

Motivating Effort with Information about Future Rewards

Chang Liu*

First version: October 8, 2021

This version: December 28, 2025

Abstract

This paper studies the optimal mechanism to motivate effort in a dynamic principal-agent model without transfers. An agent is engaged in a task with uncertain future rewards and can quit at any time. The principal knows the reward and provides information over time to motivate effort. We provide a unified framework and derive the optimal information policy in closed form across stationary and nonstationary environments. Within this framework, we identify two precise conditions, each of which guarantees that dynamic disclosure is strictly valuable. First, if the principal is *impatient* compared to the agent, she prefers the front-loaded effort schedule induced by dynamic disclosure; in a stationary environment, dynamic disclosure is beneficial if and only if the principal is less patient. Second, in an environment where the agent would become *pessimistic* over time absent any disclosure, dynamic information provision can counteract this downward trend and encourage longer effort. Notably, patience remains a crucial determinant of the structure of the optimal policy.

Keywords: Dynamic information design, delayed disclosure, informational incentives

*School of Economics, UNSW Business School; chang.liu36@unsw.edu.au. I gratefully acknowledge funding from the National Science Foundation under grant DMS-1928930 and from the Alfred P. Sloan Foundation under grant G-2021-16778 during the Fall 2023 semester. I express my deep gratitude to Benjamin Golub and Shengwu Li for their guidance and support throughout the project. An earlier version of this paper appeared as the second chapter of my Ph.D. dissertation at Harvard. I would also like to thank Yeon-Koo Che, Federico Echenique, Jeffrey Ely, Yannai Gonczarowski, Zoë Hitzig, François-Xavier Ladant, David Laibson, Annie Liang, Ziqi Lu, John Macke, Erik Madsen, Matthew Rabin, Brit Sharoni, Philipp Strack, Haoqi Tong, Kai Hao Yang, Weijie Zhong, and especially Eric Maskin and Tomasz Strzalecki for valuable comments and discussion. In addition, I thank multiple anonymous referees at ACM EC'22, as well as the audiences at ACM EC'22 and the Stony Brook International Conference on Game Theory for their insightful feedback.

1 Introduction

In long-term relationships, controlling information is a powerful tool for incentivizing behavior. This paper studies the optimal mechanism to motivate effort in a dynamic principal-agent problem without transfers, focusing on a specific question: When can the principal benefit from revealing information over time rather than all at once?

For instance, consider the relationship between a principal investigator (PI, she) and a postdoctoral researcher (postdoc, he) in a research laboratory. The monetary compensation in a postdoctoral contract usually does not include bonuses for good research results, but the postdoc does build a reputation that is beneficial to his future career path by achieving major breakthroughs in projects. The PI has worked longer in academia, so she has more information about the prospects and rewards of particular research projects. By revealing information to the postdoc, the PI can influence his attitude towards the project and indirectly motivate his effort.

How should the PI release information to motivate the postdoc's effort? In particular, is it optimal for her to provide an initial disclosure and then keep silent, or can she benefit from *dynamic disclosure*? If dynamic disclosure is beneficial, is it better to delay all disclosures to a certain date, or to gradually provide small amounts of information over time? What determines the form of the optimal information policy?¹

We study these questions in a dynamic principal-agent model without transfers (Section 2). An agent works for a principal on a task and exerts effort continuously until making an *irreversible* decision to quit. The principal aims to maintain the agent's effort, while the agent faces a trade-off between the reward from completing the task and a flow cost of exerting effort. This future reward, which we call the task *quality*, is initially only known by the principal. Task completion is stochastic and arrives according to a Poisson process, and the rate of task completion may depend on the task quality. Both players observe when the task is completed, at which point the relationship ends. The principal can commit to an arbitrary dynamic information policy, specifying how and when to disclose information to the agent about the quality of the task. The agent knows the policy and understands the principal's commitment, observes the realized disclosures, rationally updates his belief about the task quality, and at each moment chooses whether to continue working or to quit in order to maximize his expected payoff.

Our main results provide closed-form characterizations of the optimal information policy and find that dynamic disclosure is not always necessary: in some environments, a well-chosen static policy is optimal

¹Our model also applies naturally to master-apprentice relationships (where rewards represent the apprentice's future job prospect) and corporate hierarchies (where rewards correspond to promotions or contract renewals). Moreover, dynamic disclosure can be interpreted as the strategic timing of information-generating events, such as scheduling experiments in scientific research and external committee reviews in business settings.

for the principal. However, we identify two precise conditions under which dynamic information disclosure is strictly valuable: one is that the principal is *impatient* compared with the agent, and the other is that the agent would become *pessimistic* over time absent any information disclosure.

We begin with the baseline version of the problem, where the task quality is *binary* and the environment is *stationary* (Section 3). The rate of task completion does not depend on its quality, so the mere passage of time is not informative. We show that the principal benefits from dynamic disclosure if and only if she is *less patient* than the agent (the first part of Proposition 1). To see the intuition behind this result, fix some static disclosure policy, and suppose that the principal introduces a short delay before sending any messages. To maintain its persuasive effectiveness, the delayed message must give the agent greater confidence that the task quality is high; without this additional optimism, the agent will find it not worthwhile to wait for disclosure while actively working and incurring flow costs. Overall, the gain from this dynamic adjustment to the static policy is that the agent always starts working until the time of disclosure regardless of the task quality, while the loss is that the agent is more likely to quit from the low-quality task upon disclosure. Because the gain occurs earlier than the loss, the principal prefers this front-loaded effort schedule when she is less patient than the agent.

Moreover, in the binary-stationary environment, the optimal information policy for an impatient principal takes the form of a *maximally delayed disclosure*: the principal delays all disclosures up to the maximum time threshold and then fully discloses the task quality (the second part of Proposition 1). Note that an arbitrary information policy induces a lottery over the agent's quitting time, and this characterization emerges from analyzing how the principal and agent differ in their preferences over such time lotteries. Their incentives align for high quality tasks – the optimal policy never induces quitting in this case – but diverge for low-quality tasks. When the task quality is low, exponential discounting implies that the principal is risk-averse over time lotteries, while the agent is risk-seeking (since the agent's value is multiplied by a negative constant). The principal's problem reduces to choosing a time lottery that minimizes the loss from quitting (relative to eventual task completion), subject to incentive compatibility constraints. The structure of the optimal time lottery is determined by comparing the risk attitudes of the two parties, which under exponential discounting simplifies to comparing their discount rates. When the principal is less patient than the agent, her greater risk aversion leads her to prefer contracting the time lottery while maintaining the agent's value. This drives the optimal lottery toward minimal risk – becoming degenerate at a specific time threshold. Such a lottery corresponds to delayed full disclosure of the task quality. Furthermore, as long as the agent's ex ante value remains strictly positive, the principal benefits from further delay, ultimately leading to maximally delayed disclosure as the optimal policy.

Next, we analyze the optimal policy when the task quality takes more than two values (Section 4), assuming the environment remains stationary and the distribution of task quality has a full-support density. We show that, just as in the binary-stationary case, the principal benefits from dynamic disclosure if and only if she is less patient than the agent. However, the optimal information policy now exhibits richer structure compared with maximally delayed disclosure. It involves a *cutoff cascade* – a series of time-varying thresholds where at each instant, the principal reveals only whether the task quality exceeds the current cutoff. Depending on model parameters, this takes one of two forms: *immediate and gradual disclosure*, which combines immediate information release with gradual disclosure over time, or *delayed gradual disclosure*, which delays all disclosure until a specific time before beginning gradual revelation. Unlike the binary case, even when delay is optimal, the policy involves gradual rather than one-shot disclosure, as the principal now benefits from customizing disclosure times for different quality realizations. This cascade structure demonstrates the principal’s fundamental trade-off between sorting and risk-sharing. On the one hand, efficiency demands sorting effort according to task quality. Earlier disclosure achieves better sorting – the agent quits low-quality tasks and continues only with promising ones. On the other hand, if the principal withholds information, the agent remains uninformed and continues working across all quality levels, including low-quality tasks, thereby bearing the risk of wasted effort. Such opacity may increase overall effort provision, but it does so by shifting the risk onto the agent, who may work longer on a low-quality task without knowing it. The optimal information policy balances these two competing forces.

Finally, we move on to nonstationary environments (Section 5) by allowing the task completion rate to vary with the task quality, assuming the task quality is binary. The analysis boils down to two cases, depending on whether the completion rate of the high-quality task is larger than that of the low-quality task. Notably, if the high-quality task has a *shorter* expected duration than the low-quality task, then a new incentive to delay disclosure arises. Here, “no news is bad news”: the agent becomes *pessimistic* over time in absence of any information disclosure because it becomes more likely that the task has low quality. In this pessimistic case, we show that the principal always benefits from dynamic disclosure, *regardless of whether she is more patient than the agent*.² Intuitively, providing information over time can counteract this downward trend in the agent’s belief and motivate the agent to work longer. However, we find that the level of patience is still a crucial factor in determining the structure of the optimal information policy (Proposition 5). Although maximally delayed disclosure remains optimal for an impatient principal, the optimal information policy for a patient principal involves *Poisson disclosure*, where full disclosure arrives

²Conversely, in the case where the agent becomes *optimistic* over time, the results from the stationary environment directly apply. If the agent starts working, he will never find it optimal to quit early, so the optimistic case reduces to an essentially stationary problem.

at a calibrated Poisson rate in order to restore the stationarity of the environment, i.e., to ensure that “no news is no news”.

To summarize, our analysis identifies two distinct channels through which dynamic disclosure becomes strictly valuable: (i) relative impatience of the principal, which makes front-loaded effort schedules attractive, and (ii) a pessimistic environment, where dynamic disclosure counteracts the agent’s declining beliefs. Table 1 organizes the optimal information policies derived in this paper, organized by the quality distribution (binary versus general) and the environment (stationary versus nonstationary). The general nonstationary

	Stationary	Nonstationary (Pessimistic)
Binary	Static (if patient) Maximally delayed (if impatient) (Section 3)	Poisson (if patient) Maximally delayed (if impatient) (Section 5)
General	Static (if patient) Gradual (if impatient) (Section 4)	<i>Open problem</i>

Table 1: Summary of optimal information policies across environments.

case remains an open problem for future research.

Beyond identifying when dynamic disclosure is strictly valuable, our framework provides a complete characterization of optimal policies, revealing how the solution structure undergoes phase transitions as underlying parameters change. By varying players’ discount rates, the reward distribution, or task completion rates, the optimal policy shifts between distinct patterns that connect to existing literature. In some regions, the solution echoes [Ely and Szydłowski \(2020\)](#) through maximally delayed disclosure that front-loads effort. In others, it resembles [Smolin \(2021\)](#) through gradual, state-contingent revelation managing the trade-off between sorting efficiency and risk allocation. While these earlier contributions each feature one type of mechanism, our unified framework shows that these emerge as different phases of a single model, with boundaries determined by threshold conditions on patience and the environment’s informational properties.

What drives these phase transitions is which incentive constraints bind at optimality. In stationary environments, only the agent’s *ex ante* individual rationality constraint binds at optimality, and the players’ relative patience determines whether the principal uses static or dynamic disclosure through risk-shifting considerations arising from exponential discounting. In nonstationary environments, the challenge becomes identifying which of the agent’s many incentive constraints across different times must bind. In the pessimistic case, the incentive constraint at the moment the agent’s belief decays to the quitting threshold becomes crucial (our Free-Riding Lemma 4), and multiple dynamic obedience constraints may bind simultaneously. This shift in binding constraints – from static individual rationality to dynamic obedience – explains why the optimal policies range from simple delay to complex gradual rules.

On the technical side, we develop a unified proof approach that handles both stationary and nonstationary environments. Throughout, we employ a change-of-variables technique that transforms the principal's problem into an optimization over time lotteries, where comparing risk attitudes under exponential discounting yields closed-form solutions.³ For stationary environments, this works cleanly because dynamic incentive constraints reduce to a single *ex ante* condition. Our main technical contribution is extending this approach to nonstationary environments, where the natural drift of beliefs introduces a continuum of interim constraints that must be handled simultaneously.

1.1 Related Literature

The literature on information design and Bayesian persuasion, to which this paper contributes, was pioneered by the work of [Aumann, Maschler, and Stearns \(1995\)](#) and [Kamenica and Gentzkow \(2011\)](#). Within this broad category, there is a particularly relevant literature on dynamic information design, which focuses on the role of information control in dynamic environments in terms of motivating and coordinating actions.

Our model is closest to [Ely and Szydłowski \(2020\)](#), which also studies the optimal information policy for a principal to motivate effort in a dynamic moral hazard framework without transfers. In [Ely and Szydłowski \(2020\)](#), the asymmetric information is about the duration of the task: the principal can conceal task completion, strategically keeping the agent working even after the task is complete. Because the principal can always exploit this hidden information, dynamic disclosure is universally optimal. Our model, by contrast, assumes both parties observe task completion, precluding such exploitation. Dynamic disclosure then need not outperform static disclosure, and our unified framework characterizes precisely when it does: if and only if one of the two channels we identify is present.

A second closely related paper is [Smolin \(2021\)](#), which studies optimal performance evaluation design in a career-concerns model where an agent of uncertain productivity decides when to quit based on feedback. In [Smolin \(2021\)](#)'s setting, the agent learns from evaluations derived from stochastic performance signals, and the principal designs how these signals are aggregated into evaluations. The equilibrium policy produces a deterministic and typically increasing wage-tenure profile through sorting: agents who remain are increasingly likely to be high types. When dynamic disclosure is optimal in our general stationary case, the policy structure is similar – a cascade of increasing cutoffs over time – but the underlying mechanism differs. In our setting, the structure stems not only from sorting but also from front-loading effort, which benefits only an impatient principal.

This paper connects to [Au \(2015\)](#) and [Guo and Shmaya \(2018\)](#), as it identifies conditions under which

³We thank an anonymous referee for suggesting this change of variables for the stationary case.

the optimal disclosure plan in a dynamic framework is static. However, their setting involves dynamic persuasion of a privately informed receiver, where the sender reveals information sequentially to screen the receiver's type. Also related are [Ely \(2017\)](#), [Renault, Solan, and Vieille \(2017\)](#) and [Ball \(2023\)](#); in their settings, the principal privately observes a state that evolves according to some exogenous stochastic process, and the structure of the optimal information policy depends on this process. By contrast, the model in this paper characterizes the timing and form of information disclosure when the state is fixed. Moreover, in both [Ely \(2017\)](#) and [Renault, Solan, and Vieille \(2017\)](#) the agent acts myopically, whereas this paper studies the dynamic relationship between two long-lived players.

Several other recent papers study related topics but differ in their assumptions about what information the principal can control. In [Kaya \(2023\)](#), the principal can commit within each period to any signal structure but cannot commit over time to a sequence of signals.⁴ Thus, the relevant solution concept is perfect Bayesian equilibrium. In our model, the principal has full commitment power over arbitrary signal processes, yet our result that a static policy can be dynamically optimal indicates that this assumption is stronger than necessary in some cases. [Orlov, Skrzypacz, and Zryumov \(2020\)](#) and [Bizzotto, Rüdiger, and Vigier \(2021\)](#) assume instead that the principal is not the only source of information: in their models, there is an exogenous flow of payoff-relevant information beyond the principal's control. Furthermore, [Che, Kim, and Mierendorff \(2023\)](#) introduce costs of generating and processing information, and obtain a version of the folk theorem as these costs vanishes.⁵

The rest of the paper is organized as follows. Section 2 lays out the model. The first main part, Section 3, analyzes the binary-stationary case and derives one condition that makes dynamic disclosure valuable: the principal is *impatient* compared with the agent. The second main part, Section 4, extends the analysis to general reward distributions and reinforces the results from the binary case. The third main part, Section 5, analyzes the nonstationary environments and reveals another channel through which dynamic disclosure becomes valuable: the agent becomes *pessimistic* over time without any information disclosure. Section 6 concludes. Appendix A contains the proofs of all results in the main text.

⁴The commitment assumption along this style also appears in [Orlov, Skrzypacz, and Zryumov \(2020\)](#), [Bizzotto, Rüdiger, and Vigier \(2021\)](#), and [Escudé and Sinander \(2023\)](#).

