Dynamic Project Standards with Adverse Selection

Andrew McClellan* University of Chicago Booth School of Business

February 13, 2019

Abstract

We study a principal-agent relationship in which the agent has private information about the future profitability of the relationship or a currently operated project, but is biased in favor of continuing the project. When the principal retains liquidation rights over the relationship or project and must introduce distortions in the liquidation policy itself in order to elicit the agent's private information. The optimal policy consists of a threshold which, if the profitability falls below, triggers liquidation. When the agent reports a higher growth rate of the projects profitability, the optimal threshold will be either decreasing over time and approach the principal's first-best level (i.e., the distortions from eliciting the agent's information are temporary) or will be increasing and divergent over time (i.e., liquidation at later times takes place at unboundedly inefficient levels). A simple condition on the relative profitability of the project across agent types tells us when the distortions are temporary or permanent. These results are robust to the use of transfers (e.g., wage payments) provided that a limited liability condition is respected for the agent. They are also robust to the use of direct auditing methods to assess profitability. The model provides a tractable way to analyze contractual distortions in the pretense of private information, and in particular, shows that contracts simultaneously frontand back-loaded across a menu of options in the same principal-agent relationship.

^{*} Email: andrew.mcclellan@chicagobooth.edu.

I would especially like to thank Debraj Ray and Ennio Stacchetti for their support, guidance and insightful suggestions. I would also like to thank Dhruva Bhaskar for helpful conversations and suggestions as well as seminar participants at NYU.

1 Introduction

In many economic partnerships, there is an asymmetry in information about the longterm profitability of the relationship. When a company is considering the acquisition of another firm, the manager at the firm being acquired will have superior information about his company's long-term profitability. When a regulator interacts with firm using an environmentally dangerous production technology, the firm has more information about the long-term damage to the environment than the regulator does. To elicit this information, the principal can enact standards which determine when the relationship will end: if the profitability falls below the standard, the relationship is ended. In this paper we study how the nature of private information determines these standards and examine when they are becoming more stringent or lenient over time.

More formally, the goal of our paper is to understand how persistent private information about future profitability introduce distortions in a dynamic principal-agent framework. We study a situation where the flow profitability to the principal of a project run by the agent changes stochastically over time. The agent has private information about the drift of this process and would like to keep the project running as long as possible. When the principal controls liquidation rights to the project, he can elicit the agent's private information by setting profitability standards (or thresholds) which, if the project falls below, triggers a shut-down.

We show that the nature of these standards can be radically different depending on how magnitude of the drift of profitability. When the expected drift of profitability is not too high, the liquidation threshold converges to the principal's preferred leveli.e., the distortions introduced by the agent's private information are temporary. More surprisingly, when the expected profitability of the more optimistic agent is large enough, the liquidation threshold is increasing over time-i.e., the distortions are persistent and growing even though the agent only has private information at the beginning of the relationship. A simple condition gives us exactly when the threshold is increasing or decreasing.

Our model gives a clear illustration of the forces at play that drives this front- or back-loading of distortions. When determining the liquidation policy of the high type, the principal takes into account what payoffs the policy would generate for the low type. Thus, we deciding the optimal liquidation policy, the principal must consider the direct incentives (using the expectation of the agent reporting truthfully) and also the counter factual for a low type agent claiming to be a high type (which uses the expectation of the low type agent). These forces lead to different placements of distortions over time. With our continuous time framework, we are able to get a clean and tractable way to study this although the economic insights we gain are certainly not unique to our model. Briefly, the intuition is as follows. If the drift of the high type is large enough, then as time progresses the value of liquidation increases relative to the direct payoff loss the principal. The likelihood a low type ascribes to being at a particular profitability level is increasing relative to the likelihood a high type will be there. By back-loading distortions, the principal is getting the most "bang for his buck." When the drift of the high type is low, this relative likelihood (for a fixed payoff level) is decreasing as time goes on. Thus introducing distortions does relatively little to decrease the low types payoffs when compared to the direct loss of inefficient liquidation to the principal.

The model also illustrates how the specification of payoffs drastically changes the resulting mechanism. In most of the previous literature, the agent knows the profitability of the project and this profitability is static. The principal observes noisy realizations of the profitability. However, in our model the agent's information is instead about the long-term profitability of the project and current profitability is perfectly observed by the principal. This allows for low type agents to still be profitable to the principal and for high type agents to become unprofitable. In many situations, this is a more reasonable payoff framework: a firm might know that the long-term profitability of a dial-up internet access business is low, but still operates the business today because current demand is profitable. In our model, these two different payoff structures lead to very different mechanisms.

We also show that these features of the optimal thresholds are robust to a number of extensions of the model, including auditing and limited liability transfers. In all these extensions, we find the same front- and back-loading forces at play and the resulting liquidation policies are qualitatively the same and illustrate the usefulness of our model and solution methods to analyzing a number of economically important situations.

We describe the literature in Section 2 and the model in Section 3. Our main results are presented in Section 4 and various extensions are explored in Section 5. All proofs are relegated to the Appendix.

2 Literature

Our work ties into a rich literature on agency problems and continuous-time techniques. Seminal papers such as Demarzo and Sannikov (2006) and Sannikov (2007) studied agency problems in continuous time and showed how continuous time techniques add a great deal of tractability to the problem. Unlike their models, we do not allow for transfers or moral hazard, instead choosing to focus on the problem of eliciting an agent's private information. Papers by Garrett and Pavan (2012) and Pavan, Segal and Toikka (2014) have looked at the dynamic mechanism design problem in which the agent's type is private information and is non-stationary. Unlike our model, they allow for arbitrary transfers and the methods they use are quite different from our own. Our model focuses on the simpler issue of how to temporally place distortions depending on the agent's private information.

Our model most closely fits in the literature on dynamic mechanism design without transfers. Kuvalekar and Lipnowski (2016) study a situation with some similarities to ours. They model a continuous-time principal-agent game without transfers in which there is symmetric uncertainty about an agent's type and the agent tries to stay hired as long as possible. Whereas they focus on equilibrium outcomes when there is learning about the profitability of the relationship, we instead look at how to elicit private information when the current profitability of the relationship is known but its future distribution is private information of the agent.

Madsen (2017) studies a model with similar features to our own: a principal seeks to elicit private information from an agent who desires to remain employed. While our model looks at a situation in which it is common knowledge when the agent's private information is received, Madsen studies a situation in which the arrival of the agent's information is private knowledge, which yields a much richer use of transfers. Another key difference is the payoff specification: in Madsen the agent's information is about the flow profitability of the state whereas in our model the agent's information is about the drift of the flow profitability. This difference will be key to our results (and which we will discuss more in Section 3).

Looking at a delegation model in a similar spirit to ours, Guo (2016) studies a model of experimentation in which the agent has private information about a likeliehood a project is good. Aside from technical differences (our is a stopping problem, her's a bandit problem), we model different payoff structures and misalignment of the principal and agent's incentives. Fong (2007) also studies a mechanism design model without transfers with adverse selection. Her model, unlike ours, allows for moral hazard. She finds that the optimal mechanism consists of a score, which if it falls too low results in the termination of the agent. The tools available to the mechanism designer are the same as in our model. However, the difference in the payoff structure drives the difference of our results. In her model, the payoff to the principal is a function only of the agent's type and action. In our model, the agent's type will not impact the current payoffs to the principal but instead the principal's beliefs about future payoffs. This change to the payoffs separates our model from most of the previous literature and yields the qualitative difference in our mechanism from previous mechanisms in the literature.

On a technical side, our paper makes use of techniques from Peskir (2005) and finitehorizon option pricing models. We are able to use techniques where the optimization problem depends on time to solve our infinite-horizon problem with adverse selection. Our use of Lagrangian techniques in optimal stopping problems draws on constrained optimal stopping problems such as in McClellan (2017).

3 Model

We study the long-term relationship between a principal P and an agent A in an infinitehorizon continuous-time model. There is a payoff-relevant variable X_t for the principal, where X_t is a diffusion process given by

$$X_t = \mu_\theta t + \sigma B_t,$$

and B_t is a standard Brownian motion which is defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ satisfying the usual conditions¹ and X_0 is known to both P and A. The drift of X_t depends on a state of nature $\theta \in \Theta$. A learns θ at t = 0 (and we will refer to θ as A's type). It is common-knowledge that A learns θ but that P does not.

The focus of our paper is to find the optimal mechanism for P to elicit the private information of A. We study situations in which transfers are infeasible and assume that P's decision is only when to irrevocably fire the agent or liquidate the project. Formally, we define an admissible mechanism to be \mathcal{F}_t^X -measurable stopping time τ such that the game ends at time τ . We assume that the terminal payoff at (τ, X_{τ}) to both P, A is zero.² The payoff of the principal for an arbitrary stopping rule τ is given by

$$\mathbb{E}_{\theta,X_0}[\int_0^\tau e^{-rs} u_P(X_s)ds].$$

A's payoff is given by

$$\mathbb{E}_{\theta,X_0}[\int_0^\tau e^{-rs} u_A(X_s)ds],$$

¹See Karatzas and Shreve (1991).

²This assumption can be relaxed as long as terminal payoffs depend only on X_{τ} and A's terminal payoffs are positive.

where we assume that $u_A(X_s) > 0 \ \forall X_s$. For simplicity, we assume that $u_P(X_s) = X_s$ and $u_A(X_S) = \alpha \in \mathbb{R}_+$.³ In order to simplify notation slightly, we will drop the dependence of the expectations on X_0 .

We allow P full commitment power, so the revelation principle applies.⁴ This leads us to define a mechanism for P, which is a stopping time that depends on the agent's report.

Definition 1. A stopping mechanism is a menu $\{\tau_{\theta}\}_{\theta \in \Theta}$ where τ_{θ} is an \mathcal{F}_t^X -measurable stopping rule. The stopping mechanism is incentive compatible (IC) if for each type $(\theta_i, \theta_k) \in \Theta \times \Theta$, we have

$$\mathbb{E}_{\theta_i} \left[\int_0^{\tau_{\theta_i}} e^{-rt} \alpha dt \right] \ge \mathbb{E}_{\theta_i} \left[\int_0^{\tau_{\theta_k}} e^{-rt} \alpha dt \right]$$

3.1 Remarks

The restriction to only allow P to liquidate the project both simplifies the problem and is realistic modeling choice for many economic settings: in large firms, wage levels are often fixed while the termination decisions are more flexible. Additionally, in the interaction between a regulator and a firm, monetary transfers may often be limited in scope while the regulatory power to shut down production at a plant has more bite behind it. Our assumptions on the payoffs fit both of these settings well: the employee with a fixed wage desires to be employed as long as possible while the firm cares about the profitability of the employee. As we will see in Section 5, the structure of the optimal mechanism will be robust to the addition of non-negative transfers from the principal to the agent.

A's only action in the game is to report θ at t = 0. Mathematically, the problem is similar to that of a static mechanism design problem, where the set of \mathcal{F}_t^X -measurable stopping rules corresponds to the good to be allocated. A richer model in the vein of Garrett and Pavan (2016) or Madsen (2016) might allow for continual changes to θ which A must be incentivized to report; we explore this somewhat in Section 5. Instead, the goal of our model will be to narrow in on the role of perfectly persistent information. To this end, we ignore issues of moral hazard, which will introduce more distortions in the

³The results are not dependent to our particular utility functions. Our specification that A derives positive utility allows us to not be concerned with the possibility of P operating the project longer to punish A.

⁴When the P has many interactions with different agents (e.g, a manager in a large firm interacting with many employees or a government agency regulating many companies), this commitment assumption can be motivated by reputational concerns. Even when such strong commitment is not possible, we can view this problem as providing an upper-bound on the payoff P might get.

optimal contracts, in order to narrow in on the effect of persistent private information (one could interpret our model as a situation in which the agent's actions are observable to P). Given the static nature of A's private information, the persistence of distortions over the course of the mechanism is all the more surprising in the absence of moral hazard concerns.

Before moving on, it is important to note a key feature of our model: the role of X_t , which is both payoff relevant and conveys information about θ . In much of the continuous time principal-agent and delegation literature, the principal's payoffs are given by dX_t , so that X_t is equal to the principal's accumulated payoffs. Taking expectations, this means that the principal's flow payoff at any given is μ_{θ} . In contrast, our model defines the flow payoffs of the principal to be X_t . This seemingly minor change to the payoff structure leads to very different results when it comes to the optimal mechanism. If we were to model u_P as depending only on θ , then we can show that the optimal mechanism features $\tau_h = \tau_{\ell}$ -i.e., there is no use designing screening mechanisms. Thus this dual role of X_t is a driving force behind our results. Additionally, we believe that in many situations X_t rather than dX_t is the appropriate flow payoff. The dX_t payoff specification means that the agent knows what the expected profitability of the project and the profitability never changes. Instead, our X_t payoff specification means that the agent's private information is about the expected path of future payoffs, which are changing over time. This allows for the profitability of the project to change over time. For a regulator observing the level of pollution produced by a firm, their flow payoff is determined by the cumulative level of pollution, not the change in pollution. For a firm observing a projects level of demand X_t , it is the level not change in demand that determines flow payoffs. Our model allows for ℓ types to be profitable for P and for h types to become unprofitable for P.

