. For this reason, a difference between this concept and that from Farrell and Maskin (1989); Bernheim and Ray (1989)—aside from our focus on dynamic mechanisms rather than repeated games—is that we allow the principal to offer *any* mechanism provided each continuation is renegotiation-proof, whether or not the mechanism itself is renegotiation-proof. Thus, our notion is more stringent than theirs.

In addition, notice that since V is the principal's maximum payoff, the payoff frontier cannot contain a payoff profile which would make the principal strictly better off without hurting the

agent. However, it is possible to have a payoff profile that would make the agent strictly better off without hurting the principal. Additionally, notice that the domain of ΓV is $[\delta u_{min}, 1]$, since this is the set of agent utilities profiles that can be achieved when continuation utilities lie in $[u_{min}, 1]$. However, agent utility levels in $[\delta u_{min}, u_{min})$ would not arise on-path.

We note that the constraints imposed by renegotiation-proofness on the solution may be consistent with multiple different sets of mechanisms. For instance, if $u_{min} = 1$ —that is, the only continuation involves allocating to the agent with probability 1 irrespective of type—then immediate allocation is renegotiation-proof whenever $r > \max\{0, q(h - c)\}$, since in that case $\Gamma V(u) \leq r$ for all $u \in [\delta, 1]$. But in general it will be possible to find another set of renegotiation-proof mechanisms that allow the principal to do better for these parameters (e.g., if δ is sufficiently large). Intuitively, for such parameters, achieving outcomes that improve upon the single-period optimum requires access to a range of possible continuation utilities, to enable rewards for truth telling or punishments for lying. Our results in this section are focused on the maximum payoff the principal can achieve across all such potential sets of mechanisms.

While the renegotiation-proofness requirement does influence the principal’s optimal mechanism, the decreasing skepticism property remains:

Proposition 3. *The optimal renegotiation-proof mechanism takes the same form as in Theorem 1 and Theorem 2, modified so that the mechanism restarts instead of withholding allocation forever. Specifically,*

1. *It has an endogenous finite deadline T^* at which point the principal allocates irrespective of project quality and does not utilize verification.*
2. *Prior to the deadline, and in the absence of a restart, the mechanism consists of a verification phase and possibly an initial randomization phase. Furthermore, the selection and verification probabilities in each respective stage are determined by the condition that the low type is indifferent between misrepresenting and truth-telling.*

If either (a) the principal verifies a lie during the verification phase or (b) withholds allocation during the randomization phase after the an agent report of h , then the mechanism restarts from the first period, and otherwise continues as specified.

Proposition 3 shows that renegotiation-proofness can be satisfied if the principal restarts the mechanism instead of withholding it entirely. Intuitively, the fact that the agent’s utility increases over time provides scope *both* to incentivize truth-telling, as well as to punish the agent if necessary. In particular, the optimal punishment for the agent is one that ensures the principal’s

payoff is at its highest over the course of the mechanism. While this change does influence the precise probabilities of verification and selection, aside from this feature the mechanism appears qualitatively similar to the full commitment case.

The proof of Proposition 3 follows the same methodology as the derivations for the full commitment case. Appendix A shows that the tracing procedure in Section 5 applies for any set of agent utility promises, $[u_{min}, u_{max}]$; renegotiation-proofness sets $u_{max} = 1$ and u_{min} to be the smallest agent payoff less than 1 such that the value function is downward sloping everywhere.

6.3. Limited Commitment

We now examine how the principal-optimal mechanism changes when she lacks the ability to make intertemporal commitments. Formally, we study following game:

- In every period t , the principal announces a stage mechanism m_t as defined in Section 3, where message sets can in principle be arbitrary. The announced mechanism is enforced in that period.
- The agent then chooses a message from the message set, the verification decision occurs as specified by the mechanism and the chosen message m_t , and the allocation decision occurs as specified by the mechanism, chosen message, and verification outcome.

Importantly, the mechanism for a given period is only announced at the start of the period.

Continuation payoffs are therefore as specified by equilibrium. Our interest is in understanding how well the principal can do under this alternative assumption. Following the literature of dynamic games, we assume the solution concept is a PBE where the principal's strategies are functions of public histories, and the agent's strategies are functions of public histories and his type in the current period.

Our first observation is that it suffices to assume the stage mechanism in every period is a direct revelation mechanism; this follows from our version of the Revelation Principle in Online Appendix E. Given this observation, we follow the same basic approach to characterizing the payoff frontier of the equilibrium set.

Denote the minimum payoff the principal can achieve in any equilibrium by v_{min} . We can characterize the equilibrium set by applying the recursive methodology from Section 5. Specifically, we add a requirement that the principal's payoff at any history when she is called to move lies weakly above v_{min} , since it is without loss to assume off-path punishment for the principal is v_{min} . Our methodology characterizes the boundary of the payoff frontier assuming the agent's

promised utilities lie in a range $[u_{min}, u_{max}]$. Identifying the optimal mechanism thus reduces to finding the values of u_{min} and u_{max} inducing the largest set of self-generating payoffs, with $V(u_{min}) = V(u_{max}) \geq v_{min}$. In our case, the largest set of self-generating payoffs is induced by expanding the the interval $[u_{min}, u_{max}]$ until the principal's participation constraint binds.

To get specific characterizations, we first consider three benchmark strategy profiles:

- *Immediate Allocation*: At any history, the principal commits within the current period to allocate immediately. In each period, the agent truthfully reports if the principal stays on path, and reports the message that maximizes his within-period allocation probability if the principal deviates. The principal obtains r , while the agent obtains 1.
- *Always Verify*: At any history, the principal commits within the current period to verify if h is reported, with allocation only if h is verified. In each period, the agent truthfully reports if the principal stays on path, and reports the message that maximizes his within-period probability if the principal deviates. The principal's expected payoff is $\frac{q(h-c)}{1-(1-q)\delta}$ while the agent's is $\frac{q}{1-(1-q)\delta}$.
- *Never Allocate*: At any history, the principal commits within the current period not to allocate. In each period, the agent truthfully reports if the principal stays on path, and reports the message that maximizes his within-period probability if the principal deviates. Both players obtain 0.

Proposition 4. *Denote the minimum payoff the principal can achieve in any equilibrium by v_{min} . Then*

$$v_{min} \geq \max\{r, 0\}.$$

Furthermore let

$$v^* := \max\left\{r, \frac{q(h-c)}{1-(1-q)\delta}, 0\right\}.$$

Then v^* can be attained in equilibrium by one of the three strategies: *Immediate Allocation*, *Always Verify*, and *Never Allocate*. By definition, $v^* \geq v_{min}$. Specifically, if either $v^* = r$ or $v^* = 0$, we have $v^* = v_{min}$.

The first part of Proposition 4 gives a lower bound on principal's equilibrium payoff, as the principal can always choose to immediate allocate or withhold forever irrespective of the agent's strategy. The second part of Proposition 4 implies the maximum payoff among the three strategies: *Immediate Allocation*, *Always Verify*, and *Never Allocate* can always be sustained as an equilibrium outcome.

We are interested in the principal-optimal equilibrium. When immediate allocation delivers v^* (and thus v_{min}), the principal-optimal equilibrium involves a deadline at which allocation occurs with probability 1, as in our baseline model:

Corollary 1. *Suppose $r > \max\{0, \frac{q(h-c)}{1-(1-q)\delta}\}$. In the principal-optimal equilibrium:*

1. *There is an endogenous finite deadline T^* at which point the principal selects a project with probability 1, for both types, and does not utilize verification.*
2. *Before T^* , the equilibrium path involves at most two phases: a verification phase and possibly an initial randomization phase as in Theorems 1 and 2, modified so that continuation play reverts to the agent-pessimal continuation equilibrium whenever allocation does not occur after an h report. Allocation and verification probabilities in each respective phase are such that an agent with project type l is indifferent between truth-telling and misreporting.*

The conditions of Corollary 1 ensure that withholding allocation forever leads to a principal payoff lower than Immediate Allocation, and hence cannot be sustained in equilibrium. As a result, in the event of a detected lie, play restarts at a different (in particular, nonzero) promised utility for the agent. We mention that this point will typically deliver lower payoff to the agent than the mechanism outlined in Proposition 3.²⁴ Otherwise, the solution involves a decreasing verification probability on-path, culminating with project allocation, as in Theorems 1 and 2.

When $v_{min} > r$ (for instance, $r < 0$), it is not possible to allocate to the agent with probability 1. Thus, it is also not possible to ensure allocation occurs by any finite time. Still, even in these cases, principal-optimal mechanisms involve agent payoffs increasing on-path until the agent-optimal equilibrium is reached. If the agent is either caught lying or reports h and allocation does not occur, play switches to the agent-pessimal equilibrium. The main difference relative to the case covered by Corollary 1 is that the agent-optimal equilibrium involves repetition of the same stage mechanism, which does not guarantee allocation with probability 1. With this caveat, the resulting outcomes still broadly resemble the optimal mechanisms from our main results. See Appendix C for a formal description of these outcomes.

6.4. The $\Delta \rightarrow 0$ limit

This section outlines the solution in the limit as the period length becomes short, i.e., $\Delta \rightarrow 0$, where we obtain closed-form solutions and some additional insights about the form of optimal

²⁴This observation implies that, when $v_{min} = r$, the principal will do worse under the mechanism identified in Proposition 3 than in the equilibrium of Corollary 1.

verification.²⁵ This limit was alluded to in the introduction, where Figures 1,2, and 3 illustrated the time path of optimal verification. Recall that we take $q = \lambda\Delta$ and $\delta = e^{-\rho\Delta}$. Note that, since a period is length Δ , a mechanism that lasts T^* periods lasts a *length of time* equal to $T^* \cdot \Delta$. As $\Delta \rightarrow 0$, $T^* \rightarrow \infty$, but $T^*\Delta$ converges to a finite number.

Corollary 2. *Set $c \leq h$. In the $\Delta \rightarrow 0$ limit, the optimal mechanism lasts for a length of time equal to $H^* := \frac{h\lambda - l(\lambda + \rho)}{c\lambda\rho}$. The optimal probability of verification is:*

$$x_t = \frac{\rho}{\lambda + \rho} \left(1 - e^{-(\lambda + \rho)(H^* - t)}\right). \quad (7)$$

When verification costs are high, we obtain closed form expressions for the length of time of each phase, as well as a simple condition for the verification phase to have non-zero length:

Corollary 3. *Set $c > h$. In the $\Delta \rightarrow 0$ limit, the verification phase has non-zero length if and only if $c \leq \frac{(\lambda + \rho)(h - l)}{\rho}$. In this case, the randomization phase lasts a length of time equal to $\frac{\log(c/h)}{\lambda}$, and the verification phases lasts for a length of time equal to $\frac{(h - l)(\lambda + \rho) - \rho c}{c\lambda\rho}$.*

Note that the fact that the length of the randomization phase is fixed as ρ varies requires verification to be used; when $c = \infty$, so that there is no verification phase, the length of the randomization phase will depend on ρ . Nevertheless, the randomization phase becomes small as a fraction of the total interaction, for any $c < \infty$ in the $\rho \rightarrow 0$ limit. Calculating the verification probability in the verification phase and allocation probability in the randomization phase is straightforward and follows the same steps as Corollary 2.

Corollaries 2 and 3 illustrate the added tractability that the continuous time limit facilitates; e.g., enabling sharper comparative statics by avoiding integer issues that emerge when defining the deadline in Theorems 1 and 2. It is also noteworthy that the condition on c that ensures a non-zero verification phase in Corollary 3 is significantly more elegant than the corresponding discrete-time parameter (presented in Section G.1 in the Online Appendix). In Online Appendix G, we use this tractability to derive this maximum cost for renegotiation-proof mechanisms as in Section 6.2 and under limited commitment as in Section 6.3. In both cases, our derivations show that verification may be useful even when $c > h$, highlighting the generality of this possibility.

7. Conclusion

We study optimal verification of quality claims in a dynamic delegated project selection problem, where an agent samples prospects over time and does not internalize quality. One empirically

²⁵We are grateful to Eddie Dekel for suggesting the continuous-time limit.

relevant implication of our analysis is that variation in scrutiny across settings may reflect differences in the relevant dynamics: More immediate proposals should receive greater scrutiny absent exogenous constraints on the horizon. We showed how the dynamics of optimal verification depend on the costs of verification, patience, and the returns from increased quality. Our analysis sheds light on whether relaxed oversight in acquisition contexts reflects board negligence as opposed to an optimal response to costs, a question courts frequently confront when determining whether dissatisfied shareholders are entitled to legal remedies.