⁵Recent working papers, circulated after earlier versions of the present work, explore related themes in dynamic information design. For example, [Koh, Sanguanmoo, and Zhong \(2025\)](#) develop a unified framework for persuasion with optimal stopping, establishing a revelation principle under full commitment and characterizing implementable outcomes without intertemporal commitment. Our focus on the conditions under which dynamic disclosure is strictly optimal is complementary.

2 Model

The model is a continuous-time principal-agent problem without transfers. An agent works for a principal and exerts effort until he chooses to quit. We assume that the quit decision is *irreversible*. The agent exerts effort at flow cost $c > 0$, earns a reward $q \geq 0$ if he works until the task is completed, and discounts at rate $r > 0$. We refer to the reward, q , as the quality of the task. The realized quality q is the principal's private information, and the agent's prior is given by a distribution F with finite first moment: $\mathbb{E}_F[q] < \infty$.

Assume that both players observe when the task is completed, and the completion arrives with Poisson rate $\lambda_q > 0$ when the task quality is q , at which point the relationship ends. This implies that when the task quality is q , the task duration x is an exponential random variable with mean $1/\lambda_q$. Thus, given realized duration x , the agent's ex post payoff from exerting effort until time τ is

$$e^{-r\tau} q \cdot \mathbf{1}(\tau \geq x) - \int_0^\tau e^{-rt} c dt = e^{-r\tau} q \cdot \mathbf{1}(\tau \geq x) - \frac{c}{r} (1 - e^{-r\tau}).$$

The principal receives a bonus $b(q) > 0$ if the agent works until the task is completed, and discounts at rate $r_p > 0$; therefore, her ex post payoff when the agent quits at τ is

$$e^{-r_p \tau} b(q) \cdot \mathbf{1}(\tau \geq x).$$

We assume $b(\cdot)$ is a bounded measurable function with $0 < \underline{b} \leq b(q) \leq \bar{b} < \infty$ for all $q \in [0, \infty)$. These assumptions ensure that both players' expected payoffs are always well-defined.

We now compute the following expected payoff functions, which will be useful for later analysis. The agent's expected payoff from completing a task with quality q (i.e., quit at $\tau = \infty$) is

$$\begin{aligned} v(q) &\equiv \mathbb{E}_{x|q} \left[e^{-rx} q - \frac{c}{r} (1 - e^{-rx}) \middle| q \right] \\ &= \int_0^\infty \left[e^{-rx} q - \frac{c}{r} (1 - e^{-rx}) \right] \lambda_q e^{-\lambda_q x} dx = \frac{\lambda_q q - c}{r + \lambda_q}. \end{aligned}$$

The agent may also plan to quit at some time τ if the task is still incomplete. This leads to expected payoff

$$\begin{aligned} v(\tau, q) &\equiv \mathbb{E}_{x|q} \left[e^{-rx} q \cdot \mathbf{1}(x \leq \tau) - \frac{c}{r} (1 - e^{-r \min\{x, \tau\}}) \middle| q \right] \\ &= \frac{\lambda_q q - c}{r + \lambda_q} (1 - e^{-(r + \lambda_q)\tau}) = (1 - e^{-(r + \lambda_q)\tau}) v(q). \end{aligned} \tag{1}$$

Similarly, the principal's expected payoff when the agent chooses to quit at τ is

$$\begin{aligned} w(\tau, q) &\equiv \mathbb{E}_{x|q} [e^{-r_p x} b(q) \cdot \mathbf{1}(x \leq \tau) | q] \\ &= \frac{\lambda_q b(q)}{r_p + \lambda_q} (1 - e^{-(r_p + \lambda_q)\tau}) \equiv (1 - e^{-(r_p + \lambda_q)\tau}) w(q). \end{aligned} \quad (2)$$

In particular, $w(q) = \frac{\lambda_q b(q)}{r_p + \lambda_q}$ denotes the principal's expected payoff from a completed task with quality q .

Assume that the principal can commit to an arbitrary rule that specifies how and when to disclose information to the agent about the quality of the task.⁶ She can make any number of disclosures at any time during the process. The agent knows the rule and understands the principal's commitment, observes the realized disclosures, rationally updates his belief about the quality of the task, and best responds by making effort choices that maximize his expected payoff. Adapting the proof of Proposition 1 in [Kamenica and Gentzkow \(2011\)](#), one can show that it is without loss of generality to restrict attention to a particular class of rules that produce “recommended actions”, an analog to the revelation principle. Since the agent's decision to quit is irreversible, the only relevant instruction the principal can provide is the time at which the agent should stop. We refer to a rule governing these disclosures as an *information policy*, which informs the agent whether to continue ($\sigma_t = 0$) or quit ($\sigma_t = 1$) at any time.

Definition 1 (Information policy). *An information policy $\sigma = \{\sigma_t\}_{t \geq 0}$ is a $\{0, 1\}$ -valued, non-decreasing, and right-continuous stochastic process whose distribution depends on the task quality q . The policy is fully characterized by the family of conditional probabilities $\mathbb{P}(\sigma_t = 1 | q)$, which represents the probability that quitting has been recommended by time t given quality q .*

The principal's objective is to maximize her expected payoff $\mathbb{E}_{\sigma}[w(\tau, q)]$ over all possible information policies σ , where τ denotes the agent's random quitting time induced by σ .⁷

In the following sections, we derive the optimal information policy in closed form, in order to obtain the necessary and sufficient conditions for the principal to benefit from dynamic disclosure.

⁶In organizations, commitment can be sustained formally through contracts (e.g., reviewed by an external evaluator), or informally through reputation. Moreover, commitment can also be understood as using existing instruments to generate signals that cannot be hidden. In the previous example of a research lab, suppose the project is to develop a new empirical method, and the PI has control over exclusive data sets that can improve the understanding of the value of completing a breakthrough through test runs. The PI can decide which data set to provide to the postdoc and when to provide it, but once the data set is provided, the results of the test run will be publicly observed. Finally, as we show in the following sections, when the principal is patient and the environment either is stationary or makes the agent become optimistic over time, a static information policy turns out to be dynamically optimal. In these cases, full commitment power over arbitrary signal processes is stronger than necessary.

⁷The agent's quitting time is determined by comparing two payoffs at each instant of time: the continuation payoff from following the principal's recommendations versus the value of ignoring further information. We omit the formal expressions from the main text for brevity.

3 Binary-Stationary Case

We begin our analysis with the baseline version of the problem, where the task quality takes one of two possible values: low quality, $L > 0$, or high quality, $H > L$, and the environment is stationary: $\lambda_H = \lambda_L = \lambda$. The agent's prior belief assigns probability $\mu \in (0, 1)$ that the task has high quality.

Assume $v(H) > 0 > v(L)$, so that the low-quality task is not individually rational for the agent. We fully characterize the optimal dynamic information policy, and, through this analysis, reveal the first channel that makes dynamic information disclosure valuable: the principal is *impatient* compared with the agent.

The main result for the binary-stationary case is Proposition 1, which shows that the principal benefits from dynamic disclosure if and only if she is *less patient*⁸ than the agent, i.e., $r_p > r$. Moreover, if $r_p > r$, *maximally delayed disclosure* is optimal: the principal delays all disclosure up to the maximum time and then fully discloses the task quality.

3.1 Static Disclosure

Consider a one-shot information release at time zero, after which the agent decides how long to work. Under this static disclosure, once the agent receives the principal's initial message, he effectively faces a binary choice: continue working on the task until completion or quit immediately. This is because the environment is time-stationary: by the memoryless property of Poisson arrivals, the agent's continuation problem at any time $t > 0$ is identical to the one at time zero. Hence, the static version reduces to the standard Bayesian persuasion problem of Kamenica and Gentzkow (2011).

Let $\bar{\mu}$ denote the threshold belief at which the agent is indifferent between working and quitting:

$$\bar{\mu}v(H) + (1 - \bar{\mu})v(L) = 0. \quad (3)$$

For $\mu \geq \bar{\mu}$, no persuasion is necessary. When $\mu < \bar{\mu}$, the principal's objective is to design an initial signal that maximizes the probability the agent's belief remains above the $\bar{\mu}$ threshold, thereby inducing him to work. The well-known solution is a binary signal that we refer to as the KG policy.

Definition 2. KG denotes the static information policy that, for $\mu < \bar{\mu}$, satisfies:

1. The agent is told to continue ($\sigma_0 = 0$) with probability one when the task has high quality ($q = H$), and also with positive probability when the task has low quality ($q = L$).

⁸While it is common to assume that the principal is more patient than the agent, we want to point out that the principal may be relatively impatient in some cases. In the previous example of a research lab, the PI could be under pressure for a promotion review, and the postdoc may have just been hired.

2. The agent exactly holds belief $\bar{\mu}$ conditional on observing $\sigma_0 = 0$, i.e., $\mathbb{P}(q = H \mid \sigma_0 = 0) = \bar{\mu}$.
3. The policy is static: for all $t > 0$, $\sigma_t = \sigma_0$ almost surely.

The principal's value function under KG with respect to the agent's belief, $W^{\text{KG}}(\mu)$, is depicted in Figure 1. The solid line segments represent the benchmark without persuasion, and the dashed line segment, formed by the concavification of the solid line segments, represents the principal's gain obtained from persuasion.⁹

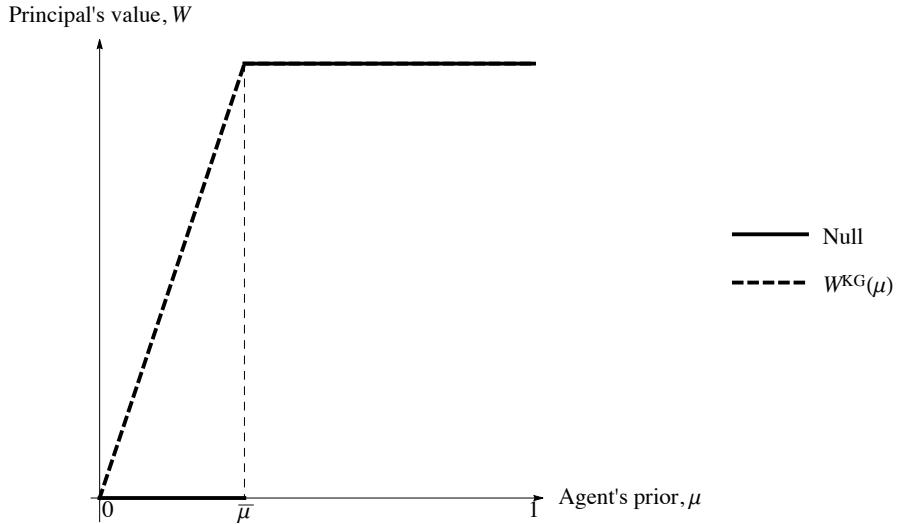


Figure 1: Optimal static information policy KG corresponds to a concavification.

The static KG policy serves as a benchmark for the dynamic policies we consider next. In fact, we will see that if the principal is relatively patient (low r_p), this static disclosure remains dynamically optimal even when more complex information policies are available.

3.2 Delayed Disclosure

We now allow the principal to spread out information over time and examine whether doing so can lead to higher payoffs. Consider the simplest dynamic policy, where the principal withholds all information until time $t > 0$, then fully discloses the task quality.¹⁰ Again, due to the memoryless property, conditional on working until time t , the agent will quit when informed of low quality and continue until completion otherwise. We call this policy *delayed disclosure* at time t , denoted by $\text{DD}(t)$.

⁹Only for illustration, Figure 1, as well as Figures 2, 3a, and 3b below, assumes $b(L) = b(H)$; this does not affect any of our results.

¹⁰Of course, the principal could employ more complex dynamic policies that do not fully reveal the state, but as we later show, such complexity is suboptimal in the binary-stationary case.

Definition 3 (Delayed disclosure). $DD(t)$ denotes the information policy of delayed disclosure at time t , where the principal tells the agent nothing prior to time t , and fully discloses the task quality at t . Formally, $\sigma_s = 0$ for all $s < t$, and for all $s \geq t$, $\sigma_s = \mathbf{1}(q = L)$.

Of course, delaying information can backfire if the agent's initial belief is too low – the agent might refuse to work from the outset unless he anticipates a sufficient chance of a high-quality task. Let $\tilde{\mu}(t)$ denote the minimum prior belief at which the agent is willing to start working under the policy $DD(t)$. This $\tilde{\mu}(t)$ is determined by the agent's ex ante incentive constraint: the solution to¹¹

$$\tilde{\mu}(t)v(H) + (1 - \tilde{\mu}(t))v(t, L) = 0. \quad (4)$$

Next, we examine whether $DD(t)$ enables the principal to increase her payoff relative to KG. The answer depends crucially on patience, i.e., the relationship between r_p and r . The gain from $DD(t)$ compared with KG is that the agent always starts working regardless of the task quality, and the loss from $DD(t)$ is that the agent quits from the low-quality task after the disclosure at time t . Since the gain occurs earlier than the loss, an impatient principal benefits from $DD(t)$. The formal result is summarized in the following lemma.

Lemma 1. Fix a time $t > 0$. $DD(t)$ increases the principal's payoff relative to KG for some range of priors if and only if the principal is less patient than the agent, i.e., $r_p > r$.

Proof. All proofs of the results in the main text are in Appendix A. □

Figure 2 illustrates this pattern: when $r_p > r$, there exists a range of priors for which $DD(t)$ yields higher payoffs than KG (Figure 2a); when $r_p < r$, $DD(t)$ is worse for the principal than KG for all priors (Figure 2b).

Taking one step further, we argue that whenever the principal gains from delaying disclosure (i.e., when $r_p > r$), she can improve her payoff by delaying as much as possible subject to keeping the agent on board. For any $\mu < \bar{\mu}$, let $\tilde{t}(\mu)$ denote the *maximum delay time*, i.e., the latest time of disclosure subject to the constraint that the agent's ex ante payoff is at least zero:¹²

$$\mu v(H) + (1 - \mu)v(\tilde{t}(\mu), L) = 0. \quad (5)$$

¹¹Since $\bar{\mu}$ is the smallest belief at which the agent would start working under no disclosure, we have $\tilde{\mu}(t) < \bar{\mu}$ for all $t > 0$. It follows from equation (1) that $v(\tau, L) = (1 - e^{-(r+\lambda)\tau})v(L)$ with $v(L) < 0$, so $v(\tau, L)$ strictly decreases from 0 to $v(L)$ as τ increases from 0 to infinity. Comparing equations (3) and (4) that define the threshold priors $\bar{\mu}$ and $\tilde{\mu}(t)$, we can see that (i) $\tilde{\mu}(t)$ is strictly increasing in t , (ii) $\lim_{t \rightarrow 0} \tilde{\mu}(t) = 0$, $\lim_{t \rightarrow \infty} \tilde{\mu}(t) = \bar{\mu}$.

¹²Note from equations (4) and (5) that $\tilde{t}(\mu)$ is the inverse of $\tilde{\mu}(t)$, so it satisfies (i) $\tilde{t}(\mu)$ is strictly increasing in μ within the range $(0, \bar{\mu})$, (ii) $\tilde{t}(0) = 0$, $\lim_{\mu \uparrow \bar{\mu}} \tilde{t}(\mu) = \infty$, (iii) $\tilde{t}(\mu) = \infty$ for all $\mu \geq \bar{\mu}$.

