

Moving the Goalposts^{*}

Jeffrey C. Ely[†]

Northwestern University

Martin Szydlowski[‡]

University of Minnesota

April 24, 2019

^{*}This research is supported by grants from the National Science Foundation. We thank the Editor and four anonymous referees for constructive feedback that helped us improve the paper. We also thank Briana Chang, Eddie Dekel, Adriana Lleras-Muney, and David Rahman for valuable discussions.

[†]Corresponding author. *E-mail Address*: jeffely@northwestern.edu.

[‡]*E-mail Address*: szydl002@umn.edu.

Abstract

We study information as an incentive device in a dynamic moral hazard framework. An agent works on a task of uncertain difficulty, modeled as the duration of required effort. The principal knows the task difficulty and provides information over time. The optimal mechanism features moving goalposts: an initial disclosure makes the agent sufficiently optimistic that the task is easy. If the task is indeed difficult the agent is told this only after working long enough to put the difficult task within reach. Then the agent completes the difficult task even though he never would have chosen to at the outset.

Keywords: goalposts, leading him on, information design.

1 Introduction

Monetary and other material rewards are standard instruments in the design of incentives. But when material rewards are either unavailable or have already been exhausted, there still remains a powerful motivational tool: information. In an organization where salaries and bonuses are set by higher-ups, a mid-level manager incentivizes his team by providing performance feedback and guiding productive activity. An athlete whose long-term motivation comes from prizes offered by the wider market nevertheless finds day-to-day motivation from a coach who offers nothing but information. And of course information is an essential currency of *intra*-personal incentives wherein a decision-maker might manage his own access to information in order to better align his impulses with longer-term goals.

We analyze a principal-agent problem in which the agent is faced with a task of uncertain difficulty. Difficult tasks require a longer duration of effort to complete. If the agent completes the task, he earns a reward that is exogenously given and outside the control of the principal. The principal is informed about the difficulty of the task and she selectively reveals information over time to the agent about how far he is from completion. The principal's objective is to induce the agent to work as much as possible whereas the agent trades off the reward from completing the task and the cost of effort.

The principal can commit to an arbitrary dynamic policy of information disclosure. The agent knows the policy, observes the realized disclosures, rationally updates his belief about the difficulty of the task, and optimally responds with effort choices to maximize his expected payoff from the reward net effort costs. We fully characterize the optimal policy.

The nature and timing of optimal disclosures vary across the lifetime of the project. In the late stages, when the agent has already made substantial progress, his goal is

within reach even if the task is difficult. At this stage the principal wants to persuade the agent that the task is difficult so that the agent will continue working longer. However, in the early stages of the project, the difficult task may be out of reach: if the agent knew the task were difficult he would prefer to quit immediately rather than spend the effort to complete it. Thus, at the early stages of the project the principal wants to persuade the agent that the task is easy.

This reversal of incentives over the duration of the project gives rise to some novel features of optimal persuasion in our dynamic setting. When the task is not too difficult (in expected terms), the optimal contract involves *leading the agent on*. Early on, the agent would like to know whether the task is indeed difficult because in that case he would rather quit immediately. Instead the principal is purposefully silent and reveals nothing about the difficulty of the task. If indeed the task is difficult the principal waits until the agent has worked long enough to complete the easy task and only then does she inform the agent that there is more work to do. Having reached that point, since the agent has already made substantial progress he completes the difficult task even though he would never have set out to do that in the first place.

This feature of the optimal mechanism utilizes a role for information as an incentive device that is novel and unique to dynamic persuasion: information as a carrot. When at the outset the agent is pessimistic that the task is difficult, he is unwilling to even begin working. One way to induce the agent to spend effort would be to disclose information right away to persuade the agent that the task is less difficult than he fears. Instead, the principal withholds this valuable information and only promises to reveal it as a reward for some minimal initial effort. Thus, delayed disclosure confers a double incentive effect: first, the carrot incentivizes the agent to commence working at the outset and later, the eventual disclosure induces him to continue working when the task turns out to be in fact difficult.

However, when the agent is very pessimistic, even promising the maximum value, i.e. full disclosure, is not enough to incentivize the agent to start working. In this case the optimal mechanism involves *moving the goalposts*. The principal begins with an initial partial disclosure that either convinces the agent the task is hard or leaves him relatively optimistic that the task is easy and asks him to begin working. Then, after the agent has made substantial progress, if the task is in fact difficult the principal provides a second disclosure which informs the agent of this and induces him to work even more.

Next we endogenize the task difficulty. We consider a principal who can choose the prior distribution over tasks and then designs an information policy. In the context of a junior partner who is working toward a promotion we are asking whether the promotion standard should be transparent and known at the outset or vague and clarified only partially and gradually over time. We show that task uncertainty, by unleashing the benefits of dynamic information disclosure can improve over the optimal deterministic mechanism, especially when the principal is patient.

1.1 Related Literature

We contribute to the literature on information design and Bayesian persuasion, see Bergemann and Morris (2016b), Bergemann and Morris (2016a), Kamenica and Gentzkow (2011), and Aumann, Maschler and Stearns (1995). We apply these tools to a dynamic setting and study how the principal should reveal information over time to motivate effort.

The most closely related paper is Smolin (2017) who studies optimal performance feedback for an agent who is learning his own ability. Ability is persistent and this is the key difference relative to our model. The goal of the principal is always to persuade the agent that he is high ability so that he will keep working. The incentive reversals

that drives the qualitative features of our optimal disclosure policy do not arise. Just as we do, [Smolin \(2017\)](#) uses efficiency arguments to simplify the characterization of optimal policies.

[Orlov, Skrzypacz and Zryumov \(2016\)](#) study a receiver who chooses when to exercise an option based on a publicly observed diffusion process and information provided by a sender. They consider Markov Perfect Equilibria. The key difference is that in our setting the principal has commitment power while their sender cannot fully make use of promised disclosures. Consequently, we obtain entirely different optimal policies.

Other closely related papers are [Renault, Solan and Vieille \(2017\)](#) and [Ely \(2017\)](#). [Ely \(2017\)](#) studies a persuasion model in which a Poisson realization arrives unknown to the agent. The principal knows about the arrival and chooses when (and how) to reveal it. In the optimal policy, the principal reveals the arrival with a fixed delay. Crucially, the agent in the setting of [Ely \(2017\)](#) acts myopically: promising future information to incentivize the agent has no use. In our model, the agent is forward-looking and promised information is the central object of our study and what drives our leading on and moving goalposts policies. [Renault, Solan and Vieille \(2017\)](#) also assume the agent is myopic. Our model studies how to structure the timing of information disclosure even when all information is available at the outset. In both [Renault, Solan and Vieille \(2017\)](#) and [Ely \(2017\)](#), the timing of information disclosure is driven by an exogenous stochastic process for information arrival.¹

Using information to motivate an agent is a theme in a number of other papers. [Varas, Marinovic and Skrzypacz \(2017\)](#) study a model in which a principal can monitor an agent with reputation concerns. They characterize optimal monitoring, but they do not consider how the principal should release information once it is ac-

¹We thank an anonymous referee for emphasizing this distinction.

quired. Hörner and Lambert (2016) study how to provide information about the effort of an agent with career concerns in order to motivate him.² Using promised information as an incentive device is reminiscent of ?. They study a model in which a principal is selling information to an agent who makes monetary payments in advance. Gradual disclosure extracts more of the value of this information. Our setting features no monetary payments.

Moving goalposts also appear in Georgiadis, Lippman and Tang (2014) but for a different reason: they arise because a principal cannot fully commit to an incentive scheme. A natural application of our model is to the traditional Master-Apprentice training relationship (and its modern variants found in many law firms and consulting practices). Our work thus naturally complements the work of Garicano and Rayo (2017) and Fudenberg and Rayo (2017) who study apprenticeships under complete information with the focus on the principal's exchange of human capital for productive services.

Our work can also be seen as a contribution to optimal incentive contracts without transfers. Ben-Porath, Dekel and Lipman (2014), Bull and Watson (2007) and Ben-Porath, Dekel and Lipman (2017) examine disclosure policies as incentive schemes in a setting of verifiable information. For a broad survey, see Dekel (2017).

Finally, our paper is related to the literature on dynamic moral hazard (e.g. Spear and Srivastava (1987) or Sannikov (2008)). There, the promise of future payments motivates current effort. In the optimal contract, payments are often backloaded. In our setting, information is optimally backloaded. When using "leading on" the principal delays valuable disclosures into the future, knowing that this is sufficient to motivate effort today.

²See also Hansen (2012) and Fang and Moscarini (2005) for earlier work.

2 Model

An agent works on behalf of a principal and spends effort continuously until he chooses to quit. At that point the relationship ends and the agent earns a reward $R > 0$ if and only if his total accumulated output exceeds a threshold $x > 0$. We refer to the threshold as the difficulty of the task. We normalize the production technology to accrue one unit of output per instant of effort at flow cost $r \cdot c$ to the agent.³ The realized threshold is unknown to the agent at the outset and the agent's prior is given by the CDF F . Thus, the total (ex ante) expected payoff to the agent from working until time $\tau > 0$ is

$$F(\tau)e^{-r\tau}R - c(1 - e^{-r\tau}).$$

The principal is the residual claimant of the agent's output until the time the agent quits. The principal wishes to maximize expected discounted output which the principal discounts at rate r_p . We normalize the principal's flow value of output to r_p so that the principal earns total discounted payoff $(1 - e^{-r_p\tau})$ when the agent works until τ . For most of our analysis we assume that the principal and agent share the same discount rate, i.e. $r_p = r$. We take up the case of unequal discounting when we turn our attention to the design of the threshold.

The principal observes the realized threshold x at the beginning of the relationship and chooses how and when to disclose information to the agent about the difficulty of the task. The principal can make any number of arbitrary disclosures at any time in the process and we refer to the rule governing these disclosures as the information policy. Crucially we assume that the principal commits to an information policy and that the agent knows the policy and understands the principal's commitment. Our goal is to

³Normalizing the flow cost by r helps simplify the notation throughout the paper. We do the same for the principal's flow value.

characterize the optimal policy.

Without information from the principal, the agent does not know whether he has already passed the threshold. The optimal policy is designed to induce the agent to work as long as possible (in expectation) in order to maximize the payoff of the principal. The time structure of disclosures engages the following tradeoff. Informing the agent that the task is easy incentivizes the agent to continue working. However, a *policy* of announcing that the task is easy entails the downside of making the agent pessimistic in the absence of the announcement, possibly inducing the agent to quit. The optimal policy balances these gains and losses at every instant while the relationship is ongoing.

We will show that two policies are optimal, depending on the agent's initial belief. *Leading on* means the principal withholds valuable information until time t . *Moving the goalposts* means the principal makes an initial disclosure at time zero aimed to decrease the agent's expectation of the threshold and future disclosures to increase it. We can describe these policies formally in terms of designing a signal $\sigma_t(x) \in \Delta(\{0, 1\})$, in the sense of Kamenica and Gentzkow (2011), which informs the agent whether to continue ($\sigma_t = 1$) or quit ($\sigma_t = 0$) at any point in time.⁴

Definition 1. A *leading on policy* consists of a random time $t > 0$ and a process $\sigma_t(x) \in \Delta(\{0, 1\})$, such that for all $\hat{t} < t$, $\sigma_{\hat{t}}(x) = 1$ for all x and the agent works at least until t .

A *moving the goalposts policy* consists of a process $\sigma_t(x) \in \Delta(\{0, 1\})$ such that $E(x|\sigma_0 = 1) < E(x)$ and at all times $t > 0$ at which the agent hasn't quit yet, (i) $E(x|\{\sigma_s\}_{s \leq t})$ is weakly increasing in t and (ii) there exist two times $\hat{t} < t$ so that the increase is strict between \hat{t} and t .

Under a *leading on policy*, the principal provides no useful information before time t .⁵ However, the information provided after this time is valuable and induces the agent

⁴In the full model in Section 5, it is more convenient to describe the principal's policies indirectly via the paths of effort they induce.

⁵Because $\sigma_{\hat{t}}(x)$ is independent of x .

to work at least until t . Under moving goalposts, she provides an initial disclosure, but then over time the agent learns that the task is hard. This learning can all occur at one time when the threshold is binary (in [Section 4](#)) or gradually when the threshold is continuous (in [Section 5](#)).

3 Interpretations of the model

In this section we describe two interpretations of the model: apprenticeships and corporate succession.

Apprenticeships Apprenticeships have historically been the main method to transfer skills and are still widely used today.⁶ In a typical apprenticeship, the apprentice works for the master at a wage which is fixed below the wage of a skilled worker.⁷ In exchange, the master trains the apprentice. It takes time for the apprentice to gain knowledge.⁸ Once the apprentice believes he has gained enough, he has no incentive to keep working at the reduced wage. Instead, he can leave and find employment as a journeyman or as a skilled laborer. But if the apprentice leaves, the master cannot recoup any costs for training the apprentice, which are often incurred at the beginning of the relationship. This moral hazard problem of the apprentice is a key friction recognized by economic historians⁹ and it has been identified as an important contributor to the decline of apprenticeships throughout the 20th century.¹⁰

⁶We are very grateful to Joel Mokyr for pointing us to this literature.