4 Baseline Mechanisms

4.1 First-Best

As a benchmark, we analyze the problem where there is no private information. Because A's only action is to report his type at t = 0, he has no role when there is no private information. The problem with no asymmetric information then simplifies to a Markov single-decision maker problem with P's preferences where the only state variable is X_t . As is standard, the optimal mechanism for P will take the form a threshold rule, in which P shuts-down at the first time $X_t \leq b$ for some $b \in \mathbb{R}$. Additionally, because of the tractable nature of Brownian motion, we can get an closed form solution for the optimal threshold.

Proposition 1. The optimal stopping policy for θ_i is given $\tau_i = \inf\{t : X_t = b_i\}$ where $b_i = \frac{\sigma^2}{\mu_i + \sqrt{\mu^2 + 2r\sigma^2}} - \frac{\mu_i}{r}$. The optimal threshold b_i is decreasing in μ_i .

It is easy to see why b_i decreasing in μ : for projects that are expected to be more profitable in the future, P is willing to accept larger losses today since the option value of continuing is higher. The solution to b_i also reveals that the optimal threshold is decreasing in r: when the discount rate is high, it is better to cut one's losses today rather than continue waiting in the hope that X_t will be profitable in the future.

The result illustrates where the tension in our agency problem arises from: P would like to keep the project running longer when μ_{θ} is high, but this creates incentives for Ato misreport μ_{θ} when he knows μ_{θ} to be low.

As is usual in one-dimensional single-decision maker stopping problems, the optimal solution takes a simple threshold form while, as we will see in the next section, threshold rules are not optimal when agency considerations come into play. The first-best contract highlights the fact that the non-standard solution to the problem with agency considerations is driven by the presence of private information and not the base model itself.

4.2 Private Information

We introduce the private information of A by looking at the case when A learns $\theta \in \Theta = \{\theta_{\ell}, \theta_h\}$, where $\theta_{\ell} < \theta_h$. We can write P's problem of eliciting A's private information as

$$\sup_{\tau_h,\tau_\ell} \sum_{i=h,\ell} \mathbb{P}(\theta_i) \cdot \mathbb{E}_{\theta_i} [\int_0^\tau e^{-rt} X_t dt]$$

subject to $\forall i, k$

$$IC(i,k): \mathbb{E}_{\theta_i}[\int_0^{\tau_i} e^{-rt} \alpha dt] \ge \mathbb{E}_{\theta_i}[\int_0^{\tau_k} e^{-rt} \alpha dt]$$

Intuitively, we expect that the high types incentive constraint will be slack: Since A only want to maximize the time until shut-down, A's incentives and P's incentives are more aligned (since P would always want to wait longer before shutting-down the high type). As we will verify that this intuition is correct later, we begin by studying the optimal τ_h subject to ℓ 's IC. Let W_ℓ be equal to the utility ℓ gets from choosing τ_ℓ (this is determined by P in his choice of τ_ℓ). Then the IC constraint for ℓ can be rewritten as

$$IC(\ell,h): \quad \mathbb{E}_{\theta_{\ell}}[\int_{0}^{\tau_{h}} e^{-rt} \alpha dt] \leq W_{\ell}.$$

In order to solve this problem, we will transform it into an optimal stopping problem of a more standard form. Our first step is to modify the payoffs. We note that for A, we have

$$\mathbb{E}_{\theta_{\ell}} \left[\int_{0}^{\tau_{h}} e^{-rt} \alpha dt \right] = \frac{\alpha}{r} (1 - \mathbb{E}_{\theta_{\ell}} [e^{-r\tau}]).$$

For P, we note that by Ito's Lemma, we have that

$$de^{-rt}X_t = -re^{-rt}X_tdt + e^{-rt}(\mu dt + \sigma dB_t).$$

By integrating from 0 to t, rearranging terms and taking expectations, we get that

$$\mathbb{E}_{\theta_h} \left[\int_0^\tau e^{-rs} X_s ds \right] = \frac{\mu_h}{r^2} + \frac{X_0}{r} - \mathbb{E}_{\theta_h} \left[e^{-r\tau} (\frac{\mu_h}{r} + X_\tau) \frac{1}{r} \right].$$

Two features separate the problem from being written in a standard optimal stopping format: the presence of IC constraints and the different expectation operators in the ICconstraint. In order to tackle the first feature, we construct a Lagrangian with Lagrange multiplier $\lambda \geq 0$ (corresponding to the IC constraint), which is written as

$$\mathcal{L} = \sup_{\tau} \frac{\mu}{r^2} - \mathbb{E}_{\theta_h} [e^{-r\tau} (\frac{\mu_h}{r} + X_\tau) \frac{1}{r}] - \lambda (\frac{\alpha}{r} (1 - \mathbb{E}_{\theta_\ell} [e^{-r\tau}]) - W_\ell) + \frac{X_0}{r} \\ = \sup_{\tau} [\lambda \alpha \mathbb{E}_{\theta_\ell} [e^{-r\tau}] - \mathbb{E}_{\theta_h} [e^{-r\tau} (\frac{\mu_h}{r} + X_\tau)] + (\frac{\mu_h}{r} + X_0 + \lambda (W_\ell - \alpha))] \frac{1}{r}.$$

From here on out we will drop the constants $\frac{\mu_h}{r} + X_0 + \lambda(W_\ell - \alpha)$ and the factor $\frac{1}{r}$ from \mathcal{L} .

We have one final step in order to transform P's problem into a standard optimal stopping problem: to transform the Lagrangian into a single expectation. The use of Brownian motion allows us to use Girsanov's theorem to change the measure for θ_{ℓ} into a measure for θ_h . By Oskendal Theorem 8.6.4, the Radon-Nikodym derivative⁵ is

$$M_t = exp(-\frac{\mu_h - \mu_\ell}{\sigma^2}(X_t - X_0) + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t).$$

This change of measure allows us to convert the θ_{ℓ} expectation into a θ_h expectation as

$$\mathbb{E}_{\theta_{\ell}}[e^{-r\tau}\lambda\alpha] = \mathbb{E}_{\theta_{h}}[e^{-r\tau}\lambda\alpha M_{t}]$$

This allows us to write P's problem as

⁵We will subsequently assume that $X_0 = 0$ in order to not carry additional notation.

$$V(t,X) := \sup_{\tau} \quad \mathbb{E}_{\theta_h}[e^{-r(\tau-t)}G(\tau,X_{\tau})]$$

where $G(\tau, X_{\tau}) := \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t} - (\frac{\mu_h}{r} + X_{\tau})$ is our gain function from stopping at (τ, X_{τ}) .

With the above manipulations, we can now view P's problem as an unconstrained optimal stopping problem. Standard optimal stopping arguments allow us to show that the space of $\mathbb{R}_+ \times \mathbb{R}$ can be partitioned into an open continuation set \mathcal{C} and a closed stopping set \mathcal{D} where

$$C = \{(t, X) : V(t, X) > G(t, X)\}$$
$$D = \{(t, X) : V(t, X) = G(t, X)\}$$

We define P's strategy with the sets C. This leads to a natural conjecture that C will take the form $C = \{(t, X) : X > b_h(t)\}$ for some function $b_h(t)$ (which we will call a threshold function). We will verify this conjecture and solve for the optimal $b_h(t)$ below.

Solving for V is technically complicated by the fact that our gain function G depends on both X_t and t. The natural formulation of the problem leads to a free-boundary problem; however this will involve solving a partial differential equation, which can often be difficult to solve. Fortunately, we can make use of a change-of-variable formula for local-times on curves from Peskir (2005) in order to solve for the optimal policy. This formula, previously used in option pricing, will prove useful in our settings.

With this formulation of P's problem, we can clearly illustrate the forces driving our main result. By quick inspection, we note that for a fixed X, the Radon-Nikodym derivative M_t is disappearing (exploding) as $t \to \infty$ if $|\mu_h| < (>)|\mu_\ell|$. This fact will turn out to lead to very different stopping rules. The intuition for our main result is quite simple and is elucidated by M_t . Let us divide the space $\mathbb{R}_+ \times \mathbb{R}$ into H and L where $H = \{(t, X) : X \geq \frac{\mu_h + \mu_\ell}{2}t\}$ and $L = \{(t, X) : X < \frac{\mu_h + \mu_\ell}{2}t\}$. For $(t, X) \in H$, we have $M_t > 1$, so that the probability that θ_ℓ reaches (t, X) is less than the probability of the same event under the measure for θ_h . Similarly, for $(t, X) \in L$, we have that $M_t < 1$, so that the probability that θ_h reaches (t, X) is less than the probability for θ_ℓ . As $-X_t + \frac{\mu_h + \mu_\ell}{2}t \to \infty$, the relative probability that ℓ has reached this point rather than h approaches infinity. This means that, for large t, by shutting-down immediately, P is reducing ℓ 's incentives more than he is harming his own (since he evaluates the shut-down policy with h's expectation). In the other direction, as $-X_t + \frac{\mu_h + \mu_\ell}{2}t \to -\infty$, the relative probability that h has reached this point rather than ℓ approaches infinity. Therefore, shutting down at this point harms P infinitely more than it decreases ℓ 's incentives. This makes shut-down at this point costly for P with no benefit.

This quick analysis indicates the role of $|\mu_h|$ and $|\mu_\ell|$. If $\mu_h + \mu_\ell < 0$, then the benefits to stopping are decreasing over time for a fixed X-i.e., after a long enough time, no ℓ types would be expected to be at X; hence, punishment here carries on a small impact ℓ 's ex-ante incentives relative to the distortion loss by early-shut down to P. This decreases the need for inefficient shut-down at this point and drives the shut-down policy towards the first-best for P as t increases. In the opposite direction, if $\mu_h + \mu_\ell > 0$, then the benefits to stopping are increasing over time.

Given the above discussion, we split the analysis up into two cases: that of $|\mu_h| > |\mu_\ell|$ and $|\mu_h| < |\mu_\ell|$. Our results below will analyze the long-term behavior of the optimal mechanism in each case. We will find that when $|\mu_h| > |\mu_\ell|$, the optimal policy features increasing distortions, in that shut-down occurs at higher X at t increases while when $|\mu_h| < |\mu_\ell|$, the shut-down policy converges to P's first best solution. The intuition behind this result follows from the previous paragraphs: P wants to shut-down when the probability of reaching X at time t from ℓ 's perspective is high relative the probability from h's perspective. How these relative probabilities move is qualitatively different depending on the relative sizes of $|\mu_h|, |\mu_\ell|$. Remarkably, whether or not $|\mu_h| > |\mu_\ell|$ is both necessary and sufficient to know whether agency distortions via inefficient shutdown are permanent or transitory. Our main result, given below, proves the intuition above to be correct.

Theorem 1. The optimal mechanism is such that:

- If |µ_h| < |µ_ℓ|, then b_h(t) is a decreasing and continuous function of t which converges to b_h^{FB} as t → ∞.
- If |µ_h| > |µ_ℓ|, then b_h(t) is an increasing and continuous function of t and b_h(t) → ∞ as t → ∞.
- If $|\mu_h| = |\mu_\ell|$, then $b_h(t) = b_\ell(t) = b^*$.

For ℓ , the stopping rule is a static threshold $\tau_{\ell} = \{t : X_t = b_{\ell}\}$ where $b_{\ell} \in \mathbb{R}_+$ and $b_{\ell} < b_{\ell}^{FB}$. If $|\mu_h \neq |\mu_{\ell}|$, then the thresholds b_{ℓ} and b_h cross exactly once.

There are two routes to satisfy the *IC* constraint for ℓ : to increase ℓ 's utility from reporting truthfully and to decrease ℓ 's utility from reporting to be *h*. Both are used in the optimal mechanism. We can see the first route being used in the fact that $b_{\ell} < b_{\ell}^{FB}$, while the second route comes in through inefficient shut-down with b_h (evident from the fact that b_h is not equal to the first-best). Interestingly, increasing the utility of a correct report by ℓ results in a constant distortion across time while decreasing the utility a misreport creates time-varying distortions. This result is tied to the fact that when increasing the utility of a correct report, P and ℓ evaluate the policy using the same expectation. However, when decreasing the utility of a misreport, P and ℓ evaluate the policy using different expectations. Because the distortions in ℓ 's mechanism are purely to provide utility to ℓ (and not to screen as in h's mechanism), ℓ 's mechanism is stationary. Thus we see that the reason for the changing standards in h's mechanism are used as a method to screen out ℓ types.

More generally, our model speaks to how to design dynamic standards in screening problems. Interestingly, our results tell us that it is not the size of μ_{θ} alone that determines the structure of the optimal stopping rule, but the relation between $|\mu_h|$ and $|\mu_{\ell}|$. Thus, whether or not h should face increasing or decreasing standards depends on the relative sizes of μ_h, μ_{ℓ} .

Note that our model can easily be recast to fit a moral-hazard story at t = 0. Suppose that P wants A to take some costly action which ensures the state to be θ_h rather than θ_{ℓ} . Thus we can replace our IC constraints with

$$\mathbb{E}_{\theta_h}[\int_0^\tau e^{-rt}\alpha] - c \ge \mathbb{E}_{\theta_\ell}[\int_0^\tau e^{-rt}\alpha]$$

so that, use the same Lagrangian approach, the problem is equivalent to solving

$$V^m(t,X) := \sup_{\tau} \quad \mathbb{E}_{\theta_h}[e^{-r(\tau-t)}G^m(\tau,X_{\tau})]$$

where $G^m(\tau, X_{\tau}) := \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t} - (\frac{\mu_h + \lambda \alpha}{r} + X_{\tau})$. Using the same arguments as in the proof of Theorem 1, we will get the optimal stopping policy is of the same form as $b_h(t)$.