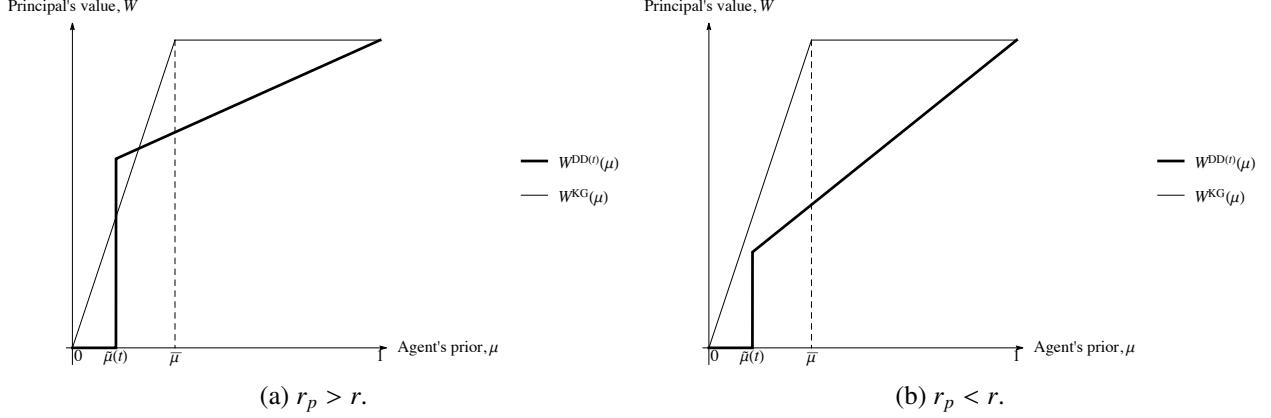


Figure 2: The principal's payoff comparison between DD(t) and KG.

Among all disclosure times t that are sufficient to motive the agent under DD(t) (i.e., $t \leq \tilde{t}(\mu)$), the principal's payoff is increasing in t . Intuitively, a longer wait means the agent will work longer on the low-quality task before quitting, which is precisely what the impatient principal values. Therefore, the optimal choice is to set $t = \tilde{t}(\mu)$, the longest delay that just barely induces the agent to participate. We refer to this extreme policy as *maximally delayed disclosure*, denoted by MDD.

Definition 4 (Maximally delayed disclosure). MDD denotes the information policy of maximally delayed disclosure. For $\mu < \bar{\mu}$, the principal provides no information prior to the time $\tilde{t}(\mu)$, and fully discloses the task quality at $\tilde{t}(\mu)$. Formally, $\sigma_s = 0$ for all $s < \tilde{t}(\mu)$, and for all $s \geq \tilde{t}(\mu)$, $\sigma_s = \mathbf{1}$ ($q = L$).

By construction, for $\mu < \bar{\mu}$, the agent is exactly indifferent to starting work under MDD (he gets zero expected payoff), so he complies with the recommendation to continue until $\tilde{t}(\mu)$. Figure 3a depicts the principal's value under MDD compared to simpler DD(t) policies and the static KG policy when $r_p > r$. In particular, for all $\mu \in (0, \bar{\mu})$, MDD strictly increases the principal's payoff relative to the optimal static

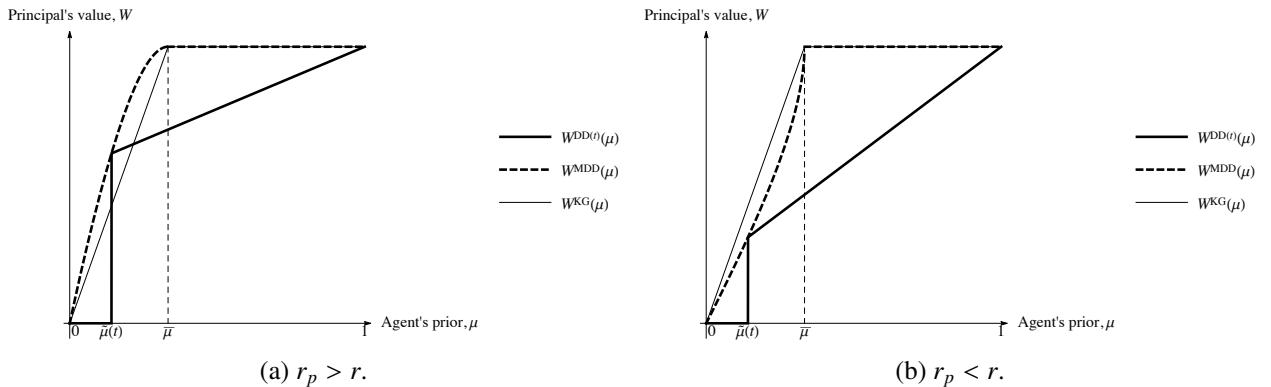


Figure 3: The principal's payoff comparison among DD(t), MDD and KG.

policy KG. Moreover, W^{MDD} is the upper envelope and concave closure of the family of value functions

$\{W^{\text{DD}(t)} : t > 0\}$. Hence, W^{MDD} itself is a concave function, and it follows that MDD yields a higher payoff than any information policy with a less informative disclosure. In other words, once we allow the disclosure time to be optimally chosen as a function of the prior, gradual or partial disclosures cannot outperform the bang-bang approach of MDD. In Proposition 1, we establish that MDD is indeed optimal within the set of all dynamic information policies.

Conversely, consider the case $r_p < r$. The relationship among the information policies DD(t), MDD and KG is shown in Figure 3b. The principal still prefers MDD to any policy in the family $\{\text{DD}(t) : t > 0\}$, but W^{MDD} is no longer concave, and its concavification is precisely W^{KG} . In this case, Proposition 1 establishes the optimality of the static policy KG within all dynamic information policies. In particular, when the principal is more patient than the agent, her commitment power over the entire signal process has no additional value.

3.3 Optimal Information Policy

We are now ready to state the main result for the binary-stationary case, Proposition 1, which establishes the optimality of the information policies analyzed above. Among all dynamic information policies, KG is optimal when $r_p \leq r$, and MDD is optimal when $r_p > r$.¹³

Proposition 1. *In the binary-stationary case,*

1. *If $r_p \leq r$, then KG is optimal among all dynamic information policies.*
2. *If $r_p > r$, then MDD is optimal among all dynamic information policies.*

The optimality is within the entire set of dynamic information policies, including those with multiple or even infinite disclosures that may induce quitting times other than zero or $\tilde{t}(\mu)$.

The intuition behind Proposition 1 is best understood by viewing the information design problem as choosing a lottery over the agent's quitting time. Any information policy, σ , induces a distribution over times at which the agent quits, τ . The principal wants to make this quitting time as “late” as possible, but at the same time she must respect the agent's incentive compatibility constraints. Because the agent never needs to be encouraged to continue when the task quality is high, the non-trivial design problem only concerns when the task quality is low.

Crucially, under exponential discounting, the principal and agent have opposing risk preferences over

¹³When $r_p = r$, a range of policies including KG and MDD are optimal, as long as they are sufficient to motivate the agent and such that the agent's ex ante expected payoff is exactly zero (e.g., non-maximally delayed partial disclosure).

time lotteries: the principal is risk-averse, while the agent is risk-seeking.¹⁴ Formally, applying Jensen's inequality to equations (1) and (2):

$$\begin{aligned}\mathbb{E}_\sigma[w(\tau, L)] &= w(L) \cdot \mathbb{E}_\sigma\left[1 - e^{-(r_p+\lambda)\tau}\right] \leq w(L)\left(1 - e^{-(r_p+\lambda)\mathbb{E}_\sigma[\tau]}\right) = w(\mathbb{E}_\sigma[\tau], L), \\ \mathbb{E}_\sigma[v(\tau, L)] &= \underbrace{v(L)}_{(-)} \cdot \mathbb{E}_\sigma\left[1 - e^{-(r+\lambda)\tau}\right] \geq \underbrace{v(L)}_{(-)} \cdot \left(1 - e^{-(r+\lambda)\mathbb{E}_\sigma[\tau]}\right) = v(\mathbb{E}_\sigma[\tau], L).\end{aligned}$$

Given this difference, the optimal policy depends on whose risk attitude dominates, and skews towards one of two extremes depending on who is more patient. If $r_p < r$, the agent's risk-seeking preference dominates, favoring an extreme lottery with maximum uncertainty: either immediate quitting or indefinite continuation, precisely what KG delivers. On the other hand, if $r_p > r$, the principal's risk-aversion dominates, favoring a deterministic quitting time, which MDD provides by concentrating all probability mass at $\tilde{t}(\mu)$.

4 General Stationary Case

We now characterize the optimal dynamic information policy for general reward distributions F , assuming that the environment is stationary ($\lambda_q \equiv \lambda$) and that F has a full-support density over $[0, \infty)$. The latter assumption implies that it is without loss of optimality to restrict attention to policies that induce a deterministic quitting time $\tau(q)$ for each quality realization.¹⁵ In addition, we assume that the agent's ex ante payoff from always completing the task is negative, i.e., $\mathbb{E}_q[v(q)] = \int_0^\infty v(q) dF(q) < 0$, as otherwise no persuasion would be necessary. Finally, we assume that the principal's bonus $b(q)$ is nondecreasing in task quality q .¹⁶

The main result for this section parallels the binary-case finding from Section 3: in a stationary environment, the principal benefits from dynamic disclosure if and only if she is less patient than the agent ($r_p > r$). If $r_p \leq r$, Proposition 2 shows that the optimal information policy is an adaptation of KG. By contrast, if $r_p > r$, the principal strictly benefits from dynamic disclosure, and the optimal policy becomes a *cutoff cascade* – a series of time-varying thresholds where at each instant, the principal reveals only whether

¹⁴The connection between exponential discounting and risk preferences over time lotteries already features in the analysis in Section IV.F of Ely and Szydłowski (2020).

¹⁵In discrete settings (e.g., Sections 3 and 5), optimal policies often require randomization to convexify the principal's feasible payoffs and satisfy individual rationality constraints with equality. With an atomless quality distribution, the principal can bind these constraints by continuously adjusting the quality cutoff, thereby eliminating the need for randomization.

¹⁶This assumption applies to many situations. One example is that the principal receives the same bonus at completion regardless of the task quality ($b(q) = b_0$). Another is that the principal and the agent have the same reward from the task ($b(q) = q$), so the difference between their payoffs is only that the principal does not incur the costs. Furthermore, the results for this section would go through even if $b(q)$ may be decreasing, as long as it does not decrease too rapidly. Formally, we need the ratio $|v(q)|/b(q)$ to be nonincreasing over $q \in (0, \bar{q})$.

the task quality exceeds the current cutoff, as shown in Propositions 3 and 4. This gradualism custom-fits timing to realized quality, generalizing MDD from a single cutoff to a filtering of low qualities over time.

The section proceeds as follows. Subsection 4.1 establishes the optimality of static disclosure when the principal is patient ($r_p \leq r$). Subsection 4.2 shows that when the principal is impatient ($r_p > r$), the optimal policy takes the form of a *cutoff cascade*. Subsections 4.3 and 4.4 then characterize the specific structure of the optimal cascade.

4.1 Static Disclosure

As in the binary case, we first consider the static version where the principal provides a one-shot disclosure and the agent chooses effort. Because F has full support and no atoms, there exists a unique task quality $q^* > 0$ such that

$$\int_{q^*}^{\infty} v(q) dF(q) = 0. \quad (6)$$

Let \bar{q} denote the *minimum individually rational quality* for the agent; that is, \bar{q} solves

$$v(\bar{q}) = \frac{\lambda\bar{q} - c}{r + \lambda} = 0, \quad \text{i.e.,} \quad \bar{q} = \frac{c}{\lambda}.$$

It follows from equation (6) that $q^* < \bar{q}$.

Consider the following static information policy, which is a counterpart of KG: the principal informs the agent whether the task has high quality ($q \geq q^*$) or low quality ($q < q^*$). By stationarity, if informed of high quality, the agent completes the task ($\tau = \infty$); otherwise, he quits immediately ($\tau = 0$). Slightly abusing notation, we also use KG to represent this information policy.

Definition 5. KG denotes the static information policy that informs the agent whether the task has high quality ($q \geq q^*$) or low quality ($q < q^*$). Formally, $\sigma_0(q) = \mathbf{1}(q < q^*)$ and $\sigma_t = \sigma_0$ for all $t \geq 0$, with induced quitting time

$$\tau^{\text{KG}}(q) = \begin{cases} 0, & q < q^*, \\ \infty, & q \geq q^*. \end{cases}$$

Indeed, the initial disclosure motivates the agent to complete some tasks with lower quality than the minimum individually rational level (those in the range $q^* \leq q < \bar{q}$). As an analog of the first part of Proposition 1, we show below that if $r_p \leq r$, this static policy is optimal among all dynamic ones.

Proposition 2. If $r_p \leq r$, then KG is optimal among all dynamic information policies.

4.2 Cutoff Cascades

Henceforth, we consider the case where the principal is less patient than the agent, $r_p > r$. As in the binary case, the principal strictly benefits from dynamic disclosure over static disclosure. However, with a continuum of qualities, the optimal dynamic policy differs from MDD: rather than a single delayed revelation, the principal now customizes disclosure times for different realized qualities q , strictly benefiting from this additional design flexibility. We define the following family of information policies:

Definition 6 (Cutoff cascade). *An information policy features a cutoff cascade if there exists a cutoff quality function $\tilde{q}(s)$ that is nondecreasing and ranges in $[0, \bar{q}]$, such that at each instant $s \geq 0$ if the task is incomplete, the principal tells the agent whether $q \geq \tilde{q}(s)$ or $q < \tilde{q}(s)$. The agent quits when informed of $q < \tilde{q}(s)$ and continues otherwise. Formally, $\sigma_s(q) = \mathbf{1}(q < \tilde{q}(s))$.*

Figure 4 shows examples of cutoff cascade policies. KG corresponds to the constant cutoff quality function $\tilde{q}(s) = q^*$, as shown in Figure 4a. Figure 4b plots one-shot delayed disclosure at time t , the direct adaptation of the information policy DD(t) (Definition 3). Figure 4c plots *immediate and gradual disclosure* with cutoff quality q_0 , IGD(q_0), which combines an immediate disclosure with cutoff quality q_0 , followed by gradual revelation. Figure 4d plots *delayed gradual disclosure* starting at time t , DGD(t), which delays all disclosure until time t before beginning gradual revelation. We show that when r_p is slightly larger than r , the optimal policy is IGD(q^{**}) for some optimal q^{**} (Proposition 3), while when r_p is much larger than r , it is DGD(\tilde{t}) for some optimal \tilde{t} (Proposition 4).

To understand the logic behind the optimality of cutoff cascade policies, we employ a two-step relaxation approach. We first consider a relaxed version of the principal's optimization problem where we temporarily set aside the interim incentive compatibility constraints. In this relaxed program, the principal can choose an arbitrary joint distribution over (q, τ) – the task quality and the quitting time – subject only to two constraints: (i) the marginal distribution of q matches the prior F , and (ii) the agent's individual rationality constraint is satisfied. Importantly, we ignore whether such a joint distribution can actually be implemented through any feasible information policy.