⁷In an apprenticeship the apprentice and master agree on a wage ex-ante, which often does not change as the apprentice gains experience. See e.g. Wallis (2008).

⁸Often, the apprentice simply learns by observing the master and there is little formal instruction. See Mokyr (2018).

⁹E.g. Epstein (1998), Wallis (2008), Mokyr (2018), and many others.

¹⁰Increased apprentice mobility is the central argument in Elbaum (1989), who examines the decline of apprenticeship in the US. Evidence indeed suggests that enforcement of apprenticeship contracts was often weak and desertion a serious problem, see e.g. Minns and Wallis (2012).

The master-apprentice relationship fundamentally relies on an information asymmetry. The apprentice enters it “by definition underinformed about the material to be taught” (see Mokyr (2018)). He does not know the skills necessary to succeed in the trade and often cannot judge the quality of instruction. Thus, at any given time, the apprentice is uncertain whether he has learned enough to strike out on his own.¹¹ But as time passes, he becomes more and more certain that he has. This is exactly what happens in our model. Without information from the principal, the probability that the agent has completed the task is monotonically increasing in time, and the agent quits when he is sufficiently certain. From the master’s perspective “it was in the interest of the master to keep the asymmetry as long as possible, since it was this asymmetry that allowed him to control his worker and thus draw rents from the apprentice’s labor” (Mokyr). This maps closely to the principal’s problem in our model. She knows whether the agent has completed the task (i.e. learned enough to leave the master) but she strategically reveals that information to the agent over time, to make him work as long as possible.

Corporate Succession As Fudenberg and Rayo (2017) note, contemporary “work-for-training” arrangements can be found in professional services firms (i.e. law, accounting, and consulting) and the sciences.¹² In these settings, employees are mainly motivated by future career progress, in the form of internal promotions or finding more desirable employment elsewhere.¹³ Attaining these goals depends on factors which

¹¹De la Croix, Doepke and Mokyr (2017) note that “The apprentice, by the very nature of the teaching process, is not in a position to assess adequately whether he has received what he has paid for until the contract is terminated.” Humphries (2003) finds that during the British Industrial Revolution, only one third of all apprentices became masters, or “freemen.” This suggests substantial uncertainty in labor market outcomes for the apprentice.

¹²For example, postdoctoral positions have been criticized for being overly long and featuring uncertain outcomes. See Stephan and Ma (2005) and Stephan (2013).

¹³Indeed, high work load and low hourly wages are characteristic at the entry level in these industries. See Fudenberg and Rayo (2017) for a list of examples.

may be difficult to observe for the employee, such as the competitiveness of the labor market, workplace politics, or his own rate of skill acquisition.¹⁴ How long the employee has to work to be eligible is therefore uncertain. A manager who can observe these factors can then use her information to motivate the employee to work.

Generally, the manager's and the employee's incentives are not aligned. Managers may prefer to retain valuable employees and thus hold them back even though they are ready to advance. The manager's and employees incentives hence map closely into our model and the manager can use "moving goalposts" to extract value from the employee. Indeed, holding back employees is a common problem in firms and managers frequently overpromise when it comes to career advancement or exit options.¹⁵

To manage employees expectations, companies build sophisticated "succession management" systems.¹⁶ These inform employees e.g. about how far away they are from building the required competencies for a particular promotion. As Conger and Fulmer (2003) note, "Succession planning systems have traditionally been shrouded in secrecy in an attempt to avoid sapping the motivation of those who aren't on the fast track. [...] If you don't know where you stand [...] you will continue to strive to climb the ladder." See also Rothwell (2010) for an extensive discussion.¹⁷

It is not difficult to imagine that these systems could be designed to guarantee commitment. Even if an individual manager might have no commitment power over

¹⁴In apprenticeships, the apprentice does not know enough to evaluate his own work performance and must rely on the master for feedback. A similar problem should arise in entry-level positions in accounting, law, the sciences, etc.

¹⁵See Behrens (2015) for holding back and Rousseau (1995) for a list of examples of overpromising.

¹⁶We are very grateful to Connie Wanberg for pointing us to this literature.

¹⁷Rothwell writes: "The most famous question in succession planning is the well-known one: Should successors be told, or not told, that they are successors? [...] Consider the disadvantages of telling. [...] individuals [...] may kick their feet up and relax, assuming that they have the promotion in the bag. [...] Worse, their peers may hear of it. Some of these otherwise excellent employees may lose hope and start looking for work elsewhere." Notice how closely this intuition maps into our model. An agent will stop working if he believes he has reached the threshold. But he will also stop if he learns the threshold is far away. In practice, managing expectations is difficult. Rothwell notes that by telling one can also "motivate successors [...] in anticipation of future promotion."

what she reveals to employees about future advancement, a system designed from the ground up, with exact rules what is revealed and to whom may accomplish just that.¹⁸

We begin our analysis with a simple example to illustrate the main ideas.

4 Binary Example

In this example the task is either easy, with threshold $x_l > 0$ or hard, with threshold $x_h > x_l$. The principal knows the task difficulty. The agent is uncertain and begins with prior probability μ that the task is easy.

The hard task is not individually rational for the agent. In particular, the maximum individually rational effort $\bar{\tau}$ is given by

$$e^{-r\bar{\tau}}R - c(1 - e^{-r\bar{\tau}}) = 0,$$

and $x_l < \bar{\tau} < x_h$. Absent any intervention by the principal, the agent would never work longer than x_l . Nevertheless the hard task is not too hard: the incremental effort is individually rational: $x_h - x_l \leq \bar{\tau}$.

The model sheds new light on optimal information design by demonstrating how dynamic disclosures can play a role even in a setting which, from the point of view of fundamentals, is essentially static in nature. In our model the state, i.e. the total amount of effort required to complete the task, is determined at the outset and never changes. The other dynamic state variables, effort and progress, are symmetrically ob-

¹⁸See Conger and Fulmer (2003) for detailed descriptions of such systems. “At Lilly, for example, people know if they are regarded as having additional potential, but they don’t know exactly how high that potential is, nor do they know about every role for which they may be considered. [...] Lilly’s Web-based succession tool is available through an icon on employees’ computer desktops. A click on the icon takes the employee to a portal on the company’s intranet, with [...] job opportunities customized for each employee.” Similarly, “At Dow Chemical, employees nominate themselves for positions online [...]. Dow’s Web tool also includes career opportunity maps that detail the sequence of jobs one can expect in a function or line of business. Some companies even show compensation ranges by level and position.”

served by principal and agent. Thus, all of the information the principal ever provides could feasibly be provided in a single disclosure at the beginning of the relationship.¹⁹ Indeed there is a natural static reduced-form of our problem in which that is the only opportunity to provide information, after which the agent chooses how much effort to spend and the game ends. We show how the optimal dynamic contract strictly improves on the best the principal could achieve in that static reduced-form.

Static Disclosure Consider as a benchmark the static reduced-form of the problem in which there is a single disclosure by the principal after which the agent chooses effort. Because $x_h > \bar{\tau}$, the agent would not set out to complete the hard task regardless of the his belief μ . Therefore there is no one-shot disclosure policy that could persuade the agent to work until x_h . The best the principal can do with such a policy is to persuade the agent that the task is easy. Indeed, if the agent is led to believe that with at least probability

$$\bar{\mu} = \frac{c}{R} \left(\frac{1 - e^{-rx_l}}{e^{-rx_l}} \right)$$

the task is easy, he will optimally choose effort $\tau = x_l$. Any more pessimistic belief will induce the agent to quit immediately and earn zero. The optimal one-shot, (or *static*) disclosure policy is designed to minimize the probability that the agent is as pessimistic as that.

Proposition 1. *No static policy can induce the agent to complete the hard task. The optimal static policy is the one that maximizes the probability that the agent has belief at least $\bar{\mu}$ and spends effort $\tau = x_l$.*

This static version of the problem is equivalent to a standard Bayesian Persuasion problem as studied by Kamenica and Gentzkow (2011) and Aumann, Maschler and

¹⁹Thus, the role for dynamic disclosures is inherently different from Ely (2017). In that paper disclosure policies are necessarily dynamic because the payoff-relevant variables privately observed by the principal are themselves changing over time.

Stearns (1995). The principal optimally sends two messages: an optimistic message which persuades the agent to believe that with exactly probability $\bar{\mu}$ the task is easy and a pessimistic message which convinces the agent the task is hard. The optimistic message is sent whenever the task is truly easy and also with positive probability conditional on the task being hard. The latter probability represents the principal's gain from persuasion. This probability is maximized subject to the constraint that the resulting pooled message leaves the agent with a belief at least $\bar{\mu}$, and this constraint binds.

Dynamic Disclosure Dynamic disclosure policies can do better for a couple of reasons. First, information is valuable to the agent and that value is often best withheld and offered only later as a carrot to incentivize early effort. Second, having reached that later date the cost of past effort is sunk and the principal may then find it possible, and indeed in her interest to persuade the agent to continue working. Based on these ideas we will show how to structure dynamic disclosure to induce the agent to complete the hard task and in fact to complete the task with probability 1 regardless of its difficulty.

Suppose the principal were to delay all disclosure until time x_l and then to fully disclose the threshold. The agent would quit upon hearing that the task was easy (and already complete) and continue working until x_h upon hearing the task was difficult (because $x_h - x_l \leq \bar{\tau}$). Of course that raises the question of how the principal could have induced the agent to reach x_l in the first place. Indeed as the static analysis revealed, when the agent begins with a prior below $\bar{\mu}$ he would quit immediately unless persuaded that the task were easy. With dynamic disclosure however, information as a carrot can substitute for information as persuasion. Consider the time-zero value to

the agent of full disclosure at time x_l :

$$V(\mu) = e^{-rx_l} \left\{ \mu R + (1 - \mu) \left[R e^{-r(x_h - x_l)} - c \left(1 - e^{-r(x_h - x_l)} \right) \right] \right\}.$$

With probability μ the agent learns that the task is already complete and therefore quits and earns the reward. With probability $1 - \mu$ the agent learns that the task is hard. The incremental effort necessary to complete the difficult task is individually rational and the agent optimally does so, incurring additional effort costs $c \left(1 - e^{-r(x_h - x_l)} \right)$ and earning the reward at date x_h .

Leading the agent on For all priors exceeding $\tilde{\mu}$ defined by $V(\tilde{\mu}) = c \left(1 - e^{-rx_l} \right)$, information provided at date x_l is of sufficient value to incentivize the agent to begin working.²⁰ No initial persuasion is necessary. Indeed full delayed disclosure eventually induces the agent to complete the task with probability 1. This includes the difficult task which, had he known at the outset, he would never have set out to complete.

We refer to this device of using the value of information as a carrot and delaying persuasion until effort costs are sunk as *leading the agent on*. Within this class of policies, the extreme of full delayed disclosure may be more incentive than necessary, i.e. when $\mu > \tilde{\mu}$. In that case further fine-tuning of this mechanism utilizes persuasion at time x_l to increase the probability that the agent works through x_h . This is illustrated in [Figure 1](#).

The left panel shows the value of information at time x_l for the agent. The two segments of the solid upper envelope show the agent's continuation value function: the payoff from quitting at x_l when his prior is high, and continuing to x_h when his prior is low. The right panel shows the value to the principal. The solid line is her

²⁰Note that $\tilde{\mu}$ is the smallest belief at which the agent would begin working without any incentive, hence $\tilde{\mu} < \bar{\mu}$.

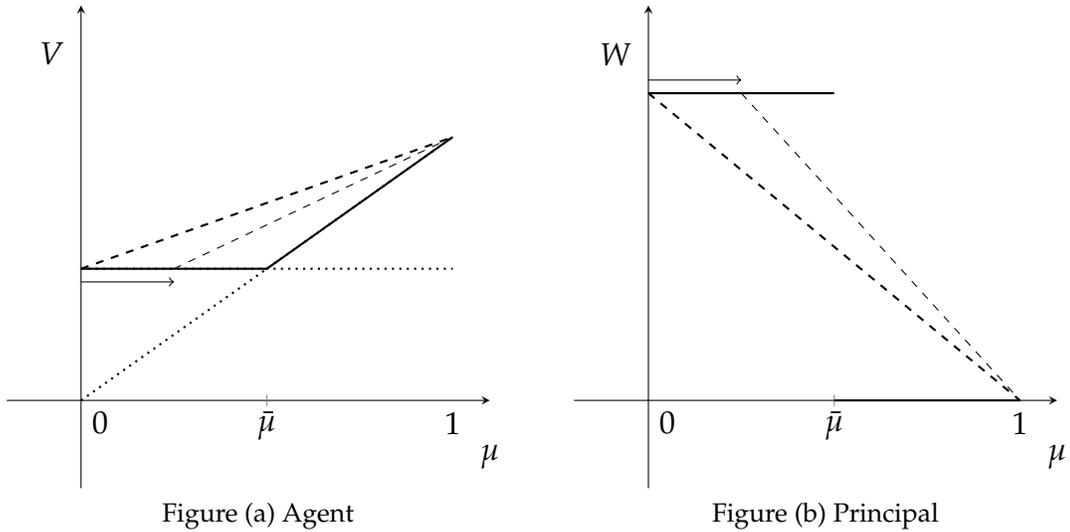


Figure 1: Extracting value. On the x-axis is the posterior belief after revealing information at x_l . At the y-axis are the values of the principal W and agent V . The solid line on the left panel is the agent's continuation value under a given posterior μ . Similarly, on the right panel, the solid lines are the principal's continuation value for given posteriors. On both panels, the dashed lines represent the expected values from different information policies.

value if the agent continues (when the prior is low) or stops immediately (when it is high). Consider a line connecting any two points on the graph of the agent's value function and a delayed disclosure policy involving two messages which induce the corresponding beliefs. By the law of total probability, that line gives for any prior belief the expected value of such a disclosure.