4.3 Front-Loaded Inefficiencies: $|\mu_h| < |\mu_\ell|$

We begin by analyzing the case when $|\mu_h| < |\mu_\ell|$. The optimal policy in this case features front-loaded inefficiencies and the shut-down policy is asymptotically P's first-best level, implying that the distortions introduced by the asymmetric information at t = 0 dissipate over time. The reason for this comes from the fact that the Radon-Nikodym derivative is decreasing in t for a fixed X. Intuitively, this the payoff to "punishing" an ℓ type that misreported (given by the term λM_t) by shutting down at X is decreasing in t for a fixed X. P would optimally like to minimize inefficiencies by shutting down at beliefs which ℓ finds more likely than h. For $|\mu_h| < |\mu_\ell|$, such only occurs at lower X as t increases; but for low enough X, P would already optimally shut-down, thereby allowing him to punish ℓ while taking efficient actions.

In order to verify the above arguments, we need to derive several properties about the optimal stopping boundary $b_h(t)$. We begin by showing that this boundary is both continuous and decreasing.

Lemma 1. If $|\mu_h| < |\mu_\ell|$, then $b_h(t)$ is a decreasing and continuous function of t.

Proof. All proofs are in the Appendix.

Next we note that $b_h(t)$ is bounded below by b_h^{FB} . If for some t', we have $b_h(t') < b_h^{FB}$, then P could modify the shut-down policy to be $b_h(t') = b_h^{FB}$ for all t' > t. This would preserve IC (since higher shut-down thresholds always decrease A's utility) and be strictly better for P by definition of b_h^{FB} .

Lemma 2. $b_h(t)$ is bounded below by b_h^{FB} .

With this in hand, we can look at the long-term properties of $b_h(t)$. We show that for large t, we have $b_h(t) \to b_h^{FB}$ -i.e., the distortions introduced by the mechanism disappear in the long-run. Looking at M_t , we can see that $M_t \to 0$ if X_t is bounded below since $\frac{\mu_h^2 - \mu_\ell^2}{2} < 0$. This means that the distortion term is vanishing as times goes on, leading us to correctly conjecture that the optimal threshold must approach the principal's firstbest.

Lemma 3. $b_h(t)$ converges to b_h^{FB} at $t \to \infty$.

This type of policy shows us that the mechanism front-loads distortions. The reasoning behind this is simple. Let μ_{ℓ} be sufficiently negative. Then for large t and $b_h(t) > b_h^{FB}$, ℓ will assign a small probability to ever making it this far. Thus the decision to shut-down above the efficient level has only a small impact on ℓ 's *IC* constraints. However, P evaluates the performance of $b_h(t)$ relative to an expectation with respect to μ_h . From the perspective of μ_h , $X_t > b_h(t)$ becomes infinitely more likely than it is from ℓ 's perspective. Thus the ratio of the benefit to slackening ℓ 's *IC* constraint becomes to the loss to P approaches zero. Thus long-run distortions have no bite and in order to satisfy *IC* constraints, distortions must be front-loaded.

In some sense, these seems intuitive. Because P must only satisfy time zero constraints, it is natural to expect that all optimal mechanisms front-load distortions. However, as we will see in the next subsection, this is not always the case.



Figure 1: The optimal thresholds when $|\mu_h| < |\mu_\ell|$. As we can see h's threshold is decreasing over time and approaches the principal's first-best solution for h while ℓ 's threshold is stationary.

4.4 Back-Loaded Inefficiencies: $|\mu_h| > |\mu_\ell|$

The more surprising result in Theorem 1 is when $|\mu_h| > |\mu_\ell|$ the optimal mechanism for h involves increasingly ex-post inefficient shut-down, implying the distortions caused by asymmetric information are increasing over time. Again, we can understand this result as coming from the optimal placement of inefficiencies. As t increases, an observation of $X_t = X$ (for some fixed X) becomes increasingly likely to have come from θ_ℓ when compared to θ_h ; hence, while shutting down at such X may be ex-post inefficient, from an ex-ante perspective they maximize the benefit of more shut-down in relaxing the IC constraint relative to inefficiency of early-shut down.

As in the previous subsection, we begin by verifying the above arguments and show that the optimal stopping boundary is both continuous and increasing.

Lemma 4. If $|\mu_h| > |\mu_\ell|$, then $b_h(t)$ is increasing and continuous in t.

This readily supplies the result that distortions will be persistent over time. Our goal for the rest of the subsection will be to study how they change over time. In order to do this, we use a change of variable. We define $Z_t = X_t - \frac{\mu_h + \mu_\ell}{2}t$, which is a monotonic transformation of M_t and therefore can be interpreted as a measure of the relative probability assigned to (t, X_t) by ℓ relative to h. Since Z_t is measurable

with respect to \mathcal{F}_t^X , we can define an equivalent problem (and transforming utilities properly) where the agent reports the drift of Z_t rather than X_t and we find a stopping rule which is measurable with \mathcal{F}_t^Z . This process is useful because the drifts of h, ℓ are symmetric around zero: h views the drift of Z_t as $\tilde{\mu} := \frac{\mu_h - \mu_\ell}{2}$ and ℓ views the drift of Z_t as $-\tilde{\mu} = \frac{\mu_\ell - \mu_h}{2}$. Our modified gain function for Z is given by

$$\widetilde{G}(Z,t) = \lambda \alpha e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z_t} - (Z_t + \frac{\mu_h}{r} + \frac{\mu_h + \mu_\ell}{2}t).$$

We define the shut-down barrier for this problem in (t, Z_t) space as $b_h^Z(t)$. We can then apply a similar argument as in Lemma 1 to show that the shut-down threshold $b_h^Z(t)$ for the Z process is decreasing in t.

Lemma 5. If $|\mu_h| > |\mu_\ell|$, then $b_h^Z(t)$ is decreasing and continuous in t.

Assuming that b_h^Z is differentiable, this implies that $\frac{\partial b_h^Z(t)}{\partial t} < 0 \iff b'_h(t) < \frac{\mu_h + \mu_\ell}{2}$. This result allows us to show that $b_h(t)$ is not growing "too" fast. However, this still allows for the possibility that $b_h(t)$ is converging to some higher stationary threshold, similar to the structure when $|\mu_h| < |\mu_\ell|$. Our next result shows that this is not the case.

Lemma 6. The threshold $b_h(t)$ is diverging: $\lim_{t\to\infty} b_h(t) = \infty$.



Figure 2: The optimal thresholds when $|\mu_h| > |\mu_\ell|$. As we can see h's threshold is increasing over time while ℓ 's threshold is stationary.

This mechanism has similar features to that of "up-or-out" contracts, in which an employee is given a set amount of time to reach certain benchmarks and is fired if they fail to meet these goals. While these types of contracts are often interpreted as a response to moral hazard considerations, our model illustrates how such contracts can actually be useful as a screening tool even when such moral hazard considerations are absent.

We might naturally wonder whether or not this persistent distortion is an artifact of the fact that X_t is both payoff-relevant and carries information about the value of θ . If instead of an agency problem R was just uncertain about the value of θ , would we get a similar result to the diverging stopping boundary? As the next proposition shows, this is not the case. The logic is simple: as $t \to \infty$, a low X_t implies that the state is almost certainly θ_{ℓ} . If it is not optimal to stop at X_t when $\theta = \theta_{\ell}$ with probability one, then it is not optimal to stop when θ is uncertain.

Proposition 2. In the single-decision learning problem, the optimal stopping boundary b satisfies $b(t) \leq b_{\ell}^{FB}$.

4.5 No Screening: $|\mu_h| = |\mu_\ell|$

We now confront the final case, when $|\mu_h| = |\mu_\ell|$ (which implies $\mu_\ell = -\mu_h$). Note that, in this knife-edge case, we have

$$M_t = exp(-\frac{2\mu_h}{\sigma^2}X_t),$$

which is independent of t. It is straightforward then to show that the solution for h's mechanism is a static threshold. Together with the fact that ℓ 's mechanism is also a static threshold, IC implies that $b_h(t) = b_\ell(t) = b^*$ for some $b^* \in \mathbb{R}$ -i.e., there is no use screening types! This surprising result comes from the fact that the time dimension no longer impacts the relative likelihood of being at a particular state X. Thus the relative loss to P and loss to ℓ from shutting down is independent of time, making time based distortions inefficient. Therefore, the best that P can do is offer h a stationary threshold.

4.6 Optimal Boundary

Up unto this point, we have left the exact nature of $b_h(t)$ unanswered. The proofs for this section will be the most technically involved, as it involves finding the solution to a free-boundary problem with a partial differential equation. However, using insights from the option pricing literature, we can solve for the optimal boundary by formulating the following free-boundary problem:

$$\begin{split} \mathbb{L}V(t,X) &= rV(t,X)\\ V(t,b_h(t)) &= G(t,b_h(t))\\ \frac{\partial V(t,X)}{\partial X}|_{X=b_h(t)} &= \frac{\partial G(t,X_t)}{\partial X}\\ V(t,X) &> G(t,X) \; \forall (t,X) \in \mathcal{C}\\ V(t,X) &= G(t,X) \; \forall (t,X) \in \mathcal{D}, \end{split}$$

where the infinitesimal generator is given by $\mathbb{L} = \frac{\partial}{\partial t} + \mu_h \frac{\partial}{\partial X} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial X^2}$. The second and third line specify that the value function must, as is standard, satisfy be continuous and satisfy smooth-pasting on the boundary.

Since our state variable is given by (t, X), we have to deal with solving a PDE for V. As is well known, this is in general a difficult problem. However, because it is a PDE involving time, we can use a change-of-variable formula for local times on a curve from Peskir (2005) to get some traction on the problem. This formula has found use in the option-pricing literature and allows for a more in depth derivation of optimal stopping boundaries for value functions which depend on time, as ours does. By using this formula and applying the smooth-pasting principles we derive in the Appendix, we can pin down $b_h(t)$ to be the solution to an integral equation.

Proposition 3. The optimal threshold function $b_h(t)$ is the solution to the Volterra integral equation

$$G(t,b_h(t)) = -\int_t^\infty e^{-r(s-t)} \mathbb{E}[\left(\mathbb{L}G(s,X_s) - rG(s,X_s)\right)\mathbb{1}(X_s < b(s))ds]$$
(1)

The function $b_h(t)$ is the unique solution to 1 in the class of functions $\{c : \mathbb{R}_+ \to \mathbb{R} : c \in C^1 \text{ and } -\mu_h - rG(t, c(t)) < 0 \ \forall t \in [0, \infty)\}.$

The integral equation 1 can also be written as

$$\lambda \alpha e^{\frac{\mu_h - \mu_\ell}{\sigma^2} (-b_h(t) + \frac{\mu_h + \mu_\ell}{2})} - \frac{\mu_h}{r} - b_h(t)$$
$$= r \int_t^\infty e^{-r(s-t)} (\mathbb{E}_{\theta_h, b_h(t)} [X_s \mathbb{1}(X_s < b_h(s))] + r \lambda \alpha \mathbb{P}_{\theta_\ell, b_h(t)} (X_s < b_h(S)) ds$$

Since we are dealing with an infinite horizon problem, it is difficult to solve this integral equation. If we were to solve a version of the equation for a finite time model, we would be able to work backwards to calculate $b_h(t)$. To be more specific, suppose that at time T the game ends (i.e., the agent is fired) and both players receive a continuation value of zero. Solving such a finite time model can be done by adapting techniques from finite horizon option pricing.

Proposition 4. Fix a $T \in \mathbb{R}_+$. Then the solution to the finite time horizon problem threshold b_h^T converges pointwise to the optimal infinite time horizon solution b_h as $T \to \infty$.

The first part of the proposition tells us that we can use these finite-time horizon solutions as valid approximations of the optimal policy. These solutions will have integral equations of a similar form to the one above. The benefit of these finite-time horizon solutions is that we can easily solve $b_h(t)$ by evaluating the optimal boundary at T and solving backwards.

5 Extensions

In order to test the robustness of our results, we now look at a number of different extensions. The model, as studied so far, is quite simple in the tools it allows the principal. As we will see, the optimal mechanisms will have qualitatively the same structure even when we allow the principal a wider range of mechanism tools with which to choose from, illustrating how the dynamics of the optimal mechanism aren't dependent on the fine details of the model.

5.1 Costly Auditing

In many real-life situations, the principal can choose to acquire costly information prior to the shut-down decision. For example, if θ describes the future profitability of a firm's division, then the principal could hire an outside auditor learn the state of the division. In order to reflect this situation, we modify the model by allowing for costly state verification. For simplicity, we move back to the case with $\Theta = \{\theta_{\ell}, \theta_h\}$. When P decides to audit, he pays a cost C and learns the type of A. Let τ be the decision to either shutdown or audit the agent and d_{τ} be a decision variable which takes value 1 if P audits at τ and value 0 otherwise. A mechanism in this extension is a menu of such pairs:

Definition 2. A mechanism with costly auditing is a menu of stopping times and decision rules $\{(\tau_{\theta}, d_{\tau}^{\theta})\}_{\theta \in \Theta}$ where τ_{θ} and d_{τ}^{θ} are \mathcal{F}_{t}^{X} -measurable.