Our analysis also highlighted subtleties toward translating insights from static models of costly verification to dynamic settings. A lesson one might draw from the optimal mechanisms derived in related static settings (e.g., Ben-Porath et al. (2014)) is that, loosely speaking, higher reports tend to be those that are subject to greater scrutiny in the sense of facing a higher probability of verification. In dynamic settings, interpreting “earlier” claims as “higher” and translating this finding would seem to suggest earlier reports should be verified with higher probability. We showed this intuition is incomplete since the conclusion is valid *only* given a sufficiently long horizon; this logic reverses with a shorter horizon, and verification is concentrated at the end.

We maintained an intentionally minimal model to isolate the essential tradeoffs driving the structure of optimal mechanisms. We hope the intuitions and tools developed here prove useful in other contexts where information asymmetries can be eliminated at a cost, and that our work provides a useful benchmark for analyses in more complex environments. While project selection with costly verification has been an active topic of recent research, many applications feature the ability to resample following a rejection. Our work sought to illustrate how this feature influences optimal verification, moving this literature beyond the rigidity of static settings.

Data Availability Statement

The data and code underlying this research is available on Zenodo at <https://dx.doi.org/10.5281/zenodo.19476913>.

References

- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.
- ADAMS, R. B. AND D. FERREIRA (2007): “A Theory of Friendly Boards,” *The Journal of Finance*, 62, 217–250.
- ADAMS, R. B., B. E. HERMALIN, AND M. S. WEISBACH (2010): “The role of boards of directors in

- corporate governance: A conceptual framework and survey,” *Journal of Economic Literature*, 48, 58–107.
- AGHION, P. AND J. TIROLE (1997): “Formal and real authority in organizations,” *Journal of Political Economy*, 105, 1–29.
- ALCHIAN, A. A. AND H. DEMSETZ (1972): “Production, information costs, and economic organization,” *American Economic Review*, 62, 777–795.
- ANTLER, Y., D. BIRD, AND S. OLIVEROS (2023): “Sequential Learning,” *American Economic Journal: Microeconomics*, 15, 399–433.
- ARMSTRONG, M. AND J. VICKERS (2010): “A model of delegated project choice,” *Econometrica*, 78, 213–244.
- BAJOORI, E. AND J. WIRTZ (2022): “Optimal delegated search with learning and no monetary transfers,” *Working Paper*.
- BALDENIUS, T., N. MELUMAD, AND X. MENG (2014): “Board composition and CEO power,” *Journal of Financial Economics*, 112, 53–68.
- BALL, I. AND J. KNOEPFLE (2023): “Should the Timing of Inspections be Predictable?” *Working Paper*.
- BEN-PORATH, E., E. DEKEL, AND B. L. LIPMAN (2014): “Optimal Allocation with Costly Verification,” *American Economic Review*, 104, 3779–3813.
- (2019): “Mechanisms with Evidence: Commitment and Robustness,” *Econometrica*, 87, 529–566.
- BERGIN, J. AND W. B. MACLEOD (1993): “Efficiency and Renegotiation in Repeated Games,” *Journal of Economic Theory*, 61, 42–73.
- BERKOVITCH, E. AND R. ISRAEL (2004): “Why the NPV Criterion Does Not Maximize NPV,” *The Review of Financial Studies*, 17, 239–255.
- BERNHEIM, B. D. AND D. RAY (1989): “Collective dynamic consistency in repeated games,” *Games and Economic Behavior*, 1, 295–326.
- BLISS, R. T. AND R. J. ROSEN (2001): “CEO compensation and bank mergers,” *Journal of Financial Economics*, 61, 107–138.

- CHADE, H. AND N. KOVRIJNYKH (2016): “Delegated information acquisition with moral hazard,” *Journal of Economic Theory*, 162, 55–92.
- CHE, Y.-K. AND J. HÖRNER (2018): “Recommender Systems as Mechanisms for Social Learning,” *The Quarterly Journal of Economics*, 133, 871–925.
- CURELLO, G. AND L. SINANDER (2025): “Screening for Breakthroughs,” *Theoretical Economics*, forthcoming.
- DÉCAIRE, P. H. AND D. SOSYURA (2024): “Self-Dealing in Corporate Investment,” Working paper.
- DEFOND, M. L., R. N. HANN, AND X. HU (2005): “Does the market value financial expertise on audit committees of boards of directors?” *Journal of Accounting Research*, 43, 153–193.
- DESSEIN, W., A. GALEOTTI, AND T. SANTOS (2016): “Rational inattention and organizational focus,” *American Economic Review*, 106, 1522–1536.
- EL-KHATIB, R., K. FOGEL, AND T. JANDIK (2015): “CEO network centrality and merger performance,” *Journal of Financial Economics*, 116, 349–382.
- EPITROPOU, M. AND R. VOHRA (2019): “Optimal On-Line Allocation Rules with Verification,” *International Symposium on Algorithmic Game Theory*, 3–17.
- ERLANSON, A. AND A. KLEINER (2020): “Costly verification in collective decisions,” *Theoretical Economics*, 15, 923–954.
- ESCOBAR, J. F. AND Q. ZHANG (2021): “Delegating Learning,” *Theoretical Economics*, 16, 571–603.
- FARRELL, J. AND E. MASKIN (1989): “Renegotiation in repeated games,” *Games and Economic Behavior*, 1, 327–360.
- FENG, F. Z., R. Y. LUO, AND B. MICHAELI (2024): “In Search of a Unicorn: Dynamic Agency with Endogenous Investment Opportunities,” *Journal of Accounting and Economics*, 78.
- FOX, T. (2021): “Mike Lynch, the HP – Autonomy Deal and M&A Due Diligence (Parts 1–4),” <https://www.jdsupra.com/legalnews/mike-lynch-the-hp-autonomy-deal-and-m-a-6865725/>, series published July 26–29, 2021. Includes Part 1: History; Part 2: Post Closing; Part 3: Anomalies; Part 4: Denouement.

- GOEL, A. M. AND A. V. THAKOR (2008): “Overconfidence, CEO selection, and corporate governance,” *The Journal of Finance*, 63, 2737–2784.
- GRAHAM, J. R., H. KIM, AND M. LEARY (2020): “CEO-Board Dynamics,” *Journal of Financial Economics*, 137, 612–636.
- GRATTON, G., R. HOLDEN, AND A. KOLOTILIN (2018): “When to Drop a Bombshell,” *Review of Economic Studies*, 85, 2139–2172, advance access publication 11 December 2017.
- GREEN, B. AND C. R. TAYLOR (2016): “Breakthroughs, Deadlines, and Self-Reported Progress: Contracting for Multistage Projects,” *American Economic Review*, 106, 3660–3699.
- GREGOR, M. AND B. MICHAELI (2024): “Board bias, information, and investment efficiency,” *Review of Accounting Studies*.
- GRENADIER, S. R., A. MALENKO, AND N. MALENKO (2016): “Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment,” *American Economic Review*, 106, 2552–2581.
- GRINSTEIN, Y. AND P. HRIBAR (2004): “CEO compensation and incentives: Evidence from M&A bonuses,” *Journal of Financial Economics*, 73, 119–143.
- GUO, Y. AND J. HÖRNER (2018): “Dynamic Allocation without Money,” *Working Paper*.
- HALAC, M. AND P. YARED (2020): “Commitment versus Flexibility with Costly Verification,” *Journal of Political Economy*, 128, 4523–4573.
- HARFORD, J. AND K. LI (2007): “Decoupling CEO Wealth and Firm Performance: The Case of Acquiring CEOs,” *The Journal of Finance*, 62, 917–949.
- HARRIS, M. AND B. HOLMSTROM (1982): “A Theory of Wage Dynamics,” *Review of Economic Studies*, 49, 315–333.
- HARRIS, M. AND A. RAVIV (1996): “The Capital Budgeting Process: Incentives and Information,” *The Journal of Finance*, 51, 1139–1174.
- JENSEN, M. C. AND W. H. MECKLING (1976): “The Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure,” *Journal of Financial Economics*, 3, 305–360.
- KHALFAN, N. (2023): “Optimal Allocation with Noisy Inspection,” *Working Paper*.

- KNOEPFLE, J. AND J. SALMI (2024): “Dynamic Evidence Disclosure: Delay the Good to Accelerate the Bad,” *Working Paper*.
- KOVAC, E., D. KRÄHMER, AND T. TATUR (2013): “Optimal stopping in a principal-agent model with hidden information and no monetary transfers,” *Working Paper*.
- LAZEAR, E. P. (1981): “Agency, Earnings Profiles, Productivity, and Hours Restrictions,” *American Economic Review*, 71, 606–620.
- LEWIS, T. R. (2012): “A theory of delegated search for the best alternatives,” *RAND Journal of Economics*, 43, 391–416.
- LEWIS, T. R. AND M. OTTAVIANI (2008): “Search Agency,” *Working Paper*.
- LIPNOWSKI, E. AND J. RAMOS (2020): “Repeated Delegation,” *Journal of Economic Theory*, 188, 105040.
- MALENKO, A. (2019): “Optimal Dynamic Capital Budgeting,” *Review of Economic Studies*, 86, 1747–1778.
- MARINOVIC, I., A. SKRZYPACZ, AND F. VARAS (2018): “Dynamic Certification and Reputation for Quality,” *American Economic Journal: Microeconomics*, 10, 58–82.
- MARINOVIC, I. AND F. VARAS (2016): “No News is Good News: Voluntary Disclosure in the Face of Litigation,” *RAND Journal of Economics*, 47, 822–856.
- MARTELLINI, P. AND G. MENZIO (2022): “The Dean and The Chair,” *Working Paper*.
- MCCALL, J. (1970): “Economics of Information and On-The-Job Search,” *Quarterly Journal of Economics*, 84, 113–126.
- MYLOVANOV, T. AND A. ZAPECHELNYUK (2017): “Optimal allocation with ex-post verification and limited penalties,” *American Economic Review*, 107, 2666–2694.
- NOCKE, V. AND M. D. WHINSTON (2010): “Dynamic Merger Review,” *Journal of Political Economy*, 118, 1201–1251.
- (2013): “Merger Policy with Merger Choice,” *American Economic Review*, 103, 1006–1033.
- PISKORSKI, T. AND M. M. WESTERFIELD (2016): “Optimal dynamic contracts with moral hazard and costly monitoring,” *Journal of Economic Theory*, 166, 242–281.

- POPOV, L. (2016): “Stochastic costly state verification and dynamic contracts,” *Journal of Economic Dynamics and Control*, 64, 1–22.
- RAY, D. (1994): “Internally Renegotiation-Proof Equilibrium Sets: Limit Behavior with Low Discounting,” *Games and Economic Behavior*, 6, 162–177.
- (2002): “The Time Structure of Self-Enforcing Agreements,” *Econometrica*, 70, 547–582.
- ROTEMBERG, J. J. AND G. SALONER (1994): “Benefits of narrow business strategies,” *American Economic Review*, 84, 1330–1349.
- (2000): “Visionaries, managers, and strategic direction,” *RAND Journal of Economics*, 31, 693–716.
- SONG, F. AND A. V. THAKOR (2006): “Information Control, Career Concerns, and Corporate Governance,” *Journal of Finance*, 61, 1845–1896.
- SUGAYA, T. AND A. WOLITZKY (2021): “The Revelation Principle in Multistage Games,” *Review of Economic Studies*, 88, 1503–1540.
- THOMAS, J. AND T. WORRALL (1988): “Self-Enforcing Wage Contracts,” *Review of Economic Studies*, 55, 541–554.
- TOWNSEND, R. M. (1979): “Optimal Contracts and Competitive Markets with Costly State Verification,” *Journal of Economic Theory*, 21, 265–293.
- ULBRICHT, R. (2016): “Optimal delegated search with adverse selection and moral hazard,” *Theoretical Economics*, 11, 253–278.
- VARAS, F. (2018): “Managerial Short-Termism, Turnover Policy, and the Dynamics of Incentives,” *Review of Financial Studies*, 31, 3409–3451.
- VARAS, F., I. MARINOVIC, AND A. SKRZYPACZ (2020): “Random Inspections and Periodic Reviews: Optimal Dynamic Monitoring,” *Review of Economic Studies*, 87, 2893–2937.
- WANG, S., P. SUN, AND F. DE VÉRICOURT (2016): “Inducing Environmental Disclosures: A Dynamic Mechanism Design Approach,” *Operations Research*, 64, 371–389.
- WATSON, J. (2021): “Theoretical Foundations of Relational Incentive Contracts,” *Annual Review of Economics*, 13, 631–659.