For example, the thick dashed line is the value function for full disclosure. By contrast, a policy such as the one represented by the thin dashed line is less informative and less valuable to the agent. In particular, it induces effort x_h with a less-informed belief. This is achieved by sending the associated message not just when the task is difficult but also sometimes when the task is easy.

The thin dashed line in the right panel shows the value of the same policy to the principal. Compared to the full-disclosure policy, the alternative policy induces the

agent to work until x_h with a larger probability and is thus better for the principal. We see that distortions of this form reduce the agent's payoffs while raising the principal's. The optimal delayed-disclosure mechanism will thus involve just enough distortion to reduce the agent's ex ante value to zero.

Proposition 2. *When the prior exceeds $\tilde{\mu}$, leading the agent on induces him to complete the task with probability 1. The optimal such mechanism maximizes the probability that the agent works through time x_h subject to the constraint that the agent's ex ante payoff is at least zero. Indeed this is the optimal mechanism among all dynamic disclosure policies.*

The last claim in the proposition is that delayed disclosure is optimal within the entire feasible set of dynamic disclosure policies, including those that involve multiple history-contingent disclosures at staggered times, possibly inducing quitting times other than x_l or x_h . This follows from two observations. First, as discussed above it holds the agent to his reservation value of zero so that no further reduction in the agent's payoff would be implementable. Second, the policy is *efficient*: there is no alternative feasible effort plan that can further raise the principal's payoff without lowering that of the agent. As we show in [Section 5](#), these are general properties of optimal dynamic disclosure policies. The proof of efficiency for this binary example is in the [Appendix](#).

Moving the goalposts Dynamic disclosure enables the principal to increase his payoff relative to static disclosure. [Figure 2](#) plots the value to the principal as a function of the agent's prior from leading the agent on.

When the agent is initially so pessimistic that $\mu < \tilde{\mu}$, even full delayed disclosure is insufficient incentive to induce the agent to start working. Thus the principal's payoff drops to zero for such priors. We can now use ideas from Bayesian persuasion

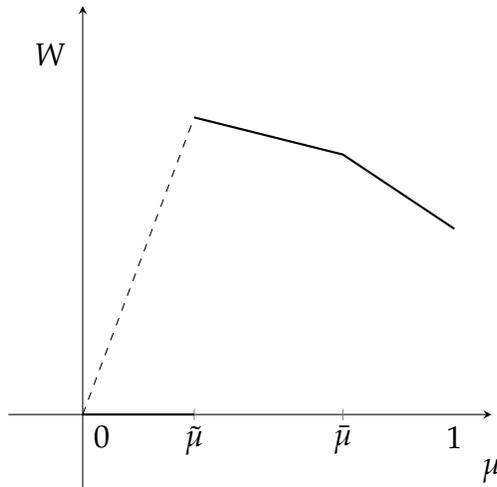


Figure 2: Moving the goalposts. The solid line is the principal's value from the optimal leading-on policy for a given prior belief μ . If the prior is too low, leading on cannot induce the agent to work and the principal's value is zero (below $\tilde{\mu}$). The dashed line characterizes the optimal information policy in this case, which randomizes between beliefs zero and $\tilde{\mu}$.

to build upon delayed disclosure and construct a new mechanism which improves for this range of priors. The dashed segment represents the concavification of the principal's value function. Following [Kamenica and Gentzkow \(2011\)](#) and [Aumann, Maschler and Stearns \(1995\)](#), the concavification shows the maximum value the principal can obtain by an initial disclosure. For priors to the left of $\tilde{\mu}$, this initial disclosure reveals one of two possibilities. The first is that the task is certainly difficult, leading the agent to quit and yielding zero for the principal. The alternative is that the task is easy with probability $\tilde{\mu}$, making the agent just optimistic enough begin working in anticipation of a second disclosure at time x_l .

Notice the pattern of disclosures in this mechanism. The goal of the initial disclosure is to persuade that the task is easy so the agent sets out working while the goal of the second disclosure is to persuade that the task is hard so that he keeps going. We refer to this mechanism as *moving the goalposts*.

Proposition 3. *When the prior is below $\tilde{\mu}$, the optimal mechanism is moving the goalposts.*

The principal provides disclosures at no more than two dates. At date zero the disclosure is designed to maximize the probability that the agent begins working and at date x_l the disclosure is designed to maximize the probability that the agent continues to x_h .

The general case The mechanisms derived above made use of the fact that the incremental effort $x_h - x_l$ was individually rational. Thus, after the sunk effort costs of reaching x_l it was possible for the principal to persuade the agent to keep working. This would not be possible if the hard task was so hard that the incremental effort exceeded $\bar{\tau}$. However, in that case, consider a third possible threshold in between x_h and x_l such that each increment was individually rational. Indeed consider adding a fourth threshold even higher than x_h with again an individually rational increment. In these extensions, it is possible to induce the agent to complete each successive increment through disclosure policies that combine the ideas of leading the agent on and moving the goalposts.

In the analysis of the general model below we characterize the optimal disclosure policies for general distributions of thresholds. In addition we consider the question of how the principal would optimally design the threshold distribution. Indeed not only can the principal use information design to manage incentives *given* the agent's initial incomplete information about the task, but in fact the principal *benefits* from that incomplete information, and strictly so when the principal is more patient than the agent, i.e. when her discount rate r_p is strictly smaller than r .

4.1 The Optimality of A Random Threshold

When the principal is more patient than the agent, we can show that the random threshold, taken for granted in the preceding sections, is in fact part of an optimal mechanism when the threshold is a choice of the principal. The key idea is that dis-

counting makes the agent effectively risk-seeking with respect to random thresholds. Increasing uncertainty therefore raises the agent's payoff and allows the principal to extract greater expected effort through disclosure. When the principal is more patient than the agent this increase in expected effort strictly raises the payoff to the principal. On the other hand, discounting makes the principal effectively risk-averse with respect to random thresholds. When the principal and agent are equally patient these forces cancel each other out so that random and deterministic policies are equally good.²¹

Let us begin by considering the optimal deterministic mechanism. In a deterministic mechanism the threshold x is known to the agent and there is hence no role for information disclosure. The agent works if and only if x is individually rational and therefore the maximal individually rational effort duration $\bar{\tau}$ is the optimal deterministic threshold for the principal.

We will now show how adding randomness and thereby enabling delayed disclosure (via leading on) expands the feasible set of implementable efforts and raises the payoff of the principal. Consider a random threshold with two possible realizations, x_l and x_h in which

$$\begin{aligned}x_l &= \bar{\tau} - \tilde{\tau} \\x_h &= \bar{\tau} + \tilde{\tau},\end{aligned}$$

where $\tilde{\tau}$ is any positive duration satisfying

$$R/2 \leq e^{-r2\tilde{\tau}}R - c(1 - e^{-r2\tilde{\tau}}). \quad (1)$$

This new distribution is a mean-preserving spread of $\bar{\tau}$ with the gap (equal to $2\tilde{\tau}$) between high and low thresholds small enough that the agent would prefer to continue

²¹Similarly, when the principal is less patient than the agent it becomes optimal to reduce uncertainty.

working from x_l to x_h in order to earn the reward with probability 1 rather than quit at x_l and earn the reward with probability 1/2. Note that $2\tilde{\tau} < \bar{\tau}$.

Without any accompanying dynamic disclosure policy, this random threshold would be strictly worse for both parties. Intuitively, the agent was just willing to work until $\bar{\tau}$ and earn the reward for sure. The easy task is not much easier than that but would now only yield the reward with probability 1/2. Faced with that option or the option of the non-individually-rational difficult task the agent would strictly prefer to quit immediately.

However, full delayed disclosure would induce the agent to complete the task with probability 1. The result is a lottery: with equal probability the agent works until either x_l or x_h whereupon he earns the reward. Due to discounting, the agent is risk-loving with respect to this lottery. To see this, re-arrange the agent's utility from earning the reward at some time t as follows:

$$(c + R)e^{-rt} - c.$$

Ignoring the constant $-c$, the agent evaluates lotteries using a strictly convex exponential utility function. He thus strictly prefers the random threshold.

The principal on the other hand is risk-averse with respect to the lottery for exactly the same reason: up to a constant his utility is $-e^{-rp^t}$. Thus, merely adding randomness is no more than a transfer of surplus from the principal to the agent.

There is however a second channel for transferring surplus from the agent back to the principal: increase the mean of the agent's effort. This is achieved by the leading on policy. Because the random mechanism is strictly individually rational (we showed it gives the agent strictly higher payoff than the just-individually-rational deterministic threshold), the principal can use persuasion at date x_l to increase the probability the

agent works until x_h . Indeed, consider the leading on policy which increases the probability of working until x_h just enough to return the agent's ex ante utility to zero (c.f. [Figure 1](#)). We show in the Proposition below that the net effect of randomization plus leading on raises the expected payoff of the principal whenever $r_p < r$.

Proposition 4. *If $r_p < r$, there exists a random threshold which, when coupled with an optimal leading on policy, strictly improves on the optimal deterministic threshold.*

For intuition, consider [Figure 3](#) where the expected payoffs of the agent (V) and principal (W) are plotted on the horizontal and vertical axes. Point A represents the optimal deterministic mechanism. The movement to point B represents the increase in the agent's payoff, call it Δ_a^1 , and decrease in the principal's payoff, call it Δ_p^1 , when the threshold is randomized. The movement to point C represents the loss to the agent Δ_a^2 and gain to the principal Δ_p^2 , from partial delayed disclosure. The proposition states that the line from B to C is steeper than the line from A to B .

For starters, consider the extreme case $r_p \rightarrow 0$ of no discounting for the principal, i.e. linear utility equal to the expected undiscounted effort duration. In this case the proposition holds because the Principal is risk-neutral with respect to the randomization and gains from persuasion; graphically the line AB is horizontal. More generally, note that exponential utility is characterized by constant absolute risk affinity/aversion with the discount rates r or r_p as the risk coefficient. The slope comparison reduces to the comparison of Δ_p^1/Δ_a^1 with Δ_p^2/Δ_a^2 or equivalently Δ_p^1/Δ_p^2 with Δ_a^1/Δ_a^2 . The latter are the ratios of the attitude toward the introduction of risk (i.e. the second derivative of utility) to the attitude toward a shift in "wealth" (the first derivative.) The comparison is thus determined by the comparison of the two parties' coefficients of absolute risk aversion, i.e. the discount rate. Specifically, if $r_p < r$, the principal is effectively less risk-averse than the agent, taking into account the higher average effort. Then, it is possible to find a policy which improves her utility compared to the deterministic

threshold without changing the utility of the agent.

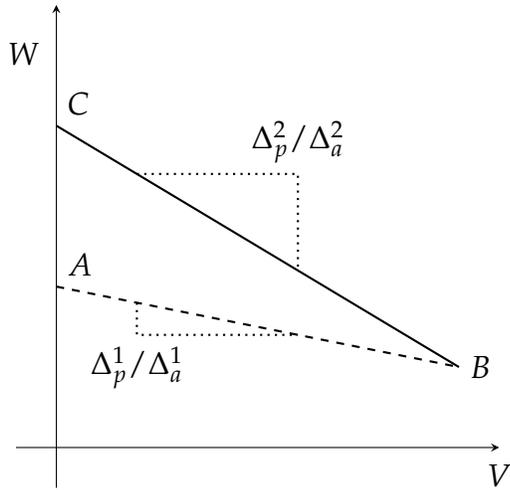


Figure 3: The value of randomness. Point A is achievable by a deterministic threshold. Choosing a random threshold, the principal can move to point B. The slope of the dashed line captures the tradeoff in principal vs. agent values in doing so. By using leading on, the principal can then move from B to C. The solid line captures the tradeoff in principal and agent values from inducing him to work more via this policy.

Given that the principal prefers a random threshold the question of the optimal distribution over thresholds naturally arises. In [Section 6](#) we show that an optimal threshold distribution exists, we characterize it, and show that it is independent of the principal's discount rate.

5 General Distributions

This section presents a general characterization of the optimal dynamic information design for continuous task distributions F . We assume that F has a continuously differentiable density, denoted f , with full support on $[0, \infty)$, and we use H to denote the

hazard rate.²² We make two assumptions on F that facilitate a clean characterization of the optimal policy:

$$\frac{f'(x)}{f(x)} < r \tag{A1}$$

and

$$H'(x) \geq 0. \tag{A2}$$

Assumption (A1) is relevant when the agent receives no information. It ensures that his marginal value of working decreases as he works longer, so that he stops the first time this value reaches zero.