The utility of P is then given by

$$\mathbb{E}\left[\int_0^\tau e^{-rt} X_t dt + e^{-r\tau} (\widetilde{V}(X_\tau) - C) d_\tau\right],$$

where $\widetilde{V}(X_{\tau})$ is the continuation value post-audit for P for some policy (we will conjecture the form of \widetilde{V} below). If P audits and discovers A be a different type than he claimed, he will immediately shut-down; if P discovers that A truthfully reported, then A will be allowed to continue and will receive some continuation value W_{τ} determined by the mechanism used after the audit. Thus the utility of A can be written as

$$\mathbb{E}_{\theta_i} \left[\int_0^{\tau_j} e^{-rt} \alpha dt + e^{-r\tau_j} W_\tau \mathbb{1}(j=i) d_{\tau_j} \right]$$

Similar to Section 5, we conjecture that h's *IC* constraint will be slack and thus we focus on studying τ_h when ℓ 's *IC* constraint binds. Since h's *IC* constraint has been dropped, we know that after auditing P will institute the first-best policy, we can replace $\tilde{V}(X_{\tau})$ with $V^{FB}(X_{\tau})$. By using the same techniques as that section, we can write P's problem as an optimal stopping problem with the gain function

$$G(\tau, X_{\tau}) = \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t} - \left(\frac{\mu_h}{r} + X_{\tau}\right) + \left(V^{FB}(X_{\tau}) - C\right) d_{\tau}$$

As before, we will have two qualitatively different contracts, depending on $|\mu_h|$ and $|\mu_\ell|$. Note that that the Radon-Nikodym derivative term has not changed with the addition of costly auditing. Our arguments on the nature of $b_h(t)$ will go through almost verbatim. Thus the only real difference is to find the optimal choice of d_{τ} . But since we are considering only the *IC* constraints of a deviating choice, the choice of shut-down or audit is equivalent to him: once *P* finds he has deviated through an audit, shut-down is immediate. Therefore, the audit choice is only dependent on the size of the continuation value relative to the audit cost. For any τ , the decision rule d_{τ} which maximizes $G(\tau, X_{\tau})$ is given by

$$d_{\tau} = \mathbb{1}(V^{FB}(X_{\tau}) > C)$$

If $b_h(t)$ is approaching P's first best level and is decreasing, then eventually the audit cost becomes too costly relative to shut-down (since the continuation value is small). The value $V^{FB}(b_h(t)) - C$ will be decreasing in t, making auditing less attractive. If $b_h(t)$ is increasing, then the value of the continuation value relative to audit cost is growing over time, making auditing more attractive. This naturally leads to cutoff times T_1, T_2



Figure 3: Optimal mechanism when $|\mu_h| > |\mu_\ell|$. We see that the optimal stopping barrier for h is still increasing, but that after t' = 7.2 P audits rather than shut-down.

such that if $b_h(t)$ is decreasing, auditing only happens before T_1 and if $b_h(t)$ is increasing, auditing only happens after T_2 .

Proposition 5. If $|\mu_h| > |\mu_\ell|$, then $b_h(t)$ is increasing and $\exists t'$ such that $d_\tau = \mathbb{1}(\tau > t')$. If $|\mu_h| < |\mu_\ell|$, then $b_h(t)$ is decreasing, approaches b_h^{FB} as $t \to \infty$, and $\exists t''$ such that $d_\tau = \mathbb{1}(\tau < t'')$. For ℓ 's menu choice, $\tau_\ell = \inf\{t : X_t = b_\ell\}$ for some $b_\ell \in \mathbb{R}$ and $d_{\tau_\ell} = 0$ always.

This result shows that in the long-run, inefficient shut-down happens only when $|\mu_{\ell}| < |\mu_h|$ and the shut-down inefficiency approaches zero. This might seem to go against our previous results. However, from *P*'s perspective, auditing is always inefficient; the mechanism has already separated the types so there is not need for auditing. So while the shut-down inefficiency vanishes in the long-run, the inefficiency from auditing vanishes if and only if $|\mu_h| < |\mu_{\ell}|$. In this way, our previous results go through with the addition of auditing. By using auditing, *P* is essentially able to put a cap on the inefficiency needed to provide punishments to a deviating ℓ type. Given the strain that the optimal mechanism may put on the plausibility of our assumption that *P* can perfectly commit, this cap helps justify the form of the optimal mechanism even when the ability of the principal to keep promises ex-post is limited.

5.2 Transfers

We might wonder if the restriction to not allowing for monetary transfers will drastically affect the model. If we allow for any kind of monetary transfers, then, because both parties are risk-neutral, P can achieve his first best by constructing Cremer-Mclean type mechanisms. Therefore, a more natural question is whether or not monetary transfers when combined with limited liability will have a large impact. We begin by assuming that A still receives a flow utility α which P can reward in addition to with transfers. Formally, transfers W_t are an \mathcal{F}_t^X -measurable control variable which leads to utility for player i of

$$\mathbb{E}\left[\int_0^\tau e^{-rt} (u_i(X_t) + \beta_i dW_t) dt + e^{-r\tau} \beta_i (W_\tau - W_{\tau-})\right]$$

where $\beta_P = -1$ and $\beta_A = 1$. Since the discount rate is the same for P, A it is without loss to only focus on wage structures with payments at $t = \tau$ (i.e., $W_{\tau-} = 0$).

In fact, wages have little role in the model. Because there is no moral hazard and agent's prefer to keep the project going, wages only serve to incentivize the reporting of information. Since ℓ 's *IC* constraint is the one that is likely to bind, wages will be used to incentivize ℓ to report truthfully. Since it is only used to make ℓ report truthfully (and assuming h's *IC* constraints aren't binding), there is no need to condition payments on incoming information. Therefore the optimal wage to ℓ can be paid out at $t = \tau$. Additionally, no wage will ever be paid to h since adding a wage only increases ℓ 's temptation value of declaring himself to be h and decreases P's payoffs.

5.3 Multiple Reports

We have assumed so far that the agent receives all private information at time t = 0. In many real life situations though, information is received randomly over time. To incorporate this into the model, we allow the state to change over time. More specifically, the state follows a Markov chain which may potentially change states at a time τ^{γ} which exponentially distributed with parameter γ . For simplicity, we assume that $\Theta = \{\theta_h, \theta_\ell\}$ and that, conditional on the state switching, the probability of staying θ_i conditional on the current state being θ_i is ρ_i and $\rho_\ell = 1$. Importantly, we assume that the time at which the state may change is observed by both A and P, although only A observes the realization of the state.

This assumption on the observability of the times at which the agent acquires information is needed to keep track of the potential deviations the agent might make. For such problems, in which the agent's potential deviations are simple, the Lagrangian approach used in the previous section works well. This assumption is plausible in many economic situations. If X_t is the profitability of a division within the firm, then τ^{γ} may correspond to a loss of personnel or some public market shift which may potentially affect the firm's profitability.

The gain function at time τ^{γ} when the type changes from *i* to *k* is given by $V_i^k(W_i^i, W_i^k, X_{\tau^{\gamma}})$ where W_i^k is the expected continuation utility given to *A* who reported to be type *k* at t = 0 and reports to be type *i* at τ^{γ} . If W_i^i, W_i^k are not incentive compatible, then we set V_i^k to be $-\infty$. We can treat W_i^j as a decision variable of *P* at time τ^{γ} . This allows us to write out the principal's problem at time t = 0 as

$$\sup_{(\tau_i, W_k^i)_{i,k=h,\ell}} \sum_{i=h,\ell} \sum_{k=h,\ell} \mathbb{E}_{\theta_i} [\int_0^{\tau \wedge \tau^{\gamma}} e^{-rt} X_t dt + \mathbb{1}(\tau^{\gamma} < \tau) e^{-r\tau^{\gamma}} V_i^k(W_k^i, X_t)] \mathbb{P}(\theta_i, \theta_k)$$

subject to $\forall i, k = h, \ell$

$$\begin{aligned} IC(\theta_i, \theta_k) : & \mathbb{E}_{\theta_i} \left[\int_0^{\tau_i \wedge \tau^{\gamma}} \alpha dt + \mathbb{1}(\tau^{\gamma} < \tau) e^{-r\tau^{\gamma}} (\rho_i W_i^i + (1 - \rho_i) W_k^i) \right] \\ & \geq \mathbb{E}_{\theta_i} \left[\int_0^{\tau_k \wedge \tau^{\gamma}} \alpha dt + \mathbb{1}(\tau^{\gamma} < \tau_k) e^{-r\tau^{\gamma}} (\rho_i W_i^k + (1 - \rho_i) W_k^k) \right] \end{aligned}$$

Let us consider the problem type-by-type and conjecture that $IC(\theta_h, \theta_\ell)$ is slack. Let w_ℓ be the utility given to ℓ in his optimal mechanism. Then P's problem can be written as

$$\sup_{(\tau,W_h^h,W_\ell^h)} \mathbb{E}_{\theta_h} \left[\int_0^{\tau \wedge \tau^{\gamma}} e^{-rt} X_t dt + \mathbb{1}(\tau^{\gamma} < \tau) e^{-r\tau^{\gamma}} V(W_\ell^h, W_h^h, X_{\tau^{\gamma}}) \right]$$

subject to

$$IC(\theta_{\ell},\theta_{h}): \mathbb{E}_{\theta_{\ell}}[\int_{0}^{\tau \wedge \tau^{\gamma}} e^{-rt} \alpha dt + \mathbb{1}(\tau^{\gamma} < \tau)e^{-r\tau^{\gamma}}W_{\ell}^{h}] \le w_{\ell}$$

where $V(W_h^h, W_\ell^h, X_{\tau^{\gamma}}) := \rho_h V_h^h(W_h^h, W_\ell^h, X_{\tau^{\gamma}}) + (1 - \rho_h) V_h^\ell(W_h^h, W_\ell^h, X_{\tau^{\gamma}})$. Define $\widetilde{V}(W_\ell^h, X_t) := \sup_{W_h^h} V(W_h^h, W_\ell^h, X_t)$. Since W_h^h doesn't enter the constraint set of the problem for ℓ , we can replace V with \widetilde{V} , thereby simplify our problem to be the choice of τ and W_ℓ^h . Using the same arguments as in Section 4.2, the solution to above problem will also solve (for some Lagrange multiplier $\lambda \geq 0$)

$$\sup_{(\tau,W_{\ell}^{h})} \quad \mathbb{E}_{\theta_{h}}[e^{-r\tau} \big(\lambda \alpha M_{t} - (\frac{\mu_{h}}{r} + X_{t}) + \int_{0}^{\tau} e^{-(r+\gamma)t} \gamma (\lambda W_{\ell}^{h} M_{t} + \widetilde{V}(W_{\ell}^{h}, X_{t})) dt],$$

where M_t is the Radon-Nikodym derivative. Using the same techniques as in Section 4.2, we see that the optimal stopping boundary for h follows the same increasing/decreasing pattern before and after τ^{γ} .

Proposition 6. If the agent reports the state to be *i* at t = 0, then $\tau = \inf\{t : X_t \leq b_i(t)\}$ where $b_h(t)$ is increasing if and only if $\frac{\mu_h + \mu_\ell}{2} > 0$ while $b_\ell(t)$ is a constant threshold. After the first reporting opportunity, the optimal stopping boundary is again increasing after a report of θ_h if and only if $\frac{\mu_h + \mu_\ell}{2} > 0$. The optimal stopping boundary after a report of ℓ is constant.

Thus we see that the qualitative features of the optimal mechanism are robust to the public arrival of new information. The key assumption is that the time of the potential state change is observable by both P and A. If this doesn't hold, then it is difficult to know what the agent's best deviation on when to report a state change is. While Madsen (2016) explored this in a setting where P's flow payoff is equal to dX_t , further analysis in our model is beyond the technical means so far used.

6 Conclusion

We introduce a simple principal-agent model in which the agent has private information about the expected profitability of a project. Our model illustrates how the principal places distortions in order to elicit the agent's information and how the nature of these distortions depends on the parameters of the problem. When high types' drifts are high enough, back-loading distortions minimizes the expected probability that they are realized but leads to increasing ex-post inefficiency. On the other hand, high types' drifts are low enough, the threshold for h is decreasing over time as later termination provides less and less incentives to prevent a ℓ type from imitating h (and so the threshold approaches the principal's first best). The use of Lagrangian techniques and Girsanov's Theorem gives us a clear picture of the tensions in the model. We find that high types are offered one of two types of contracts: up-or-out or asymptotically stationary. Our framework allows us see what features of persistent private information drive the characteristics of the optimal mechanism and a clear picture as to why private information may create persistent distortions.

Additionally, we find that the qualitative structure of the optimal mechanism is robust to a number of extensions. While the up-or-out mechanisms might strain the commitment assumption, we believe that it illustrates a natural force in agency problems: that the principal wants to introduce distortions where they provide the most bang for their buck. We believe that generally this characterization of when distortions are front- or backloaded would carry over when we limit the commitment power of the principal. This finding might help explain why some professions feature up-or-out structures (e.g., limit on time to reach promotion or tenure) while some feature less stringent standards as time goes on.