ZHAO, R. R. (2006): “Renegotiation-proof contract in repeated agency,” *Journal of Economic Theory*, 131, 263–281.

ZHOU, B. (2024): “Optimal Disclosure Windows,” Working paper.

A. General Methodology for the Infinite Horizon Case

A.1. Full Program

We present a description of the Bellman equation, prior to the simplifications which would emerge using our results. Given some specific promised utility u , this yields the following objective:

$$V(u) = \max q[hx_h p_{hh}^* + x_h(1 - p_{hh}^*)\delta V(u_{hh}^*) + h(1 - x_h)\hat{p}_h + (1 - x_h)(1 - \hat{p}_h)\delta V(\hat{u}_h) - x_h c] \\ + (1 - q)[lx_l p_{ll}^* + x_l(1 - p_{ll}^*)\delta V(u_{ll}^*) + l(1 - x_l)\hat{p}_l + (1 - x_l)(1 - \hat{p}_l)\delta V(\hat{u}_l) - x_l c].$$

The incentive compatibility conditions are: For the high type,

$$x_h p_{hh}^* + x_h(1 - p_{hh}^*)\delta u_{hh}^* + (1 - x_h)\hat{p}_h + (1 - x_h)(1 - \hat{p}_h)\delta \hat{u}_h \\ \geq x_l p_{lh}^* + x_l(1 - p_{lh}^*)\delta u_{lh}^* + (1 - x_l)\hat{p}_l + (1 - x_l)(1 - \hat{p}_l)\delta \hat{u}_l,$$

and for the low type:

$$x_l p_{ll}^* + x_l(1 - p_{ll}^*)\delta u_{ll}^* + (1 - x_l)\hat{p}_l + (1 - x_l)(1 - \hat{p}_l)\delta \hat{u}_l \\ \geq x_h p_{hl}^* + x_h(1 - p_{hl}^*)\delta u_{hl}^* + (1 - x_h)\hat{p}_h + (1 - x_h)(1 - \hat{p}_h)\delta \hat{u}_h.$$

Meanwhile, the promise-keeping constraint is:

$$u = q[x_h p_{hh}^* + x_h(1 - p_{hh}^*)\delta u_{hh}^* + (1 - x_h)\hat{p}_h + (1 - x_h)(1 - \hat{p}_h)\delta \hat{u}_h] \\ + (1 - q)[x_l p_{ll}^* + x_l(1 - p_{ll}^*)\delta u_{ll}^* + (1 - x_l)\hat{p}_l + (1 - x_l)(1 - \hat{p}_l)\delta \hat{u}_l].$$

A.2. General Simplification

The following Lemma is a version of the results stated in Section 5. We generalize them to allow for arbitrary values for the lowest and highest possible agent utility levels. The analysis in Section 5 corresponds to the case of $u_{min} = 0$ and $u_{max} = 1$.

Lemma 5. *The following statements hold for the main model presented in Section 3, as well as the extensions relaxing commitment in Sections 6.2 and 6.3.*

1. $V(\cdot)$ is a concave function supported on $[u_{min}, u_{max}]$.
2. $u_{lh}^* = u_{hl}^* = u_{min}$; $p_{hl}^* = p_{lh}^* = 0$.

3. (a) $\hat{u}_h = u_{hh}^* = u_{min}$.

(b) Let a denote the required utility value (representing either $\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l$ or $\hat{p}_l^* + (1 - \hat{p}_l^*)\delta u_{ll}^*$). The optimal controls are given by:

$$(p, u) = \begin{cases} (0, \frac{a}{\delta}) & \text{if } a \leq \delta u_{max}, \\ \left(\frac{a - \delta u_{max}}{1 - \delta u_{max}}, u_{max}\right) & \text{if } a > \delta u_{max}. \end{cases}$$

where (p, u) corresponds to (\hat{p}_l, \hat{u}_l) or (\hat{p}_l^*, u_{ll}^*) respectively.

4. $p_{hh}^* = 1$.

5. $x_l = 0$.

6. The low type's incentive compatibility constraint binds; the high type's incentive constraint holds even if dropped.

Proof of Lemma 5. $V(\cdot)$ is concave due to the public randomization device; specifically, given any two achievable levels of utility for the agent, u_1 and u_2 , the principal can choose a mechanism with probability α that delivers u_1 to the agent and $V(u_1)$ to the principal, and with probability $1 - \alpha$ a mechanism that delivers u_2 to the agent and $V(u_2)$ to the principal. The principal's payoff under this mechanism is $\alpha V(u_1) + (1 - \alpha)V(u_2)$, so that the optimal mechanism delivering $\alpha u_1 + (1 - \alpha)u_2$ to the agent must lead to an objective at least as large.

The second part follows from inspection of the objective, since minimizing these only makes incentive compatibility easier to satisfy.

For the third part, note a fixed amount of promised utility is equivalent to the future discounted probability of allocation. Thus, it is optimal for the principal to put more weights on high type projects and less weights on low type projects without changing the agent's incentive constraint.

For the first half, suppose $\hat{p}_h + (1 - \hat{p}_h)\delta\hat{u}_h = a$. For the sake of contradiction suppose $\hat{u}_h > u_{min}$, then there exists $\hat{p}'_h > \hat{p}_h$ such that

$$\hat{p}'_h + (1 - \hat{p}'_h)\delta u_{min} = a.$$

In other words, when h is reported and verification doesn't occur, instead of allocating with probability \hat{p}_h and promising \hat{u}_h , the principal could 'deviate' to allocating with probability \hat{p}'_h and promising u_{min} . Because the total promised utility is the same as before, no incentive constraint

will be violated. Choose $0 < b < 1$ such that

$$(1 - \hat{p}_h)b = 1 - \hat{p}'_h.$$

Considering the payoff for the principal, we have

$$\begin{aligned} & \hat{p}'_h h + (1 - \hat{p}'_h)\delta V(u_{min}) \\ &= \hat{p}'_h h + (1 - \hat{p}_h)b\delta V(u_{min}) \\ &= \hat{p}'_h h + (1 - \hat{p}_h)\delta V(\hat{u}_h + u_{min} - \hat{u}_h) - (1 - \hat{p}_h)(1 - b)\delta V(u_{min}) \\ &\geq \hat{p}'_h h + (1 - \hat{p}_h)\delta V(\hat{u}_h) - (1 - \hat{p}_h)\delta(\hat{u}_h - u_{min})h - (1 - \hat{p}_h)(1 - b)\delta u_{min}h \\ &= \hat{p}_h h + (1 - \hat{p}_h)\delta V(\hat{u}_h), \end{aligned}$$

where the inequality comes from the fact that $V'(u) \leq h$ which directly follows from the fact that the derivative is bounded by h at $u = 0$ and $V(\cdot)$ is concave. Thus, the principal would have a weakly better alternative, and it is optimal to have $\hat{u}_h = u_{min}$, same for u_{hh}^* .

For the second half, suppose $\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = a$, for the sake of contradiction suppose $\hat{p}_l > 0$ and $\hat{u}_l < u_{max}$. Then there exist $\hat{p}'_l < \hat{p}_l$ and $\hat{u}'_l > \hat{u}_l$ such that

$$\hat{p}'_l + (1 - \hat{p}'_l)\delta\hat{u}'_l = a.$$

Then choosing $0 < b < 1$ such that

$$(1 - \hat{p}'_l)b = 1 - \hat{p}_l,$$

we have

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = \hat{p}'_l + (1 - \hat{p}'_l)b\delta\hat{u}_l + (1 - \hat{p}'_l)(1 - b)\delta\frac{\hat{u}'_l - b\hat{u}_l}{1 - b}.$$

Note because $\hat{u}_l < u_{max}$ we can find $\hat{u}'_l < u_{max}$ and \hat{u}_l and \hat{u}'_l are close enough such that $\frac{\hat{u}'_l - b\hat{u}_l}{1 - b} < u_{max}$. Consider an alternative dynamic mechanism where, after an l report, the principal:

1. Allocates with probability \hat{p}'_l .
2. If not allocating, with probability b plays the original mechanism which delivers utility \hat{u}_l .
3. If not allocating, with probability $1 - b$ plays a mechanism which delivers utility $\frac{\hat{u}'_l - b\hat{u}_l}{1 - b}$.

Since the total promised utility is the same as before, no incentive constraint will be violated. Now

we consider the trade-off for the principal. Note the payoff is

$$\begin{aligned}
& \hat{p}'_l l + (1 - \hat{p}'_l) b \delta V(\hat{u}_l) + (1 - \hat{p}'_l)(1 - b) \delta V\left(\frac{\hat{u}'_l - b \hat{u}_l}{1 - b}\right) \\
& \geq \hat{p}'_l l + (1 - \hat{p}'_l) \delta V(\hat{u}_l) + (1 - \hat{p}'_l)(1 - b) \delta\left(\frac{\hat{u}'_l - b \hat{u}_l}{1 - b}\right) \cdot l \\
& = \hat{p}'_l l + (1 - \hat{p}'_l) \delta V(\hat{u}_l),
\end{aligned}$$

where the second inequality comes from $V(u) \geq lu$. The same argument works for p_{ll}^* and u_{ll}^* .

For the fourth part, note $\hat{u}_h = u_{hh}^* = u_{min}$. Given x_h, \hat{p}_h , and p_{hh}^* , there are two possible cases. If $x_h p_{hh}^* + (1 - x_h) \hat{p}_h \geq x_h$, let $p_{hh}' = 1$ and $\hat{p}'_h = [x_h p_{hh}^* + (1 - x_h) \hat{p}_h - x_h] / (1 - x_h)$. Clearly we have

$$x'_h p_{hh}' + (1 - x'_h) \hat{p}'_h = x_h p_{hh}^* + (1 - x_h) \hat{p}_h,$$

and

$$\begin{aligned}
& x_h (p_{hh}' + (1 - p_{hh}') u_{min}) + (1 - x_h) (\hat{p}'_h + (1 - \hat{p}'_h) u_{min}) \\
& = x_h (p_{hh}^* + (1 - p_{hh}^*) u_{min}) + (1 - x_h) (\hat{p}_h + (1 - \hat{p}_h) u_{min}).
\end{aligned}$$

This tuning only decreases the right-hand side of IC_l and makes incentive compatibility easier to satisfy.

If $x_h p_{hh}^* + (1 - x_h) \hat{p}_h < x_h$, let $p_{hh}' = 1, \hat{p}'_h = 0$, and $x'_h = x_h p_{hh}^* + (1 - x_h) \hat{p}_h$. Then:

$$x'_h p_{hh}' + (1 - x'_h) \hat{p}'_h = x_h p_{hh}^* + (1 - x_h) \hat{p}_h.$$

This tuning decreases the right-hand side of IC_l as well as the total cost of verification. Thus it is always optimal to set $p_{hh}^* = 1$.

For the fifth part, given $x_l, \hat{u}_l, u_{ll}^*, \hat{p}_l$, and p_{ll}^* , let

$$\hat{u}'_l = u_{ll}' = \frac{x_l (1 - p_{ll}^*) u_{ll}^* + (1 - x_l) (1 - \hat{p}_l) \hat{u}_l}{x_l (1 - p_{ll}^*) + (1 - x_l) (1 - \hat{p}_l)}.$$

Since $V(u)$ is concave in u (because of the public randomization device), we have

$$x_l (1 - p_{ll}^*) \delta V(u_{ll}^*) + (1 - x_l) (1 - \hat{p}_l) \delta V(\hat{u}_l) \leq x_l (1 - p_{ll}^*) \delta V(u_{ll}') + (1 - x_l) (1 - \hat{p}_l) \delta V(\hat{u}'_l).$$

Thus, by setting $\hat{u}_l = u_{ll}^*$, we can increase the principal's utility without violating any constraints. Now given $\hat{u}_l = u_{ll}^*$, we can further let $x'_l = 0$ and $\hat{p}'_l = x_l p_{ll}^* + (1 - x_l) \hat{p}_l$ to minimize the cost of

verification without violating any constraints.