Assumption (A2) captures the main intuition of our introductory example in Section 4. There, if the agent reaches the low threshold but learns that the true threshold is high, he will continue working. In fact, as time passes, he only becomes more willing to work because the high threshold moves closer. Assumption (A2), which says that the hazard rate is increasing in t , implies the same. If at a positive time the agent learns that he has not reached the threshold, he is more willing to continue than before.

In this section we assume that principal and agent have the same discount factor r .

Effort Schedules The main tool of our analysis will be an *effort schedule*, defined as a joint probability of effort levels and thresholds. It describes in probabilistic terms how long the agent works conditional on each possible realization of the threshold. In particular we will consider a probability measure G over \mathbf{R}^2 describing the joint probability of the (exogenous) threshold x and the (endogenous) duration of effort τ . The marginal of G on x coincides with F , and for any fixed threshold, the conditional distribution $G(\cdot \mid x)$ describes the CDF of the agent's effort duration conditional on x being the true difficulty of the task.

²²That is, $H(x) = \frac{f(x)}{1-F(x)}$.

The values to the agent and principal of an effort schedule G are simply the expectations of their corresponding ex-post payoffs with respect to G . The agent earns ex-post payoff

$$v(\tau, x) = \begin{cases} -c(1 - e^{-r\tau}), & \text{if } \tau < x \\ e^{-r\tau}R - c(1 - e^{-r\tau}) & \text{if } \tau \geq x \end{cases}$$

when he works for a duration τ and the realized threshold is x . Thus, the schedule G provides the agent with expected payoff

$$V(G) = \mathbf{E}_G v.$$

The principal's payoff depends only on effort and is given by $w(\tau, x) = 1 - e^{-r\tau}$, so that the expected value to the principal from a schedule G is

$$W(G) = \mathbf{E}_G w.$$

Because we assumed that the distribution of the threshold has no atoms, the relevant effort schedules will be *pure*, i.e. $G(\cdot | x)$ assigns probability one to a single effort duration.²³ In the case of pure effort schedules we will use the notation $g(x)$ to refer to the effort level chosen when the threshold is x . With slight abuse of notation, we denote the expected values under a pure schedule as $\mathbf{E}_g v$ and $\mathbf{E}_g w$.

Not every effort schedule is implementable by an information policy. Information

²³We have assumed that the pdf f is continuously differentiable, which implies that F is continuous and hence atomless. In the binary threshold case of Section 4, moving the goalposts does not induce a pure effort schedule. If $x = x_h$ then there is a positive probability that the agent quits immediately and a positive probability that he works until x_h . Intuitively, revealing information which induces quitting at time zero provides the agent with enough utility so that he is willing to work if the information is favorable. With discrete states, mixing over quitting and continuing is necessary to make the agent's participation constraint bind, which is necessary for an optimal mechanism. That is, without mixing, the value of information and hence the agent's utility cannot be adjusted continuously. With continuous states, no such mixing is necessary. One can instead continuously increase the interval of states in which the agent quits right away.

is a low-powered incentive instrument. The agent is never prevented from simply ignoring all information provided by the principal and making an uninformed optimal effort choice. Thus, a necessary condition for an effort schedule to be implementable is that it provides no less than what we refer to as the agent's *no-information* value.

No-Information Benchmark Consider the decision problem of the agent if the principal were not present or if she provided no information, or if the agent simply ignored any information provided. Since the agent does not observe when he passes the threshold unless he is told by the principal, he simply chooses an optimal effort duration τ , works for that long, and then quits. The agent then earns

$$F(\tau)e^{-r\tau}R - c(1 - e^{-r\tau}),$$

and his optimal strategy would choose τ to maximize the above expression.

The effort schedule associated with the no-information benchmark is the constant pure schedule $g(x) = \tau$ for all x . Denote by V_{ni} the associated expected payoff for the agent. This is the agent's *no-information value*.

Full-Information Benchmark At the opposite extreme we may consider the policy in which the agent learns the realized threshold at the outset and chooses his effort accordingly. It is represented by the pure effort schedule \bar{g} :

$$\bar{g}(x) = \begin{cases} x, & \text{if } e^{-rx}R - c(1 - e^{-rx}) \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

The agent spends just enough effort to complete the task whenever it is *individually rational* to do so, and otherwise he quits immediately. Note that there exists a maximum

individually rational effort level which we denote by $\bar{\tau}$. It is the most difficult task that the agent would knowingly set out to complete and it is given by the equation

$$e^{-r\bar{\tau}}R - c(1 - e^{-r\bar{\tau}}) = 0.$$

Efficiency A schedule G is (ex ante, Pareto) *efficient* if there is no other schedule G' such that $W(G') > W(G)$ and $V(G') \geq V(G)$.²⁴ The full-information schedule is efficient because any schedule that differs with positive probability must give the agent a strictly lower expected payoff. Indeed the full-information policy maximizes the agent's expected payoff.

The no-information schedule, on the other hand, is typically not efficient. This can be seen from the binary example in [Section 4](#). When the agent was pessimistic he quit immediately and earned zero. We constructed an implementable schedule which provided the same utility to the agent but strictly higher payoff to the principal.

Generally, there are two reasons a schedule may be inefficient. First, as above if the threshold is individually rational but the agent never starts working there is a Pareto-improving information policy. Secondly, regardless of whether the threshold is individually rational it would be inefficient for the agent to begin working but quit before completing the task. We show below that these two conditions exhaust all possible inefficiencies.

Proposition 5. *A schedule is efficient if and only if it puts zero probability on the events $x > \tau > 0$ and $\bar{\tau} \geq x > \tau = 0$.*

Individually Rational and Implementable Schedules Individual rationality, the condition that $V(G) \geq V_{ni}$, is a necessary condition for a schedule G to be implementable

²⁴Note that we do not require the dominating schedule G' to be implementable by an information policy.

through an information policy. Going one step further, we may consider the incentives of the agent at any point in the course of an information policy. For each date t , there is an associated *continuation schedule* $G_t = G(\cdot \mid \tau > t)$. This is the conditional distribution of task difficulties and effort durations given that the agent does not quit at or prior to t . It affords the agent a continuation no-information value $V_{ni,t}$ given by the expected payoff from ignoring any further messages from the principal and choosing an optimal time to quit. A necessary condition for G to be implementable is that at every t , doing so would be no better than the continuation value promised by G :

$$V_t(G) := \mathbf{E}_{G_t} v(\tau - t, x - t).$$

In fact, these conditions are also sufficient.

Lemma 1. *Given a schedule G and a date t , let G_t be the continuation schedule. Let $V_{ni,t}$ denote the optimal continuation value for the agent in the absence of any additional information from the principal. If for every date $t \geq 0$,*

$$V_t(G) \geq V_{ni,t},$$

then G is implementable.

The logic behind [Lemma 1](#) is simple. The principal is using information to incentivize the agent. The principal can threaten to withhold all future planned information disclosures, leaving the agent to his own devices. If this is enough to dissuade the agent, then the principal can implement G .

5.1 The Schedule g^∞

Even though the threshold distribution has unbounded support, it is often possible to persuade the agent to continue working arbitrarily long and complete the task with probability 1. This is represented by the pure schedule g^∞ defined by $g^\infty(x) = x$. Falling short of that, the principal may try to implement the schedule g^t :

$$g^t(x) = \begin{cases} x, & \text{if } x \leq t \\ 0 & \text{otherwise.} \end{cases}$$

According to this schedule the agent quits immediately when the task is too difficult ($x > t$) but otherwise works just long enough to complete the task.

These schedules are the main elements of our analysis. They have the convenient property that, under Assumption (A2), individual rationality is not only necessary but sufficient for implementability.

Lemma 2. *For $t \in [0, \infty]$, the schedule g^t is implementable if and only if it is individually rational.*

To understand the lemma, take the schedule g^∞ .²⁵ Fix some positive time s and suppose that without information from s onward, the agent works for an additional time τ^0 . The agent's continuation value under g^∞ at time s consists of two parts: his value until τ^0 and his continuation value from working past τ^0 . The first part is larger than his no-information value, because he receives the reward with the same ex-ante probability,²⁶ but stops working earlier on average. Thus, g^∞ is implementable at s if the agent's continuation value at $s + \tau^0$ is positive. Under the increasing hazard rate assumption (A2) we can prove that the continuation value is non-decreasing over time.

²⁵The intuition for finite t is similar.

²⁶This probability is $\frac{F(\tau^0) - F(s)}{1 - F(s)}$.

Therefore if the schedule was individually rational at time zero, then its continuation value is non-negative at all future dates.

The principal prefers g^∞ to any finite g^t , since it extracts strictly more effort from the agent. However, for exactly that reason g^∞ may not be implementable. Hence the analysis divides into two cases depending on whether g^∞ is implementable.

5.2 Leading The Agent On

When g^∞ is implementable the principal can persuade the agent to work arbitrarily long to complete the task. Still, the schedule may provide the agent with some surplus, i.e. $V(g^\infty) > V_{ni}$, in which case it is possible for the principal to improve further.

Consider a different schedule in which the agent works with probability 1 for an initial duration t , after which he follows some implementable continuation schedule G . The total value to the agent is

$$e^{-rt}V(G) - c(1 - e^{-rt}) \quad (2)$$

which decreases from $V(G)$ to $-c$ as t increases. By choosing t we can design a schedule whose total value is exactly V_{ni} . As we show below this schedule is implementable and, when choosing G to be an initial disclosure followed by g^∞ , also optimal.

Proposition 6. *When g^∞ is implementable, the optimal mechanism is leading the agent on. There exists a t^* such that the following pure schedule is implementable and optimal for the principal.*

$$g^*(x) = \begin{cases} t^*, & \text{if } x \leq t^* \\ x & \text{otherwise.} \end{cases}$$

In particular the agent works until at least t^ which is (weakly) larger than the no-information*

effort level.

The mechanism that implements g^* is *leading the agent on*. The principal tells the agent nothing unless he first works for the duration t^* . At time t^* if the agent has worked throughout then the principal discloses whether the task has been completed (i.e. if $x \leq t^*$). If so the agent immediately quits. If not the principal proceeds to implement the continuation schedule g^∞ . At each subsequent point in time the agent is told only whether or not he has just reached the threshold, providing no information about any remaining effort required if not.²⁷

Here is why the policy is implementable. First of all, the schedule g^* reduces to the continuation schedule g^∞ beginning at time t^* , and as we have already shown in [Lemma 2](#), if g^∞ is individually rational it is also implementable.

All that remains is to show that the policy is implementable prior to time t^* . By [Lemma 1](#) this amounts to showing that its continuation value $V_t(G)$ exceeds the no-information continuation value $V_{ni,t}$ at each such time. By choosing t^* so that the expression in [Equation 2](#) is equal to V_{ni} , this is satisfied for $t = 0$. The main step of the proof involves showing that $V_t(G)$ and $V_{ni,t}$ cannot cross at any date in between. Intuitively this is true because $V_t(G)$ increases faster as time progresses because a larger reward is on the horizon.

The proof that g^* is optimal follows the lines from the binary example. By construction g^* provides the agent with exactly his no-information value. It is also efficient, implying that any schedule which yields a higher payoff for the principal cannot be individually rational for the agent. Here is a brief sketch of the proof that g^* is efficient. To incentivize effort the principal has but two instruments at his disposal. First the design of the schedule allocates effort across different realizations of the threshold. Any

²⁷This means that the agent's expectation about x is increasing in time. This is similar to bandit problems with perfect good news. See e.g. Keller, Rady and Cripps (2005). We are grateful to a referee for pointing out this connection.

given amount of effort is least costly for the agent, and hence most efficient, when it is allocated in a way that maximizes the probability of surpassing the threshold. With the schedule g^* , the agent completes the task with probability 1. However, g^* also induces the agent to spend additional, wasteful, effort that does not lead to additional rewards and indeed delays them. This is most efficiently incentivized by the second incentive instrument: the promise of valuable future disclosures. The schedule g^* delays any disclosure until after the agent has completed the initial t^* period of effort.

5.3 Moving The Goalposts

When $V(g^\infty)$ is negative, it is no longer individually rational for the agent to follow a schedule such as g^* in which he completes the task with probability 1. However, the schedule $V(g^{\bar{t}})$ is always implementable: it is the schedule in which the agent exactly completes all individually rational tasks. Indeed, as t ranges from \bar{t} to infinity, the payoff $V(g^t)$ is continuous and strictly decreasing (as more and more non-individually rational tasks are added to the schedule). Thus, there exists a finite $t^{**} > \bar{t}$ such that $V(g^{t^{**}}) = 0$.