We argue that this relaxed program admits an optimal solution that is *positively assortative* – higher quality tasks are assigned longer working times. This follows from an optimal assignment principle: since the principal's bonus $b(q)$ is nondecreasing in task quality q , each additional moment of work generates higher returns on high-quality tasks than on low-quality tasks. If the principal mistakenly assigned a high-quality task a short working time and a low-quality task a long working time, she could improve her payoff by swapping these assignments without affecting the agent's expected payoff. This positive assortativity

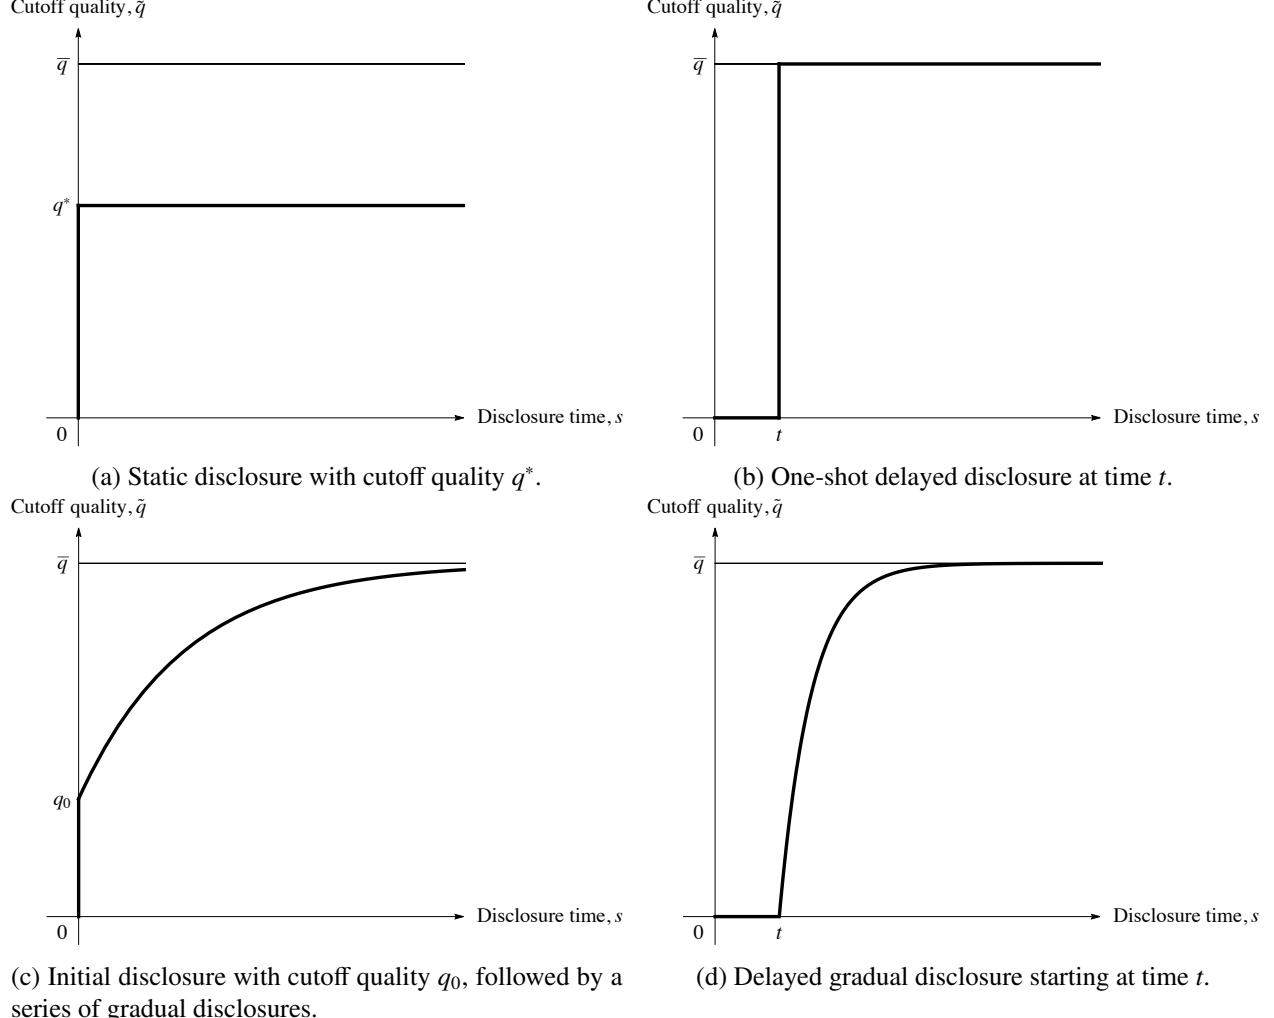


Figure 4: Examples of cutoff cascades.

naturally leads to the structure of a cutoff cascade: tasks are filtered out in order of increasing quality, with the lowest-quality tasks revealed (and quit) first, while higher-quality tasks continue.

This insight is not merely intuitive – it plays a central role in the formal analysis. By establishing positive assortativity in the relaxed program, we reduce the optimization to a tractable family of cutoff cascades. The second step of our approach shows that a positively assortative solution to the relaxed program can indeed be implemented through a cutoff cascade policy, thereby restoring incentive compatibility. This implementability result is formalized in the following Lemma 2. Individual rationality is clearly necessary: the agent must obtain at least zero ex ante expected payoff, since otherwise he can simply ignore all information, quit immediately, and ensure a zero payoff. Lemma 2 shows that individual rationality is both necessary and sufficient for the agent to follow cutoff cascade policies.

Lemma 2. *Consider any cutoff cascade policy σ . The agent's best response is to follow σ if and only if σ*

is individually rational; that is, following σ yields a nonnegative *ex ante* payoff for the agent.

The intuition behind Lemma 2 relies on the positive assortativity: higher quality leads to later quitting under cutoff cascades. As time passes without being told to quit, the agent's beliefs about task quality improve – after all, if the task were low quality, it likely would have been revealed by now. This belief updating makes cutoff cascade policies incentive compatible, as being asked to continue is good news about quality.

Having established that cutoff cascades are without loss of optimality, we now characterize the specific form of the optimal cascade. A key question is whether the principal should begin disclosing immediately or delay. The analysis divides into two cases based on whether a baseline policy – *immediate and gradual disclosure* with no initial information, IGD(0) (formally defined in Subsection 4.3) – satisfies the agent's individual rationality constraint. When the principal is moderately impatient, IGD(0) is not individually rational, and the principal optimally provides an initial disclosure to boost the agent's belief before beginning gradual revelation (Proposition 3). When the principal is highly impatient, IGD(0) is individually rational, and the principal can profitably delay before beginning gradual revelation (Proposition 4).

4.3 Immediate and Gradual Disclosure

When the principal is moderately impatient (r_p is slightly larger than r), she optimally provides an initial disclosure to raise the agent's belief and then begins gradual revelation. To formalize this, we first define *immediate and gradual disclosure* starting from an initial cutoff q_0 , denoted IGD(q_0), and then characterize the optimal initial cutoff.

We begin by defining the ratio between the agent's loss and the principal's gain for any task $q \in [0, \bar{q}]$ that is not individually rational: $u(q) \equiv |v(q)|/w(q)$. Since the principal's bonus $b(q)$ is assumed to be nondecreasing, so is $w(q) = \frac{\lambda b(q)}{r_p + \lambda}$. Moreover, since $|v(q)|$ decreases from $|v(0)|$ to zero as q rises from zero to \bar{q} , the ratio $u(q)$ is strictly decreasing and ranges from the positive value $u(0)$ to zero. Using this ratio, we define a specific cutoff cascade policy as follows.

Definition 7 (Immediate and gradual disclosure). IGD(q_0) denotes the cutoff cascade policy of immediate and gradual disclosure with initial cutoff quality $q_0 \in (0, \bar{q})$, specified by the cutoff quality function

$$\tilde{q}^{\text{IGD}(q_0)}(s) = u^{-1}(u(q_0) e^{-(r_p - r)s}), \quad (7)$$

where $u^{-1}(\cdot)$ is the inverse function of u .

Under this policy, the cutoff quality rises continuously from q_0 to \bar{q} as time s increases. Solving for the time at which each quality is revealed yields the agent's quitting time:

$$\tau^{\text{IGD}(q_0)}(q) = \begin{cases} 0, & q < q_0, \\ \frac{1}{r_p - r} \log\left(\frac{u(q_0)}{u(q)}\right), & q_0 \leq q < \bar{q}, \\ \infty, & q \geq \bar{q}. \end{cases} \quad (8)$$

Equation (8) is obtained by inverting the cutoff quality function (7); these equivalent representations characterize the policy through the timing of disclosure and the schedule of quitting, respectively.

Note that if the principal sets $q_0 = q^*$ (the cutoff in KG), $\text{IGD}(q^*)$ reveals strictly more information than KG, thereby granting the agent a strictly positive ex ante payoff. The principal can improve her payoff by lowering the initial cutoff q_0 to extract this surplus, but q_0 is bounded below by zero.

The analysis therefore divides into two cases, depending on whether $\text{IGD}(0)$ is individually rational. From equation (8), the quitting time under $\text{IGD}(0)$ for any quality $q \in (0, \bar{q})$ is $\tau^{\text{IGD}(0)}(q) = \frac{1}{r_p - r} \log\left(\frac{u(0)}{u(q)}\right)$. As r_p increases, this quitting time decreases, making the policy more favorable to the agent. When the principal is moderately impatient (r_p is slightly larger than r), the quitting times under $\text{IGD}(0)$ are long and the policy violates individual rationality. When the principal is highly impatient (r_p is much larger than r), the quitting times are shorter and $\text{IGD}(0)$ becomes individually rational.

If $\text{IGD}(0)$ violates individual rationality, the principal chooses an optimal initial cutoff quality $q^{**} \in (0, q^*)$ such that the agent's payoff from $\text{IGD}(q^{**})$ equals exactly zero.

Definition 8 (Optimal immediate and gradual disclosure). OIGD denotes the cutoff cascade policy of optimal immediate and gradual disclosure, specified by the cutoff quality function

$$\tilde{q}^{\text{OIGD}}(s) = \tilde{q}^{\text{IGD}(q^{**})}(s) = u^{-1}\left(u(q^{**})e^{-(r_p - r)s}\right),$$

where $u^{-1}(\cdot)$ is the inverse function of u , and $q^{**} \in (0, \bar{q})$ is chosen such that the agent's ex ante expected payoff equals zero.

For such a moderately impatient principal, OIGD is optimal among all dynamic information policies.

Proposition 3. If $r_p > r$ and $\text{IGD}(0)$ is not individually rational, then OIGD is optimal among all dynamic information policies.

The relationship between the information policies KG and OIGD is shown in Figure 5a, which corresponds to the case where r_p is slightly larger than r , and OIGD is optimal. As r_p approaches r from above,

the initial cutoff quality q^{**} approaches q^* from below. Moreover, for any $q \in (q^*, \bar{q})$, the quitting time $\tau^{\text{OIGD}}(q) = \frac{1}{r_p - r} \log\left(\frac{u(q^{**})}{u(q)}\right)$ tends to infinity. Therefore, OIGD converges to KG, demonstrating continuity in the optimal information policy as the principal's discount rate varies.

4.4 Delayed Gradual Disclosure

We have shown that when IGD(0) is not individually rational, the principal's optimal policy is OIGD – immediate disclosure followed by gradual revelation. We now turn to the complementary case: when IGD(0) is individually rational. This occurs when the principal is sufficiently impatient (r_p is much larger than r), and in this case IGD(0) yields a strictly positive ex ante payoff for the agent. The principal can extract this surplus by delaying disclosure. To formalize this, we first define *delayed gradual disclosure* starting at time t , denoted DGD(t), and then characterize the optimal delay time.

Definition 9 (Delayed gradual disclosure). *DGD(t) denotes the cutoff cascade policy of delayed gradual disclosure starting at time t , specified by the cutoff quality function*

$$\tilde{q}^{\text{DGD}(t)}(s) = \begin{cases} 0, & s < t, \\ u^{-1}(u(0) e^{-(r_p - r)(s-t)}), & s \geq t, \end{cases} \quad (9)$$

where $u^{-1}(\cdot)$ is the inverse function of u .

Under this policy, the cutoff quality remains zero until time t , after which it rises continuously to \bar{q} . Solving for the time at which each quality is revealed yields the agent's quitting time:

$$\tau^{\text{DGD}(t)}(q) = \begin{cases} t + \frac{1}{r_p - r} \log\left(\frac{u(0)}{u(q)}\right), & q < \bar{q}, \\ \infty, & q \geq \bar{q}. \end{cases} \quad (10)$$

Equation (10) is obtained by inverting the cutoff quality function (9); as q rises from 0 to \bar{q} , the agent's quitting time continuously rises from t to infinity.

Let \tilde{t} denote the *maximum delay time*, i.e., the latest time such that DGD(\tilde{t}) is individually rational for the agent. Since the agent's expected payoff decreases continuously in the delay time, \tilde{t} is unique and finite.

Definition 10 (Maximally delayed gradual disclosure). *MDGD denotes the cutoff cascade policy of maxi-*

mally delayed gradual disclosure, specified by the cutoff quality function

$$\tilde{q}^{\text{MDGD}}(s) = \tilde{q}^{\text{DGD}(\tilde{t})}(s) = \begin{cases} 0, & s < \tilde{t}, \\ u^{-1}(u(0)e^{-(r_p-r)(s-\tilde{t})}), & s \geq \tilde{t}, \end{cases}$$

where $u^{-1}(\cdot)$ is the inverse function of u , and $\tilde{t} \geq 0$ is chosen such that the agent's ex ante expected payoff equals zero.

For a highly impatient principal, MDGD is optimal among all dynamic information policies.

Proposition 4. *If $r_p > r$ and $\text{IGD}(0)$ is individually rational, then MDGD is optimal among all dynamic information policies.*

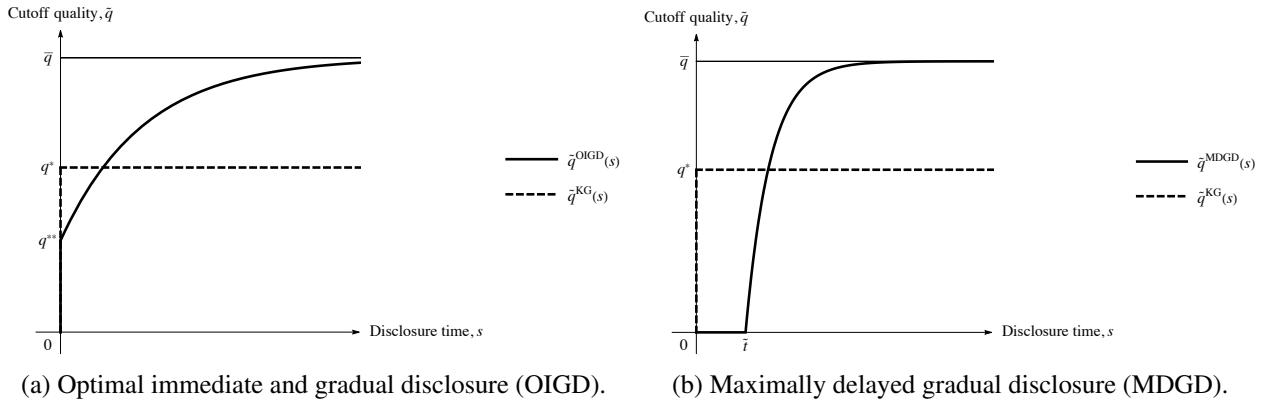


Figure 5: Relationship among KG, OIGD and MDGD.

The relationship between KG and MDGD is shown in Figure 5b, corresponding to the case where r_p is much larger than r and MDGD is optimal. As r_p increases beyond the range where OIGD is optimal, the principal shifts from lowering the initial cutoff quality to introducing a delay before disclosure begins. The notable difference from the binary case (Section 3) is that even when delay is optimal, the policy involves gradual rather than one-shot disclosure. The principal customizes disclosure times for different quality realizations, extracting more value than any one-shot delayed policy.

5 Binary-Nonstationary Case

We now move to nonstationary environments by allowing the task completion rate to vary with task quality, while maintaining the binary quality assumption from Section 3. In this setting, the mere passage of time conveys information about the task's quality.

The analysis divides into two cases, depending on whether the completion rate of the high-quality task is larger or smaller than that of the low-quality task. We focus on the *pessimistic case* ($\lambda_H > \lambda_L$), where the high-quality task has a shorter expected duration than the low-quality task. In this case, “no news is bad news” – as time passes without completion, the agent becomes increasingly pessimistic about the task’s quality. The analysis reveals a new channel that makes dynamic disclosures valuable: the agent becomes *pessimistic* over time.¹⁷

In the pessimistic case, the principal benefits from dynamic disclosure *regardless of the relationship between r_p and r* . However, Proposition 5 shows that the structure of the optimal information policy still depends on the patience levels. While MDD remains optimal for an impatient principal ($r_p > r$), a patient principal ($r_p \leq r$) optimally employs *Poisson disclosure*, where full disclosure arrives at the calibrated Poisson rate ($\lambda_H - \lambda_L$) when the task has low quality. This calibrated rate ensures that the absence of disclosure neither raises nor lowers the agent’s beliefs – effectively restoring stationarity.

The section proceeds as follows. Subsection 5.1 introduces the critical belief threshold $\hat{\mu}$ by analyzing the agent’s behavior in the absence of information provision; this threshold shapes the structure of optimal policy throughout the section. Subsection 5.2 establishes the Free-Riding Lemma (Lemma 4), showing that for high priors ($\mu > \hat{\mu}$), the principal optimally remains silent until the belief decays to $\hat{\mu}$. Subsection 5.3 introduces Poisson disclosure, which stabilizes the agent’s belief by offsetting the natural pessimistic drift. Subsection 5.4 states the main optimality result: the form of optimal policy still depends crucially on relative patience.