References

- Balzer, Thomas, and Klaus Janßen. "A duality approach to problems of combined stopping and deciding under constraints." Mathematical methods of operations research 55, no. 3 (2002): 431-446.
- [2] DeMarzo, Peter M., and Yuliy Sannikov. "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model." The Journal of Finance 61.6 (2006): 2681-2724.
- [3] Fong, Kyna. "Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons." Working Paper (2007).
- [4] Garrett, Daniel F., and Alessandro Pavan. "Managerial turnover in a changing world." Journal of Political Economy 120.5 (2012): 879-925.
- [5] Guo, Yingni. "Dynamic delegation of experimentation." The American Economic Review 106.8 (2016): 1969-2008.
- [6] Karatzas, Ioannis, and Steven Shreve. Brownian motion and stochastic calculus. Vol. 113. Springer Science and Business Media, 2012.
- [7] Kuvalekar, Aditya V., and Elliot Lipnowski. "Job Insecurity." Working Paper (2016).
- [8] Madsen, Erik. "Optimal project termination with an informed agent." Working Paper (2016).
- [9] McClellan, Andrew. "Experimentation and Approval Mechanisms." Working Paper (2017).
- [10] Pavan, Alessandro, Ilya Segal, and Juuso Toikka. "Dynamic mechanism design: A myersonian approach." Econometrica 82.2 (2014): 601-653.

- [11] Peskir, Goran. "A change-of-variable formula with local time on curves." Journal of Theoretical Probability 18.3 (2005): 499-535.
- [12] Sannikov, Yuliy. "A continuous-time version of the principal-agent problem." The Review of Economic Studies 75.3 (2008): 957-984.
- [13] Williams, Noah. "Persistent private information." Econometrica 79.4 (2011): 1233-1275.

7 Appendix

7.1 First-Best

Proof of Proposition 1. With symmetric information, P's problem is a single-decision maker problem given by

$$\sup_{\tau} (\frac{\mu_i}{r} + X_0 - \mathbb{E}_{\theta_i, X_0}[e^{-r\tau_i}(\frac{\mu_i}{r} + X_{\tau_i})])\frac{1}{r}$$

It is easy to see that τ_i will take the form of a static threshold $\tau_i = \inf\{t : X_t \leq b_i\}$. The expected discounted time until τ when starting at X_0 is given by

$$e^{-R(b-X_0)}$$
 where $R = \frac{-\mu_i - \sqrt{\mu_i^2 + 2r\sigma^2}}{\sigma^2}$

Taking first-order conditions and solving gives us the optimal b_i . To verify that $\frac{\partial b_i}{\partial \mu_i} < 0$, we note that

$$\frac{\partial b_i}{\partial \mu_i} = -\frac{1}{r} - \sigma^2 (1 + \frac{\mu_i}{\sqrt{\mu_i^2 + 2r\sigma^2}}) < 0$$

7.2 Front-Loaded Distortions

Lemma 7. V(t, X) is convex in x.

Proof. Starting at (t, X), it is straightforward to verify that the gain function $e^{-r\tau}G(\tau, X_{\tau})$ is convex in x. Therefore, we know that $\mathbb{E}_x[e^{-r\tau}G(\tau, X_{\tau})]$ is convex in x. Moreover, $V(t, X) = \sup_{\tau} \mathbb{E}_x[e^{-r\tau}G(\tau, X_{\tau})]$ is convex in x since it is the sup over a set of convex functions.

Lemma 8. Smooth pasting holds in X: $\frac{\partial V(t,X)}{\partial X}|_{X=b_h(t)} = G_x(t,X).$

Proof. Let τ_{ϵ} be the optimal stopping rule when at $(t, X + \epsilon)$ and $X = b_h(t)$. Then we know that

$$V(t, X + \epsilon) - V(t, X) \leq \mathbb{E}_{X+\epsilon} [e^{-r(\tau_{\epsilon} - t)} (G(\tau_{\epsilon}, X_{\tau_{\epsilon}} + \epsilon))] - \mathbb{E}_{X} [e^{-r(\tau_{\epsilon} - t)} (G(\tau_{\epsilon}, X_{\tau_{\epsilon}}))]$$

$$= \mathbb{E}_{X+\epsilon} [e^{-r(\tau_{\epsilon} - t)} (\lambda \alpha (e^{\frac{\mu_{h} - \mu_{\ell}}{\sigma^{2}} (\frac{\mu_{h} + \mu_{\ell}}{2} \tau_{\epsilon} - X_{\tau_{\epsilon}} - \epsilon)} - e^{\frac{\mu_{h} - \mu_{\ell}}{\sigma^{2}} (\frac{\mu_{h} + \mu_{\ell}}{2} \tau_{\epsilon} - X_{\tau_{\epsilon}})}) - \epsilon)]$$

$$= \mathbb{E}_{X+\epsilon} [e^{-r(\tau_{\epsilon} - t)} (\epsilon (\lambda \alpha e^{\frac{\mu_{h} - \mu_{\ell}}{\sigma^{2}} (\frac{\mu_{h} + \mu_{\ell}}{2} \tau_{\epsilon} - X_{\tau_{\epsilon}})} \frac{\mu_{h} - \mu_{\ell}}{\sigma^{2}} - 1) + o(\epsilon^{2}))]$$

$$\Rightarrow \frac{V(t, X + \epsilon) - V(t, X)}{\epsilon} \leq \mathbb{E}_{X+\epsilon} [e^{-r(\tau_{\epsilon} - t)} G_{x}(\tau_{\epsilon}, X) + \frac{o(\epsilon^{2})}{\epsilon}]$$

Taking $\epsilon \to 0$, we have $\frac{\partial_+ V(t,X)}{\partial X} \leq \frac{\partial G(t,X)}{\partial X}$ (convexity of V implies that the right derivative exists).

To get the reverse inequality, we use the fact that $V(t, X + \epsilon) > G(t, X + \epsilon)$ and V(t, X) = G(t, X) to see that

$$\frac{V(t,X+\epsilon)-V(t,X)}{\epsilon} \geq \frac{G(t,X+\epsilon)-G(t,X)}{\epsilon}$$

Taking $\epsilon \to 0$, we have $\frac{\partial_+ V(t,X)}{\partial X} \ge \frac{\partial G(t,X)}{\partial X}$, completing the proof $\left(\frac{\partial_- V(t,X)}{\partial X} = G_x(t,X) \right)$ follows immediately from the definition of \mathcal{D} .).

Lemma 9. Smooth-pasting holds in t: $\frac{\partial V(t,X)}{\partial t}|_{X=b_h(t)} = G_t(t,X).$

Proof. Let $(t, X) \in \partial \mathcal{C}$. Let τ_{δ} be the optimal stopping rule at $(t + \delta, X)$ (as we will see in the proof of Lemma 1, $(t + \delta, X) \in \mathcal{C}$) and define $\tau = \tau_{\delta} - \delta$. Then we know that

$$V(t+\delta,X) - V(t,X) \leq \mathbb{E}[e^{-r(\tau+\delta-(t+\delta))}G(\tau+\delta,X_{\tau+\delta})] - \mathbb{E}[e^{-r(\tau-t)}G(\tau,X_{\tau})]$$

= $\mathbb{E}[e^{-r(\tau-t)}(\lambda\alpha e^{-\frac{\mu_h-\mu_\ell}{\sigma^2}X_{\tau}+\frac{\mu_h^2-\mu_\ell^2}{2\sigma^2}(\tau+\delta)} - (\frac{\mu_h}{r}+X_{\tau}))]$
- $\mathbb{E}[e^{-r(\tau-t)}(\lambda\alpha e^{-\frac{\mu_h-\mu_\ell}{\sigma^2}X_{\tau}+\frac{\mu_h^2-\mu_\ell^2}{2\sigma^2}\tau} - (\frac{\mu_h}{r}+X_{\tau}))]$
= $\mathbb{E}[e^{-r(\tau-t)}\lambda\alpha e^{-\frac{\mu_h-\mu_\ell}{\sigma^2}X_{\tau}+\frac{\mu_h^2-\mu_\ell^2}{2\sigma^2}\tau}](e^{\frac{\mu_h^2-\mu_\ell^2}{2\sigma^2}\delta} - 1).$

Dividing both sides by δ and taking $\delta \to 0$, we get that

$$\frac{\partial_+ V(t,X)}{\partial t} \leq \mathbb{E}[e^{-r(\tau-t)}\lambda\alpha e^{-\frac{\mu_h-\mu_\ell}{\sigma^2}X_\tau + \frac{\mu_h^2-\mu_\ell^2}{2\sigma^2}\tau}]\frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}$$

Because $e^{-\frac{\mu_h-\mu_\ell}{\sigma^2}X_\tau+\frac{\mu_h^2-\mu_\ell^2}{2\sigma^2}\tau}$ is always positive and is a martingale, we have that

$$\frac{\partial_+ V(t,X)}{\partial t} \leq \lambda \alpha \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t} = G_t(t,X).$$

For the reverse inequality, we note that V(t, X) = G(t, X) and $V(t + \delta, X) > G(t + \delta, X)$ (by definition of C), which implies that

$$\frac{V(t+\delta,X)-V(t,X)}{\delta} \geq \frac{G(t+\delta,X)-G(t,X)}{\delta}$$

Taking $\delta \to 0$ gives us that $\frac{\partial_+ V(t,X)}{\partial t} \ge G(t,X)$ and implies that $\frac{\partial_+ V(t,X)}{\partial t} = G_t(t,X)$. Checking that $\frac{\partial_- V(t,X)}{\partial t} = G_t(t,X)$ follows immediately from the definition of \mathcal{D} .

Proof of Lemma 1. We will show that $V(t', X) \in \mathcal{C} \Rightarrow V(t, X) \in \mathcal{C}$ for t' < t. We know that the optimal stopping rule given that the current state is (t, X) will be measurable with respect to the process starting at (t, X) (and ignoring all past events). This allows us to create an isomorphism from such stopping rules starting at t and t' by the mapping $\tau = \tau' + (t - t')$. Take a large T and define $\tau_T = \tau \wedge T$ and $\tau'_T = \tau_T + (t' - t)$.

We claim that $V(t, X) - \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t}$ is increasing in t. To see this, note that

$$\begin{split} & \mathbb{E}[e^{-r(\tau_{T}-t)}(\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{\tau}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}\tau}-(\frac{\mu_{h}}{r}+X_{\tau}))]-\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{t}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}t} \\ & = \mathbb{E}[e^{-r(\tau_{T}-t)}(-(\frac{\mu_{h}}{r}+X_{\tau_{T}})))+\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{\tau_{T}}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}\tau_{T}}(e^{-r(\tau_{T}-t)}-1)] \\ & = \mathbb{E}[e^{-r(\tau_{T}'-t')}(-(\frac{\mu_{h}}{r}+X_{\tau'})))+\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{\tau'_{T}}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}(\tau'_{T}+(t-t'))}(e^{-r(\tau'_{T}-t')}-1)] \\ & > \mathbb{E}[e^{-r(\tau'_{T}-t')}(-(\frac{\mu_{h}}{r}+X_{\tau'_{T}})))+\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{\tau'_{T}}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}\tau'_{T}}(e^{-r(\tau'_{T}-t')}-1)] \\ & = \mathbb{E}[e^{-r(\tau'_{T}-t')}(\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{\tau'_{T}}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}\tau'_{T}}-(\frac{\mu_{h}}{r}+X_{\tau'_{T}}))]-\lambda\alpha e^{-\frac{\mu_{h}-\mu_{\ell}}{\sigma^{2}}X_{t'}+\frac{\mu_{h}^{2}-\mu_{\ell}^{2}}{2\sigma^{2}}t'}} \end{split}$$

where equality holds by the fact that $M_t = e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t}$ is a martingale and τ_T, τ_T' are bounded by T so the Optimal Sampling Theorem implies that

$$e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t} = \mathbb{E}[e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_\tau + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}\tau}]$$

By taking the sup over τ and τ' and $T \to \infty$, we get our desired result.

Thus we have that

$$V(t,X) - G(t,X) = V(t,X) - \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t} + \frac{\mu_h}{r} + X$$

$$\geq V(t',X) - \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t'} + \frac{\mu_h}{r} + X$$

$$= V(t',X) - G(t',X) > 0$$

So $V(t', X) \in \mathcal{C} \Rightarrow V(t, X) \in \mathcal{C}$, implying that b_h is decreasing.