Consider the following information policy. Before any effort decision the principal informs the agent whether the task is difficult ($x > t^{**}$) or easy ($x \leq t^{**}$). If the task is difficult the agent quits immediately. If the task is easy the agent commences working and at each future instant is informed only when he has completed the task. We show now that in response to this information policy the agent continues working until he completes the task and that the associated effort schedule

$$g^{**}(x) = \begin{cases} x, & \text{if } x \leq t^{**} \\ 0 & \text{otherwise} \end{cases}$$

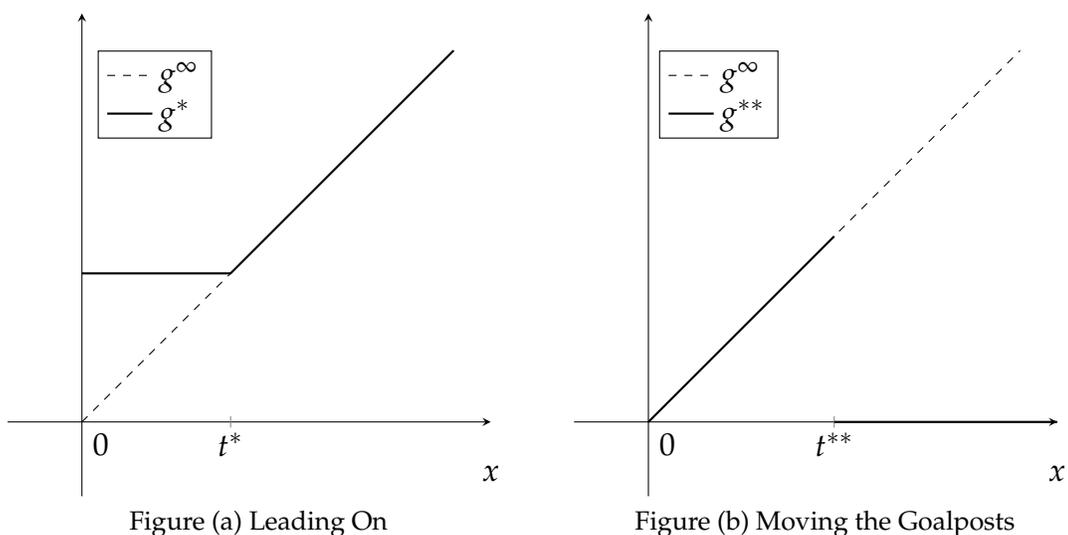


Figure 4: Optimal Policies

is optimal for the principal. We refer to this policy as *moving the goalposts* because of the way the principal uses it to shift around the agent's expectations of the difficulty of the task. It first drops expectations discretely (from $\mathbf{E}_F(x)$ to $\mathbf{E}_F(x \mid x \leq t^{**})$) when the principal suggests that the task is easy. The purpose is to encourage the agent to embark with the task rather than walk away as he would do in the absence of this disclosure. Indeed, it ultimately induces the agent to complete tasks even more difficult than the maximum individually rational task ($t^{**} > \bar{\tau}$). But once the agent begins working his expectations become more and more pessimistic as the principal repeatedly breaks the bad news that the task is not yet complete (i.e. $\mathbf{E}_F(x \mid t < x \leq t^{**})$).

Proposition 7. *When g^∞ is not implementable, there exists a $t^{**} > \bar{\tau}$ such that the schedule g^{**} is implementable and optimal and is achieved by the policy of moving the goalposts.*

We illustrate the leading on and moving the goalposts policies in [Figure 4](#).

6 Designing the Task

The binary example in [Section 4.1](#) shows that the principal benefits from a random threshold. In particular, by “splitting” a deterministic threshold into an easier task and a harder task, the principal can extract more effort on average from the agent. This requires that the incremental effort is individually rational so that the agent can be persuaded to keep working when the task turns out to be the harder one. Taking this one step further, each of these extreme tasks can be further split, resulting in a sequence of individually rational increments leading to higher overall expected effort. In this section we show how to exploit this possibility to the extreme by designing the optimal distribution over task difficulties. We do this in full generality: the principal can use any distribution and is not restricted to mean preserving spreads as in [Section 4.1](#). As we have seen in [Section 4.1](#), a random threshold is only beneficial when the principal is more patient than the agent, i.e. $r_p \leq r$. We maintain this condition throughout this section.

When the principal can design the distribution of thresholds, she has a two-stage optimization. First she chooses the task distribution F and then she chooses the optimal information policy, leading to the induced schedule G . However, we can simplify the problem by focusing on effort schedules in which with probability 1 the agent works just long enough to complete the task. That is, the schedule g^∞ of [Section 5.1](#) is optimal and $F = G$. This corresponds to using the leading on policy in [Proposition 6](#), with $t^* = 0$.

Since only the distribution over effort durations matters for the payoffs of principal and agent, we denote the choice of the principal by G .

Lemma 3. *A schedule G is implementable if and only if $V_t(G) \geq 0$ at all dates $t \geq 0$ such that $1 - F(t) > 0$. Without loss of generality, the optimal schedule is g^∞ and $F = G$.*

The following lemma formalizes the key intuition behind the design of random tasks. A direct schedule is a lottery over quitting times. Due to discounting and costly effort the agent is risk-loving with respect to these lotteries. Intuitively the agent is willing to increase the difficulty of already difficult tasks in exchange for reducing the difficulty of tasks that are already easy. On the other hand, the principal values effort and she therefore has the opposite preferences. She is effectively risk averse. By building riskier lotteries into the schedule the principal can increase the agent's value, then extract that value by asking the agent to work longer on average. As in [Section 4.1](#), this improves her value whenever $r_p < r$.

Lemma 4. *Suppose $r_p \leq r$ and that G is an individually rational schedule assigning positive mass to an interval (x_l, x_h) . Then there exists an individually rational schedule H in which all of the mass in the interval is moved to the endpoints and which is better for the principal, $W(H) \geq W(G)$. The inequality is strict when $r_p < r$.*

The lemma states that it is always possible to increase the principal's payoff by adding risk to the schedule without violating individual rationality for the agent: the agent will still be willing to begin working. However, if the resulting increment between successive tasks is too large, the schedule will not be implementable: it may not be possible to persuade the agent to continue. Thus it is the constraint of implementability that limits the extent to which [Lemma 4](#) can be exploited and ensures that an optimal schedule exists.²⁸

Lemma 5. *An optimal direct schedule exists.*

Having established existence we can complete the derivation of the optimal schedule using necessary conditions. By [Lemma 4](#) when the principal is strictly more

²⁸The complication is that the implementability constraint is a conditional expectation which is not continuous in G , raising doubts about compactness. However we can show that if $G_k \rightarrow G$ (in the weak topology) and $V_t(G) < 0$ then for some k there must be some (possibly different) time s such that $V_s(G_k) < 0$, so that the set of schedules satisfying $V_t(G) \geq 0$ at all dates $t \geq 0$ is indeed compact.

patient than the agent, adding risk to the effort schedule enables an improvement for the principal. Thus, if $r_p < r$, a distribution of task difficulties can be optimal only if there is no room for further increases in risk without violating implementability. We show that this implies that the agent's continuation value must be identically zero at every point in time. In particular, the task distribution cannot have atoms, it must have unbounded support, and it must have a constant hazard rate. With these conditions, we can find the optimal distribution.

Proposition 8. *The exponential distribution with hazard rate rc/R is optimal, uniquely so when $r_p < r$.*

Because in a direct schedule the agent earns the reward at time x , in order to provide a constant continuation value the distribution of x must have a constant hazard rate, i.e. F must be distributed exponentially. For an exponential distribution with parameter λ , the continuation value is easily calculated:

$$\frac{\lambda R - rc}{\lambda + r}$$

and setting $\lambda = rc/R$ yields a zero continuation value, hence the optimal F .

The gains identified in [Lemma 4](#) are strict when $r_p < r$.²⁹ difference in discount rates translates to a difference in tolerance to riskiness in effort schedules. In the case of equal discounting, the exponential distribution is among many optimal distributions including the deterministic one.

²⁹The increase in effort from using a random task difficulty can be large. The optimal deterministic threshold is the maximum individually rational effort, which equals $\frac{1}{r} \log\left(\frac{R}{c} + 1\right)$, whereas the expected effort under the exponential distribution is the inverse of the hazard rate, i.e. R/rc . The increase in expected effort is thus $\frac{R}{rc} - \frac{1}{r} \log\left(\frac{R}{c} + 1\right)$, which increases without bound as R becomes large.

References

- Aliprantis, C, and K Border.** 1999. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin.
- Aumann, Robert J, Michael Maschler, and Richard E Stearns.** 1995. *Repeated Games with Incomplete Information*. MIT press.
- Behrens, Alfredo.** 2015. "Are Bad Managers Holding Back Your Best Talent?" *Harvard Business Review*.
- Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman.** 2014. "Optimal Allocation with Costly Verification." *American Economic Review*, 104(12): 3779–3813.
- Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman.** 2017. "Mechanisms with Evidence: Commitment and Robustness." Working Paper.
- Bergemann, Dirk, and Stephen Morris.** 2016a. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics*, 11(2): 487–522.
- Bergemann, Dirk, and Stephen Morris.** 2016b. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *American Economic Review*, 106(5): 586–591.
- Bull, Jesse, and Joel Watson.** 2007. "Hard Evidence and Mechanism Design." *Games and Economic Behavior*, 58(1): 75–93.
- Conger, Jay A, and Robert M Fulmer.** 2003. "Developing Your Leadership Pipeline." *Harvard Business Review*, 81(12): 76–85.
- Dekel, Eddie.** 2017. "On Evidence Disclosure and Mechanism Design."

- De la Croix, David, Matthias Doepke, and Joel Mokyr.** 2017. "Clans, Guilds, and Markets: Apprenticeship Institutions and Growth in the Preindustrial Economy." *The Quarterly Journal of Economics*, 133(1): 1–70.
- Elbaum, Bernard.** 1989. "Why Apprenticeship Persisted in Britain but not in the United States." *The Journal of Economic History*, 49(2): 337–349.
- Ely, Jeffrey C.** 2017. "Beeps." *American Economic Review*, 107(1): 31–53.
- Epstein, Stephan R.** 1998. "Craft Guilds, Apprenticeship, and Technological Change in Preindustrial Europe." *The Journal of Economic History*, 58(3): 684–713.
- Fang, Hanming, and Giuseppe Moscarini.** 2005. "Morale Hazard." *Journal of Monetary Economics*, 52(4): 749–777.
- Fudenberg, Drew, and Luis Rayo.** 2017. "Training and Effort Dynamics in Apprenticeship."
- Garicano, Luis, and Luis Rayo.** 2017. "Relational Knowledge Transfers." *American Economic Review*, 107(9): 2695–2730.
- Georgiadis, George, Steven A Lippman, and Christopher S Tang.** 2014. "Project Design with Limited Commitment and Teams." *RAND Journal of Economics*, 45(3): 598–623.
- Hansen, Stephen E.** 2012. "Performance Feedback with Career Concerns." *Journal of Law, Economics, & Organization*, 29(6): 1279–1316.
- Hartman, Philip.** 2002. *Ordinary Differential Equations. Classics in Applied Mathematics*, SIAM.
- Hörner, Johannes, and Nicolas S. Lambert.** 2016. "Motivational Ratings."

- Humphries, Jane.** 2003. "English Apprenticeship: A Neglected Factor in the First Industrial Revolution." In *The Economic Future in Historical Perspective.* , ed. Paul A. David and Mark Thomas, 73–102. Oxford University Press.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.
- Keller, Godfrey, Sven Rady, and Martin Cripps.** 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73(1): 39–68.
- Minns, Chris, and Patrick Wallis.** 2012. "Rules and Reality: Quantifying the Practice of Apprenticeship in Early Modern England." *The Economic History Review*, 65(2): 556–579.
- Mokyr, Joel.** 2018. "The Economics of Apprenticeship." In *Apprenticeship through the Ages (prelim. title).* , ed. Maarten Prak and Patrick Wallis, forthcoming. Cambridge University Press.
- Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov.** 2016. "Persuading the Regulator To Wait." *Available at SSRN 2767234.*
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille.** 2017. "Optimal Dynamic Information Provision." *Games and Economic Behavior*, 104: 329–349.
- Rothwell, William J.** 2010. *Effective succession planning: Ensuring leadership continuity and building talent from within.* . 4th ed., AMACOM American Management Association.
- Rousseau, Denise.** 1995. *Psychological Contracts in Organizations: Understanding Written and Unwritten Agreements.* Sage Publications.

- Sannikov, Yuliy.** 2008. "A Continuous-Time Version of the Principal-Agent Problem." *The Review of Economic Studies*, 75(3): 957–984.
- Smolin, Alex.** 2017. "Dynamic Evaluation Design."
- Spear, Stephen E, and Sanjay Srivastava.** 1987. "On Repeated Moral Hazard with Discounting." *The Review of Economic Studies*, 54(4): 599–617.
- Stephan, Paula.** 2013. "How to Exploit Postdocs." *BioScience*, 63(4): 245–246.
- Stephan, Paula, and Jennifer Ma.** 2005. "The Increased Frequency and Duration of the Postdoctorate Career Stage." *American Economic Review*, 95(2): 71–75.
- Varas, Felipe, Iván Marinovic, and Andrzej Skrzypacz.** 2017. "Random Inspections and Periodic Reviews: Optimal Dynamic Monitoring."
- Wallis, Patrick.** 2008. "Apprenticeship and Training in Premodern England." *The Journal of Economic History*, 68(3): 832–861.