5.1 Agent Behavior and the Critical Belief Threshold

In the pessimistic case, when a task remains unfinished for an extended period, the agent infers that it is likely of low quality. This creates a fundamental challenge: even an agent who begins with high beliefs about the task may eventually lose faith and quit before completion, creating a dynamic problem that the principal must solve.

We first consider the decision problem of an *uninformed* agent – one who receives no disclosure from the principal. This benchmark characterizes behavior in the absence of any information policy and will be essential for understanding the value of disclosure. Bayes’ rule implies that if the task is not completed at

¹⁷By contrast, in the *optimistic case* ($\lambda_H < \lambda_L$) where “no news is good news”, all the results from Section 3 carry over directly: the principal benefits from dynamic disclosure if and only if she is less patient than the agent ($r_p > r$), and MDD is optimal for an impatient principal. In fact, if the agent starts working, he will never find it optimal to quit early, so the optimistic case reduces to an essentially stationary problem.

time $t \geq 0$, the agent's belief μ_t is given by

$$\begin{aligned}\mu_t \equiv \mathbb{P}(q = H \mid x > t) &= \frac{\mathbb{P}(x > t \mid q = H)\mathbb{P}(q = H)}{\mathbb{P}(x > t \mid q = H)\mathbb{P}(q = H) + \mathbb{P}(x > t \mid q = L)\mathbb{P}(q = L)} \\ &= \frac{\mu}{\mu + e^{(\lambda_H - \lambda_L)t}(1 - \mu)},\end{aligned}\tag{11}$$

which decreases over time since $\lambda_H > \lambda_L$. Let $\hat{\mu}$ denote threshold belief where

$$\hat{\mu}(r + \lambda_H)v(H) + (1 - \hat{\mu})(r + \lambda_L)v(L) = 0.\tag{12}$$

The threshold $\hat{\mu}$ plays a central role in all subsequent results, determining both the agent's behavior and the structure of optimal policy throughout this section. If the agent's prior is below $\hat{\mu}$, he quits immediately. Otherwise, he begins working, but with a contingency plan: if the task remains incomplete for too long, his belief will drop to $\hat{\mu}$, at which point it becomes optimal to quit.

Let $\bar{t}(\mu)$ denote the time when the agent's belief falls to $\hat{\mu}$ starting from a prior $\mu > \hat{\mu}$. Setting $\mu_t = \hat{\mu}$ in equation (11) and solving yields:

$$\frac{\mu}{\mu + e^{(\lambda_H - \lambda_L)\bar{t}(\mu)}(1 - \mu)} = \hat{\mu} \iff \bar{t}(\mu) = \frac{1}{\lambda_H - \lambda_L} \left(\log \frac{\mu}{1 - \mu} - \log \frac{\hat{\mu}}{1 - \hat{\mu}} \right).\tag{13}$$

Lemma 3 characterizes the agent's behavior in the absence of information provision.

Lemma 3. *Suppose that the principal does not provide any information to the agent.*

1. *If $\mu \leq \hat{\mu}$, then the agent quits immediately.*
2. *If $\mu > \hat{\mu}$, then the agent starts working but plans to quit at time $\bar{t}(\mu)$ if the task remains incomplete.*

Beyond characterizing when the uninformed agent quits, $\hat{\mu}$ determines the structure of optimal disclosure – in particular, whether the optimal policy involves delay. As we formalize in the Free-Riding Lemma (Lemma 4), when the agent's prior exceeds $\hat{\mu}$, the principal optimally remains silent until the belief decays to $\hat{\mu}$, free-riding on the agent's willingness to continue without information provision.

The threshold $\hat{\mu}$ also reveals the limitations of static disclosure, demonstrating the need for dynamic information provision. Under a static policy, the agent may still quit the high-quality task before completion. For example, Figure 6 illustrates the optimal static information policy (which we still call KG). The solid curve represents the principal's value without persuasion. For $\mu \leq \hat{\mu}$, the agent quits immediately, and the principal's value is zero. As μ rises from $\hat{\mu}$ to one, the agent works increasingly longer, so the principal's

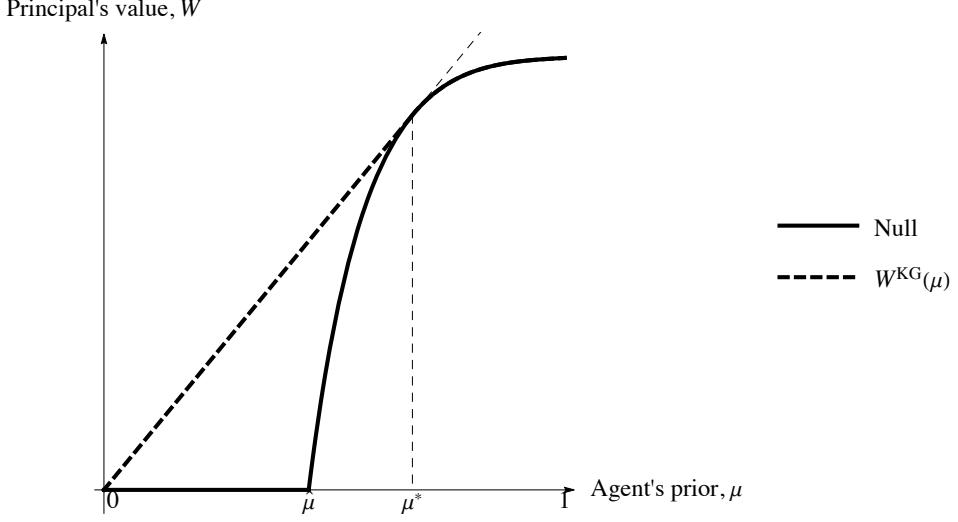


Figure 6: Optimal static information policy KG corresponds to a concavification.

value also rises. The dashed line segment, formed by the concavification of the solid curve, represents the principal's gain from persuasion. In particular, KG raises the agent's belief to some $\mu^* > \hat{\mu}$ to induce working, but the agent still plans to quit at time $\bar{t}(\mu^*)$ if the task has not been completed by then. Under any static information policy, the only way to guarantee completion of the high-quality task is initial full disclosure (to push the belief to one), which is too costly.

Dynamic disclosure overcomes this limitation. Continuing the above example, suppose that in addition to the initial disclosure that motivates the agent to start working, the principal also commits to fully disclose the task quality at time $\bar{t}(\mu^*)$ (the agent's planned quitting time) if the task is still incomplete. If the task has high quality, the agent will learn this at $\bar{t}(\mu^*)$ and thus continue working until eventual completion. If the task has low quality, the principal's disclosure at $\bar{t}(\mu^*)$ simply causes the agent to quit (which he was going to do anyway). Crucially, the agent knows this disclosure plan from the start, so he is willing to work until time $\bar{t}(\mu^*)$ since the dynamic plan only strengthens his incentives to continue compared to the static policy. For the same reason, following this information policy yields a strictly positive ex ante payoff for the agent, giving the principal room for improvement.

5.2 The Free-Riding Lemma

We now analyze the optimal structure of dynamic information disclosure. The critical step in this nonstationary setting is identifying the binding incentive constraints within the continuum of constraints the agent faces over time. A key insight emerges: if the agent has a high prior ($\mu > \hat{\mu}$), the incentive constraint at time $\bar{t}(\mu)$ implies all earlier constraints, regardless of the information policy. Therefore, it is optimal for the

principal to provide no information and “free-ride” until the belief falls to $\hat{\mu}$. This observation is formalized in the Lemma 4 below.

Lemma 4 (Free-Riding). *If $\mu > \hat{\mu}$, any optimal information policy provides no information at all $t < \bar{t}(\mu)$. That is, the agent never quits and is never told to quit before the belief naturally decays to $\hat{\mu}$.*

In light of this Free-Riding Lemma, once we obtain the optimal information policy for prior $\mu = \hat{\mu}$, we can solve the problem for any $\mu > \hat{\mu}$ by translating the time axis: relabel time $\bar{t}(\mu)$ as time zero and adapt the optimal policy accordingly. We therefore focus on characterizing the optimal policy at or below the critical prior $\hat{\mu}$.

5.3 Poisson Disclosure

To sustain the agent’s effort at the critical belief $\hat{\mu}$, the principal must prevent the belief from drifting downward. Consider the following policy: instead of setting a deterministic disclosure time, the principal discloses the task quality at Poisson rate $(\lambda_H - \lambda_L)$ when the task has low quality. We first show that this *Poisson disclosure* stabilizes the agent’s belief over time.

Suppose that the agent neither receives disclosure nor completes the task during a small time interval Δt . This event occurs with probability $(1 - \lambda_H \Delta t)$ if the task quality is high, and with probability

$$(1 - \lambda_L \Delta t) (1 - (\lambda_H - \lambda_L) \Delta t) = 1 - \lambda_H \Delta t$$

if the task quality is low. Since the event occurs is equally likely under both task qualities, Bayes’ rule implies that the agent does not update his belief. In this way, “no news is no news”, and the stationarity of the problem is restored. Thus, Poisson disclosure plays the same role as keeping silence in the stationary case analyzed in Section 3.

If $\mu < \hat{\mu}$, Poisson disclosure at rate $(\lambda_H - \lambda_L)$ alone cannot induce the agent to work. The principal addresses this by combining Poisson disclosure with an initial announcement that raises the agent’s belief to exactly $\hat{\mu}$ – the threshold at which he is just willing to start. This initial announcement follows the same logic as KG: it pools high-quality with low-quality to boost the agent’s belief. Once the agent begins working, his belief remains constant at $\hat{\mu}$ as long as no disclosure arrives, so he continues. He quits only when informed that the task has low quality. We refer to this information policy as *immediate and Poisson disclosure*, denoted by IPD.

Definition 11 (Immediate and Poisson disclosure). *For $\mu \leq \hat{\mu}$, IPD denotes the information policy of immediate and Poisson disclosure where*

1. At time zero, the agent is told to continue ($\sigma_0 = 0$) with probability one when $q = H$, and also with positive probability when $q = L$, such that $\mathbb{P}(q = H | \sigma_0 = 0) = \hat{\mu}$.
2. For $s > t$, conditional on $\sigma_t = 0$, the recommendation switches to quit ($\sigma_s = 1$) according to a Poisson arrival with rate $(\lambda_H - \lambda_L)$ when $q = L$. Formally, $\mathbb{P}(\sigma_s = 1 | \sigma_t = 0) = (1 - e^{-(\lambda_H - \lambda_L)(s-t)}) \cdot \mathbf{1}(q = L)$.

Under IPD, once the agent starts working, the agent's quitting time follows an exponential distribution with rate $(\lambda_H - \lambda_L)$ when the task has low quality, whereas he never quits when the task has high quality. Part (1a) of Proposition 5 below shows that when the principal is more patient than the agent ($r_p \leq r$), IPD is optimal among all dynamic information policies for low priors ($\mu \leq \hat{\mu}$).

For high priors ($\mu > \hat{\mu}$), the optimal policy for a patient principal is constructed by combining this result with the Free-Riding Lemma (Lemma 4): the principal provides no information until time $\bar{t}(\mu)$ and switches to IPD thereafter. This constitutes *delayed Poisson disclosure*, denoted by DPD, where Poisson disclosure begins only after $\bar{t}(\mu)$.

Definition 12 (Delayed Poisson disclosure). *For $\mu > \hat{\mu}$, DPD denotes the information policy of delayed Poisson disclosure. The principal provides no information until $\bar{t}(\mu)$, and then discloses at Poisson rate $(\lambda_H - \lambda_L)$ if $q = L$. Formally, $\sigma_s = 0$ for all $s \leq \bar{t}(\mu)$, and for any $s > t \geq \bar{t}(\mu)$, $\mathbb{P}(\sigma_s = 1 | \sigma_t = 0) = (1 - e^{-(\lambda_H - \lambda_L)(s-t)}) \cdot \mathbf{1}(q = L)$.*

5.4 Optimal Information Policy

We now state the main result characterizing the optimal policy, Proposition 5. While Poisson disclosure achieves belief stationarity, the principal is not restricted to stochastic disclosure; deterministic policies (such as MDD from Section 3) remain available. Part 2 of Proposition 5 establishes that for an impatient principal ($r_p > r$), MDD is optimal, paralleling Proposition 1. This highlights that while dynamic disclosure is always valuable in a pessimistic environment, the *form* of optimal disclosure still depends crucially on relative patience.

For any $\mu \leq \hat{\mu}$, the maximum delay time $\bar{t}(\mu)$ remains defined by equation (5). For $\mu > \hat{\mu}$, the Free-Riding Lemma (Lemma 4) implies that persuasion only takes effect from time $\bar{t}(\mu)$ onward. Therefore, the maximum delay time for a high prior is the free-riding duration plus the standard maximum delay at $\hat{\mu}$:

$$\tilde{t}(\mu) = \bar{t}(\mu) + \tilde{t}(\hat{\mu}) \quad \text{for } \mu > \hat{\mu}. \tag{14}$$

With $\tilde{t}(\mu)$ now defined for all $\mu \in (0, 1)$, we extend the Definition 4 of MDD to this nonstationary setting:

Definition 4' (Maximally delayed disclosure). MDD denotes the information policy of maximally delayed disclosure, where for any $\mu \in (0, 1)$, the principal provides no information prior to the time $\tilde{t}(\mu)$, and fully discloses the task quality at $\tilde{t}(\mu)$. Formally, $\sigma_s = 0$ for all $s < \tilde{t}(\mu)$, and for all $s \geq \tilde{t}(\mu)$, $\sigma_s = \mathbf{1}$ ($q = L$).

Proposition 5 establishes the optimality of these information policies.¹⁸ When $r_p \leq r$, either IPD or DPD is optimal depending on the prior. When $r_p > r$, MDD is optimal.

Proposition 5. In the pessimistic case with $\lambda_H > \lambda_L$,

1. If $r_p \leq r$, then the optimal policy involves Poisson disclosure.

(a) If $\mu \leq \hat{\mu}$, then IPD is optimal among all dynamic information policies.

(b) If $\mu > \hat{\mu}$, then DPD is optimal among all dynamic information policies.

2. If $r_p > r$, then MDD is optimal among all dynamic information policies.

The intuition mirrors the stationary case (Subsection 3.3): the optimal policy is determined by the parties' risk attitudes towards time lotteries. The key difference here is the natural pessimistic drift in the agent's belief absent disclosure. Any information policy that ensures completion of the high-quality task must reveal low quality at a rate of at least $(\lambda_H - \lambda_L)$ to offset this downward drift (i.e., to make “no news no news”). The principal can choose to disclose at a faster rate, with full disclosure at a deterministic time representing the limiting case of an infinite revelation rate.

The structure of binding constraints differs between the two cases. Under Poisson disclosure (IPD or DPD), multiple dynamic obedience constraints bind simultaneously: throughout the Poisson phase, the agent's belief remains constant at $\hat{\mu}$, so his continuation value equals zero at every instant. By contrast, under MDD, only a single constraint binds: the individual rationality constraint when $\mu \leq \hat{\mu}$, or the time- $\tilde{t}(\mu)$ constraint when $\mu > \hat{\mu}$ (as implied by the Free-Riding Lemma). In either case, the remaining constraints hold with slack. This difference shapes the proof strategies: establishing optimality of Poisson disclosure requires working with a continuum of incentive constraints and showing that all of them bind at optimality, whereas for MDD, it suffices to identify a single binding constraint and verify that all others are slack.

Figure 7a depicts the relationship among IPD/DPD, MDD and KG when $r_p < r$. Both the Poisson

¹⁸We thank an anonymous referee for pointing out the connection between Proposition 5 (part 1) and Proposition 8 in Section VI of Ely and Szydłowski (2020), as both results feature exponentially distributed time lotteries with calibrated parameter. However, neither result directly nests the other. Section VI of Ely and Szydłowski (2020) considers the problem where the principal can both *design* the distribution of task durations and *observe* its realization before choosing an optimal information policy. In our model, the principal can neither choose nor ex ante observe the task duration. Nonetheless, in the special case of $\lambda_L = 0$ and $\mu = \hat{\mu}$, one can deduce from Lemma 3 and Proposition 8 in Ely and Szydłowski (2020) that if $r_p \leq r$, then the time lottery induced by Poisson disclosure at rate λ_H (which equals $\lambda_H - \lambda_L$ when $\lambda_L = 0$) is optimal in our model.