We now want to show that b_h is continuous. For the sake of contradiction, suppose that b_h had a jump at t. Fix an $X \in (b_h(t+), b_h(t-))$ and consider some t' > t. Note that we can rewrite the difference between V and G as

$$V(t',X) - G(t',X) = \int_{b_h(t'+)}^X \int_{b_h(t'+)}^u [V_{XX}(t',w) - G_{XX}(t',w)] dw du.$$
(2)

By Ito's Lemma, we know that

$$V_{XX}(t',X) - G_{XX}(t',X) = \frac{2}{\sigma^2} [rV(t',X) - V_t(t',X) - \mu_h V_X(t',X)] - \lambda \alpha \frac{(\mu_h - \mu_\ell)^2}{\sigma^4} e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t'}$$

Using Lemmas 9 and 10, we have that for t' close to t

$$V_X(t',X) \approx G_X(t,X) = -\lambda \alpha \frac{\mu_h - \mu_\ell}{\sigma^2} e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t'} - 1,$$

$$V_t(t',X) \approx G_t(t',X) = \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} X + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t'}.$$

Substituting these into equation 2, we have

$$\begin{split} V(t',X) - G(t',X) \approx & \int_{b_h(t'+)}^X \int_{b_h(t'+)}^u \left[\frac{2}{\sigma^2} (rV(t',w) - \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t'} (\frac{\mu_h^2 - \mu_\ell^2}{\sigma^2} - \mu_h \frac{\mu_h - \mu_\ell}{\sigma^2}) \right] \\ & - \lambda \alpha \frac{(\mu_h - \mu_\ell)^2}{\sigma^4} e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t'} dw du \\ & = \int_{b_h(t'+)}^Z \int_{b_h(t'+)}^u \left[\frac{2}{\sigma^2} (rV(t',w) + \mu_h) \right] dw du, \end{split}$$

Note that for it be optimal to stop at X, it must be that $de^{-rt}G(t, X) < 0$, which implies (by Ito's Lemma) that

$$-rG(t,X) + G_t(t,X) + \mu_h G_X(t,X) + \frac{\sigma^2}{2} G_{XX}(t,X) < 0.$$

Using our functional form for G, this simplifies to $\mu_h + rG(t, X) > 0$. Since V(t, X) > 0, this implies that $\mu_h + rV(t, X) > 0$. Since G(t, X) is decreasing in X and $\mu_h + rG(t, b_h(t-)) > 0$, there exists a $c, \delta > 0$ such that $\mu_h + rG(t, X) > c$ for all $X \in [b_h(t+), b_h(t+) + \delta]$. Therefore, we can conclude that

$$\begin{split} V(t',X) - G(t',X) &\approx \int_{b_{h}(t'+)}^{Z} \int_{b_{h}(t'-)}^{u} \left[\frac{2}{\sigma^{2}} (rV(t',w) + \mu_{h}) \right] dw du \\ &\geq \int_{b_{h}(t'+)}^{b_{h}(t'+)+\delta} \int_{b_{h}(t'+)}^{b_{h}(t'+)+\delta} \left[\frac{2}{\sigma^{2}} (rV(t',w) + \mu_{h}) \right] dw du \\ &\geq \int_{b_{h}(t'+)}^{b_{h}(t'+)+\delta} \int_{b_{h}(t'+)}^{b_{h}(t'+)+\delta} c dw du \\ &\geq c \frac{\delta^{2}}{2} > 0 \end{split}$$

Therefore, we can conclude that (given a small enough ϵ) for all $t' \in [t, t + \epsilon]$, we have

$$V(t',X)-G(t',X)>c\frac{\delta^2}{3}>0.$$

Taking $t' \to t$, since $X > b_h(t+)$, we have that V(t, X) > G(t, X), a contradiction of $(t, X) \in \mathcal{D}$. Therefore we cannot have a jump in $b_h(t)$.

Proof of Lemma 2. Suppose that b_h crosses b_h^{FB} . Since b_h is decreasing, it will never rise above b_h^{FB} again. Now consider the alternative stopping rule \hat{b}_h which is defined as

$$\hat{b}_{h}(t) = \begin{cases} b_{h}(t) & \text{if } t \leq sup\{s : \ b(s) > b_{h}^{FB}\}, \\ b_{h}^{FB} & \text{if } t > sup\{s : \ b(s) > b_{h}^{FB}\}. \end{cases}$$

Because this shut-down rule leads to quicker shutdown, \hat{b}_h is less attractive for ℓ than b_h was. Additionally, from time $sup(s: b_h(s) > b_h^{FB})$ onward, \hat{b}_h leads to a higher utility for the principal since it delivers the first best outcome. Therefore, b_h cannot have been optimal.

Proof of Lemma 3. We want to show that for $\widetilde{G}^{\epsilon}(x) = \epsilon - (\frac{\mu_h}{r} + x)$, the function $\widetilde{V}(x) \approx V(t, X)$ for t large enough.

Because $b_h(t)$ is decreasing and bounded below by b_h^{FB} , we know that $b_h(t)$ is converging to some constant line $\underline{b} = \sup\{b : b \leq b_h(t) \ \forall t\}$. Let $\tau(\underline{b}) = \inf\{s \geq t : X_t = \underline{b}\}$. For large t, we will show that $\underline{b} > b_h^{FB}$ is strictly sub-optimal.

For large t, the Radon-Nikodym derivative $M_t = e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t} \to 0$ unless $X_t \to -\infty$ and thus the distortion term becomes negligible. Since X_t is bounded below, we can conclude that $M_t \to 0$ and hence we have

$$V(t,X) = \sup_{\tau \ge t} \mathbb{E}[e^{-r(\tau-t)}(\lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_\tau + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}\tau} - (\frac{\mu_h}{r} + X_\tau))]$$

$$\approx \mathbb{E}[e^{-r(\tau(\underline{b}) - t)}(\lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_{\tau(\underline{b})} + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}\tau(\underline{b})} - (\frac{\mu_h}{r} + X_{\tau(\underline{b})}))].$$

Let t be large enough and $\epsilon = \lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} \underline{b} + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} t}$. Then we have

$$V(t,X) = \sup_{\tau \ge t} \mathbb{E}[e^{-r(\tau-t)}(\lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}X_\tau + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}\tau} - (\frac{\mu_h}{r} + X_\tau))]$$

$$\leq \sup_{\tau \ge t} \mathbb{E}[e^{-r(\tau-t)}(\lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2}\underline{b} + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t} - (\frac{\mu_h}{r} + X_\tau))]$$

$$= \sup_{\tau \ge t} \mathbb{E}[e^{-r(\tau-t)}(\epsilon - (\frac{\mu_h}{r} + X_\tau))].$$

Let $\tau_{\epsilon} = argsup_{\tau \geq t} \mathbb{E}[e^{-r\tau^*}(\epsilon - (\frac{\mu_h}{r} + X_{\tau^*}))]$. Since this problem is time-homogeneous, the optimal stopping policy will take a threshold form, $\tau_{\epsilon} = inf\{s \geq t : X_s = b_h^{\epsilon}\}$ for some b_h^{ϵ} . As $t \to \infty$, we have $\epsilon \to 0$; it is easy to show that $b_h^{\epsilon} \to b_h^{FB}$. Let $\tau(b) := inf\{s \geq t : X_t \leq b\}$. Then, since $\underline{b} > b^{FB}$, for small ϵ we know that

$$\begin{split} V(t,X) &< \mathbb{E}[e^{-r(\tau(\underline{b})-t)} \left(\epsilon - \left(\frac{\mu_h}{r} + \underline{b}\right)\right)] \\ &< \mathbb{E}[-e^{-r(\tau(b^{FB})-t)} \left(\frac{\mu_h}{r} + b^{FB}\right)] \\ &< \mathbb{E}[e^{-r(\tau(b^{FB})-t)} (\lambda \alpha e^{-\frac{\mu_h - \mu_\ell}{\sigma^2} b^{FB} + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} \tau(b^{FB})} - \left(\frac{\mu_h}{r} + b^{FB}\right))], \end{split}$$

where the final line is the payoff that P would get from using the stopping rule $\tau(b^{FB})$, contradicting the the optimality of V. Therefore, we cannot have $\underline{b} > b^{FB}$.

7.3 Back-Loaded Distortions

Proof of Lemma 4. We claim that $\frac{V(t,X)}{e^{-rt}} - \lambda \alpha e^{\frac{-(\mu_h - \mu_\ell)}{\sigma^2}X_t + \frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2}t}$ is decreasing in t. This follows from a slight modification of the proof of Lemma 1 and noting that $\frac{\mu_h^2 - \mu_\ell^2}{2\sigma^2} > 0$. Suppose that V(t,X) - G(t,X) > 0. We want to show that V(t',X) - G(t',X) > 0 for t' < t. This follows from a similar argument as in Lemma 1.

Proof of Lemma 5. We claim that $\frac{\tilde{V}(Z,t)}{e^{-rt}} + \frac{\mu_h + \mu_\ell}{2}t$ is increasing in t. Let t' < t and define $\tau' = \tau + (t' - t)$. Then we know that

$$\begin{split} & \mathbb{E}[e^{-r(\tau-t)}(\lambda\alpha e^{-\frac{2\tilde{\mu}}{\sigma^2}Z_{\tau}} - (\frac{\mu}{r} + Z_{\tau} + \frac{\mu_h + \mu_\ell}{2}\tau)] + \frac{\mu_h + \mu_\ell}{2}t \\ &= \mathbb{E}[e^{-r(\tau-t)}(\lambda\alpha e^{-\frac{2\tilde{\mu}}{\sigma^2}Z_{\tau}} - (\frac{\mu}{r} + Z_{\tau} + \frac{\mu_h + \mu_\ell}{2}(\tau-t)) - t(e^{-r(\tau-t)} - 1)\frac{\mu_h + \mu_\ell}{2}] \\ &= \mathbb{E}[e^{-r(\tau'-t')}(\lambda\alpha e^{-\frac{2\tilde{\mu}}{\sigma^2}Z_{\tau'}} - (\frac{\mu}{r} + Z_{\tau'} + \frac{\mu_h + \mu_\ell}{2}(\tau'-t')) - t(e^{-r(\tau'-t')} - 1)\frac{\mu_h + \mu_\ell}{2}] \\ &\geq \mathbb{E}[e^{-r(\tau'-t')}(\lambda\alpha e^{-\frac{2\tilde{\mu}}{\sigma^2}Z_{\tau'}} - (\frac{\mu}{r} + Z_{\tau'} + \frac{\mu_h + \mu_\ell}{2}(\tau'-t')) - t'(e^{-r(\tau'-t')} - 1)\frac{\mu_h + \mu_\ell}{2}] \\ &= \mathbb{E}[e^{-r(\tau'-t')}(\lambda\alpha e^{-\frac{2\tilde{\mu}}{\sigma^2}Z_{\tau'}} - (\frac{\mu}{r} + Z_{\tau'} + \frac{\mu_h + \mu_\ell}{2}(\tau'-t')) - t'(e^{-r(\tau'-t')} - 1)\frac{\mu_h + \mu_\ell}{2}] \\ &= \mathbb{E}[e^{-r(\tau'-t')}(\lambda\alpha e^{-\frac{2\tilde{\mu}}{\sigma^2}Z_{\tau'}} - (\frac{\mu}{r} + Z_{\tau'} + \frac{\mu_h + \mu_\ell}{2}(\tau'-t')) - t'(e^{-r(\tau'-t')} - 1)\frac{\mu_h + \mu_\ell}{2}] \end{split}$$

Taking the sup over τ , we get our desired result. Therefore, we have that

$$\begin{split} \widetilde{V}(t,Z) - \widetilde{G}(t,Z) &= \frac{\widetilde{V}(Z,t)}{e^{-rt}} + \frac{\mu_h + \mu_\ell}{2} t - (\lambda \alpha e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z} - (\frac{\mu}{r} + Z)) \\ &\geq \widetilde{V}(t',Z) + \frac{\mu_h + \mu_\ell}{2} t' - (\lambda \alpha e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z} - (\frac{\mu}{r} + Z)) \\ &= \widetilde{V}(t',Z) - \widetilde{G}(t',Z) > 0. \end{split}$$

Therefore, $\widetilde{V}(t', Z) \in \mathcal{C} \Rightarrow \widetilde{V}(t, Z) \in \mathcal{C}$.

We now want to show that $b_h^Z(t)$ is continuous in t. For the sake of contradiction, suppose that there was a jump at t. Let our value function be $\widetilde{V}(t, Z)$. It follows from the same arguments as in Lemma 9 that at the boundary of C, we have $\frac{\partial \widetilde{V}(t,Z)}{\partial Z} = \frac{\partial \widetilde{G}(t,Z)}{\partial Z} = \frac{-2\widetilde{\mu}}{\sigma^2}\lambda\alpha e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z} - 1$. We note that $\widetilde{V}(t,Z) \ge 0$ (since otherwise the principal to could wait never stop and guarantee himself a payoff of zero) and $\widetilde{V}_t(t,X) \le 0$ (since $\frac{\mu_h + \mu_\ell}{2} > 0$).

Fix $Z \in (b_h(t-), b_h(t+))$.Note that

$$\widetilde{V}(t',Z) - \widetilde{G}(t',Z) = e^{rt} \int_{b(t'-)}^{Z} \int_{b_h(t'-)}^{u} (\widetilde{V}_{ZZ}(t',w) - \widetilde{G}_{ZZ}(t',w)) dw du$$

By Ito's Lemma and our previous claims, we have that for t' close to t,

$$\begin{split} \widetilde{V}_{ZZ}(t,Z) &- \widetilde{G}_{ZZ}(t,Z) = \frac{2}{\sigma^2} [r \widetilde{V}(t,Z) - \widetilde{V}_t(t,Z) - \widetilde{\mu} V_Z(t,Z)] - \lambda \alpha \frac{4\widetilde{\mu}^2}{\sigma^4} e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z} \\ &= \frac{2}{\sigma^2} [r \widetilde{V}(t,Z) - \widetilde{V}_t(t,Z) - \widetilde{\mu} (\frac{-2\widetilde{\mu}}{\sigma^2} \lambda \alpha e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z} - 1)] - \lambda \alpha \frac{4\widetilde{\mu}^2}{\sigma^4} e^{-\frac{2\widetilde{\mu}}{\sigma^2}Z} \\ &\geq \frac{2}{\sigma^2} [r \widetilde{V}(t,Z) + \mu_h] \\ &> 0 \end{split}$$

since $\mu_h > 0$ and which implies that $\widetilde{V}(t', Z) - \widetilde{G}(t', Z) > 2\frac{\widetilde{\mu}}{\sigma^2} \frac{(Z-b_h(t'-))^2}{2} > 0$; taking $t' \to t$, we have $\widetilde{V}(t, Z) - \widetilde{G}(t, Z) > 0$ a contradiction of \mathcal{D} . Therefore, there cannot have been a jump and $b_h^Z(t)$ must be continuous.