Lastly, note first that by previous lemmas, the problem can be reduced to

$$V(u) = \max q[hx_h + h(1-x_h)\hat{p}_h - x_h c + (1-x_h)(1-\hat{p}_h)\delta V(u_{min})] + (1-q)[l\hat{p}_l + (1-\hat{p}_l)\delta V(\hat{u}_l)],$$

subject to IC_h :

$$x_h + (1-x_h)\hat{p}_h + (1-x_h)(1-\hat{p}_h)\delta u_{min} \geq \hat{p}_l + (1-\hat{p}_l)\delta \hat{u}_l,$$

IC_l :

$$\hat{p}_l + (1-\hat{p}_l)\delta \hat{u}_l \geq x_h \delta u_{min} + (1-x_h)\hat{p}_h + (1-x_h)(1-\hat{p}_h)\delta u_{min},$$

and

$$u = q[x_h + (1-x_h)\hat{p}_h + (1-x_h)(1-\hat{p}_h)\delta u_{min}] + (1-q)[\hat{p}_l + (1-\hat{p}_l)\delta \hat{u}_l].$$

Suppose $x_h = 0$, then both IC_h and IC_l are binding because

$$\hat{p}_l + (1-\hat{p}_l)\delta \hat{u}_l = \hat{p}_h + (1-\hat{p}_h)\delta u_{min},$$

thus there is nothing to prove. Now suppose IC_l is not binding; if $\hat{p}_h = 1$, $x_h > 0$, then by decreasing x_h , the principal can be strictly better off while maintaining reservation utility u . If $\hat{p}_h < 1$, $x_h > 0$, then by decreasing x_h and increasing \hat{p}_h while keeping $x_h + (1-x_h)\hat{p}_h$ unchanged, the principal can be strictly better off while maintaining reservation utility u . Thus the low type incentive constraint binds, and the high type incentive constraint satisfies automatically, completing the proof of the Lemma. \square

A.2.1. Restated, Simplified Program

We summarize the implications of Lemma 5 by restating the optimization problem presented in Appendix A.1. Of course, when the principal has full commitment, we have $u_{min} = 0$ and $u_{max} = 1$, which gives the formulation in Section 5. It may not be the case when commitment is relaxed. In general, we study the following:

$$V(u) = \max q[hx_h + h(1-x_h)\hat{p}_h - x_h c + (1-x_h)(1-\hat{p}_h)\delta V(u_{min})] + (1-q)[l\hat{p}_l + (1-\hat{p}_l)\delta V(\hat{u}_l)],$$

subject to the low type's (binding) incentive compatibility constraint:

$$\hat{p}_l + (1-\hat{p}_l)\delta \hat{u}_l = x_h \delta u_{min} + (1-x_h)\hat{p}_h + (1-x_h)(1-\hat{p}_h)\delta u_{min},$$

where:

1. $\hat{p}_l = 0$ and

$$\hat{u}_l = [x_h \delta u_{min} + (1 - x_h) \hat{p}_h + (1 - x_h)(1 - \hat{p}_h) \delta u_{min}] / \delta,$$

if $x_h \delta u_{min} + (1 - x_h) \hat{p}_h + (1 - x_h)(1 - \hat{p}_h) \delta u_{min} \leq \delta u_{max}$.

2.

$$\hat{p}_l = \frac{x_h \delta u_{min} + (1 - x_h) \hat{p}_h + (1 - x_h)(1 - \hat{p}_h) \delta u_{min} - \delta u_{max}}{1 - \delta u_{max}},$$

and $\hat{u}_l = u_{max}$ if $x_h \delta u_{min} + (1 - x_h) \hat{p}_h + (1 - x_h)(1 - \hat{p}_h) \delta u_{min} > \delta u_{max}$.

and promise keeping:

$$u = q[x_h + (1 - x_h) \hat{p}_h + (1 - x_h)(1 - \hat{p}_h) \delta u_{min}] + (1 - q)[\hat{p}_l + (1 - \hat{p}_l) \delta \hat{u}_l].$$

To describe our tracing procedure, we perform the change of variables $y = \hat{p}_h(1 - x_h)$, as alluded to in Section 5. Under this change, the problem can be further re-written as:

$$V(u) = \max q[(h - c)x_h + hy + (1 - x_h - y) \delta V(u_{min})] + (1 - q)[l \hat{p}_l + (1 - \hat{p}_l) \delta V(\hat{u}_l)],$$

subject to the low type's incentive compatibility constraint:

$$\hat{p}_l + (1 - \hat{p}_l) \delta \hat{u}_l = (1 - y) \delta u_{min} + y,$$

where

1. $\hat{p}_l = 0$ and $\hat{u}_l = [(1 - y) \delta u_{min} + y] / \delta$ if $(1 - y) \delta u_{min} + y \leq \delta u_{max}$.

2. $\hat{p}_l = [(1 - y) \delta u_{min} + y - \delta u_{max}] / (1 - \delta u_{max})$ and $\hat{u}_l = u_{max}$ if $(1 - y) \delta u_{min} + y > \delta u_{max}$.

and promise keeping:

$$u = y + (1 - y) \delta u_{min} + q x_h (1 - \delta u_{min}).$$

Note that the only two free variables in this optimization problem are x_h and y , which are subject to the constraints

$$\begin{aligned} x_h &\geq 0 \\ y &\geq 0 \\ x_h + y &\leq 1. \end{aligned}$$

In particular, given any x_h, y satisfying these constraints, \hat{p}_h can be uniquely determined; \hat{p}_l and \hat{u}_l are pinned down by the low-types incentive compatibility constraint, together with the two enumerated conditions. Furthermore, we can define each x_h and y as functions of u , say $x_h(u)$ and $y(u)$ —by promise keeping, changes in u must also generate changes in these variables.

A.3. Characterizing V : Three-Region Solutions

The general idea is to understand the boundaries of $V(\cdot)$ and the (left/right) derivatives of $V(\cdot)$. Note that a concave function defined on a compact interval has non-increasing left/right derivatives defined everywhere. The following generalizes Lemma 4 to arbitrary utility promise intervals:

Lemma 6. *There exist $u_{low} < u_{high}$ such that:*

1. $x_h(u) = 0$ when $u \in [u_{min}, u_{low}]$,
2. $\frac{dV(u)}{du} = \frac{(h-c)-\delta V(u_{min})}{1-\delta u_{min}}$ when $u \in [u_{low}, u_{high}]$, and
3. $x_h(u) + y(u) = 1$ when $u \in [u_{high}, u_{max}]$.

Proof of Lemma 6. We define the principal's expected payoff function $W(x_h, y)$ as:

$$W(x_h, y) = q[(h-c)x_h + hy + (1-x_h-y)\delta V(u_{min})] + (1-q)[l\hat{p}_l(y) + (1-\hat{p}_l(y))\delta V(\hat{u}_l(y))],$$

where the incentive compatibility constraints are substituted in (using Lemma 5) to define \hat{p}_l and \hat{u}_l as functions of y :

1. If $(1-y)\delta u_{min} + y \leq \delta u_{max}$, then $\hat{p}_l(y) = 0$ and $\hat{u}_l(y) = [(1-y)\delta u_{min} + y]/\delta$.
2. If $(1-y)\delta u_{min} + y > \delta u_{max}$, then $\hat{p}_l(y) = [(1-y)\delta u_{min} + y - \delta u_{max}]/(1-\delta u_{max})$ and $\hat{u}_l(y) = u_{max}$.

The principal's optimization problem is:

$$V(u) = \max_{x_h, y} W(x_h, y)$$

subject to the promise-keeping constraint:

$$u(x_h, y) \equiv y + (1-y)\delta u_{min} + qx_h(1-\delta u_{min}) = u,$$

and the feasibility constraints:

$$\begin{aligned}x_h &\geq 0, \\y &\geq 0, \\x_h + y &\leq 1.\end{aligned}$$

We consider the marginal effect of variations in the choice variables. Consider the variation with respect to x_h while holding y fixed. The key observation is that x_h enters linearly into both the objective function $W(x_h, y)$ and the constraint $u(x_h, y)$. The marginal rate of providing utility via x_h is determined by:

$$\frac{\partial u(x_h, y)}{\partial x_h} = q(1 - \delta u_{min}),$$

and the marginal rate to the principal is:

$$\frac{\partial W(x_h, y)}{\partial x_h} = q(h - c - \delta V(u_{min})).$$

Thus, if the principal increases u solely by increasing x_h , the derivative of the value function is constant:

$$\left. \frac{dV(u)}{du} \right|_{dy=0} = \frac{\partial W / \partial x_h}{\partial u / \partial x_h} = \frac{(h - c) - \delta V(u_{min})}{1 - \delta u_{min}}.$$

This value is independent of x_h and y .

The concavity of the value function $V(\cdot)$, combined with the constant marginal return of x_h , dictates the structure of the solution. This is because if at some utility level u it is optimal to increase x_h to deliver additional utility (i.e., x_h is interior), then $V'(u)$ must equal the constant derived above. Because $V'(\cdot)$ is non-increasing, if the optimal solution involves increasing x_h over an interval, it must continue to do so until the feasibility constraint binds ($x_h + y = 1$). It cannot be optimal to use x_h at some u , switch to a different instrument (implying a different slope), and then return to using x_h at a higher u , as this would imply $V'(u)$ decreases and then increases back to the constant level, violating concavity.

By this argument, there are only three cases of (x_h, y) as u increases.

1. $x_h = 0$ and y increases.
2. x_h increases and y is fixed until $x_h + y = 1$. Note this corresponds to $\frac{dV(u)}{du} = \frac{(h-c) - \delta V(u_{min})}{1 - \delta u_{min}}$.
3. y increases and x_h decreases while $x_h + y = 1$ holds.

And if $x_h(u) + y(u) = 1$ then $x_h(u') + y(u') = 1$ for all $u' \in [u, u_{max}]$; if $x_h(u) = 0$ then $x_h(u') = 0$ for all $u' \in [u_{min}, u]$. \square

A.4. Tracing Procedure

We are now in position to describe our tracing procedure. Here, we present all details behind how the derivative of the value function changes as u decreases from $u = u_{max}$. We present the computations for the general case, but describe the implications on the corresponding restrictions when these methods are applied subsequently.

With this in mind, first suppose there exists u with $x_h(u) + y(u) = 1$. Define

$$u_1 = \inf\{u \mid y(u) + x_h(u) = 1 \text{ and } \hat{u}_l = u_{max}\}.$$

Note $u_1 = u_{max}$ implies $x_h(u) = 0$ for any u . In other words, this is the degenerate case when the cost of verification is too high. Thus on $[u_1, u_{max}]$, we have

$$V(u) = \max q[(h - c)(1 - y) + hy] + (1 - q)[l\hat{p}_l + (1 - \hat{p}_l)\delta V(u_{max})],$$

subject to

$$\hat{p}_l + (1 - \hat{p}_l)\delta u_{max} = (1 - y)\delta u_{min} + y,$$

and

$$u = y + (1 - y)\delta u_{min} + q(1 - y)(1 - \delta u_{min}).$$

Now define

$$W(y) := q[(h - c)(1 - y) + hy] + (1 - q)[l\hat{p}_l(y) + (1 - \hat{p}_l(y))\delta V(\hat{u}_{max})],$$

where the incentive compatibility constraints are substituted in (using Lemma 5) to define \hat{p}_l as function of y :

1. If $(1 - y)\delta u_{min} + y \leq \delta u_{max}$, then $\hat{p}_l(y) = 0$ and $\hat{u}_l(y) = [(1 - y)\delta u_{min} + y]/\delta$.
2. If $(1 - y)\delta u_{min} + y > \delta u_{max}$, then $\hat{p}_l(y) = [(1 - y)\delta u_{min} + y - \delta u_{max}]/(1 - \delta u_{max})$ and $\hat{u}_l(y) = u_{max}$.