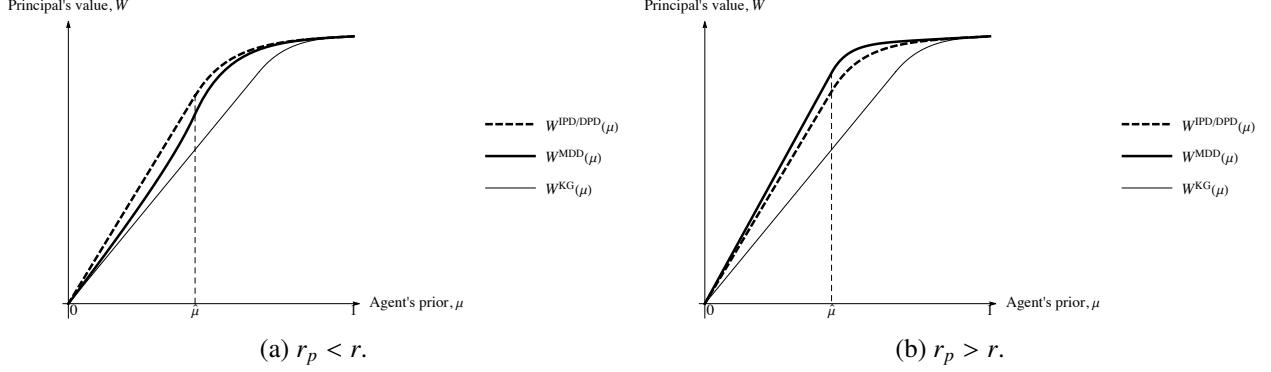


Figure 7: The principal’s payoff comparison among IPD/DPD, MDD and KG.

policy and MDD strictly outperform KG for all $\mu \in (0, 1)$. When $r_p < r$, the agent’s risk attitude dominates, so the optimal lottery has the highest risk. Among policies ensuring incentive compatibility, the Poisson policy induces the highest-risk time lottery in the sense of second-order stochastic dominance. Conversely, when $r_p > r$, the principal’s risk attitude dominates, making the optimal lottery degenerate at a fixed time $\tilde{t}(\mu)$, induced by MDD. As Figure 7b shows, both dynamic policies still improve upon KG, but Poisson policy becomes suboptimal here due to the excessive risk it imposes.

6 Conclusion

In this paper, we study the optimal mechanism to motivate effort in a dynamic moral hazard model without transfers. We develop a unified framework and characterize the optimal policy in closed form, identifying two precise conditions under which dynamic disclosure is strictly valuable to the principal: one is that the principal is impatient compared with the agent, and the other is that the agent would become pessimistic over time absent any disclosure.

Our results reveal distinct optimal policy structures across different environments. In binary-stationary settings, an impatient principal optimally employs maximally delayed disclosure – withholding all information until a specific threshold before complete revelation – while a patient principal prefers static disclosure. With general quality distributions, the optimal policy for an impatient principal becomes a cutoff cascade, where the principal gradually reveals whether the reward exceeds a time-varying threshold. In pessimistic environments, maximally delayed disclosure remains optimal for an impatient principal, while a patient principal chooses Poisson disclosure – releasing information at a calibrated rate that exactly offsets the downward drift in the agent’s belief. Crucially, our unified framework synthesizes these policy forms by identifying which constraints bind: static individual rationality in stationary environments versus dynamic

obedience in pessimistic ones.

Our analysis also emphasizes how players' risk attitudes toward time lotteries shape optimal information policies. Under exponential discounting, these risk attitudes reduce to simple comparisons of discount rates. With other time preferences, however, this relationship becomes considerably richer. This connection between dynamic information design and time preferences suggests a promising direction for future research, particularly in environments where standard exponential discounting may not hold.

Beyond time preferences, Table 1 points to another natural direction for future research: the general nonstationary case, where the quality is continuously distributed and the task completion rate varies with quality. This setting combines the challenges of continuous quality (requiring gradual disclosure) with time-varying beliefs (requiring dynamic obedience constraints). Characterizing closed-form solutions in such environments remains open.

References

Au, Pak Hung, (2015). "Dynamic Information Disclosure." *The RAND Journal of Economics* 46 (4):791–823.

Aumann, Robert J., Michael B. Maschler, and Richard E. Stearns, (1995). *Repeated Games with Incomplete Information*. MIT Press.

Ball, Ian, (2023). "Dynamic Information Provision: Rewarding the Past and Guiding the Future." *Econometrica* 91 (4):1363–1391.

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier, (2021). "Dynamic Persuasion with Outside Information." *American Economic Journal: Microeconomics* 13 (1):179–94.

Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff, (2023). "Keeping the Listener Engaged: A Dynamic Model of Bayesian Persuasion." *Journal of Political Economy* 131 (7):1797–1844.

Ely, Jeffrey C., (2017). "Beeps." *American Economic Review* 107 (1):31–53.

Ely, Jeffrey C. and Martin Szydłowski, (2020). "Moving the Goalposts." *Journal of Political Economy* 128 (2):468–506.

Escudé, Matteo and Ludvig Sinander, (2023). "Slow Persuasion." *Theoretical Economics* 18 (1):129–162.

Guo, Yingni and Eran Shmaya, (2018). "The Value of Multistage Persuasion." Working Paper.

Kamenica, Emir and Matthew Gentzkow, (2011). “Bayesian Persuasion.” *American Economic Review* 101 (6):2590–2615.

Kaya, Ayça, (2023). “Paying with information.” *Theoretical Economics* 18 (2):669–706.

Koh, Andrew, Sivakorn Sanguanmoo, and Weijie Zhong, (2025). “Persuasion and Optimal Stopping.” Working paper. Available at arXiv: <https://arxiv.org/abs/2406.12278>.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov, (2020). “Persuading the Principal to Wait.” *Journal of Political Economy* 128 (7):2542–2578.

Renault, Jérôme, Eilon Solan, and Nicolas Vieille, (2017). “Optimal Dynamic Information Provision.” *Games and Economic Behavior* 104:329–349.

Smolin, Alex, (2021). “Dynamic Evaluation Design.” *American Economic Journal: Microeconomics* 13 (4):300–331.

A Proofs of Results in the Main Text

For any information policy σ , let τ denote the induced random quitting time, and let $\tau(q)$ denote the conditional distribution of τ given task quality q .

A.1 Proofs for Section 3

Proof of Lemma 1. We first express the principal's payoff from KG as a function of the agent's prior μ . For $\mu < \bar{\mu}$, the KG policy induces a posterior of exactly $\bar{\mu}$ with maximum probability, and a posterior of 0 otherwise. By the standard concavification result, the probability of inducing posterior $\bar{\mu}$ equals $\mu/\bar{\mu}$. Thus,

$$W^{\text{KG}}(\mu) = (\mu/\bar{\mu}) \cdot (\bar{\mu}w(H) + (1 - \bar{\mu})w(L)) = \mu w(H) + \mu \cdot \frac{1 - \bar{\mu}}{\bar{\mu}} w(L) \quad \text{if } \mu < \bar{\mu}. \quad (\text{A.1})$$

For $\mu \geq \bar{\mu}$, the agent always starts out completing the task, yielding

$$W^{\text{KG}}(\mu) = \mu w(H) + (1 - \mu) w(L) \quad \text{if } \mu \geq \bar{\mu}.$$

Now we consider the principal's payoff from $\text{DD}(t)$. If the agent's prior is too low (below $\tilde{\mu}(t)$), he quits immediately and the principal's payoff is zero, strictly worse than KG. If the prior exceeds $\tilde{\mu}(t)$, $\text{DD}(t)$ induces $\tau(H) = \infty$ and $\tau(L) = t$. Thus, the principal's payoff in this range is

$$W^{\text{DD}(t)}(\mu) = \mu w(H) + (1 - \mu) w(t, L) \quad \text{if } \mu \geq \tilde{\mu}(t), \quad (\text{A.2})$$

where $\tilde{\mu}(t) < \bar{\mu}$. Since $w(t, L) < w(L)$ for all finite t , we have $W^{\text{DD}(t)}(\mu) < W^{\text{KG}}(\mu)$ for $\mu \geq \bar{\mu}$. The only range where $\text{DD}(t)$ can outperform KG is $[\tilde{\mu}(t), \bar{\mu}]$. Comparing equation (A.2) with equation (A.1), for any $\mu \in [\tilde{\mu}(t), \bar{\mu}]$, we have

$$\begin{aligned} W^{\text{DD}(t)}(\mu) > W^{\text{KG}}(\mu) &\iff (1 - \mu) w(t, L) > \mu \cdot \frac{1 - \bar{\mu}}{\bar{\mu}} w(L) \\ &\iff (1 - \mu) (1 - e^{-(r_p + \lambda)t}) > \mu \cdot \frac{1 - \bar{\mu}}{\bar{\mu}}, \end{aligned}$$

where the last equivalence follows from the principal's payoff function, equation (2). The last inequality imposes an upper bound on μ , so we need the strict inequality to hold at $\mu = \tilde{\mu}(t)$, i.e.,

$$(1 - \tilde{\mu}(t)) (1 - e^{-(r_p + \lambda)t}) > \tilde{\mu}(t) \cdot \frac{1 - \bar{\mu}}{\bar{\mu}}. \quad (\text{A.3})$$

Comparing equations (3) and (4), which define the threshold priors $\bar{\mu}$ and $\tilde{\mu}(t)$, we get

$$(1 - \tilde{\mu}(t)) (1 - e^{-(r_p + \lambda)t}) = \tilde{\mu}(t) \frac{v(H)}{[-v(L)]} = \tilde{\mu}(t) \cdot \frac{1 - \bar{\mu}}{\bar{\mu}}.$$

Therefore, strict inequality (A.3) holds if and only if

$$1 - e^{-(r_p + \lambda)t} > 1 - e^{-(r_p + \lambda)t} \iff r_p > r.$$

This completes the proof. \square

Proof of Proposition 1. For $\mu \geq \bar{\mu}$, both KG and MDD provide no information, and the agent completes the task with probability one without any persuasion. This is clearly optimal. Henceforth, we assume $\mu < \bar{\mu}$.

We introduce a change of variables to linearize the principal's objective function.¹⁹ Since the principal's value from quitting time τ is proportional to $1 - e^{-(r_p + \lambda)\tau}$, we define the transformed variable ζ as $\zeta \equiv 1 - e^{-(r_p + \lambda)\tau} \in [0, 1]$. In this transformed space, $\zeta = 0$ corresponds to immediate quitting ($\tau = 0$), while $\zeta = 1$ corresponds to never quitting ($\tau = \infty$). The principal's value then becomes $\mathbb{E}_\sigma[w(\tau, q)] = \mathbb{E}_\sigma[w(q) \cdot \zeta]$.

Next we express the agent's value in terms of ζ . The agent's value from quitting time τ is proportional to $1 - e^{-(r + \lambda)\tau} = 1 - (1 - \zeta)^{\frac{r+\lambda}{r_p+\lambda}}$. Define $\rho \equiv \frac{r_p + \lambda}{r + \lambda}$ as the ratio of the principal's effective discount rate to the agent's, and let $\phi(z) \equiv 1 - (1 - z)^{1/\rho}$. The agent's value can then be written as $\mathbb{E}_\sigma[v(\tau, q)] = \mathbb{E}_\sigma[v(q) \cdot \phi(\zeta)]$. Note that $\phi'(z) = \frac{1}{\rho}(1 - z)^{1/\rho - 1} > 0$, so $\phi(z)$ is strictly increasing, and $\phi''(z) = \frac{\rho-1}{\rho^2}(1 - z)^{1/\rho - 2}$, whose sign depends entirely on $\rho - 1$, i.e., whether the principal is more patient than the agent ($r_p > r$).

Because the agent never needs encouragement to continue when the task quality is high, the non-trivial design problem concerns only the low-quality task. We therefore set $\zeta = 1$ when the task quality is high and focus entirely on determining the optimal distribution of ζ for the low-quality task. Let G denote this distribution of ζ conditional on low quality. The principal's expected payoff is then

$$W(G) \equiv \mathbb{E}_\sigma[w(q) \cdot \zeta] = \mu w(H) + (1 - \mu)w(L)\mathbb{E}_G[\zeta].$$

Since the first term is constant and $w(L) > 0$, the principal's problem reduces to maximizing $\mathbb{E}_G[\zeta]$.

The agent's expected payoff is

$$V(G) \equiv \mathbb{E}_\sigma[v(q) \cdot \phi(\zeta)] = \mu v(H) + (1 - \mu)v(L)\mathbb{E}_G[\phi(\zeta)].$$

The principal's optimization requires individual rationality $V(G) \geq 0$. Since $v(L) < 0$, we can rewrite this as

$$\mathbb{E}_G[\phi(\zeta)] \leq \frac{\mu v(H)}{(1 - \mu)[-v(L)]} \equiv C_1.$$

To find the optimal information policy, we adopt the following approach: we first solve a relaxed problem that temporarily omits all interim incentive compatibility constraints,

$$\begin{aligned} & \max_{G \in \Delta([0,1])} \quad \mathbb{E}_G[\zeta], \\ & \text{s.t.} \quad \mathbb{E}_G[\phi(\zeta)] \leq C_1, \end{aligned} \tag{A.4}$$

where the principal chooses the optimal distribution G subject only to individual rationality. We then verify that the solution to (A.4) can be implemented by an information policy, thereby establishing its optimality.

If $r_p < r$, then $\rho < 1$ and $\phi(z)$ is strictly concave. By Jensen's inequality, any interior point $\zeta \in (0, 1)$ yields a higher $\phi(\zeta)$ than the equivalent mean-preserving combination of 0 and 1, so replacing interior points with such randomizations relaxes the constraint in (A.4). Therefore, the optimal distribution G must be bang-bang. This is precisely the structure of KG: $\zeta \in \{0, 1\}$ with probabilities that exactly bind individual rationality, implying that KG is optimal among all dynamic information policies.

If $r_p > r$, then $\rho > 1$ and $\phi(z)$ is strictly convex. By Jensen's inequality, $\mathbb{E}_G[\phi(\zeta)] \geq \phi(\mathbb{E}_G[\zeta])$, and equality holds if and only if G is degenerate at a specific point. It follows that $\mathbb{E}_G[\zeta] \leq \phi^{-1}(\mathbb{E}_G[\phi(\zeta)]) \leq \phi^{-1}(C_1)$, with equality if and only if G is degenerate at $z = \phi^{-1}(C_1)$. This is precisely the structure of MDD: a deterministic value $\zeta = z$ that exactly binds individual rationality, implying that MDD is optimal among all dynamic information policies.

If $r_p = r$, then $\rho = 1$ and $\phi(z) = z$, making the objective and constraint in (A.4) identical; thus any policy sufficient to motivate the agent while exactly binding individual rationality is optimal, including both KG and MDD. \square

¹⁹We thank an anonymous referee for suggesting this change of variables that streamlines the proof.

A.2 Proofs for Section 4

As discussed in the main text, because we assume that the task quality follows a full-support and atomless distribution, it is without loss to assume that each $\tau(q)$ is pure; that is, $\tau(q)$ assigns probability one to a deterministic quitting time.

Adopting the same change of variables as in the proof of Proposition 1, we define $\zeta(q) \equiv 1 - e^{-(r_p + \lambda)\tau(q)} \in [0, 1]$, where $\zeta = 0$ corresponds to immediate quitting ($\tau = 0$) and $\zeta = 1$ corresponds to never quitting ($\tau = \infty$). With this transformation, the principal's and agent's values become

$$W(\zeta) = \mathbb{E}_\sigma[w(\tau, q)] = \int_0^\infty w(q)\zeta(q) dF(q),$$

$$V(\zeta) = \mathbb{E}_\sigma[v(\tau, q)] = \int_0^\infty v(q)\phi(\zeta(q)) dF(q),$$

where $\phi(z) \equiv 1 - (1 - z)^{1/\rho}$ and $\rho \equiv \frac{r_p + \lambda}{r + \lambda}$ as before.