Proof of Lemma 6. Suppose for the sake of contradiction that $\exists \bar{b}$ such that $\bar{b} > b_h(t) \forall t$. This implies that for any fixed X_t , we have that $\lambda \alpha e^{\frac{\mu_h - \mu_\ell}{\sigma^2}(-X_\tau + \frac{\mu_h + \mu_\ell}{2}\tau)} \to \infty$ as $t \to \infty$. Since μ_h is positive, with positive probability it will never cross \bar{b} and as we increase X_t , we will have that the probability it will never cross \bar{b} is $e^{-\frac{2\mu_h}{\sigma^2}(X_t - \bar{b})}$, which approaches 1 as $X_t \to \infty$.

Given a > 0, b, the density of the first-passage time of a Brownian motion starting at zero at time t is given by

$$f(t') = \frac{a}{\sqrt{2\pi(t'-t)^3}} e^{-\frac{(a+b(t'-t))^2}{t'-t}}$$

Therefore, the density of the first-crossing time of \bar{b} when starting at time t and X is given by

$$f(t'|t,X) = \frac{X_t - \bar{b}}{\sqrt{2\pi(t'-t)^3}} e^{-\frac{(X_t - \bar{b} - \mu_h(t'-t))^2}{\sigma^2(t'-t)}}$$

Fix a small $\delta > 0$. Since the gain function G is decreasing in X, for large t the value function V(t, X) will be bounded above by the expected utility of stopping at \bar{b} , when the payoff of stopping at \bar{b} is $G(t, \bar{b} - \delta)$ (since the process is stopping sooner and at a higher payoff).

$$V(t,X) \leq \int_{t}^{\infty} e^{-r(t'-t)} (\lambda \alpha e^{\frac{\mu_{h} - \mu_{\ell}}{\sigma^{2}} (-(\bar{b}-\delta) + \frac{\mu_{h} + \mu_{\ell}}{2} t')} - (\frac{\mu_{h}}{r} + \bar{b} - \delta)) f(t'|X) dt'$$
(3)
$$= \int_{0}^{\infty} e^{-rs} \lambda \alpha e^{\frac{\mu_{h} - \mu_{\ell}}{\sigma^{2}} (-(\bar{b}-\delta) + \frac{\mu_{h} + \mu_{\ell}}{2} (t+s)} \frac{X_{t} - \bar{b}}{\sqrt{2\pi s^{3}}} e^{-\frac{(X_{t} - \bar{b} - \mu_{h} s)^{2}}{\sigma^{2} s}} ds$$
$$- \int_{0}^{\infty} e^{-rs} (\frac{\mu_{h}}{r} + \bar{b} - \delta) \frac{X_{t} - \bar{b}}{\sqrt{2\pi s^{3}}} e^{-\frac{(X_{t} - \bar{b} - \mu_{h} s)^{2}}{\sigma^{2} s}} ds$$

If we take $X = \frac{\mu_h + \mu_\ell}{3}t + \bar{b}$, we can see that the exponential terms in the second line of 3 are equal to

$$\frac{\mu_h - \mu_\ell}{\sigma^2} (-\bar{b} + \delta) + \frac{\mu^2 - \mu_\ell^2}{2\sigma^2} (t + s) - \frac{(\mu_h - \mu_\ell)^2}{9\sigma^2 s} t^2 + 2\mu_h \frac{\mu_h + \mu_\ell}{3\sigma^2} t - \frac{\mu_h^2}{\sigma^2} s^2$$

Dropping the constant $\frac{\mu_h - \mu_\ell}{\sigma^2} (-\bar{b} + \delta)$ and factoring out $\frac{1}{\sigma^2}$, we are left with

$$\frac{\mu^2 - \mu_\ell^2}{2}(t+s) - \frac{(\mu_h - \mu_\ell)^2}{9s}t^2 + 2\mu_h \frac{\mu_h + \mu_\ell}{3}t - \mu_h^2 s^2$$

which can be made arbitrarily negative for all s by taking $t \to \infty$. Using a similar argument for the third line of 3, we can see that the upper bound on V(t, X) can be made arbitrarily small by taking t large enough. However, for large t, we have that $G(t, \bar{b} + \frac{\mu_h + \mu_\ell}{2}t) \to \infty$, contradicting the fact that V(t, X) > G(t, X). Therefore, b_h cannot be bounded above.

7.4 Binding Constraint

The previous Lemmas completely characterize the solution to the relaxed problem in which $IC(h, \ell)$ is dropped. In order to complete the proof of Theorem 1 we only need to verify that the solution to the relaxed problem doesn't violate $IC(h, \ell)$.

Proposition 7. The solution to the relaxed problem satisfies $IC(h, \ell)$.

Proof. It is clear that $IC(\ell, h)$ must bind in our relaxed problem; otherwise, P could get his first-best which violates $IC(\ell, h)$.

Suppose that the solution to the relaxed problem did violate $IC(h, \ell)$ and let J_h^h be the utility that h gets from $b_h(t)$ in the relaxed problem. Now lets define a relaxed version of the relaxed problem: lets find the optimal $\hat{b}_h(t)$ only under the constraint that it delivers J_h^h to h. Using Lagrangian techniques, we get

$$\sup_{\tau} \mathbb{E}_{\theta_h} \left[e^{-r(\tau-t)} (\lambda \alpha_h^h - (\frac{\mu_h}{r} + X_{\tau})) \right].$$

which has a solution of a static threshold-i.e., $\hat{b}_h(t) = \hat{b}_h \in \mathbb{R}$. Because this delivers J_h^h to h and J_h^h is less than the utility he would get from $b_\ell(t)$ (which is also a static threshold b_ℓ), we know that $\hat{b}_h > b_\ell$. This shows that $IC(\ell, h)$ is satisfied and \hat{b}_h is allowable in our original relaxed problem with only $IC(\ell, h)$. Moreover, it yields at least as much utility as $b_h(t)$ and $IC(\ell, h)$ is slack, a contradiction. Therefore, we cannot have $IC(h, \ell)$ violated.

7.5 Optimal $b_h(t)$

Proof of Proposition 3. Since $b_h(t)$ is continuous, we can use the change of variable formula from Peskir (2005) to show that

$$e^{-r(T-t)}V(T, X_T) = V(t, X) + \int_t^T e^{-r(s-t)} (\mathbb{L}V(s, X_s) - rV(s, X_s)) \mathbb{1}(X_s \neq b_h(s)) ds + Y_T + \frac{1}{2} \int_t^T \Delta_X V(s, X_s) \mathbb{1}(X_s = b_h(s)) d\ell_s,$$

where $\Delta_s V_Z(s, b_h(s)) = V_Z(s, b_h(s)+) - V_Z(s, b_h(s)-)$, Y_T is a martingale and ℓ_s is the local time of X_s which is given by

$$\ell_s = P_{s,X} - \lim_{\epsilon \downarrow 0} \frac{1}{2\epsilon} \int_t^s \mathbb{1}(b_h(u) - \epsilon < X_u < b_h(u) + \epsilon).$$

By the principle of smooth fit, the last integrand is equal to zero. If we take $P_{t,X}$ expectations of both sides, we have

$$\mathbb{E}[e^{-r(T-t)}V(T,X_T)] = V(t,X) + \int_t^T e^{-r(s-t)} \mathbb{E}[(\mathbb{L}V(s,X_s) - rV(s,X_s))\mathbb{1}(X_s \neq b_h(s))]ds.$$

Evaluating V(t, X) at $b_h(t)$, we have that

$$\mathbb{E}e^{-r(T-t)}[V(T,X_T)] = G(t,b_h(t)) + \int_t^T e^{-rs} \mathbb{E}[(\mathbb{L}G(s,X_s) - rG(s,X_s))] \mathbb{1}_{X_s < b_h(s)}]ds$$

Taking $T \to \infty$, we get our desired solution. To complete the proof, we only need to check the uniqueness of $b_h(t)$ within an appropriate class of functions is verified in Lemma 11.

Lemma 10. The function $b_h(t)$ is the unique solution to 1 in the class of functions $\{c : \mathbb{R}_+ \to \mathbb{R} : c \in C^1 \text{ and } -\mu_h - rG(t, c(t)) < 0 \ \forall t \in [0, \infty)\}.$

Proof. Let c(t) be another function in the relevant class which solves the integral equation such that $\forall t \ G(t, c(t)) \ge 0$. We will show that c and b coincide. Let us define U(t, X) as

$$U(t, X_t) = \int_t^\infty e^{-r(s-t)} \mathbb{E}[(\mathbb{L}G(s, X_s) - rG(s, X_s))\mathbb{1}(X_s < c(s)))]ds.$$

and $V^c(t, X)$ as

$$V^{c}(t, X_{t}) = \begin{cases} U(t, X_{t}) & \text{if } X_{t} > c(t), \\ G(t, X_{t}) & \text{if } X_{t} \le c(t). \end{cases}$$

By standard arguments, we know that $\mathbb{L}V^{c}(t, X_{t}) - rV^{c}(t, X_{t}) = 0$ for $X_{t} \geq c(t)$. Using the change-of-variable formula, we know that

$$V^{c}(t, X_{t}) = \int_{t}^{\infty} e^{-r(s-t)} \mathbb{E}[(\mathbb{L}G(s, X_{s}) - rG(s, X_{s}))\mathbb{1}(X_{s} < c(s)))]ds$$
$$-\frac{1}{2} \int_{t}^{\infty} e^{-r(s-t)} \Delta_{x} V_{x}^{c}(t+s, c(t+s)) \mathbb{E}_{x}[\ell_{s}^{c}(X_{s})].$$

By definition of V^c and U, we have that

$$\begin{split} \mathbb{1}(X_t \le c(t))G(t, X_t) + \mathbb{1}(X_t > c(t))U(t, X_t) = &U(t, X_t) - \frac{1}{2} \int_t^\infty e^{-rs} \Delta_x V_x^c(t+s, c(t+s)) \mathbb{E}_x[\ell_s^c(X_s)] \\ \Rightarrow 2\mathbb{1}(X_t \le c(t))(U(t, X_t) - G(t, X_t)) = \int_t^\infty e^{-rs} \Delta_x V_x^c(t+s, c(t+s)) \mathbb{E}_x[\ell_s^c(X_s)]. \end{split}$$

Thus, if $U(t, X_t) = G(t, X_t)$ for all $X_t \leq c(t)$, then it must be that V^c is C^1 at $X_t = c(t)$.

Define a stopping time $\tau_c^+ = \inf\{s \ge t : X_s \ge c(s)\}$ and lets consider an $X_t < c(t)$. We know that, since c(t) solves the integral equation, $U(\tau_c, X_{\tau_c}) = G(\tau_c, X_{\tau_c})$. Moreover, it is straightforward to show that both U, G are C^1 and hence, satisfy

$$e^{-rT}U(T, X_T) = U(t, X) + \int_t^T e^{-r(s-t)} \left(\mathbb{L}U(s, X_s) - rU(s, X_s) \right) \mathbb{1}(X_s \neq c(s)) ds + Y_t^U$$
(4)

$$e^{-rT}G(T, X_T) = G(t, X) + \int_t^T e^{-r(s-t)} \left(\mathbb{L}G(s, X_s) - rG(s, X_s) \right) \mathbb{1}(X_s \neq c(s)) ds + Y_t^G$$
(5)

for some martingales Y^U_t, Y^G_t . Taking $T \to \infty$ and expectations of both sides when evaluated at X = c(t), we get

$$0 = U(t, X) + \mathbb{E}\left[\int_{t}^{\infty} e^{-r(s-t)} \left(\mathbb{L}U(s, X_{s}) - rU(s, X_{s})\right) \mathbb{1}(X_{s} \neq c(s))ds\right]$$

= $U(t, X) + \mathbb{E}\left[\int_{t}^{\infty} e^{-r(s-t)} \left(\mathbb{L}U(s, X_{s}) - rU(s, X_{s})\right) \mathbb{1}(X_{s} < c(s))ds\right]$
= $U(t, X) - G(t, X)$

where the first equality follows from definition of U and the second equality follows from the fact that c(t) solves equation 1; therefore, at X = c(t), we have U(t, X) = G(t, X).