A Appendix

A.1 Proofs for Section 4

For the binary example we prove the results for the case of equal discounting $r = r_p$ because the specific method of proof previews the techniques for the general model of Section 5 where equal discounting is assumed. When $r_p < r$ we can use splitting arguments such as in the proof of Proposition 4 and Lemma 4 to show that an optimal policy must induce effort levels in $\{0, x_l, x_h\}$ and then proceed with the arguments below.

Proof of Proposition 1. This is a standard static Bayesian persuasion problem (see Kamenica and Gentzkow (2011) and Aumann, Maschler and Stearns (1995)) with effectively two actions for the agent, $\tau = 0$ and $\tau = x_l$. No other effort level would be optimal for the agent against any belief. In particular the agent's optimal choice rule is $\tau = 0$ if $\mu < \bar{\mu}$ and $\tau = x_l$ if $\mu \geq \bar{\mu}$. The principal's value function with respect to the agent's belief is depicted in Figure 5 below as the solid line segments. The concavification is depicted as the dashed line.

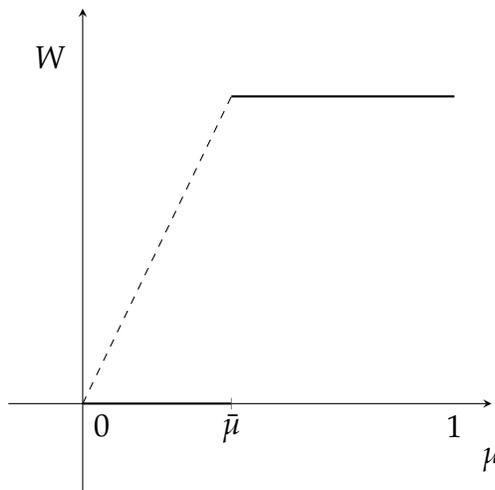


Figure 5: Static Bayesian Persuasion

As is standard for these problems, the concavification illustrates the principal's optimal value from maximizing the probability the agent chooses the principal's preferred action, here $\tau = x_l$. When the prior exceeds $\bar{\mu}$ the agent chooses x_l with probability 1 without any persuasion. When the prior is less than $\bar{\mu}$ the optimal persuasion sends two messages inducing the two beliefs $\mu = 0$ and $\mu = \bar{\mu}$, the latter message inducing the response $\tau = x_l$. \square

Proof of Proposition 2. The mechanism described in the text leads to the following *effort schedule*, i.e. joint distribution over thresholds x and effort durations τ :

	$x = x_l$	$x = x_h$
$\tau = 0$	0	0
$\tau = x_l$	α	0
$\tau = x_h$	β	$1 - \mu$

for some $\alpha, \beta > 0$ with $\alpha + \beta = \mu$. Generally, a feasible effort schedule is some joint distribution whose total probability of x_l equals μ . Any implementable policy induces such a schedule (possibly involving more than just these three effort levels). Any schedule which is better for the principal must have a larger expected discounted effort duration. Since under the schedule above the agent completes the task with probability 1, increasing expected discounted effort only adds cost to the agent and delays rewards. Therefore the schedule above is efficient among all feasible schedules. Since it was constructed to yield zero expected payoff for the agent, no individually rational policy can improve for the principal. \square

Proof of Proposition 3. Moving the goalposts implements the following effort schedule:

	$x = x_l$	$x = x_h$
$\tau = 0$	0	α
$\tau = x_l$	μ	0
$\tau = x_h$	0	β

where $\alpha, \beta > 0, \alpha + \beta = 1 - \mu$.

To show that the schedule is efficient, it suffices to show that there is no alternative schedule with the same expected discounted effort but higher expected payoff for the agent. To see why, consider any schedule with a strictly greater payoff for the principal (i.e. expected discounted effort) and weakly higher expected payoff for the agent. Then there must be some excess effort: i.e. a positive probability of working for a positive duration not equal to the threshold. Eliminate enough excess effort to reduce the total expected effort to that of the moving the goalposts schedule, raising the agent's expected payoff.

Now note that the implemented schedule completes the task with probability 1 conditional on the task being individually rational. This means that the effort is efficiently allocated. That is, any alternative schedule with the same expected discounted effort must give the agent a lower expected payoff either because of delayed rewards or a lower probability of completing the task, or both. Thus, the schedule is efficient.

Moreover it gives the agent zero expected utility implying that the policy is optimal. This was already shown for priors greater than $\tilde{\mu}$. For priors less than $\tilde{\mu}$ note that the moving the goalposts mechanism leads to two possible initial beliefs: $\tilde{\mu}$ and 0. Quitting immediately is an optimal strategy for the agent at either of these beliefs and therefore the agent's expected payoff is zero. □

A.2 Proofs for Section 4.1

Proof of Proposition 4. When the threshold is randomized and full delayed disclosure is used, the agent's expected payoff increases by

$$(c + R) \left[\frac{(e^{-rx_l} + e^{-rx_h})}{2} - e^{-r\bar{t}} \right] := (c + R) \Delta_a^1,$$

whereas the principal's payoff decreases by

$$[1 - e^{-rp\bar{t}}] - \left[1 - \frac{1}{2} (e^{-rp x_l} + e^{-rp x_h}) \right] = \frac{(e^{-rp x_l} + e^{-rp x_h})}{2} - e^{-rp\bar{t}} := \Delta_p^1.$$

When partial delayed disclosure is added, the agent shifts effort probability from x_l to x_h . Due to increased effort and delayed rewards, the agent's payoff declines at rate

$$(c + R) [e^{-rx_l} - e^{-rx_h}] := (c + R) \Delta_a^2.$$

Whereas the principal's payoff increases at rate

$$(1 - e^{-rp x_h}) - (1 - e^{-rp x_l}) = e^{-rp x_l} - e^{-rp x_h} := \Delta_p^2.$$

Note that partial delayed disclosure ranges from full disclosure to no disclosure. Under the extreme of no disclosure the agent arrives at date x_l assigning probability 1/2 to threshold x_h . By Equation 1 it would then be optimal for the agent to continue working (with probability 1) to x_h . Since x_h is not individually rational this would give the agent a negative ex ante payoff. Thus, in the range between full delayed disclosure and no disclosure there exists a partial delayed disclosure policy which gives the agent exactly zero ex ante payoff. See point C in Figure 3.

Therefore, to prove the Proposition it is enough to show that that $\Delta_p^1 / (c + R) \Delta_a^1 \leq$

$\Delta_p^2 / (c + R)\Delta_a^2$, or equivalently that $\Delta_p^1 / \Delta_p^2 \leq \Delta_a^1 / \Delta_a^2$ with strict inequality when $r_p < r$.

For any discount rate r , define

$$\Delta^1 = \left(\frac{e^{-rx_l} + e^{-rx_h}}{2} \right) - e^{-r\tilde{\tau}}$$

$$\Delta^2 = e^{-rx_l} - e^{-rx_h}$$

Using the definition of $\tilde{\tau}$, these may be rewritten as follows

$$\Delta^1 = e^{-r\tilde{\tau}} \left[\left(\frac{e^{r\tilde{\tau}} + e^{-r\tilde{\tau}}}{2} \right) - 1 \right]$$

$$\Delta^2 = e^{-r\tilde{\tau}} (e^{r\tilde{\tau}} - e^{-r\tilde{\tau}})$$

so that the ratio equals

$$\frac{\Delta^1}{\Delta^2} = \frac{(e^{r\tilde{\tau}} + e^{-r\tilde{\tau}}) - 2}{2(e^{r\tilde{\tau}} - e^{-r\tilde{\tau}})}.$$

The derivative with respect to r has the same sign as

$$(e^{r\tilde{\tau}} - e^{-r\tilde{\tau}})^2 - (e^{r\tilde{\tau}} + e^{-r\tilde{\tau}}) [(e^{r\tilde{\tau}} + e^{-r\tilde{\tau}}) - 2]$$

or

$$\left[(e^{r\tilde{\tau}} - e^{-r\tilde{\tau}})^2 - (e^{r\tilde{\tau}} + e^{-r\tilde{\tau}})^2 \right] + 2(e^{r\tilde{\tau}} + e^{-r\tilde{\tau}}).$$

Notice that the term in square brackets equals negative four, hence the derivative has the same sign as $(e^{r\tilde{\tau}} + e^{-r\tilde{\tau}}) - 2$ which is positive for all $r > 0$. This shows that $\Delta_p^1 / \Delta_p^2 \leq \Delta_a^1 / \Delta_a^2$ with a strict inequality when $r_p < r$. \square

B Proofs for Section 5

Proof of Proposition 5. The “only if” part is immediate. If a policy puts positive probability on $x > \tau > 0$ we can take all such x , and then split the probability mass between $\tau = 0$ and $\tau = x$. Doing so increases the surplus, since now x is reached with a larger probability. If a policy puts positive probability on $\bar{\tau} \geq x > \tau = 0$, we can pick $\tau = x$ instead for any such pair. To show the “if” part, consider a policy G which either makes the agent quit immediately, but only if $x > \bar{\tau}$, or puts probability one on $\tau \geq x$. Using the definitions of the principal’s and agent’s values, we can write $V(G) = \mathbf{E}_G [e^{-r\tau}R - c(1 - e^{-r\tau})]$ and $W(G) = \mathbf{E}_G [1 - e^{-r\tau}]$. Rearranging these expressions yields $V(G) = R - (R + c)W(G)$. Thus, any such schedule is efficient. \square

To prove the next propositions, we need to begin with some preliminaries. Define for any t , $v(t) = v(t, t) = e^{-rt}R - c(1 - e^{-rt})$, i.e. the value of exactly completing the task when the level of difficulty is t . Note that $v(t)$ is positive for $t = 0$, decreasing in t , and crosses zero at $\bar{\tau}$, the maximum individually-rational threshold. We have for any t , including $t = \infty$,

$$V(g^t) = \frac{1}{F(t)} \left[\int_0^t v(\hat{t})f(\hat{t})d\hat{t} \right].$$

As a function of t , $V(g^t)$ is strictly increasing up to $\bar{\tau}$ then decreasing. Similarly, for $s < t$ the agent’s continuation value satisfies

$$V_s(g^t) = \frac{1}{F(t) - F(s)} \left[\int_s^t v(\hat{t} - s)f(\hat{t})d\hat{t} \right].$$

Lemma 6. $V_s(g^t)$ is weakly increasing in s .

Proof. First, consider $V_s(g^t)$ for some $s < t < \infty$. Taking derivatives, $\frac{\partial}{\partial s} V_s(g^t) = V_s(g^t)(r + H(s)) - (H(s)R - rc)$ and $\frac{\partial^2}{\partial s^2} V_s(g^t) = \frac{\partial}{\partial s} V_s(g^t)(r + H(s)) - (R - V_s(g^t))H'(s)$. We know that $V_t(g^t) = R$ and that $V_s(g^t) < R$, since the agent incurs effort cost be-

tween s and t . Suppose that $\frac{\partial}{\partial s} V_s(g^t) < 0$. Then, since H is increasing by Assumption A2, $\frac{\partial^2}{\partial s^2} V_s(g^t) < 0$ and hence for all times $u \in [s, t]$, $V_u(g^t) < V_s(g^t) < 0$. But then $V_t(g^t) < V_s(g^t) < R$, which is a contradiction. Similarly, if $\frac{\partial}{\partial s} V_s(g^\infty) < 0$ then for any $u > s$ $V_u(g^\infty) < (u - s) \frac{\partial}{\partial s} V_s(g^\infty) + V_s(g^\infty)$, which crosses $-c$ as u becomes sufficiently large. This is a contradiction since $V_u(g^\infty) \geq -c$, i.e. the agent cannot do worse than to work forever and never get the reward.³⁰ \square

Proof of Lemma 2. Consider any date u . For any $u < \tau^0 \leq t$ we can decompose $V_u(g^t)$ into

$$V_u(g^t) = \int_u^{\tau^0} v(\hat{t})f(\hat{t})d\hat{t} + \frac{F(t) - F(\tau^0)}{F(t) - F(u)} V_{\tau^0}(g^t).$$

Suppose that at date u the agent's optimal no-information continuation strategy is to work until τ^0 and then quit. The first term above is the value of exactly completing the task when its difficulty is less than τ^0 . This is larger than the agent's no-information continuation value $V_{ni,u}$ since the latter completes the task with the same probability but later on average.

Thus, a sufficient condition for $V_u(g^t) \geq V_{ni,u}$ is that $V_{\tau^0}(g^t)$ is positive. If g^t is individually rational then $V(g^t) = V_0(g^t) \geq V_{ni} \geq 0$, so this follows directly from Lemma 6. The argument for $V_s(g^\infty)$ is identical and hence omitted. \square

Proof of Proposition 6. Let t^* equate $e^{-rt^*} [F(t^*)R + (1 - F(t^*))V_{t^*}(g^\infty)] - c(1 - e^{-rt^*}) = V_{ni}$, and consider the pure schedule g^* defined in the statement of the Proposition.³¹ Note that t^* exceeds the no-information optimal effort τ^0 . This is because the left-hand

³⁰Note here that we have normalized the flow cost of effort to be rc . The cost of working forever is thus $\frac{rc}{r} = c$.