To find the optimal information policy, we adopt the following approach: we first solve a relaxed problem that temporarily omits all interim incentive compatibility constraints,

$$\begin{aligned} \max_{\zeta \in (0, \infty) \rightarrow [0, 1]} \quad & W(\zeta), \\ \text{s.t.} \quad & V(\zeta) \geq 0, \end{aligned} \tag{A.5}$$

where the principal chooses the optimal ζ subject only to individual rationality. We then verify that the solution to (A.5) can be implemented by an information policy, thereby establishing its optimality.

We set up the Lagrangian functional \mathcal{L} with a multiplier $\eta \geq 0$ for the individual rationality constraint:

$$\mathcal{L}(\zeta, \eta) \equiv \int_0^\infty (w(q)\zeta(q) + \eta v(q)\phi(\zeta(q))) dF(q).$$

By separability, we can maximize the Lagrangian pointwise for each q :

$$\max_{\zeta \in [0, 1]} H(z, q, \eta) \equiv w(q)z + \eta v(q)\phi(z). \tag{A.6}$$

For an individually rational task $q \geq \bar{q}$ where $v(q) \geq 0$, both terms are increasing in z , so $z = 1$ is optimal. Henceforth we focus on tasks that are not individually rational, $q \in (0, \bar{q})$, where we can write $H(z, q, \eta) = w(q)z - \eta|v(q)|\phi(z)$.

Proof of Proposition 2. Suppose $r_p \leq r$, then $\rho \leq 1$ and $\phi(z)$ is concave, implying that $H(z, q, \eta) = w(q)z - \eta|v(q)|\phi(z)$ is convex in z for all $q \in (0, \bar{q})$. Since a convex function on $[0, 1]$ is maximized at the boundaries, the optimal solution is bang-bang: either $z = 0$ or $z = 1$. As a result,

$$\max_{z \in [0, 1]} H(z, q, \eta) = \max \{w(q) - \eta|v(q)|, 0\}.$$

Since $b(q)$ is nondecreasing and $|v(q)|$ is strictly decreasing, $w(q) = \frac{\lambda b(q)}{r_p + \lambda}$ is nondecreasing and the difference $w(q) - \eta|v(q)|$ is strictly increasing. Define $y \equiv \inf \{q \in (0, \bar{q}) : w(q) \geq \eta|v(q)|\}$. It follows that $\max \{w(q) - \eta|v(q)|, 0\} = (w(q) - \eta|v(q)|) \cdot \mathbf{1}(q \geq y)$. Thus, the optimal ζ to (A.5) is the step function $\zeta(q) = \mathbf{1}(q \geq y)$. This is precisely the structure of KG: $\zeta(q) \in \{0, 1\}$ with a cutoff y that exactly binds individual rationality, implying that KG is optimal among all dynamic information policies. \square

Proof of Lemma 2. Denote by τ the agent's quitting schedule under a cutoff cascade policy σ with cutoff

quality function $\tilde{q}(s)$. At any instant $s \geq 0$, if the task is incomplete and the agent has not been informed to quit, he infers that $q \geq \tilde{q}(s)$. Thus, $\tau(q)$ is the inverse of $\tilde{q}(s)$:

$$q = \tilde{q}(s) \iff s = \tau(q).$$

If the individual rationality constraint is violated, i.e., $\mathbb{E}_\sigma[v(\tau, q)] < 0$, the agent ignores all information and quits immediately. Therefore, individual rationality is necessary for the agent to follow σ .

Now suppose $\mathbb{E}_\sigma[v(\tau, q)] \geq 0$. We show that the agent's best response is to follow σ at any time $s \geq 0$ while the task remains incomplete. If the agent is informed that $q < \tilde{q}(s)$, his expected continuation value is bounded above by $v(\tilde{q}(s)) \leq v(\bar{q}) = 0$, so quitting is optimal. If the agent is informed that $q \geq \tilde{q}(s)$, his expected continuation value promised by σ is $\mathbb{E}_\sigma[v(\tau - s, q) \mid q \geq \tilde{q}(s)]$. We need to show that this value is nonnegative, thereby establishing that it is optimal for the agent to continue when he is informed to do so.

Define

$$h_1(s) \equiv \mathbb{E}_\sigma[v(\tau - s, q) \mid q \geq \tilde{q}(s)] \cdot \mathbb{P}_\sigma(q \geq \tilde{q}(s)) = \int_{\tilde{q}(s)}^{\infty} v(\tau(q) - s, q) dF(q).$$

Since τ is the inverse of \tilde{q} , we have $\tau(\tilde{q}(s)) = s$. Using the Leibniz integral rule and the fact that $v(0, q) = 0$, we obtain

$$h'_1(s) = \int_{\tilde{q}(s)}^{\infty} \left(-\frac{\partial v}{\partial t}(\tau(q) - s, q) \right) dF(q).$$

Since $v(t, q) = (1 - e^{-(r+\lambda)t})v(q)$, we have $\frac{\partial v}{\partial t}(t, q) = (r + \lambda)e^{-(r+\lambda)t}v(q)$. For $q \leq \bar{q}$, we have $v(q) \leq 0$, which implies that $-\frac{\partial v}{\partial t}(\tau(q) - s, q) \geq 0$. For $q > \bar{q}$, the agent is never asked to quit (as $\tilde{q}(s) \in [0, \bar{q}]$), so $\tau(q) = \infty$ and consequently $-\frac{\partial v}{\partial t}(\tau(q) - s, q) = 0$. Combining both cases, we have $-\frac{\partial v}{\partial t}(\tau(q) - s, q) \geq 0$ for all q , implying that $h'_1(s) \geq 0$. Since $h(0) = \mathbb{E}_\sigma[v(\tau, q)] \geq 0$ by assumption, we conclude that $h_1(s) \geq 0$ for all $s \geq 0$. This completes the proof that the agent will follow σ at each time $s \geq 0$. \square

Proof of Propositions 3 and 4. When the principal is less patient than the agent ($r_p > r$), we have $\rho > 1$ and $\phi(z)$ is strictly convex. This implies that $H(z, q, \eta) = w(q)z - \eta|v(q)|\phi(z)$ is strictly concave in z for all $q \in (0, \bar{q})$. Any interior optimal solution to (A.6) is determined by the first-order condition:

$$\frac{\partial H}{\partial z} = w(q) - \eta|v(q)|\phi'(z) = 0.$$

Substituting $\phi'(z) = \frac{1}{\rho}(1-z)^{1/\rho-1}$ and rearranging yields

$$(1 - \zeta(q))^{1/\rho-1} = \frac{\rho w(q)}{\eta|v(q)|} = \frac{\rho}{\eta u(q)}. \quad (\text{A.7})$$

We now examine the boundary conditions for the interior solution. Since $1 - \zeta(q) \in [0, 1]$ and $1/\rho - 1 < 0$, the left-hand side of equation (A.7) is bounded below by 1, which requires that $\rho \geq \eta u(q)$. Conversely, if $\rho < \eta u(q)$, the first-order condition would imply $\zeta(q) < 0$. The KKT conditions then imply the corner solution $\zeta(q) = 0$.

As q increases from zero to \bar{q} , the ratio $u(q)$ is strictly decreasing and ranges from the positive value $u(0)$ to zero. Consequently, the analysis distinguishes between two cases, depending on whether $\rho \geq \eta u(0)$.

Case 1: If $\rho < \eta u(0)$, then by the intermediate value theorem, there exists a unique $q^{**} \in (0, \bar{q})$ satisfying $\rho = \eta u(q^{**})$. For all $q < q^{**}$, we have $\rho < \eta u(q)$, so $\zeta(q) = 0$. For all $q \in [q^{**}, \bar{q}]$, the solution is interior and characterized by equation (A.7).

Substituting $\zeta(q) = 1 - e^{-(r_p + \lambda)\tau(q)}$ and $\rho = \frac{r_p + \lambda}{r + \lambda}$ into equation (A.7) and simplifying, we obtain

$$e^{(r_p - r)\tau(q)} = \frac{\rho}{\eta u(q)} = \frac{u(q^{**})}{u(q)} \implies \tau(q) = \frac{1}{r_p - r} \log \left(\frac{u(q^{**})}{u(q)} \right).$$

This is precisely the structure of OIGD: a cutoff cascade policy IGD (q^{**}) with initial cutoff quality q^{**} that exactly binds individual rationality. By Lemma 2, the quitting time τ^{OIGD} is indeed induced by the information policy OIGD, which establishes that OIGD is optimal among all dynamic information policies. This completes the proof of Proposition 3.

Case 2: If $\rho \geq \eta u(0)$, the monotonicity of $u(q)$ ensures that $\rho \geq \eta u(q)$ for all $q \in (0, \bar{q})$. Consequently, the solution is interior for all $q \in (0, \bar{q})$ and characterized by equation (A.7).

Substituting $\zeta(q) = 1 - e^{-(r_p + \lambda)\tau(q)}$ and $\rho = \frac{r_p + \lambda}{r + \lambda}$ into equation (A.7) and simplifying, we obtain

$$e^{(r_p - r)\tau(q)} = \frac{\rho}{\eta u(q)} = \frac{\rho}{\eta u(0)} \cdot \frac{u(0)}{u(q)} \implies \tau(q) = \frac{1}{r_p - r} \left(\log \left(\frac{\rho}{\eta u(0)} \right) + \log \left(\frac{u(0)}{u(q)} \right) \right).$$

Let $\tilde{t} \equiv \frac{1}{r_p - r} \log \left(\frac{\rho}{\eta u(0)} \right)$. This is precisely the structure of MDGD: a cutoff cascade policy DGD (\tilde{t}) with delay time \tilde{t} that exactly binds individual rationality. By Lemma 2, the quitting time τ^{MDGD} is indeed induced by the information policy MDGD, which establishes that MDGD is optimal among all dynamic information policies. This completes the proof of Proposition 4. \square

A.3 Proofs for Section 5

Proof of Lemma 3. Let τ be the agent's planned quitting time. His expected payoff $h(\tau)$ is given by $h_2(\tau) \equiv \mu v(\tau, H) + (1 - \mu)v(\tau, L)$. Substituting the functional forms (1) leads to:

$$h_2(\tau) = \mu \left(1 - e^{-(r + \lambda_H)\tau} \right) v(H) + (1 - \mu) \left(1 - e^{-(r + \lambda_L)\tau} \right) v(L).$$

To find the optimal τ , consider the first-order derivative:

$$h'_2(\tau) = \mu(r + \lambda_H)e^{-(r + \lambda_H)\tau}v(H) + (1 - \mu)(r + \lambda_L)e^{-(r + \lambda_L)\tau}v(L). \quad (\text{A.8})$$

First assume $\mu \leq \hat{\mu}$. Since $\lambda_H > \lambda_L$, we have $e^{-(r + \lambda_H)\tau} < e^{-(r + \lambda_L)\tau}$ for any $\tau > 0$. We can therefore bound the derivative $h'_2(\tau)$ as follows:

$$\begin{aligned} h'_2(\tau) &< \mu e^{-(r + \lambda_L)\tau} (r + \lambda_H)v(H) + (1 - \mu) e^{-(r + \lambda_L)\tau} (r + \lambda_L)v(L) \\ &= e^{-(r + \lambda_L)\tau} (\mu(r + \lambda_H)v(H) + (1 - \mu)(r + \lambda_L)v(L)) \leq 0, \end{aligned}$$

where the final inequality holds due to the definition of $\hat{\mu}$ in equation (12). This shows that $h_2(\tau)$ is strictly decreasing in τ , and hence the optimal quitting time is $\tau = 0$.

Now assume $\mu > \hat{\mu}$. We argue that $h_2(\tau)$ attains its maximum at $\tau = \bar{t}(\mu)$. Using the expression for $h'_2(\tau)$ in equation (A.8) and the definition of $\hat{\mu}$ in equation (12), we have:

$$\begin{aligned} h'_2(\tau) > 0 &\iff e^{(\lambda_H - \lambda_L)\tau} < \frac{\mu(r + \lambda_H)v(H)}{(1 - \mu)(r + \lambda_L)[-v(L)]} = \frac{\mu}{1 - \mu} \cdot \frac{1 - \hat{\mu}}{\hat{\mu}} \\ &\iff \tau < \frac{1}{\lambda_H - \lambda_L} \left(\log \frac{\mu}{1 - \mu} - \log \frac{\hat{\mu}}{1 - \hat{\mu}} \right) = \bar{t}(\mu), \end{aligned}$$

where the last equality follows from the definition of $\bar{t}(\mu)$ in equation (13). Therefore, $h_2(\tau)$ is strictly increasing on $(0, \bar{t}(\mu))$, attains its maximum at $\tau = \bar{t}(\mu)$, and is strictly decreasing on $(\bar{t}(\mu), \infty)$. The optimal strategy for the agent is to keep working until time $\bar{t}(\mu)$ and then quit if the task is not yet completed. \square

Proof of Lemma 4. For any information policy σ , let α denote the CDF of $\tau(H)$, the agent's quitting time conditional on the task being high quality, and let β denote the CDF of $\tau(L)$.

Assume $\mu > \hat{\mu}$. Consider an optimal information policy σ . Since the agent never needs encouragement to continue when the task quality is high, we have $\alpha(t) = 0$ for all $t \in [0, \infty)$. To complete the proof, it remains to show that $\beta(t) = 0$ for all $t \in [0, \bar{t}(\mu))$.

Suppose to the contrary that $\beta(t) > 0$ for some $t \in [0, \bar{t}(\mu))$. Consider the perturbation $\tilde{\beta}$ given by

$$\tilde{\beta}(t) = \begin{cases} 0, & t \in [0, \bar{t}(\mu)), \\ \beta(t), & t \in [\bar{t}(\mu), \infty). \end{cases}$$

Note that $\tilde{\beta}$ first-order stochastically dominates β , as it shifts the mass accumulated on $[0, \bar{t}(\mu))$ to a discrete jump at $t = \bar{t}(\mu)$. The principal's payoff can be written as:

$$W(\beta) \equiv \mathbb{E}_\sigma[w(\tau, q)] = \mu w(H) + (1 - \mu) w(L) \int_0^\infty (1 - e^{-(r_p + \lambda_L)t}) d\beta(t).$$

Since the integrand $1 - e^{-(r_p + \lambda_L)t}$ is strictly increasing in t , the principal strictly prefers $\tilde{\beta}$ to β . To complete the contradiction, it remains only to verify that $\tilde{\beta}$ satisfies the agent's incentive compatibility constraints.

For any time $t < \bar{t}(\mu)$, under $\tilde{\beta}$, the principal provides no information. The agent's belief μ_t follows the natural decay path given by equation (11), and by definition of $\bar{t}(\mu)$ in equation (13), we have $\mu_t > \hat{\mu}$, so the agent is willing to work at all $t < \bar{t}(\mu)$. At any time $t \geq \bar{t}(\mu)$, the conditional distribution of quitting times under $\tilde{\beta}$ is identical to that under β . Since β was incentive compatible, the constraints at $t \geq \bar{t}(\mu)$ remain satisfied. This contradicts the optimality of β , completing the proof. \square

Proof of Proposition 5. For $\mu > \hat{\mu}$, optimality follows directly from the Free-Riding Lemma (Lemma 4). We therefore assume $\mu \leq \hat{\mu}$ for the remainder of the proof.

Since the agent never needs encouragement to continue when the task quality is high, the non-trivial design problem concerns only the low-quality task. Adopting the same change of variables as in the proof of Proposition 1, we define $\zeta \equiv 1 - e^{-(r_p + \lambda_L)\tau(L)} \in [0, 1]$, where $\zeta = 0$ corresponds to immediate quitting ($\tau(L) = 0$) and $\zeta = 1$ corresponds to never quitting ($\tau(L) = \infty$), and let G denote this distribution of ζ . The principal's expected payoff is then

$$W(G) = \mathbb{E}_\sigma[w(\tau, q)] = \mu w(H) + (1 - \mu) w(L) \mathbb{E}_G[\zeta].$$

Since the first term is constant and $w(L) > 0$, the principal's problem reduces to maximizing $\mathbb{E}_G[\zeta]$.