Then we have

$$\frac{dW(y)}{dy} = qc + (1 - q)\frac{1 - \delta u_{min}}{1 - \delta u_{max}}[l - \delta V(u_{max})],$$




Figure 6: The first interval in the tracing procedure. The arrow points from an intermediate value $u \in [u_1, u_{max}]$ to $\hat{u}_l(u) = u_{max}$. The dashed line depicts the value function on this interval to the left of u . This plot considers commitment, where $u_{min} = 0$ and $u_{max} = 1$.

and

$$\frac{du(y)}{dy} = 1 - \delta u_{min} - q(1 - \delta u_{min}) = (1 - q)(1 - \delta u_{min}),$$

as

$$u(y) = y + (1 - y)\delta u_{min} + q(1 - y)(1 - \delta u_{min}).$$

Thus,

$$\frac{dV}{du}(u) = \frac{q}{(1 - q)(1 - \delta u_{min})}c + \frac{1}{1 - \delta u_{max}}[l - \delta V(u_{max})],$$

on $[u_1, u_{max}]$. In particular, a kink in V arises at the value of $u(y)$ such that $\hat{p}_l = 0$.

We derive the solution by tracing the value function $V(u)$ starting from the highest possible utility u_{max} and working backwards to lower utilities. We define an “iteration” as the range of utility u such that the continuation promise for the low type, $\hat{u}_l(u)$, falls within the range of the previous iteration. This recursive structure mirrors the backward induction of a finite horizon model, as the process terminates within a finite number of steps, where each iteration corresponds to one additional period of interaction. Formally, however, we are not analyzing different time horizons; rather, we are iteratively characterizing the value function across different partitions of its domain.

Now consider future iterations when $\hat{u}_l < u_{max}$, note here we have $\hat{p}_l = 0$:

1. If $x_h(u) + y(u) = 1$, we have:

$$\frac{dW(y)}{dy} = qc + (1 - q)(1 - \delta u_{min})\frac{dV}{du}(\hat{u}_l),$$




Figure 7: The second interval in the tracing procedure. The arrow points from an intermediate value $u \in [u_2, u_1]$ to $\hat{u}_l(u) \in [u_1, u_{max}]$. The dashed line depicts the value function on this interval to the left of u . This plot considers commitment, where $u_{min} = 0$ and $u_{max} = 1$.

thus,

$$\frac{dV}{du}(u) = \frac{q}{(1-q)(1-\delta u_{min})}c + \frac{dV}{du}(\hat{u}_l).$$

Also, plugging $\hat{p}_l = 0$ to equation

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = (1 - y)\delta u_{min} + y,$$

we have

$$\delta\hat{u}_l = (1 - y)\delta u_{min} + y.$$

Combing this expression with

$$u(y) = y + (1 - y)\delta u_{min} + q(1 - y)(1 - \delta u_{min}),$$

we have

$$u(y) = \delta\hat{u}_l + q(1 - \delta\hat{u}_l) = (1 - q)\delta\hat{u}_l + q.$$

In particular, if we set \hat{u}_l to be a cutoff (kinks) where the derivative jump, this calculation implies that the kinks, say u_1, u_2, \dots from right to left follow (with $u_0 = u_{max}$):

$$u_{i+1} = (1 - q)\delta u_i + q.$$

2. Now after some number of iterations, the derivative reaches $\frac{h-c-\delta V(u_{min})}{1-\delta u_{min}}$ (as outlined in Lemma 6). Let $u_{\tilde{i}}$ to be the last kink where $x_h(u) + y(u) = 1$, which is exactly the right

endpoint of the segment with derivative $\frac{h-c-\delta V(u_{min})}{1-\delta u_{min}}$. Using the expression for the promised utility, we have:

$$u_{\tilde{i}+1} = \frac{u_{\tilde{i}} - q}{1 - q}.$$

3. If $x_h(u) = 0$, we have

$$V(u) = \max q[hy + (1 - y)\delta V(u_{min})] + (1 - q)[l\hat{p}_l + (1 - \hat{p}_l)\delta V(\hat{u}_l)],$$

subject to

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = (1 - y)\delta u_{min} + y,$$

and

$$u = y + (1 - y)\delta u_{min}.$$

Then define

$$W(y) := q[hy + (1 - y)\delta V(u_{min})] + (1 - q)[l\hat{p}_l(y) + (1 - \hat{p}_l(y))\delta V(\hat{u}_l(y))],$$

where the incentive compatibility constraints are substituted in (using Lemma 5) to define \hat{p}_l and \hat{u}_l as functions of y :

- (a) If $(1 - y)\delta u_{min} + y \leq \delta u_{max}$, then $\hat{p}_l(y) = 0$ and $\hat{u}_l(y) = [(1 - y)\delta u_{min} + y]/\delta$.
- (b) If $(1 - y)\delta u_{min} + y > \delta u_{max}$, then $\hat{p}_l(y) = [(1 - y)\delta u_{min} + y - \delta u_{max}]/(1 - \delta u_{max})$ and $\hat{u}_l(y) = u_{max}$.

Then we have

$$\frac{dW(y)}{dy} = q[h - \delta V(u_{min})] + (1 - q)(1 - \delta u_{min})\frac{dV}{du}(\hat{u}_l).$$

and

$$\frac{du(y)}{dy} = 1 - \delta u_{min},$$

as

$$u(y) = y + (1 - y)\delta u_{min}.$$

Thus,

$$\frac{dV}{du}(u) = \frac{q[h - \delta V(u_{min})]}{1 - \delta u_{min}} + (1 - q)\frac{dV}{du}(\hat{u}_l).$$




Figure 8: An interval in the tracing procedure for which $x_h(u) = 0$. The arrow points from an intermediate $u \in [u_{i+1}, u_i]$ to the continuation utility $\hat{u}_l(u)$. The dashed line depicts the value function on this interval to the left of u . This plot considers commitment, where $u_{min} = 0$ and $u_{max} = 1$

Also, plugging $\hat{p}_l = 0$ to equation

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = (1 - y)\delta u_{min} + y,$$

we have

$$\delta\hat{u}_l = (1 - y)\delta u_{min} + y.$$

Combing it with

$$u(y) = y + (1 - y)\delta u_{min},$$

we have

$$u(y) = \delta\hat{u}_l.$$

In particular, if there is a cutoff (kink) at some \tilde{u} , then there will be another jump in the derivative at $\delta\tilde{u}$, since $\hat{u}_l > u$ and $\frac{dV}{du}$ is either linear or has a discontinuous jump in the derivative at all values to the right of u . Given that there are kinks at $u_{\tilde{i}}$ and $u_{\tilde{i}+1}$, all other kinks will be at $\{\delta^s u_{\tilde{i}}\}_{s=1}^{\infty}$ and $\{\delta^s u_{\tilde{i}+1}\}_{s=1}^{\infty}$.

We know $V(\cdot)$ is concave and only the Pareto frontier can appear on-path of the optimal mechanism (unless allocation does not occur following a report of h , in which case continuation play delivers u_{min}). Thus, the optimal mechanism always increases the agent's continuation utility following a report of l until it hits u_{max} . Since the above implies $\frac{dV}{du}(u)$ will be positive in finite iterations, optimal mechanisms will take the form of an endogenous finite deadline of the kinds described in the main text.

B. Proofs of the Results in Section 4

Our proofs utilize the methodology presented in Appendix A, which in particular identify the Bellman equation as piecewise linear and identify its derivatives. Each particular setting we study requires determining the restrictions on the corresponding solution. In the full-commitment case, we have $u_{min} = 0$ and $u_{max} = 1$ because the principal can always choose to withhold the allocation forever, or allocate immediately irrespective of type.

Proof of Theorem 1. Since $h - c \geq 0$ and $V(u_{min}) = u_{min} = 0$, we have

$$\frac{(h - c) - \delta V(u_{min})}{1 - \delta u_{min}} = h - c \geq 0.$$

Note that optimal mechanisms on-path remain on the Pareto frontier following reports of l —following reports of h , allocation either occurs or the continuation utility is 0. However, since on the Pareto frontier we always have $\frac{dV(u)}{du} \leq 0 \leq h - c$, the optimal mechanism features $x_h(u) + y(u) = 1$ by our analysis in subsection A.3 (in particular, Lemma 6). According to subsection A.3 (setting $V(u_{max}) = r$), we have

$$\frac{dV}{du}(u) = \frac{q}{(1 - q)}c + \frac{1}{1 - \delta}[l - \delta r].$$

for the rightmost derivative and

$$\frac{dV}{du}(u) = \frac{q}{(1 - q)}c + \frac{dV}{du}(\hat{u}_l).$$

Note here $\frac{dV}{du}(u)$ will increase for a fixed amount $\frac{q}{(1 - q)}c$ after each iteration. Thus, the optimum of $V(\cdot)$ must be achieved at

$$u_{T^*-1} = [(1 - q)\delta]^{T^*-1} \left(1 - \frac{q}{1 - (1 - q)\delta}\right) + \frac{q}{1 - (1 - q)\delta},$$

where the left derivative is

$$\frac{T^*qc}{1 - q} + \frac{l - \delta r}{1 - \delta} \geq 0,$$

and the right derivative is

$$\frac{(T^* - 1)qc}{1 - q} + \frac{l - \delta r}{1 - \delta} < 0.$$

Note this is exactly the kink of the value function $V(\cdot)$ where increasing at the left and decreasing

at the right. Thus, $T^* = \lceil \frac{(1-q)(\delta r - l)}{qc(1-\delta)} \rceil$ is the endogenous optimal deadline. To identify the other components of the optimal mechanism, we note that the each successive iteration puts the agent's promised utility at a new kink, until the terminal period where $u = 1$ and hence project allocation occurs. Since $y(u) = 1 - x_h(u)$ in this region, we have $\hat{p}_h = 1$ as well, and $1 - x_h(u)$ is set to be equal to the agent's promised utility. Thus, the optimal mechanism takes the form as claimed in the Theorem. \square

Proof of Theorem 2. Since $h - c < 0$ and $V(u_{min}) = 0$, we have

$$\frac{(h - c) - \delta V(u_{min})}{1 - \delta u_{min}} = h - c < 0.$$

According to subsection A.3 (now $V(u_{max}) = r$), we have

$$\frac{dV}{du}(u) = \frac{q}{(1-q)}c + \frac{1}{1-\delta}[l - \delta r],$$

for the rightmost derivative, and for all other derivatives,

$$\frac{dV}{du}(u) = \frac{q}{(1-q)}c + \frac{dV}{du}(\hat{u}_l).$$

Thus, there exists k such that:

$$\frac{kq}{(1-q)}c + \frac{1}{1-\delta}[l - \delta r] \leq h - c,$$

and

$$\frac{(k+1)q}{(1-q)}c + \frac{1}{1-\delta}[l - \delta r] > h - c.$$

This means after k iterations $\frac{dV}{du}(u) = h - c$. This is exactly the point that the principal will start to use verification—at all values of u lower than this value, $x_h(u) = 0$ in the principal's optimization problem. At all values of u lower than this kink, by subsection A.3 we have

$$\frac{dV}{du}(u) = \frac{q[h - \delta V(u_{min})]}{1 - \delta u_{min}} + (1-q)\frac{dV}{du}(\hat{u}_l) = qh + (1-q)\frac{dV}{du}(\hat{u}_l).$$

Since we start at a kink at which $\frac{dV}{du}(u) = h - c$, further iterations yield a derivative of the form $h - (1-q)^k c$, which follows from inspection of the formula for the iterations; since $c > h$, at some point we must have $h - (1-q)^k c \leq 0$ and $h - (1-q)^{k+1} c > 0$. By identifying the optimal mechanism, this corresponds to initial randomization phase because $x_h = 0$. Using the

above characterization for the length of the randomization phase, algebra reveals²⁶ that the randomization phase lasts for $\lfloor \frac{\log(h/c)}{\log(1-q)} \rfloor + 1$. In this region, we have $y = \hat{p}_h$, so that the probability of allocating following a report of h is equal to the utility of the low type, yielding the form of the mechanism.