³¹The LHS of the above equation is strictly decreasing in t , so the value t^* at which it equals V_{ni} is unique. We can see this by rewriting it as

$$(R + c) \left(e^{-rt} F(t) + \int_t^\infty e^{-rs} f(s) \right) ds - c$$

and calculating the derivative.

side evaluated at any $t^* \leq \tau^0$ is strictly larger than V_{ni} .

We first show that g^* is implementable. Note that after time t^* , the policy g^* is identical to g^∞ . Since g^∞ is implementable, $V_t(g^*) \geq V_{ni,t}$ for all $t \geq t^*$. Next consider dates t earlier than t^* , and note that g^* provides no information prior to t^* . The agent's no-information continuation value is therefore

$$V_{ni,t} = \begin{cases} e^{-r(\tau^0-t)} F(\tau^0)R - c \left(1 - e^{-r(\tau^0-t)}\right) & \text{if } t \leq \tau^0 \\ F(t)R & \text{if } \tau^0 \leq t < t^* \end{cases}$$

because prior to τ^0 the agent plans to continue working until τ^0 , and after τ^0 the agent will quit immediately. (By Assumption (A1), past τ^0 the marginal payoff from additional work is strictly negative.)

As for $V_t(g^*)$, the continuation value provided by the schedule g^* , we have for $t < t^*$

$$V_t(g^*) = e^{-r(t^*-t)} [F(t^*)R + (1 - F(t^*))V_{t^*}(g^\infty)] - c \left(1 - e^{-r(t^*-t)}\right).$$

The analysis now divides into two cases: when τ^0 is positive and when τ^0 is zero. First, suppose that $\tau^0 > 0$. Recall that t^* is chosen such that $V(g^*) = V_{ni}$. For $t < \tau^0$, we calculate the following derivatives: $\frac{\partial}{\partial t} V_t(g^*) = r(V_t(g^*) + c)$ and $\frac{\partial}{\partial t} V_{ni,t} = r(V_{ni,t} + c)$. Since by construction, $V(g^*) = V_{ni}$, this ensures that $V_t(g^*) = V_{ni,t}$ for all $t \leq \tau^0$. We now extend the argument to $t \geq \tau^0$. We can express $\frac{\partial}{\partial t} V_t(g^*)$ as $\frac{\partial}{\partial t} V_t(g^*) = re^{rt} (V(g^*) + c) = re^{rt} (V_{ni} + c)$, while the derivative of the no-information value is simply $\frac{\partial}{\partial t} V_{ni,t} = f(t)R$. From the definition of V_{ni} , it follows that $\frac{\partial}{\partial t} V_t(g^*) = re^{r(t-\tau^0)} (F(\tau^0)R + c)$. By Assumption (A1), the no-information effort τ^0 is the first

time marginal value of additional effort is zero, i.e.

$$f(\tau^0)R - r \left(F(\tau^0)R + c \right) = 0.$$

Plugging this in, we obtain $\frac{\partial}{\partial t} V_t(g^*) = e^{r(t-\tau^0)} f(\tau^0)R$. Thus, $\frac{\partial}{\partial t} V_t(g^*) \geq \frac{\partial}{\partial t} V_{ni,t}$ whenever $e^{r(t-\tau^0)} f(\tau^0) \geq f(t)$. This inequality holds because of Assumption (A1) and Grönwall's Lemma.³² Therefore $V_t(g^*) \geq V_{ni,t}$ on $[\tau^0, t^*]$. Thus, g^* is implementable if $\tau^0 > 0$.

If $\tau^0 = 0$, we can use a similar proof as above. Specifically, we have $V_{ni,t} = F(t)R$ for all $t < t^*$ and as above $\frac{\partial}{\partial t} V_t(g^*) = re^{rt} (V_{ni} + c) = re^{rt}c$, because $V_{ni} = 0$. By Assumption (A1), we have $\frac{\partial}{\partial t} V_{ni,t} = Rf(t) \leq Re^{rt}f(0)$. Note that $\tau^0 = 0$ only if $rc > Rf(0)$ (the marginal value of effort at time zero is negative). We again have $\frac{\partial}{\partial t} V_t(g^*) \geq \frac{\partial}{\partial t} V_{ni,t}$.

Having established that g^* is implementable, it remains to show that g^* is also optimal. Since g^* is exactly individually rational, we can establish that g^* is optimal by showing it is efficient. Suppose on the contrary that there was a schedule G such that $W(G) > W(g^*)$ and $V(G) \geq V(g^*)$. Since $W(G) = \mathbf{E}_G(1 - e^{-rt})$ it must be that $\mathbf{E}_G e^{-rt} < \mathbf{E}_{g^*} e^{-rt}$. Consider the pure schedule \tilde{g} defined by $\tilde{g}(x) = g^*(x) + z$ where the constant $z > 0$ is chosen so that $\mathbf{E}_{\tilde{g}} e^{-rt} = \mathbf{E}_G e^{-rt}$, and in particular $W(\tilde{g}) = W(G)$. The new schedule \tilde{g} completes the task with probability 1 since g^* does and therefore $V(\tilde{g}) = \mathbf{E}_{\tilde{g}} [e^{-rt}R - c(1 - e^{-rt})] = \mathbf{E}_G [e^{-rt}R - c(1 - e^{-rt})]$. Moreover, $V(G) = \mathbf{E}_G [1_{t \geq x} e^{-rt}R - c(1 - e^{-rt})] \leq \mathbf{E}_G [e^{-rt}R - c(1 - e^{-rt})]$ and therefore \tilde{g} dominates g^* as well. But that's clearly impossible by the definition of \tilde{g} . \square

Proof of Proposition 7. We begin by proving that g^{**} is implementable. Using Lemma 1 we know that it is sufficient for its continuation value to always exceed the continuation no-information value. Since t^{**} is chosen such that $V(g^{**}) = V_{ni}$, the result follows

³²See e.g. Hartman (2002), p.24. The Lemma states that whenever $f'(s) \leq rf(s)$ on some interval $[a, b]$, then $f(s) \leq f(a)e^{rs}$. Taking $a = \tau^0$ and $s = t - \tau^0$ yields the result.

from Lemma 2.

Next we prove that g^{**} is efficient which will establish that it is optimal (since it is exactly individually rational, i.e. $V(g^{**}) = 0$). Note that for any schedule G , $V(G) = \mathbf{E}_G [1_{t \geq x} e^{-rt} R - c(1 - e^{-rt})] \leq \mathbf{E}_G [e^{-rx} R - c(1 - e^{-rt})]$, where $1_{t \geq x}$ is the indicator function for having completed the task. For the particular schedule g^{**} , with probability one, either $t = x$ or $t = 0$. Therefore $V(g^{**}) = \mathbf{E}_{g^{**}} [e^{-rx} R - c(1 - e^{-rt})]$. Now if G is any schedule for which $W(G) > W(g^{**})$ then $\mathbf{E}_G e^{-rt} < \mathbf{E}_{g^{**}} e^{-rt}$ and therefore $V(G) \leq \mathbf{E}_G [e^{-rx} R - c(1 - e^{-rt})] < \mathbf{E}_{g^{**}} [e^{-rx} R - c(1 - e^{-rt})] = V(g^{**})$. This proves that g^{**} is efficient. Finally, $t^{**} > \bar{\tau}$ because $V_t(g^{**})$ is strictly increasing in t over $[0, \bar{\tau}]$ and $V(g^{**}) = 0$. Since t^{**} is defined by $V(t^{**}) = 0$ we must have $t^{**} > \bar{\tau}$. \square

C Proofs for Section 6

Proof of Lemma 3. The only-if part is obvious. Suppose that G is a direct schedule satisfying the conditions given. Consider the information policy consisting of two signals q (for “quit”) and s (for “stay”) such that at each date t , the principal sends the signal q whenever $x \leq t$ and otherwise the signal s . The strategy for the agent of quitting immediately in response to q and otherwise continuing results in the schedule G and thus yields non-negative continuation value at every date. Indeed this strategy is optimal. There can be no improvement after receiving the signal q , because conditional on q the agent knows that the task is complete. The only other deviation would be to quit after receiving s but this would yield a continuation payoff of zero which cannot improve upon the non-negative continuation value from continuing.

To prove the second claim, let \tilde{G} be any implementable schedule and consider the direct schedule G whose marginal over effort duration is the same as \tilde{G} . Since the

principal's payoff depends only on effort we have $W(G) = W(\tilde{G})$. Moreover at every date t , $V_t(G) \geq V_t(\tilde{G}) \geq 0$. We have the first inequality because G entails the same effort costs but provides the reward with probability 1 and with no delay. We have the second inequality because \tilde{G} is implementable and therefore provides at least the no-information continuation value at each date. It now follows that G is implementable. \square

We prove the following more specific version of [Lemma 4](#).

Lemma 7. *If G is a schedule assigning positive mass to an interval (x_l, x_h) then there exists another schedule H which is identical to G outside the interval $[x_l, x_h]$ and for which $V(H) = V(G)$ and $W(H) \geq W(G)$ with a strict inequality when $r_p < r$.*

Proof. The agent's expected payoff from any schedule G is $V(G) = \mathbf{E}_G [e^{-rt}R - c(1 - e^{-rt})] = \mathbf{E}_G [e^{-rt} (R + c)] - c$ while the principal's is $W(G) = \mathbf{E}_G(1 - e^{-r_p t})$. Applying positive affine transformations we can represent these preferences equivalently by

$$\tilde{V}(G) = \mathbf{E}_G \tilde{v}(t)$$

$$\tilde{W}(G) = -\mathbf{E}_G \tilde{w}(t)$$

where $\tilde{v}(t) = \frac{e^{-rt} - e^{-rx_h}}{e^{-rx_l} - e^{-rx_h}}$ and $\tilde{w}(t) = \frac{e^{-r_p t} - e^{-r_p x_h}}{e^{-r_p x_l} - e^{-r_p x_h}}$. In particular $\tilde{V}(G) \geq \tilde{V}(G')$ if and only if $V(G) \geq V(G')$ and likewise for the principal.

Suppose G attaches positive probability to the interval (x_l, x_h) . Consider schedules that are identical to G except that all of the mass inside the interval (x_l, x_h) is moved to atoms at the endpoints $\{x_l, x_h\}$. Among such schedules, the one with all of the mass at x_h is strictly worse for the agent than G and the one with all of the mass at the low end is strictly better. Moreover the agent's payoff increases continuously as this mass moves from x_h to x_l and therefore there exists a schedule H such that the agent is indifferent between G and H .

Similarly there exists among these a schedule J such that $\mathbf{E}J = \mathbf{E}G$. We have the following identity.

$$\begin{aligned}\tilde{V}(H) &= \overbrace{[\tilde{V}(H) - \tilde{V}(J)]}^{\Delta_a^2} + \overbrace{[\tilde{V}(J) - \tilde{V}(G)]}^{\Delta_a^1} + \tilde{V}(G) \\ \tilde{W}(H) &= \underbrace{[\tilde{W}(H) - \tilde{W}(J)]}_{\Delta_p^2} + \underbrace{[\tilde{W}(J) - \tilde{W}(G)]}_{\Delta_p^1} + \tilde{W}(G).\end{aligned}$$

We will show that because $r_p \leq r$, $\Delta_p^2 + \Delta_p^1 \geq \Delta_a^2 + \Delta_a^1$ and therefore since $\tilde{V}(H) = \tilde{V}(G)$ we also have $\tilde{W}(H) \geq \tilde{W}(G)$ and the inequality will be strict when $r_p < r$.

Note that Δ_a^1 is positive because \tilde{v} is strictly convex. That means Δ_a^2 is negative. In particular $\tilde{V}(H) < \tilde{V}(J)$. Since H and J differ only in how the mass is divided between the points x_l and x_h , it follows that H has a larger mass at x_h . It follows that Δ_p^2 is positive and indeed by our normalization of payoffs, $\Delta_p^2 = -\Delta_a^2$ since according to the normalized payoff functions \tilde{v} and \tilde{w} any shift of mass from x_l to x_h is a one-for-one transfer of utility from agent to principal.

Thus, to show that $\tilde{W}(H) \geq \tilde{W}(G)$ it is now enough to show that $|\Delta_p^1| \leq \Delta_a^1$, i.e. the principal's loss from the mean-preserving spread is smaller than the agent's gain with a strict inequality when $r_p < r$. The inequality expands to $|\tilde{W}(J) - \tilde{W}(G)| \leq \tilde{V}(J) - \tilde{V}(G)$ or $|\mathbf{E}_G \tilde{w} - \mathbf{E}_J \tilde{w}| \leq \mathbf{E}_J \tilde{v} - \mathbf{E}_G \tilde{v}$.

Since \tilde{w} is convex, $\mathbf{E}_J \tilde{w} \geq \mathbf{E}_G \tilde{w}$ and so the inequality is equivalent to $\mathbf{E}_J \tilde{w} - \mathbf{E}_G \tilde{w} \leq \mathbf{E}_J \tilde{v} - \mathbf{E}_G \tilde{v}$ or $\mathbf{E}_J (\tilde{w} - \tilde{v}) \leq \mathbf{E}_G (\tilde{w} - \tilde{v})$. Since the schedules J and G are identical outside of the interval $[x_l, x_h]$, this reduces to $\mathbf{E}_J (\tilde{w} - \tilde{v} \mid x_l \leq x \leq x_h) \leq \mathbf{E}_G (\tilde{w} - \tilde{v} \mid x_l \leq x \leq x_h)$. Furthermore, conditional on the interval $[x_l, x_h]$, the schedule J assigns all of its mass to the points x_l and x_h where the two functions \tilde{w} and \tilde{v} are equal. Thus, the left-hand side is zero and we need only to show that $\mathbf{E}_G (\tilde{w} - \tilde{v} \mid [x_l, x_h]) \geq 0$. In fact \tilde{w} pointwise dominates \tilde{v} on the interval as we now show.