The agent's expected payoff is

$$V(G) = \mathbb{E}_\sigma[v(\tau, q)] = \mu v(H) + (1 - \mu) v(L) \mathbb{E}_G[\phi(\zeta)],$$

where $\phi(z) \equiv 1 - (1 - z)^{1/\rho}$ and $\rho \equiv \frac{r_p + \lambda_L}{r + \lambda_L}$. Since $v(L) < 0$, the individual rationality constraint $V(G) \geq 0$ rearranges to $\mathbb{E}_G[\phi(\zeta)] \leq C_1$ where $C_1 \equiv \frac{\mu v(H)}{(1 - \mu)(-v(L))}$.

Next, we simplify the interim incentive compatibility constraints. At any time $s > 0$, the agent's continuation value under σ can be expanded using the law of iterated expectations, conditioning on the

realization of task quality q :

$$\begin{aligned}\mathbb{E}_{\sigma} [v(\tau - s, q) \mid x > s, \tau > s] &= \underbrace{\mathbb{P}_{\sigma}(q = H \mid x > s, \tau > s)}_{\mu_s} \cdot \mathbb{E}_{\sigma} [v(\tau - s, H) \mid x > s, \tau > s, q = H] \\ &\quad + \underbrace{\mathbb{P}_{\sigma}(q = L \mid x > s, \tau > s)}_{1 - \mu_s} \cdot \mathbb{E}_{\sigma} [v(\tau - s, L) \mid x > s, \tau > s, q = L],\end{aligned}$$

where the posterior belief is μ_s is given by

$$\mu_s \equiv \mathbb{P}_{\sigma}(q = H \mid x > s, \tau > s) = \frac{\mu}{\mu + (1 - \mu)e^{(\lambda_H - \lambda_L)s}\mathbb{P}_{\sigma}(\tau > s \mid q = L)}. \quad (\text{A.9})$$

Equation (A.9) is derived from Bayes' rule and nests equation (11) as the special case with no disclosure.

When the task quality is high, the agent never quits, so $\mathbb{E}_{\sigma} [v(\tau - s, H) \mid x > s, \tau > s, q = H] = v(H)$. When the task quality is low, the event $\tau > s$ corresponds to $\zeta > z_s$ in the transformed domain, where $z_s \equiv 1 - e^{-(r_p + \lambda_L)s}$. It follows that $\mathbb{P}_{\sigma}(\tau > s \mid q = L) = \mathbb{P}_G(\zeta > z_s)$. The agent's continuation value at time s , $v(\tau - s, L)$, is proportional to $1 - e^{-(r + \lambda_L)(\tau - s)}$, which can be expressed as

$$1 - e^{-(r + \lambda_L)(\tau - s)} = 1 - \frac{e^{-(r + \lambda_L)\tau}}{e^{-(r + \lambda_L)s}} = 1 - \frac{1 - \phi(\zeta)}{1 - \phi(z_s)} = \frac{\phi(\zeta) - \phi(z_s)}{1 - \phi(z_s)}.$$

Therefore,

$$\mathbb{E}_{\sigma} [v(\tau - s, L) \mid x > s, \tau > s, q = L] = v(L) \frac{\mathbb{E}_G [\phi(\zeta) - \phi(z_s) \mid \zeta > z_s]}{1 - \phi(z_s)}.$$

Substituting these values and the belief μ_s from equation (A.9) into the interim incentive compatibility constraint at time $s > 0$, $\mathbb{E}_{\sigma} [v(\tau - s, q) \mid x > s, \tau > s] \geq 0$, and clearing the strictly positive denominator, yields:

$$\mu v(H) + (1 - \mu)e^{(\lambda_H - \lambda_L)s}\mathbb{P}_G(\zeta > z_s) v(L) \frac{\mathbb{E}_G [\phi(\zeta) - \phi(z_s) \mid \zeta > z_s]}{1 - \phi(z_s)} \geq 0.$$

Combining the probability $\mathbb{P}_G(\zeta > z_s)$ with the conditional expectation:

$$\mu v(H) + (1 - \mu)e^{(\lambda_H - \lambda_L)s}v(L) \frac{\mathbb{E}_G [(\phi(\zeta) - \phi(z_s)) \mathbf{1}(\zeta > z_s)]}{1 - \phi(z_s)} \geq 0.$$

Since $v(L) < 0$, rearranging the inequality to isolate the expectation term yields:

$$\mathbb{E}_G [(\phi(\zeta) - \phi(z_s)) \mathbf{1}(\zeta > z_s)] \leq \frac{\mu v(H)}{(1 - \mu)[-v(L)]} \cdot e^{-(\lambda_H - \lambda_L)s} (1 - \phi(z_s)) = C_1 e^{-(\lambda_H - \lambda_L)s} (1 - \phi(z_s)).$$

Expressing the right-hand side in terms of z_s using the identities $e^{-(\lambda_H - \lambda_L)s} = (e^{-(r_p + \lambda_L)s})^{\frac{\lambda_H - \lambda_L}{r_p + \lambda_L}} = (1 - z_s)^{\frac{\lambda_H - \lambda_L}{r_p + \lambda_L}}$ and $1 - \phi(z_s) = (1 - z_s)^{1/\rho} = (1 - z_s)^{\frac{r + \lambda_L}{r_p + \lambda_L}}$:

$$\mathbb{E}_G [(\phi(\zeta) - \phi(z_s)) \mathbf{1}(\zeta > z_s)] \leq C_1 (1 - z_s)^{\frac{\lambda_H - \lambda_L}{r_p + \lambda_L}} (1 - z_s)^{\frac{r + \lambda_L}{r_p + \lambda_L}} = C_1 (1 - z_s)^{\frac{r + \lambda_H}{r_p + \lambda_L}}.$$

The transformation $z_s = 1 - e^{-(r_p + \lambda_L)s}$ maps $s \in (0, \infty)$ bijectively to $z \in (0, 1)$; thus, satisfying the constraint for all $s > 0$ is equivalent to satisfying it for all $z \in (0, 1)$. The boundary condition at $z = 0$ represents the individual rationality constraint $\mathbb{E}_G [\phi(\zeta)] \leq C_1$, while the condition at $z = 1$ is satisfied trivially since $\zeta \leq 1$.

Therefore, to find the optimal information policy, we can rewrite the principal's program as

$$\begin{aligned} \max_{G \in \Delta([0,1])} \quad & \mathbb{E}_G[\zeta], \\ \text{s.t.} \quad & \mathbb{E}_G[(\phi(\zeta) - \phi(z)) \mathbf{1}(\zeta > z)] \leq C_1(1 - z)^{\frac{r+\lambda_H}{r_p+\lambda_L}}, \quad \forall z \in [0, 1], \end{aligned} \quad (\text{A.10})$$

If $r_p < r$, then $\rho < 1$ and $\phi(z)$ is strictly concave. Let

$$h_3(z) \equiv \mathbb{E}_G[(\phi(\zeta) - \phi(z)) \mathbf{1}(\zeta > z)] = \int_z^1 (\phi(\zeta) - \phi(z)) dG(\zeta).$$

By integration by parts,

$$h_3(z) = (\phi(\zeta) - \phi(z))(G(\zeta) - 1) \Big|_{\zeta=z}^{\zeta=1} - \int_z^1 (G(\zeta) - 1)\phi'(\zeta) d\zeta = \int_z^1 (1 - G(\zeta))\phi'(\zeta) d\zeta.$$

Differentiating with respect to z using the Leibniz integral rule yields

$$h_3'(z) = -(1 - G(z))\phi'(z). \quad (\text{A.11})$$

The principal's problem (A.10) is to maximize $\mathbb{E}_G[\zeta] = \int_0^1 (1 - G(z)) dz$ subject to $h_3(z) \leq C_1(1 - z)^{\frac{r+\lambda_H}{r_p+\lambda_L}}$ for all $z \in [0, 1]$. Using equation (A.11), we can express the objective in terms of $h_3(z)$:

$$\int_0^1 (1 - G(z)) dz = \int_0^1 \frac{-h_3'(z)}{\phi'(z)} dz.$$

Apply integration by parts again, and note that $h_3(1) = 0$:

$$\mathbb{E}_G[\zeta] = \int_0^1 \frac{-h_3'(z)}{\phi'(z)} dz = \frac{-h_3(z)}{\phi'(z)} \Big|_0^1 - \int_0^1 h_3(z) \frac{d}{dz} \left(\frac{-1}{\phi'(z)} \right) dz = \frac{h_3(0)}{\phi'(0)} + \int_0^1 h_3(z) \left(\frac{-\phi''(z)}{\phi'(z)^2} \right) dz.$$

Since $\phi(z)$ is strictly concave, $-\phi''(z) > 0$, so the objective in (A.10) is strictly increasing in $h_3(z)$ pointwise. To maximize the objective, the principal must set $h_3(z)$ to its maximum feasible value, which means the constraint binds everywhere: $h_3(z) = C_1(1 - z)^{\frac{r+\lambda_H}{r_p+\lambda_L}}$ for all $z \in [0, 1]$. Since $\phi'(z) = \frac{1}{\rho}(1 - z)^{1/\rho - 1} = \frac{r+\lambda_L}{r_p+\lambda_L}(1 - z)^{\frac{r+\lambda_L}{r_p+\lambda_L} - 1}$, we can solve for $1 - G(z)$ using equation (A.11):

$$1 - G(z) = \frac{-h_3'(z)}{\phi'(z)} = \frac{C_1 \cdot \frac{r+\lambda_H}{r_p+\lambda_L} \cdot (1 - z)^{\frac{r+\lambda_H}{r_p+\lambda_L} - 1}}{\frac{r+\lambda_L}{r_p+\lambda_L} \cdot (1 - z)^{\frac{r+\lambda_L}{r_p+\lambda_L} - 1}} = C_1 \cdot \frac{r + \lambda_H}{r + \lambda_L} \cdot (1 - z)^{\frac{\lambda_H - \lambda_L}{r_p + \lambda_L}}.$$

From the definition of $\hat{\mu}$ in equation (12), the condition $\mu \leq \hat{\mu}$ is equivalent to $C_1 = \frac{\mu v(H)}{(1-\mu)[-v(L)]} \leq \frac{r+\lambda_L}{r+\lambda_H}$. This ensures that $1 - G(0) \leq 1$, confirming that the solution constitutes a valid probability distribution. Substituting $\zeta = 1 - e^{-(r_p+\lambda_L)\tau}$, which implies that $(1 - \zeta)^{\frac{\lambda_H - \lambda_L}{r_p + \lambda_L}} = e^{-(\lambda_H - \lambda_L)\tau}$, yields

$$\mathbb{P}_{\sigma}(\tau > s \mid q = L) = C_1 \cdot \frac{r + \lambda_H}{r + \lambda_L} \cdot e^{-(\lambda_H - \lambda_L)s}$$

This is precisely the structure of IPD: immediate disclosure with probability $1 - C_1 \cdot \frac{r+\lambda_H}{r+\lambda_L}$ followed by Poisson disclosure at rate $\lambda_H - \lambda_L$, implying that IPD is optimal among all dynamic information policies.

If $r_p > r$, then $\rho > 1$ and $\phi(z)$ is strictly convex. By Jensen's inequality, $\mathbb{E}_G[\phi(z)] \geq \phi(\mathbb{E}_G[z])$, and equality holds if and only if G is degenerate at a specific point. It follows that $\mathbb{E}_G[z] \leq \phi^{-1}(\mathbb{E}_G[\phi(z)]) \leq \phi^{-1}(C_1)$, with equality if and only if G is degenerate at $z^* = \phi^{-1}(C_1)$. This is precisely the structure of MDD: a deterministic value $\zeta = z^*$ that exactly binds individual rationality. It remains to verify that the agent will follow the recommendations given by MDD; that is, the constraints in (A.10) hold for all z . For $z \geq z^*$, we have $\mathbf{1}(\zeta > z) = 0$, so the left-hand side equals zero and the constraint is trivially satisfied. For $z < z^*$, the constraint becomes: $\phi(z^*) - \phi(z) \leq C_1(1 - z)^{\frac{r+\lambda_H}{r_p+\lambda_L}}$. Since $\phi(z^*) = C_1$, this is equivalent to $\phi(z) \geq C_1 \left(1 - (1 - z)^{\frac{r+\lambda_H}{r_p+\lambda_L}}\right)$. Note that $\phi(z) = 1 - (1 - z)^{1/\rho}$, so the inequality holds trivially at $z = 0$. Let $\kappa \equiv \frac{r_p+\lambda_L}{r+\lambda_H}$. For $z \in (0, z^*)$, it suffices to show

$$h_4(z) \equiv \frac{1 - (1 - z)^{1/\rho}}{1 - (1 - z)^{1/\kappa}} \geq C_1. \quad (\text{A.12})$$

Since $\lambda_H > \lambda_L$, the exponents satisfy $1/\rho < 1/\kappa$. We first show that $h_4(z)$ is strictly increasing. Differentiating with respect to z yields:

$$h'_4(z) = \frac{\frac{1}{\rho}(1 - z)^{1/\rho-1} \left(1 - (1 - z)^{1/\kappa}\right) - \frac{1}{\kappa}(1 - z)^{1/\kappa-1} \left(1 - (1 - z)^{1/\rho}\right)}{(1 - (1 - z)^{1/\kappa})^2}.$$

The sign of $h'_4(z)$ is determined by the numerator. Factoring out $(1 - z)^{1/\rho-1}$, we define the remaining term:

$$\begin{aligned} h_5(z) &\equiv \frac{1}{\rho} \left[1 - (1 - z)^{1/\kappa}\right] - \frac{1}{\kappa} (1 - z)^{1/\kappa-1/\rho} \left[1 - (1 - z)^{1/\rho}\right] \\ &= \frac{1}{\rho} - \frac{1}{\kappa} (1 - z)^{1/\kappa-1/\rho} + \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) (1 - z)^{1/\kappa}. \end{aligned}$$

We have $h_5(0) = \frac{1}{\rho} - \frac{1}{\kappa} + \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) = 0$, and

$$\begin{aligned} h'_5(z) &= \frac{1}{\kappa} \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) (1 - z)^{1/\kappa-1/\rho-1} - \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) \frac{1}{\kappa} (1 - z)^{1/\kappa-1} \\ &= \frac{1}{\kappa} \left(\frac{1}{\kappa} - \frac{1}{\rho}\right) (1 - z)^{1/\kappa-1/\rho-1} \left(1 - (1 - z)^{1/\rho}\right) > 0. \end{aligned}$$

Thus, $h_5(z) > 0$ for all $z \in (0, 1)$, implying $h'_4(z) > 0$. Since $h_4(z)$ is strictly increasing, the inequality (A.12) is most restrictive as $z \rightarrow 0^+$. Applying l'Hôpital's rule:

$$\lim_{z \rightarrow 0} h_4(z) = \lim_{z \rightarrow 0} \frac{\frac{1}{\rho}(1 - z)^{1/\rho-1}}{\frac{1}{\kappa}(1 - z)^{1/\kappa-1}} = \frac{\kappa}{\rho} = \frac{r + \lambda_L}{r + \lambda_H}.$$

Therefore, $h_4(z) \geq C_1$ holds for all $z \in (0, z^*)$ if and only if $\frac{r+\lambda_L}{r+\lambda_H} \geq C_1$. From the definition of $\hat{\mu}$ in equation (12), this is equivalent to $\mu \leq \hat{\mu}$, which is satisfied by assumption. Thus, the constraints in (A.10) hold for all z , confirming that the agent will follow the recommendations given by MDD. Therefore, MDD is optimal among all dynamic information policies.

If $r_p = r$, then $\rho = 1$ and $\phi(z) = z$, so the objective and constraint in (A.10) have the same functional form. Thus, any policy that satisfies the interim incentive compatibility constraints while exactly binding the individual rationality constraint is optimal, including both IPD and MDD. \square