To determine the number of periods after the randomization phase, we examine the formulas for the value function kinks using the calculations in Subsection A.3 to identify which kink on the interval where $\frac{dV}{du} = h - c$ is reached in the optimal mechanism. Let $e = \lfloor \frac{(h-c-\frac{l-\delta r}{1-\delta})(1-q)}{qc} \rfloor \frac{qc}{1-q} + \frac{l-\delta r}{1-\delta}$. If $h - (1 - q)^{\lfloor \frac{\log(h/c)}{\log(1-q)} \rfloor + 1} (h - e) \leq 0$, the number of periods after the randomization phase (i.e., verification phase plus the terminal period) is $\lfloor \frac{(h-c-\frac{l-\delta r}{1-\delta})(1-q)}{qc} \rfloor + 1$; otherwise if $h - (1 - q)^{\lfloor \frac{\log(h/c)}{\log(1-q)} \rfloor + 1} (h - e) > 0$, the number of periods after the randomization phase is $\lfloor \frac{(h-c-\frac{l-\delta r}{1-\delta})(1-q)}{qc} \rfloor$. In particular, by the calculations from subsection A.3, these conditions determine which kink is the first one which produces a positive slope of the value function, and hence the length of the verification phase. \square

C. Additional Results for Limited Commitment

In this Appendix, we describe principal-optimal equilibrium outcomes for parameter values not covered by Corollary 1. We first describe the resulting outcome when Never Allocate yields higher payoff than Always Verify and Immediate Allocation:

Corollary 4. *Suppose $r < 0$ and $c > h$. There exists \bar{c} such that verification is part of the principal optimal equilibrium whenever $c < \bar{c}$,²⁷ which can be implemented as follows:*

- *There is an endogenous time T^* at which point the agent-optimal stage outcome is played in every period prior to allocation, where the stage mechanism involves allocation after a report of h if no verification occurs or h is verified. In this stage:

 - *Allocation of l is chosen so that the principal's payoff under this mechanism is 0.*
 - *Reports of h are verified with probability making the low-type indifferent between truth-telling and misreporting, where continuation play following a verified lie is Never Allocate.**
- *Before T^* , the equilibrium path involves two phases as described in Theorem 2. If the agent reports h in the randomization phase but allocation does not occur, the equilibrium switches to*

²⁶Note that the main text also presents a derivation of this expression, under the assumption that the mechanism takes a two-phase form.

²⁷The condition on verification costs ensures that verification is part of the principal optimal equilibrium. While this result would be vacuous if $\bar{c} = h$, we show in Online Appendix G.3 that as $\Delta \rightarrow 0$, \bar{c} converges to $\frac{\rho+q}{\rho}h > h$.

Never Allocate. Allocation and verification probabilities in each respective phase are such that an agent with project type l is indifferent between truth-telling and misreporting.

In the above case, we have $v^* = v_{min} = 0$. Here, the distinction from the commitment optimal mechanism stems from what happens at the deadline. Since Immediate Allocation is not an equilibrium strategy when $r < 0$, the principal still utilizes verification to some extent to discriminate on quality once the deadline is reached. Once the deadline is reached, continuation play is constant and maximizes the agent's payoff.

When $v_{min} > \max\{0, r\}$, both the payoff at the deadline as well after a punishment become endogenous. As implied by Proposition 4, one necessary condition for $v_{min} > \max\{0, r\}$ is $\frac{q(h-c)}{1-(1-q)\delta} > \max\{0, r\}$.

Corollary 5. *Suppose $v_{min} > \max\{0, r\}$. In the principal-optimal equilibrium:*

- *There is an endogenous time T^* at which point the agent-optimal stage outcome is played in every period prior to allocation, where the stage mechanism involves allocation after a report of h if no verification occurs or h is verified. In this stage:*
 - *Allocation of l is such that the agent obtains u_{max} from repeatedly playing this stage mechanism every period prior to allocation.*
 - *Reports of h are verified with probability making the low-type agent indifferent between truth-telling and misreporting. Following a verified lie, the agent's continuation payoff is u_{min} , where $u_{min} < u_{max}$ satisfies $V(u_{min}) = V(u_{max}) \leq \frac{q(h-c)}{1-(1-q)\delta}$.*
- *Before T^* , the equilibrium path coincides with Theorem 1, with verification probabilities modified to maintain incentive compatibility, assuming that if the agent is verified to have lied, his continuation payoff is u_{min} .*

In this case, neither the optimal commitment mechanism's final reward (allocation) nor its optimal punishment (withholding forever) is implementable—in contrast to the other two cases described, where one of the two is. Hence, both u_{min} and u_{max} are endogenously determined. One subtlety here is that $\frac{q(h-c)}{1-(1-q)\delta} > \max\{0, r\}$ only implies Always Verify can be implemented in equilibrium but it doesn't imply that it delivers v_{min} .

Recall that $v^* := \max\{r, \frac{q(h-c)}{1-(1-q)\delta}, 0\}$. One reason principal-optimal equilibria can improve beyond v^* stems from the fact that a higher probability of allocation following l reports can allow the principal to lower the verification probability. The intuition for these gains are analogous to

the commitment case. However, it is *also* possible to construct equilibria improving the principal's payoff beyond v^* without increasing the probability of allocating to type l .

We illustrate this idea by considering the $l = -\infty$ case, so that any positive probability of allocation to the low type yields a payoff less than v^* . When Never Allocate delivers v^* , the threat of reverting to this outcome can provide scope for the principal to verify with probability less than 1, potentially enabling the use of verification in equilibrium even when $c > h$ (provided c is not too high). Interestingly, even when Always Verify delivers v^* , it may be possible to improve beyond v^* by similarly creating a wedge between the agent's payoff across continuation equilibria, as the following example shows:

Example 1. *Consider two mechanisms:*

(M1) *When l is reported, allocate with probability 0, play M2 in the next period. When h is reported, no verification occurs while allocation occurs with probability $\frac{q\delta}{1+q\delta}$; if no allocation following a report of h , play M1 in the next round.*

(M2) *When l is reported, allocate with probability 0, and play M2 in the next round. When h is reported, verify with probability $\frac{1}{1+q\delta}$. If the type is verified as h , allocate with probability 1; if the type is verified as l , allocate with probability 0 and play \mathcal{M}_1 in the next round. Following no verification after h , allocate with probability 1.*

Let u_i denote the agent's reservation utility in each mechanism i . We have

$$\begin{aligned} u_1 &= q\left(\frac{q\delta}{1+q\delta} + \frac{1}{1+q\delta}\delta u_1\right) + (1-q)\delta u_2, & u_2 &= q + (1-q)\delta u_2, \\ & \Rightarrow u_1 = \frac{q\delta}{1-(1-q)\delta}, & u_2 &= \frac{q}{1-(1-q)\delta}. \end{aligned}$$

It is straightforward to show that truth-telling is incentive compatible. Denoting the principal's reservation utility in each mechanism i by v_i , we have

$$\begin{aligned} v_1 &= q\left(\frac{hq\delta}{1+q\delta} + \frac{1}{1+q\delta}\delta v_1\right) + (1-q)\delta v_2, & v_2 &= q\left(h - \frac{c}{1+q\delta}\right) + (1-q)\delta v_2, \\ & \Rightarrow v_1 = \frac{q(\delta h - (1-q)\delta c)}{1-(1-q)\delta}, & v_2 &= \frac{q\left(h - \frac{c}{1+q\delta}\right)}{1-(1-q)\delta}. \end{aligned}$$

Note that $v_2 > \frac{q(h-c)}{1-(1-q)\delta}$. Thus, provided $\delta h - (1-q)\delta c \geq h - c$ (which we can rewrite $h \leq \frac{1-(1-q)\delta}{1-\delta}c$), then both M1 and M2 will yield higher principal payoff than Always Verify. Thus if Always Verify achieves v^* , these strategies form an equilibrium.

This example shows that even when Always Verify is an equilibrium—so that the verification costs are not prohibitive—the principal can do better by rewarding truthful reports and punishing lies.

D. Proofs for Results in Section 6 and Appendix C

D.1. Exogenous Horizon

Calculation Summary If the horizon is finite, we can use the same method in A.3 and directly compute the Bellman Equation Solution, and subsequently use this solution to identify the optimal mechanism directly. In particular, the proof of Lemma 5 applies unchanged to the case of finite horizon Bellman operators. As stated in the main text, this yields the following formulation:

$$V^T(u) = \max q[(h - c)x_h + hy] + (1 - q)[l\hat{p}_l + (1 - \hat{p}_l)\delta V^{T-1}(\hat{u}_l)],$$

subject to the low type's incentive compatibility constraint:

$$\hat{p}_l + (1 - \hat{p}_l)\delta\hat{u}_l = y,$$

where:

1. $\hat{p}_l = 0$ and $\hat{u}_l = y/\delta$ if $y \leq \delta$.
2. $\hat{p}_l = (y - \delta)/(1 - \delta)$ and $\hat{u}_l = 1$ if $y > \delta$.

and promise keeping:

$$u = y + qx_h.$$

Using $V^0(u) = 0$, the solution for V^{T-1} determines the solution for V^T for $T \in \{1, \dots, T^*\}$. As the solution method is similar to our original tracing procedure, we provide a summary of these calculations assuming $h - c > r$ (which, as stated in the main text, applies to both Propositions).

Define $d = \frac{q}{1-(1-q)\delta}$, $a_i = [(1 - q)\delta]^i(q - d) + d$, and $b_i = [(1 - q)\delta]^i(1 - d) + d$. Defining $\gamma_T = \frac{Tqc}{1-q} + l$, let T'' be the first positive integer that satisfies $qh + (1 - q)\gamma_{T''} \geq h - c$. Then $T \leq T''$, $V^T(u)$ is a piece-wise linear function with $2T - 1$ kinks (not including the two endpoints of the interval, i.e., 0 and 1), with kinks at $\delta^{T-1}a_0, \delta^{T-2}a_1, \dots, a_{T-1}, b_{T-1}, \dots, b_1$ respectively.²⁸ In

²⁸To interpret these formulas, a_{T-1} represents a mechanism whereby the principal utilizes the verification technology in every period, including the final period, with the verification probability set so that type- l is indifferent between lying and truthful reporting; $\delta^{T-1-i}a_i$ represents a length $T - 1 - i$ randomization phase and utilizing verification in the remaining periods; b_{T-i} represents utilizing verification following h reports with intermediate probability for $T - i$ periods, and subsequently allocating with probability 1 irrespective of type.

particular, a kink at $\delta^{T-i}a_i$ for $i < T - 1$ is generated when continuation utility is set to $\delta^{T-i-1}a_i$, while the kink at a_{T-1} is generated when continuation utility is set to a_{T-2} and a kink at b_{T-i} is generated when continuation utility is set to b_{T-i-1} . Set $e_i = -(1-q)^{i-1}c + h$ and $s_i = \frac{iqc}{1-q} + \frac{l-\delta r}{1-\delta}$. We have:

$$\begin{aligned}\frac{dV^T(u)}{du} &= e_T, \quad \text{for } u \in [0, \delta^{T-1}a_0], \\ \frac{dV^T(u)}{du} &= e_i, \quad \text{for } u \in [\delta^i a_{T-1-i}, \delta^{i-1} a_{T-i}] \text{ and } i = T-1, \dots, 1, \\ \frac{dV^T(u)}{du} &= \gamma_T \quad \text{for } u \in [a_{T-1}, b_{T-1}], \\ \frac{dV^T(u)}{du} &= s_i, \quad \text{for } u \in [b_i, b_{i-1}] \text{ and } i = T-1, \dots, 1.\end{aligned}$$

Let T' be the first integer that satisfies $qh + (1-q)s_{T'-1} \geq h - c$, (noting that, since $h - c > r$, we also have $h - c > qh + (1-q)\frac{l-\delta r}{1-\delta}$, implying that $T' \geq 2$). For $T'' < T \leq T'$, $V^T(u)$ is a piece-wise linear function with $2T - 1$ kinks (aside from the two endpoints, 0 and 1), with kinks:

$$\delta^{T-1}a_0, \delta^{T-2}a_1, \dots, \delta^{T-T''}a_{T''-1}, \delta^{T-T''}b_{T''-1}, \delta^{T-T''-1}b_{T''}, \dots, \delta b_{T-2}, b_{T-1}, \dots, b_1$$

respectively. Define $k_i = (1-q)^i(\gamma_{T''} - h) + h$, with e_i and s_i as above. We have:

$$\begin{aligned}\frac{dV^T(u)}{du} &= e_T, \quad \text{for } u \in [0, \delta^{T-1}a_0], \\ \frac{dV^T(u)}{du} &= e_i, \quad \text{for } u \in [\delta^i a_{T-1-i}, \delta^{i-1} a_{T-i}] \text{ and } i = T-1, \dots, T-T''+1, \\ \frac{dV^T(u)}{du} &= k_{T-T''}, \quad \text{for } u \in [\delta^{T-T''}a_{T''-1}, \delta^{T-T''}b_{T''-1}], \\ \frac{dV^T(u)}{du} &= e_i, \quad \text{for } u \in [\delta^i b_{T-1-i}, \delta^{i-1} b_{T-i}] \text{ and } i = T-T'', \dots, 1, \\ \frac{dV^T(u)}{du} &= s_i, \quad \text{for } u \in [b_i, b_{i-1}] \text{ and } i = T-1, \dots, 1.\end{aligned}$$