Since $\tilde{w}(x_l) = \tilde{v}(x_l)$ and $\tilde{w}(x_h) = \tilde{v}(x_h)$ and both functions are decreasing and continuously differentiable, we can show that $\tilde{w} \geq \tilde{v}$ on the entire interval by showing that $\tilde{w}'(\cdot) - \tilde{v}'(\cdot)$ is decreasing. Because if $\tilde{v}(t) > \tilde{w}(t)$ for some $t \in [x_l, x_h]$ then $\tilde{v}'(s) \geq \tilde{w}'(s)$ for some $s \leq t$ but if $\tilde{w}'(\cdot) - \tilde{v}'(\cdot)$ is continuous and decreasing then $\tilde{v}'(s') \geq \tilde{w}'(s')$ for all $s' \geq t$ implying $\tilde{v}(x_h) > \tilde{w}(x_h)$, a contradiction.

Now $\tilde{w}'(\cdot) - \tilde{v}'(\cdot)$ is decreasing if $\tilde{w}''(\cdot) \leq \tilde{v}''(\cdot)$ for all $t \in [x_l, x_h]$ and this inequality follows immediately because $r \geq r_p$ and the inequality is strict when $r > r_p$. \square

Proof of Lemma 5. A direct schedule has $x = \tau$ with probability 1, therefore the set of all direct schedules can be identified with the set of probability measures on the non-negative real numbers. The principal's value for a direct schedule is $W(G)$ which, viewed as function of G , is continuous in the weak topology. An optimum exists provided the feasible set, i.e. the subset of implementable schedules, is compact in that topology.

A schedule is implementable only if it is individually rational $V(G) \geq 0$. We begin by showing that the subset of individually rational schedules G is compact. Consider any $\varepsilon > 0$. Define $t(\varepsilon)$ to satisfy $R + \varepsilon \cdot e^{-rt(\varepsilon)}v(t(\varepsilon)) = 0$. There is no individually rational schedule assigning greater than ε probability to thresholds $t(\varepsilon)$ or higher. This shows that the set of individually rational schedules is *tight* and therefore relatively compact.³³ It is also closed (and hence compact) as the set of measures giving a non-negative expected value of a continuous payoff function.

To show that the subset of implementable schedules is compact it is now enough to show that it is closed. Consider a sequence G_k of implementable schedules converging to a schedule G . Let t be any date such that $1 - G(t) > 0$. Then also³⁴ for all k sufficiently large and for all $\varepsilon > 0$, sufficiently small, $1 - G_k(t + \varepsilon) > 0$. Since G_k is

³³See (Aliprantis and Border, 1999, Section 12.5).

³⁴See (Aliprantis and Border, 1999, Section 14.5)

implementable, $V_{t+\varepsilon}(G_k) \geq 0$. Since $V_{t+\varepsilon}(G_k) = \mathbf{E}_{G_k}(v(x) \mid x > t + \varepsilon)$ we have³⁵

$$(1 - G_k(t + \varepsilon)) V_{t+\varepsilon}(G_k) + G_k(\{t + \varepsilon\}) \cdot R = \int_{t+\varepsilon}^{\infty} v(s - t - \varepsilon) dG_k(s)$$

and therefore the right-hand side is non-negative. Moreover, the integral of a continuous function over a closed interval is upper-semicontinuous³⁶ in G_k and thus

$$\int_{t+\varepsilon}^{\infty} v(s - t - \varepsilon) dG(s) \geq 0.$$

Now

$$V_t(G) \geq \frac{1}{1 - G(t)} \left[\int_{t+\varepsilon}^{\infty} e^{-r(s-t)} R - c(1 - e^{-r(s-t-\varepsilon)}) dG(s) - c(1 - e^{-r\varepsilon}) \right]$$

or

$$V_t(G) \geq \frac{1}{1 - G(t)} \int_{t+\varepsilon}^{\infty} v(s - t - \varepsilon) dG(s) - \delta(\varepsilon)$$

where $\delta(\varepsilon) = c(1 - e^{-r\varepsilon}) + R \int_{t+\varepsilon}^{\infty} e^{-r(s-t-\varepsilon)} - e^{-r(s-t)} dG(s)$. Since the integral is non-negative, $V_t(G) \geq -\delta(\varepsilon)$. This is true for all sufficiently small ε and moreover $\delta(\varepsilon)$ vanishes as ε does, hence $V_t(G) \geq 0$ for all t such that $1 - G(t) > 0$ proving that G is implementable and the set of implementable schedules is closed and hence compact. \square

Lemma 8. *When $r_p < r$ an optimal schedule must be atomless.*

Proof. Suppose G is implementable and has an atom at time t . For any $\varepsilon > 0$, the continuation value at time $t - \varepsilon$ is at least $\delta e^{-r\varepsilon} R - c(1 - e^{-r\varepsilon})$ where δ is the size of the atom. This is bounded away from zero for ε small enough. Pick such an ε and consider

³⁵Recall that we have defined the notation $v(x) = e^{-rx} R - c(1 - e^{-rx})$ as the payoff from exactly completing the task. Also, $G_k(\{t + \varepsilon\})$ is the mass of a potential atom at $t + \varepsilon$.

³⁶See (Aliprantis and Border, 1999, Section 14.5)

the interval $[t - \varepsilon, t + \varepsilon]$. We split a fraction ν of the atom between $t - \varepsilon$ and $t + \varepsilon$ so that we make the agent indifferent at $t - \varepsilon$. That is, if νq is the fraction of mass moved to $t - \varepsilon$, we pick q such that $e^{-r\varepsilon}R - c(1 - e^{-r\varepsilon}) = qR + (1 - q)(e^{-r2\varepsilon}R - c(1 - e^{-r2\varepsilon}))$. At this q , the principal is weakly better off at $t - \varepsilon$. This is because the above equation is equivalent to $1 - e^{-r\varepsilon} = (1 - q)(1 - e^{-r2\varepsilon})$ so that

$$1 - q = \frac{1 - e^{-r\varepsilon}}{1 - e^{-r2\varepsilon}}.$$

The change in the principal's value is then proportional to $(1 - q)(1 - e^{-r2\varepsilon}) - (1 - e^{-rp\varepsilon})$. Plugging in the expression for q we can see that this value is positive because the function $(1 - e^{-r\varepsilon}) / (1 - e^{-r2\varepsilon})$ is increasing in r for sufficiently small ε .

For small enough $\nu > 0$ and $\varepsilon > 0$ the new schedule G' is implementable. For dates $s < t - \varepsilon$ this follows because G' is identical to G prior to $t - \varepsilon$. Any change after $t - \varepsilon$ which leaves the continuation value unchanged at $t - \varepsilon$ must also leave the continuation value unchanged at prior dates. For dates $s \geq t + \varepsilon$ the continuation schedule under G' is the same as under G and since G was implementable, we have $V_s(G') \geq 0$.

Finally for $s \in (t - \varepsilon, t + \varepsilon)$, the agent's value under G' is bounded from below by $\nu(1 - q)e^{-r2\varepsilon}R - c(1 - e^{-r2\varepsilon})$. This value is positive for ε sufficiently small. Hence G' is implementable. \square

Lemma 9. *Suppose G has no atoms and $V_t(G) > 0$ for some t . Then there exists $B > 0$ such that for all $\varepsilon > 0$ for all sufficiently small intervals $[t, u]$ we have*

1. $V_s(G) > B$ for all $s \in [t, u]$
2. $\frac{1 - G(u)}{1 - G(t)} > 1 - \varepsilon$.

Proof. Since G is atomless, it is continuous, which implies that $V_t(G)$ is continuous at

t . The first point in the Lemma then follows from continuity of $V_t(G)$, and the second from continuity of G and the fact that $1 - G(u) \leq 1 - G(t)$. \square

Lemma 10. *When $r_p < r$, an optimal schedule must have $V_t(G) = 0$ for all $t \geq 0$.*

Proof. By Lemma 8 an optimal schedule G has no atoms. Suppose $V_t(G) > 0$ for some t . Then consider an interval $[t, u]$ and the schedule H obtained from G by applying Lemma 4. Consider any $s \in [t, u]$. When the interval is small enough, $V_s(H) \geq \left(\frac{1-H(u)}{1-H(s)}\right) e^{-r(u-t)} V_u(H)$. The continuation value $V_u(H)$ is equal to $V_u(G)$ because H and G are identical after u . And since $V_s(G)$ cannot be larger than R , we have $V_s(G) - V_s(H) \leq R \left(1 - \frac{1-H(u)}{1-H(s)}\right)$. Now

$$\frac{1 - H(u)}{1 - H(s)} = \frac{1 - H(u)}{1 - H(t)} \geq \frac{1 - G(u)}{1 - G(t)}$$

where the equality follows from the fact that H has no mass in the interval (t, u) and the inequality from the fact that H was obtained from G by moving mass in that interval to the endpoints.

We now apply Lemma 9 to choose for any $\varepsilon > 0$, the endpoint u close enough to t to obtain the bounds $V_s(G) > B > 0$ for all $s \in [t, u]$ and $V_s(G) - V_s(H) \leq \varepsilon R$. Therefore the schedule H has a positive continuation value on the interval $[t, u]$. The continuation value at later dates is the same as under the original schedule G and therefore non-negative. The continuation value at earlier dates is the same as that of G since $V(H) = V(G)$ and H is identical to G at all dates earlier than t . The schedule H is thus implementable and strictly better for the principal, hence G cannot be optimal. \square

Lemma 11. *If $V_t(G) = 0$ for all $t \geq 0$ then G is absolutely continuous.*

Proof. If $V_t(G) = 0$ for all $t \geq 0$ then for any $\Delta > 0$,

$$\frac{1}{1 - G(t)} \int_0^\Delta R e^{-rx} - c(1 - e^{-rx}) dG(t + x) + \frac{1 - G(t + \Delta)}{1 - G(t)} e^{-r\Delta} V_{t+\Delta}(G) = 0$$

and since $V_{t+\Delta}(G) = 0$ we have $\int_0^\Delta Re^{-rx} - c(1 - e^{-rx})dG(t+x) = 0$. Now the left-hand side is larger than $[G(t+\Delta) - G(t)]Re^{-r\Delta} - c(1 - e^{-r\Delta})$, hence $G(t+\Delta) - G(t) \leq \frac{c(1 - e^{-r\Delta})}{Re^{-r\Delta}} = \frac{c}{R}(e^{r\Delta} - 1)$. Note that the bounding function, call it $\psi(\Delta)$, is convex and increasing. Let Δ^* satisfy $\psi(\Delta^*) = 1$, and let $K = 1/\Delta^*$. We will show that $G(t+\Delta) - G(t) \leq K\Delta$ for all $t \geq 0$ and $\Delta > 0$. This is immediate for $\Delta \geq \Delta^*$ because $G(t+\Delta) - G(t) \leq 1$. For $\Delta < \Delta^*$, we have $G(t+\Delta) - G(t) \leq \psi(\Delta) = \psi((\frac{\Delta}{\Delta^*})\Delta^* + (1 - \frac{\Delta}{\Delta^*})0) \leq (\frac{\Delta}{\Delta^*})\psi(\Delta^*) + (1 - \frac{\Delta}{\Delta^*})\psi(0) = (\frac{\Delta}{\Delta^*})\Delta^*K + (1 - \frac{\Delta}{\Delta^*})0 = \Delta K$.

Thus G is Lipschitz with constant K and therefore absolutely continuous. \square

Proof of Proposition 8. By Lemma 5 an optimal schedule G exists. If $r_p < r$ then by Lemma 10, it must hold the agent to a continuation value of zero at all points in time. To prove that the exponential distribution is optimal, it thus suffices to show that it is the unique distribution with that property. By Lemma 11, such a schedule G has a density γ . We can express the continuation value at time t as

$$\frac{1}{1 - G(t)} \int_t^\infty Re^{-r(x-t)} - c(1 - e^{-r(x-t)})\gamma(x)dx = \frac{e^{rt}}{1 - G(t)} \int_t^\infty (R + c)e^{-rx}\gamma(x)dx - c.$$

Since the continuation value is constant and equal to zero, $\int_t^\infty (R + c)e^{-rx}\gamma(x)dx = e^{-rt}c(1 - G(t))$ and both sides are differentiable with the following derivatives with respect to t : $-(R + c)e^{-rt}\gamma(t) = -ce^{-rt}[r(1 - G(t)) + \gamma(t)]$. Re-arranging yields $\frac{\gamma(t)}{1 - G(t)} = \frac{rc}{R}$. An optimal schedule has a constant hazard rate equal to rc/R and the exponential distribution with parameter rc/R is the unique such distribution. \square