Replacing T with τ_c^+ and taking expectations on 4, we have

$$\mathbb{E}[e^{-r\tau_{c}^{+}}U(\tau_{c}^{+},X_{\tau_{c}^{+}})] = U(t,X) + \mathbb{E}[\int_{t}^{\tau_{c}} e^{-r(s-t)} \big(\mathbb{L}U(s,X_{s}) - rU(s,X_{s})\big)\mathbb{1}(X_{s} < c(s))ds]$$
$$\mathbb{E}[e^{-r\tau_{c}}G(\tau_{c}^{+},X_{\tau_{c}^{+}})] = G(t,X) + \mathbb{E}[\int_{t}^{\tau_{c}^{+}} e^{-r(s-t)} \big(\mathbb{L}G(s,X_{s}) - rG(s,X_{s})\big)\mathbb{1}(X_{s} \neq c(s))ds]$$

Using our change-of-variable formula, we have

$$\begin{aligned} U(t, X_t) &= \mathbb{E}_{X_t} [e^{-r(\tau_c^+ - t)} U(\tau_c^+, X_{\tau_c^+})] - \mathbb{E}_{X_t} [\int_t^{\tau_c^+} e^{-r(s-t)} \big(\mathbb{L}G(s, X_s) - rG(s, X_s) \big) \mathbb{1}(X_s < c(s)) ds] \\ &= \mathbb{E}_{X_t} [e^{-r(\tau_c^+ - t)} G(\tau_c^+, X_{\tau_c^+})] - \mathbb{E}_{X_t} [\int_t^{\tau_c^+} e^{-r(s-t)} \big(\mathbb{L}G(s, X_s) - rG(s, X_s) \big) \mathbb{1}(X_s < c(s)) ds] \\ &= G(t, X_t) + \mathbb{E}_{X_t} [\int_t^{\tau_c^+} e^{-r(s-t)} \big(\mathbb{L}G(s, X_s) - rG(s, X_s) \big) ds] \\ &- \mathbb{E}_{X_t} [\int_t^{\tau_c^+} e^{-r(s-t)} \big(\mathbb{L}G(s, X_s) - rG(s, X_s) \big) \mathbb{1}(X_s < c(s)) ds] \\ &= G(t, X_t). \end{aligned}$$

Thus, we have our desired result and so V^c is C^1 at c(t).

Now we want to show that $V(t, X) \ge V^c(t, X)$. To do this, we define $\tau_c^- = \inf\{s \ge t : X_s \le c(s)\}$. Again using the change-of-variable formula, we have

$$e^{-r(s-t)}V^{c}(s,X_{s}) = V^{c}(t,X_{t}) + \int_{t}^{\infty} e^{-r(s-t)} [\mathbb{L}G(s,X_{s}) - rG(s,X_{s})] \mathbb{1}(X_{s} \le c(s))ds + M_{t}^{c},$$

where M_t^c is a martingale. Then replacing s with τ_c and taking expectations, we have

$$V^c(t, X_t) = \mathbb{E}[e^{-r\tau_c^-} V^c(\tau_c^-, X_{\tau_c^-})]$$

Hence by definition of V, we have $V(t, X) \ge V^{c}(t, X)$.

Let us now show that $c(t) \ge b(t)$. Pick an $X < \min\{b(t), c(t)\}$ such that $\mathbb{L}G(t, X) - rG(t, X) < 0$ (remember $\mathbb{L}G(s, X_s) - rG(s, X_s) < 0$ at $X_s = b(s)$) and define a stopping rule $\tau_b^+ = \inf\{s \ge t : X_s \ge b(s)\}$ and $\tau_- = \inf\{s \ge t : \mathbb{L}G(s, X_s) - rG(s, X_s) \ge 0\}$. Using the change of variable formula, we have

$$\mathbb{E}[e^{-r(\tau_{b}^{+}\wedge\tau_{-}-t)}V^{c}(\tau_{b}^{+}\wedge\tau_{-},X_{\tau_{b}^{+}\wedge\tau_{-}})] = G(t,X_{t}) + \mathbb{E}[\int_{t}^{\tau_{b}^{+}\wedge\tau_{-}}e^{-r(s-t)}\left(\mathbb{L}G(s,X_{s}) - rG(s,X_{s})\right)\mathbb{1}(X_{s} < c(s))ds]$$
$$\mathbb{E}[e^{-r(\tau_{b}^{+}\wedge\tau_{-}-t)}V(\tau_{b}^{+}\wedge\tau_{-},X_{\tau_{b}^{+}\wedge\tau_{-}})] = G(t,X_{t}) + \mathbb{E}[\int_{t}^{\tau_{b}^{+}\wedge\tau_{-}}e^{-r(s-t)}\left(\mathbb{L}G(s,X_{s}) - rG(s,X_{s})\right)ds]$$

With the fact that $V(t, X) \ge V^c(t, X)$, we have that

$$\mathbb{E}\left[\int_{t}^{\tau_{b}^{+} \wedge \tau_{-}} e^{-r(s-t)} \left[\mathbb{L}G(s, X_{s}) - rG(s, X_{s})\right] ds\right] \geq \mathbb{E}\left[\int_{t}^{\tau_{b}^{+} \wedge \tau_{-}} e^{-r(s-t)} \left(\mathbb{L}G(s, X_{s}) - rG(s, X_{s})\right) \mathbb{1}(X_{s} < c(s))\right] ds$$

By the continuity of c, b, this implies that $c(t) \ge b(t)$ for all t. Now suppose that c(t) > b(t) for some t. Pick an $x \in (b(t), c(t))$. By using the change-of-variable formula, we have that

$$\begin{split} \mathbb{E}[e^{-r(\tau_b^--t)}G(\tau_b^-, X_{\tau_b^-})] &= V^c(t, X) + \mathbb{E}[\int_t^{\tau_b^-} e^{-r(s-t)}(\mathbb{L}G(s, X_s) - rG(s, X_s))\mathbb{1}(X_s < c(s))],\\ \mathbb{E}[e^{-r(\tau_b^--t)}G(\tau_b^-, X_{\tau_b^-})] &= V(t, X). \end{split}$$

This, together with the fact that $V(t, X) \ge V^c(t, X)$, implies that

$$\mathbb{E}[\int_{t}^{\tau_{b}^{-}} e^{-r(s-t)} (\mathbb{L}G(s, X_{s}) - rG(s, X_{s})) \mathbb{1}(X_{s} < c(s))] \ge 0.$$

Since $\mathbb{L}G(s, X_s) - rG(s, X_s) < 0$ for $X_s < c(s)$ by assumption, the continuity of c(t), b(t) implies that c(t) = b(t).

Lemma 11. Let $b^T(t)$ be the optimal stopping threshold when the end date is T. Take at $S \in \mathbb{R}_+$. Then for a large enough T, the policy $b^T(t)$ approximates the optimal policy $b^{\infty}(t)$ over [0, S].

Proof. The stopping policy for time horizon S is a valid stopping rule in the finite horizon case.

Let τ_S be the optimal stopping rule when the horizon is S and τ_{∞} be the optimal stopping rule when the horizon is infinite.

$$\mathbb{E}[e^{-r\tau_S}G(\tau_S, X_{\tau_S})] < \mathbb{E}[e^{-r\tau_\infty}G(\tau_\infty, Z_{\tau_\infty})]$$
$$\mathbb{E}[e^{-r\tau_S}G(\tau_S, X_{\tau_S})] > \mathbb{E}[e^{-r(\tau_\infty \wedge S)}G(\tau_\infty \wedge S, Z_{\tau_\infty \wedge S})]$$

Taking the limit as $S \to \infty$, we have that $\lim_{S \to \infty} \mathbb{E}[e^{-r\tau_S}G(\tau_S, X_{\tau_S})] = \mathbb{E}[e^{-r\tau_\infty}G(\tau_\infty, Z_{\tau_\infty})]$. Suppose that τ_S doesn't converge pointwise to τ_∞ . Let t be such that $\lim_{S \to \infty} b^S(t) \neq b^\infty(t)$.

Using the change of variable formula, we know that

$$V^{S}(t,X) = \mathbb{E}[e^{-r(\tau_{\infty}\wedge S)}G(\tau_{\infty}\wedge S, Z_{\tau_{\infty}\wedge S})] - \mathbb{E}_{t,X}[\int_{t}^{\tau_{\infty}\wedge S} e^{-rs}r(Z_{s} + \frac{\mu_{h} + \mu_{\ell}}{2}s - e^{-2\frac{\tilde{\mu}}{\sigma^{2}}Z_{s}})\mathbb{1}(X_{s} < b^{S}(s))]$$

Taking the limit as $S \to \infty$ and using $\lim_{S \to \infty} V^S(t, X) = V^{\infty}(t, X)$, we have that

$$\lim_{S \to \infty} \mathbb{E}_{t,X}\left[\int_t^{\tau_\infty \wedge S} e^{-rs} r(Z_s + \frac{\mu_h + \mu_\ell}{2}s - \lambda \alpha e^{-2\frac{\tilde{\mu}}{\sigma^2}Z_s}) \mathbb{1}(X_s < b^S(s))\right] = 0$$

Therefore (using the continuity of b and the positivity of G), we must have for each s, $\lim_{S\to\infty} b^S(s) \ge b^{\infty}(s)$. Suppose that for some large S, we have $b^S(s) > b^{\infty}(s) + \epsilon$. Pick an $x \in (b^{\infty}(s) + \epsilon, b^S(s))$. By $G \ge 0$ at $b^{\infty}(t)$, we must have

$$Z_s + \frac{\mu_h + \mu_\ell}{2} s - \lambda \alpha e^{-2\frac{\tilde{\mu}}{\sigma^2} Z_s} < -\frac{\mu_h}{r}$$

7.6 Extensions

Proof of Lemma 7. Take any realization ω of X_t for θ_{N-1} such that $\tau = s$. Then θ_i will assign the same probability to ω' , which at each point t is equal to $X(\omega'_t) = X(\omega_t) - (\mu_{\theta_{N-1}} - \mu_{\theta_i})t$. Since this is strictly lower, this means that the $X(\omega')$ will cross b before τ . Since earlier stopping lowers A's utility, it must be that the expected utility of θ_i is lower than that of θ_{N-1} .

Proof of Proposition 5. The proof that $b_h(t)$ is increasing if and only if $\frac{\mu_h + \mu_\ell}{2}$ follows directly from the proofs of Lemmas 1 and 4 (adding a $d_\tau(V^{FB}(X_\tau) - C)$ to the direct payoffs of P). Thus, we only need to argue that there exists cutoff times which determine whether or not auditing takes place. To see, this note that for an arbitrary τ , the optimal d_τ is equal to 1 if and only if $V^{FB}(X_\tau) > C$. Since $V^{FB}(X_\tau)$ is increasing in X_τ , there exists a cutoff X_c such that $d_\tau = 1 \iff X_\tau > X_c$. When $b_h(t)$ is increasing, there is a t' such that auditing only happens after t'; when $b_h(t)$ is decreasing, this means that there exists a time t'' such that audity only happens before t''.

Proof of Proposition 6. We proceed similar to how we did in Section 4.2. Let us begin by arguing that $b_h(t)$ must be increasing if and only if $\frac{\mu_h + \mu_\ell}{2} > 0$. We claim that $V(t, X) - \lambda \alpha M_t$ is increasing in t if $\frac{\mu_h + \mu_\ell}{2} < 0$. Let t' < t and fix an arbitrary W_ℓ^h policy as a function of s, X_s which is *IC* after the report. Note that

$$\mathbb{E}[e^{-r(\tau_{T}-t)}(\lambda\alpha M_{\tau_{T}}-(\frac{\mu_{h}}{r}+X_{\tau_{T}}))-\int_{0}^{\tau_{T}-t}e^{-rs}\gamma(\lambda W_{\ell}^{h}(s,X_{s})M_{s}+\widetilde{V}(W_{\ell}^{h}(s,X_{s}),X_{s}))ds]-\alpha M_{t}$$

$$=\mathbb{E}[e^{-r(\tau_{T}-t)}(\lambda\alpha M_{\tau}(1-e^{r(\tau_{T}-t)})-(\frac{\mu_{h}}{r}+X_{\tau}))-\int_{0}^{\tau_{t}-t}e^{-rs}\gamma(\lambda W_{\ell}^{h}(s,X_{s})M_{\tau}+\widetilde{V}(W_{\ell}^{h}(s,X_{s}),X_{t}))dt])$$

$$>\mathbb{E}[e^{-r(\tau_{T}'-t')}(\lambda\alpha M_{\tau_{T}'}(1-e^{r(\tau_{T}'-t')})-(\frac{\mu_{h}}{r}+X_{\tau}))-\int_{0}^{\tau_{T}'-t'}e^{-rs}\gamma(\lambda W_{\ell}^{h}(s,X_{s})M_{s}+\widetilde{V}(W_{\ell}^{h}(s,X_{s}),X_{s}))ds]$$

$$-\mathbb{E}[e^{-r(\tau_{T}'-t')}(\lambda\alpha M_{\tau_{T}'}-(\frac{\mu_{h}}{r}+X_{\tau_{T}}))-\int_{0}^{\tau_{T}'-t'}e^{-rs}\gamma(\lambda W_{\ell}^{h}(s,X_{s})M_{s}+\widetilde{V}(W_{\ell}^{h}(s,X_{s}),X_{s}))ds]-\alpha M_{t'}$$

Taking $T \to \infty$ and making the same arguments as in Lemma 1, we can conclude that $b_h(t)$ is decreasing. Similar arguments as in Lemma 4 lead us to conclude that $b_h(t)$ is increasing when $\frac{\mu_h + \mu_\ell}{2} > 0$. The rest of the proposition follows from the same arguments as in Section 4.2.