Moreover, for $T > T'$, $V^T(u)$ is a piece-wise linear function with $2T - 1$ kinks, with kinks $\delta^{T-1}a_0, \delta^{T-2}a_1, \dots, \delta^{T-T''}a_{T''-1}, \delta^{T-T''}b_{T''-1}, \delta^{T-T''-1}b_{T''}, \dots, \delta^{T-T'+1}b_{T'-2}, \delta^{T-T'}b_{T'-1}, \delta^{T-T'}b_{T'-2}, \dots, \delta b_{T'-1}, \delta b_{T'-2}, b_{T'-1}, \dots, b_1$. Define $g_i = (1-q)^i(s_{T'-1} - h) + h$. We have:

$$\frac{dV^T(u)}{du} = e_T, \quad \text{for } u \in [0, \delta^{T-1}a_0],$$

$$\begin{aligned}
\frac{dV^T(u)}{du} &= e_i, \quad \text{for } u \in [\delta^i a_{T-1-i}, \delta^{i-1} a_{T-i}] \text{ and } i = T-1, \dots, T-T''+1, \\
\frac{dV^T(u)}{du} &= k_{T-T''}, \quad \text{for } u \in [\delta^{T-T''} a_{T''-1}, \delta^{T-T''} b_{T''-1}], \\
\frac{dV^T(u)}{du} &= e_i, \quad \text{for } u \in [\delta^i b_{T-1-i}, \delta^{i-1} b_{T-i}] \text{ and } i = T-T'', \dots, T-T'+1, \\
\frac{dV^T(u)}{du} &= g_i, \quad \text{for } u \in [\delta^i b_{T'-1}, \delta^i b_{T'-2}] \text{ and } i = T-T', \dots, 1, \\
\frac{dV^T(u)}{du} &= e_i, \quad \text{for } u \in [\delta^i b_{T'-2}, \delta^{i-1} b_{T'-1}] \text{ and } i = T-T', \dots, 1, \\
\frac{dV^T(u)}{du} &= s_i, \quad \text{for } u \in [b_i, b_{i-1}] \text{ and } i = T'-1, \dots, 1.
\end{aligned}$$

Proof of Proposition 1. Suppose $T \leq \lfloor \frac{-(1-q)l}{qc} \rfloor$. Since $h - c > 0$, we have $e_i > 0$ for all $i \in \{1, \dots, T-1\}$, but $\gamma_T = \frac{Tqc}{1-q} + l \leq 0$. Therefore, the optimal mechanism sets the agent's initial utility equal to a_{T-1} , as this kink corresponds to where the slope switches from positive to negative. Furthermore, as mentioned above, the next period's promised utility is a_{T-2} , which we note is the maximizer over u of $V^{T-1}(u)$ by the same argument. Using the formula for a_i , we note that:

$$((1-q)\delta)^i(q-d) + d > ((1-q)\delta)^{i-1}(q-d) + d,$$

since $q - d < 0$. Thus, the agent's promised utility decreases over time. Using this observation, since the verification probability is equal to $1 - \delta a_{T-i-1}$ for all $i < T$ and 1 when $i = T$, we have that the verification probability increases over time. In particular, note that when the utility promise is equal to a_0 , the principal verifies reports of h with probability 1 and allocating only if verified to be truthful. Other properties of the optimal mechanism follow from the inspection of the solution given these observations.

Note that if $\gamma_T > 0$, the initial utility promise is b_{T-1} ; furthermore, if the agent is at a kink b_{T-i} in the i th period, then the next period kink is b_{T-i-1} . The formula for b_t implies that the continuation utility promise is increasing over time, implying that the verification probability decreases. As before, other properties of the optimal mechanism follow from the inspection of the solution given these observations. \square

Proof of Proposition 2. First, if $h - c(1-q)^{T-1} \leq 0$, then $V(u)$ is decreasing at 0. Concavity of V thus implies that there is no value of u such that $V(u) > 0$. It follows that the optimal mechanism is to never allocate; in turn, T satisfies this condition if and only if $T \leq T_1$.

Thus suppose $T > T_1$, so that $V(u) > 0$ for some $u > 0$ implying that the optimal mechanism must allocate with positive probability. Let T_2 be the largest value of $T > T''$ such that $k_{T-T''} < 0$.

In this case, we have the slope of V^T is negative in $[\delta^{T-T''} a_{T''-1}, \delta^{T-T''} b_{T''-1}]$; thus, the kink at which the slope of the value function switches from positive to negative must be at some $u \leq \delta^{T-T''} a_{T''-1}$. Hence, the initial utility promise to the agent is $\delta^{T-i} a_{i-1}$ for some $i \in \{1, \dots, T''\}$.

To complete the proof, we next determine the evolution of the agent's utility promises for this case. Consider first the case where $k_{T-T''} < 0$, and let \tilde{i} denote the smallest i such that $e_i > 0$, so that the initial utility promise is $\delta^{i-1} a_{T-i}$. In particular, \tilde{i} is equal to T_1 independently of T . Furthermore, in the j th period for $j \in \{2, \dots, T_1\}$, we have the agent's promised utility is $\delta^{T_1-j} a_{T-T_1}$, as mentioned prior to the proof of Proposition 1. From inspection, we see that this utility path and principal payoffs can be generated by implementing a randomization phase as claimed in the statement of the Proposition. In the j th period for $j > T_1$, the agent's promised utility is equal to $a_{T-T_1-(j-T_1)}$; in particular, the evolution of the agent's promised utility coincides with the description in Proposition 1. Note that in particular the optimal mechanism in these periods can be implemented in the same way as described in Proposition 1—i.e., by verifying reports of h with probability sufficient to deter an l type from misreporting. By the argument from the proof of Proposition 1, the verification probability is increasing in the verification phase, as claimed. \square

D.2. Renegotiation

Proof of Proposition 3. According to Definition 6.1 and subsection A.1, $\Gamma V : [\delta u_{min}, 1] \rightarrow \mathbb{R}$ is generated by $V(\cdot)$ defined on $[u_{min}, 1]$ if for every $u \in [\delta u_{min}, 1]$, $\Gamma V(u)$ is given by:

$$\begin{aligned} \Gamma V(u) = \max q [& h x_h p_{hh}^* + x_h (1 - p_{hh}^*) \delta V(u_{hh}^*) + h (1 - x_h) \hat{p}_h + (1 - x_h) (1 - \hat{p}_h) \delta V(\hat{u}_h) - x_h c] \\ & + (1 - q) [l x_l p_{ll}^* + x_l (1 - p_{ll}^*) \delta V(u_{ll}^*) + l (1 - x_l) \hat{p}_l + (1 - x_l) (1 - \hat{p}_l) \delta V(\hat{u}_l) - x_l c]. \end{aligned}$$

The incentive compatibility conditions are: For the high type,

$$\begin{aligned} x_h p_{hh}^* + x_h (1 - p_{hh}^*) \delta u_{hh}^* + (1 - x_h) \hat{p}_h + (1 - x_h) (1 - \hat{p}_h) \delta \hat{u}_h \\ \geq x_l p_{lh}^* + x_l (1 - p_{lh}^*) \delta u_{lh}^* + (1 - x_l) \hat{p}_l + (1 - x_l) (1 - \hat{p}_l) \delta \hat{u}_l, \end{aligned}$$

and for the low type:

$$\begin{aligned} x_l p_{ll}^* + x_l (1 - p_{ll}^*) \delta u_{ll}^* + (1 - x_l) \hat{p}_l + (1 - x_l) (1 - \hat{p}_l) \delta \hat{u}_l \\ \geq x_h p_{hl}^* + x_h (1 - p_{hl}^*) \delta u_{hl}^* + (1 - x_h) \hat{p}_h + (1 - x_h) (1 - \hat{p}_h) \delta \hat{u}_h. \end{aligned}$$

Meanwhile, the promise-keeping constraint is:

$$u = q[x_h p_{hh}^* + x_h(1 - p_{hh}^*)\delta u_{hh}^* + (1 - x_h)\hat{p}_h + (1 - x_h)(1 - \hat{p}_h)\delta \hat{u}_h] \\ + (1 - q)[x_l p_{ll}^* + x_l(1 - p_{ll}^*)\delta u_{ll}^* + (1 - x_l)\hat{p}_l + (1 - x_l)(1 - \hat{p}_l)\delta \hat{u}_l].$$

Renegotiation-proofness, as defined in Definition 6.1, require that the Pareto frontier of the value function is self-generating. We note that we can directly apply the tracing methodology derived in subsection A.3. The tracing procedure in subsection A.3 determines the value of the function $V(u)$ on $[u_{min}, 1]$. Now, each $(u, V(u))$ corresponds to a renegotiation-proof mechanism.

Existence follows similar steps as Zhao (2006) applied to our particular setting; the details for our setting are as follows. Specifically, we need to show that there exists u_{min} such that Definition 6.1 is satisfied. By the tracing methodology derived in subsection A.3, we know this is equivalent to that the left derivative of $V(\cdot)$ at u_{min} being positive and the right derivative of $V(\cdot)$ being nonpositive. In other words, u_{min} is the local optimal (and hence, due to concavity, also the global optimal) on $[0, 1]$ when u_{min} is the smallest reservation utility.

First, consider the set of S of \underline{u} such that $V(\cdot, \underline{u})$ is non-increasing on $[\underline{u}, 1]$. This set is non-empty because $\underline{u} = 1$ trivially satisfies. Now consider $u_{min} = \inf S$. Since $V(\cdot, \underline{u}_n)$ is non-increasing on $[u_{min}, 1]$ for each n and $V(\cdot, \underline{u}_n) \rightarrow V(\cdot, u_{min})$ uniformly on $[u_{min}, 1]$, the limit function $V(\cdot, u_{min})$ is also non-decreasing on $[u_{min}, 1]$. Uniform convergence follows from observing that, extending the domain of $V(\cdot, \underline{u})$ to $[u_{min}, 1]$ using values implied by the tracing procedure in Section A.3, on $[u_{min}, 1]$, the family $\{V(\cdot, \underline{u})\}$ is uniformly Lipschitz: there exists $L < \infty$ depending only on primitives such that

$$|V(u, \underline{u}) - V(u', \underline{u})| \leq L|u - u'|$$

for all $u, u' \in [u_{min}, 1]$ and all \underline{u} . As $V(\cdot, \underline{u}) \rightarrow V(\cdot, u_{min})$ pointwise, the Arzelà–Ascoli Theorem implies uniform convergence, passing to a subsequence if necessary. Now, note that $V(\cdot, u_{min})$ must be non-decreasing on $[0, u_{min}]$, since otherwise we would be able to find an even smaller $\underline{u} < u_{min}$ with $V(\cdot, \underline{u})$ non-increasing on $[\underline{u}, 1]$ by the continuity of $V(\cdot, \underline{u})$ in \underline{u} , contradicting that $u_{min} = \inf S$.

We know if there exist u_1 and u_2 both satisfying the previous characterization such that $u_1 < u_2$, then $V(u_1) > V(u_2)$ since $V(\cdot)$ gives the largest utility the principal can obtain under commitment with the constraint to promise at least u_1 (resp. u_2) to the agent at any continuation history. Thus, $u_{min} = \inf S$ corresponds to the optimal set of renegotiation-proof mechanisms, which contains the principal-optimal renegotiation-proof mechanism.

It follows that the optimal mechanism $(u_{min}, V(u_{min}))$ implements the initial mechanism when no allocation is made. The features of the optimal mechanism in the Proposition follow from the reductions in Lemma 5 and the characterizations implied by the tracing procedure in Section A.3. Specifically, the endogenous finite deadline follows from the observation that this point must be reached after a finite number of changes in the derivatives (i.e., kinks), following the steps outlined in the Tracing Procedure in Section A.4. Given the evolution of the agent's utility, the optimal verification and allocation probabilities satisfy the claimed properties; since u_l increases over time, each must be set so that the low type's IC constraint holds with equality. \square

D.3. The Limited Commitment Case

Proof of Proposition 4. It is immediate that the principal can implement Immediate Allocation and Never Allocation independently of the agent's strategy, so that a lower bound on her equilibrium payoff is $\max\{0, r\}$. This proves the first part.

To show the second part, we show that both players have no incentives to deviate from the corresponding strategy profile. We denote the principal's payoff by v^* , as indicated in the statement of the Proposition. We use the one-shot deviation principle to verify that this is indeed an equilibrium. First, note that the agent has no incentive to deviate on-path, since in the case of Never Allocate and Immediate Allocation, the agent's payoff does not change as a function of his actions; while in the case of Always Verify, misreporting from l to h doesn't change expected payoff, and misreporting from h to l strictly decreases agent's payoff.

Now suppose the principal deviates to offer an any mechanism in a particular period. Since within the current period the agent chooses to maximize his within-period allocation probability, the principal's within-period payoff is dominated by the optimal static mechanism. Recall that $v^* := \max\{r, \frac{q(h-c)}{1-q(1-\delta)}, 0\}$. In particular, if the agent's ex-ante expected within-period allocation probability before drawing the type is p^* , then the principal's expected payoff must be weakly less than

$$p^* \cdot \max\{r, h - c, 0\} + (1 - p^*)\delta v^* \leq v^*$$

where the first term comes from the maximum slope (with respect to the agent's promised utility) of the static value function and the second term comes from both player reverting to the strategy profile. Hence, no profitable deviation exists for the principal, and the proposed strategy profile supports v^* as equilibrium payoff for the principal. \square

Proof of Corollary 1. First, note that the equilibrium payoff set is a compact convex subset of $[0, 1] \times [r, \frac{hq}{1-q(1-\delta)}]$. That this bounds the set of possible equilibrium payoffs holds since $r \geq$

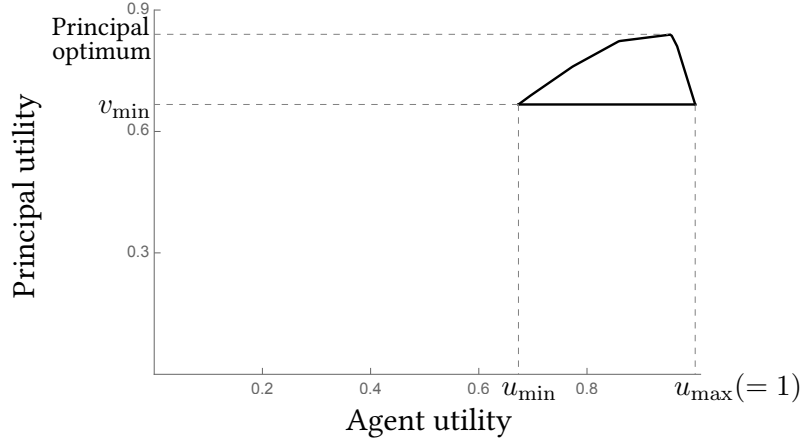


Figure 9: Equilibrium payoff set without across-period commitment, when immediate allocation delivers v_{min} . Parameters are $h = 1, c = 1/3, q = 2/3, l = 0, \delta = 9/10$.

$\max\{0, \frac{q(h-c)}{1-(1-q)\delta}\}$ and since $\frac{hq}{1-q(1-\delta)}$ denotes the principal's payoff when $c = 0$, which bounds any feasible payoff. The claim regarding compactness of the equilibrium payoff set formally follows from the arguments outlined in Sections 8.1 to 8.3 in Lipnowski and Ramos (2020). Specifically, their Lemma 4 shows that the Abreu et al. (1990) algorithm applied to our stage game can reduce to a promise-keeping constraint, a binding agent incentive compatibility constraint, and a principal participation constraint. Since Proposition 4 characterizes the principal participation constraint, the left end-point is the smallest u_{min} such that $V(u_{min}) = r$. Note that this must be achieved at some $u_{min} > 0$, since $V(0) = 0 < r$, by continuity of the value function. Since immediate allocation delivers v_{min} , it follows that $u_{max} = 1$.

Thus, the tracing procedure applied on the domain $[u_{min}, 1]$ characterized $V(u)$ on its domain, where u_{min} is endogenously determined. Lemma 6 and the calculations in Section A.4 show that $V(u)$ is a piecewise-linear function, with the derivative increasing discretely at each kink (when u decreases from 1). Given that the principal-optimal initial utility promise u^* must correspond to a kink of the value function, the optimal mechanism involves continuation promises that eventually lead to immediate allocation if allocation has not occurred after finitely many periods. The decreasing verification probability follows from the observation that y increases over time and satisfies $y = 1 - x_h$. In particular, verification probabilities and allocation probabilities in each phase are set so that the low type's IC constraint binds. Hence, the mechanism takes the form stated in the Corollary. \square

Proof of Corollary 4. By the conditions in the Corollary, since $r < 0$ and $c > h$, Proposition 4 implies that the $v_{min} = 0$ and is achieved by Never Allocate. Therefore, we obtain that $u_{min} = 0$. Furthermore, since $r < 0$ and since immediate allocation is the unique mechanism that achieves




Figure 10: Equilibrium payoff set without across-period commitment, when withholding delivers v_{min} . Parameters are $h = 1, c = 5/4, q = 1/2, l = -2, \delta = 9/10$.

$u = 1$, we have that $u = 1$ cannot be part of equilibrium and hence $u_{max} < 1$. Lemma 5 implies that u_{max} is supported by an equilibrium in which the principal repeatedly offers a stage mechanism delivering expected utility u_{max} , and the continuation following a verified lie is Never Allocate.

Thus, the set of equilibrium payoffs is a compact convex subset of $[0, 1] \times [0, \frac{hq}{1-q(1-\delta)}]$, where the left endpoint is 0 and the right endpoint is given by the condition that $V(u_{max}) = 0$; the argument for this is identical to the Proof of Corollary 1. As in the proof of Corollary 1, the tracing procedure implies that the optimum is attained after finitely many iterations, using the conditions that $u_{min} = 0$ and $u_{max} < 1$, with the optimum being achieved by setting the agent's promised utility to be at a kink of the value function. This shows that the principal-optimal equilibrium takes the form described in the statement of the Corollary. \square

Proof of Corollary 5. Since $v_{min} > \max\{0, r\}$, we have that $u_{min} > 0$ and $u_{max} < 1$. Lemma 5 implies that u_{max} is supported by an equilibrium in which the principal repeatedly offers a stage mechanism delivering expected utility u_{max} , and the continuation following a verified lie is u_{min} . We thus have that the set of equilibrium payoffs is a compact convex subset of $[0, 1] \times [0, \frac{hq}{1-q(1-\delta)}]$, with u_{min} and u_{max} determined by the condition that $V(u_{min}) = V(u_{max}) = v_{min} \leq v^*$.

We show that there is no randomization phase on-path, as in the mechanism described in Theorem 1. Note that Lemma 4 shows that $x_h + y < 1$ only in a region where the derivative of $V(u)$ is greater than or equal to:

$$\frac{h - c - \delta V(u_{min})}{1 - \delta u_{min}} \geq \frac{h - c - \delta v^*}{1 - \delta u_{min}} \propto h - c - \delta \frac{q(h - c)}{1 - (1 - q)\delta} \propto (h - c)(1 - \delta) > 0.$$

Thus, only case 3 of Lemma 4 can occur at any kink that is achieved on-path in the principal optimal mechanism, so that $x_h + y = 1$.

Thus, since the derivative of $V(u)$ increases by a discrete amount bounded from below as u moves to progressively lower kinks, the tracing procedure implies that the slope of the frontier switches from negative to positive after finitely many iterations—thus, the optimum is achieved after finitely many iterations, as in the proofs of Corollaries 1 and 4. In particular, the optimum is achieved at a kink of the value function. Recalling the structure of the stationary equilibrium that supports u_{max} , we conclude that the principal-optimal equilibrium takes the form described in the Corollary by tracing the evolution of the utility profiles under this equilibrium. \square

Proof of Corollary 2. We first consider the mechanism length. Note that if a mechanism lasts for \hat{T} length of time, then the number of periods is $\frac{\hat{T}}{\Delta}$. By Theorem 1,

$$\frac{\hat{T}}{\Delta} = \frac{(1 - \lambda\Delta)(e^{-\rho\Delta}(\lambda\Delta h + (1 - \lambda\Delta)l) - l)}{\lambda\Delta c(1 - e^{-\rho\Delta})}.$$

We use the approximation that $e^{-\rho\Delta} \approx 1 - \rho\Delta$. Substituting and simplifying, we have:

$$\hat{T} = \frac{(1 - \lambda\Delta)((1 - \rho\Delta)(\lambda\Delta h + (1 - \lambda\Delta)l) - l)}{\lambda c \rho \Delta}.$$

Algebra shows that the numerator is $(h\lambda - l(\lambda + \rho))\Delta + o(\Delta)$, so that the length of time approaches $\frac{h\lambda - l(\lambda + \rho)}{c\lambda\rho}$, as claimed. For the verification path, we derive an ODE for the agent's continuation payoff. Write $u_t = \lambda\Delta + (1 - \lambda\Delta)e^{-\rho\Delta}u_{t+\Delta}$. Using the same approximation for $e^{-\rho\Delta}$ we have:

$$\frac{u_t - u_{t+\Delta}}{\Delta} = \lambda - (\lambda + \rho)u_{t+\Delta} - o(1).$$

Taking $\Delta \rightarrow 0$ and using continuity of u_t implied the optimal mechanism in Theorem 1 yields $\dot{u}_t = (\lambda + \rho)u_t - \lambda$. This is a first-order ODE with boundary condition for the utility at time $u_{\hat{T}}$, with unique solution:

$$u_t = e^{-(\lambda+\rho)(\hat{T}-t)}u_{\hat{T}} + \frac{\lambda}{\lambda + \rho} \left(1 - e^{-(\lambda+\rho)(\hat{T}-t)}\right).$$

Note that $x_t = 1 - u_t$; indeed, with x_t set as this value, the agent is indifferent between lying about project quality when it is l and and telling the truth. \square

Proof of Corollary 3. We first consider the verification length. Note that if verification lasts for T'

length of time, then the number of periods is $\frac{T'}{\Delta}$. Because

$$\lambda\Delta h + (1 - \lambda\Delta)\left(\frac{(T'/\Delta - 1)\lambda\Delta c}{1 - \lambda\Delta} + \frac{l - \delta r}{1 - \delta}\right) = h - c$$

as $\Delta \rightarrow 0$, we have

$$\frac{T'}{\Delta} = \frac{(1 - \lambda\Delta)(e^{-\rho\Delta}(\lambda\Delta h + (1 - \lambda\Delta)l) - l)}{\lambda\Delta c(1 - e^{-\rho\Delta})} + \frac{(1 - \lambda\Delta)(h - c)}{\lambda\Delta c}.$$

We use the approximation that $e^{-\rho\Delta} \approx 1 - \rho\Delta$. Substituting and simplifying, algebra shows that the length of time approaches $\frac{(h-l)(\lambda+\rho)-\rho c}{c\lambda\rho}$, as claimed.

For the randomization length, because as $\Delta \rightarrow 0$,

$$(1 - \lambda\Delta)^{(T^* - T')/\Delta} c = h \Rightarrow \frac{T^* - T'}{\Delta} = \frac{\log(h/c)}{\log(1 - \lambda\Delta)}.$$

Algebra shows that the length of time approaches $\frac{\log(c/h)}{\lambda}$. For the verification phase to have non-zero length, we need to have $\frac{(h-l)(\lambda+\rho)-\rho c}{c\lambda\rho} > 0$, equivalent to $c < \frac{(\lambda+\rho)(h-l)}{\rho}$, as claimed. \square