Experimentation and Approval Mechanisms

Andrew McClellan*

January 28, 2022

Abstract

We study the design of approval rules when costly experimentation must be delegated to an agent with misaligned preferences. When the agent has the option to end experimentation, we show that in contrast to standard stopping problems, the optimal approval rule must be history-dependent. We characterize the optimal rule and show the approval threshold moves downward over the course of experimentation. We find that private information may qualitatively change the optimal mechanism: an agent can choose a fast-track option in the form of an initially depressed approval threshold. On expiry of the fast track, the threshold jumps up, introducing more stringent standards. Our results provide a theoretical foundation for both the loosening of approval standards on longer experimentation paths and fast-track mechanisms.

Keywords: Dynamic mechanism design, experimentation, approval rules.

1 Introduction

In many real-world economic situations, decision-makers rely on information generated by other parties. Often the party, or agent, generating such information

* Email: andrew.mcclellan@chicagobooth.edu.

I am deeply indebted to and would especially like to thank Debraj Ray and Ennio Stacchetti for their constant support, guidance and insightful suggestions throughout this project. I would also like to thank Simon Board, Dhruva Bhaskar, Eric Budish, Drew Fudenberg, Boyan Jovanovic, Emir Kamenica, Alessandro Lizzeri, Laurent Mathevet, Marco Ottaviani, David Pearce, Jacopo Perego, Ariel Rubinstein, Andrzej Skrzypacz and Can Urgun for helpful conversations and suggestions, and the seminar participants at NYU, Princeton, UPenn, USC, Minnesota, Chicago Booth, Yale SOM, UCL, UCSD, Columbia, UC Davis, Berkeley and Stanford. I am particularly grateful to three anonymous referees for comments and suggestions which have improved the paper. This work is supported by the Liew Family Junior Faculty Fellowship at the University of Chicago Booth School of Business.

will have misaligned incentives and private information. For example, when deciding whether to approve a drug, the FDA bases its decision on data generated by clinical trials run by drug companies. The drug company, which bears the cost of experimentation and prefers earlier approval than the FDA, might also have private information about the quality of the drug generated during the R&D process before experimentation begins. The FDA can commit to an approval rule, which describes what clinical trial results are necessary for approval. This approval rule will impact both how long the company is willing to experiment and whether the company will truthfully report its private information. A misalignment of incentives will prevent straightforward elicitation: for example, the company, which wants the drug to be approved more quickly, may have an incentive to exaggerate their optimism about the drug's quality.

In this paper, we revisit the canonical Wald hypothesis-testing problem with the new feature that approval and experimentation are controlled by separate players. We study how a regulator can design stopping and decision rules (without monetary transfers) that incentivize an agent to perform experimentation and truthfully reveal any private information they have about the state of nature that determines the efficacy of a project. Experimentation generates evidence of efficacy, which is captured by a one-dimensional variable X. The players have misaligned incentives: the agent is biased toward approval and pays higher experimentation costs.

We define a novel class of approval rules and prove their optimality. These rules are history-dependent but still quite simple. They give us new insights into how the agency problem adds a rich set of dynamics and generates interesting implications (e.g., longer length of experimentation is associated with more erroneous approvals). We also study the effects of private information and find it adds new qualitative features and dynamics to the optimal approval rule. Agents who report a higher initial belief that the project is good may be given a *fast-track mechanism*, where they have a chance for quick approval but are punished upon expiration of the fast track with a higher approval standard.

We start by studying the problem in which the agent has no private information (*symmetric information*), and focus on how the regulator provides incentives for the agent to experiment. A standard result in the optimal stopping literature with a single decision-maker is the optimality of *stationary threshold* strategies, in which the decision-maker stops whenever the state crosses a fixed, history-independent,

threshold. However, when the regulator relies on the agent to generate information through experimentation, such static threshold rules are no longer optimal: once the agent is about to quit, the regulator may be better off changing the approval rule to provide additional incentives for the agent to continue experimentation. The regulator could change the approval rule in many ways to give the agent incentives to experiment. Given the richness of the set of approval rules, solving for the optimal rule can appear quite daunting. Nevertheless, we find a tractable way to relax the problem and find the solution to be history-dependent and non-stationary but still remarkably simple.

We show how the optimal mechanism can be written as a function of only of the current level of evidence of efficacy X_t and the minimum over the realized path of X up to the current time t. The regulator uses an approval threshold that is initially stationary. When the evidence drifts low enough that the agent is tempted to quit if the threshold were to remain fixed, the regulator begins to lower the threshold just enough to incentivize the agent to continue experimentation. When the evidence moves higher than the current minimum, the threshold stays fixed, never increasing, and will only decrease when the evidence again reaches a new low. This downward drift of the approval threshold is bounded; if the evidence reaches a fixed lower threshold, the regulator allows the agent to quit. Unlike in the case of a single decision-maker, the probability of Type I error is not constant over time. We also find that the optimal approval threshold, when written in terms of the regulator's beliefs, is independent of her initial beliefs. This independence would not arise if we were to restrict the regulator to consider only stationary threshold rules.

Having explored the tension caused by the need to incentivize experimentation, we introduce private information about the state, which we call the agent's type, and look at the new features private information adds to the problem. The regulator faces a tradeoff: giving one type a lower approval threshold gives the other type an incentive to misreport his type. We find that the optimal mechanism may take the form of a fast track given to agents who report a high initial belief. In such a fast track, the agent is initially given a low approval threshold, but also faces a stationary "failure" threshold. If the failure threshold is reached, the project is not rejected, but the approval threshold takes a discrete jump upward (the agent is thrown out of the fast track). They are allowed to continue to experiment but

face a more stringent standard. Introducing the threat of being thrown out of the fast track allows the regulator to lower the initial approval threshold while still preserving the incentive compatibility constraint for the agent to truthfully report his initial belief. By throwing the agent out of the fast track, the regulator hurts both his and the agent's payoffs. However, an agent who has misreported his type will view this distortion as more likely, allowing the regulator to separate types without transfers and increase the probability of quicker approval for the high type. We then discuss the application of our model to the drug approval process and how the two most notable features of our optimal mechanisms— the changing approval standard and fast-track—match features of the real-world drug approval process that cannot be explained by standard single-decision-maker models.

The outline of the paper is as follows: We review the literature in Section 2 and then introduce our baseline model with symmetric information in Section 3. We solve for the optimal mechanism under symmetric information in Section 4. In Section 5 we introduce asymmetric information and study how it affects the optimal mechanism, as well as look at some comparative statics before discussing applications to the drug approval example in Section 6. We conclude in Section 7. Proofs of the our main theorems can be found in the Appendix; all other proofs are in the Supplementary Material.

2 Literature

The setting of our paper ties into a large literature on the problem of dynamic hypothesis testing building on the seminal model of Wald (1947). Peskir and Shiryaev (2006) provide a textbook summary and history of the problem.

A growing literature has studied the strategic forces in experimentation. Papers such as Bergemann and Hege (2005), Halac et al. (2016) and Halac et al. (2017) have studied the design of mechanisms with transfers to incentivize experimentation. Kruse and Strack (2015) study which rules are implementable in a general optimal-stopping principal-agent problem with transfers. Georgiadis and Szentes (2020) use a Brownian learning model to study optimal costly information acquisition when monitoring an agent's action. Guo (2016), one of the closest papers to our own, looks at a bandit problem in a principal-agent model where the agent possesses private information about the probability that the bandit is

"good." Similar to our model, Guo (2016) solves for the optimal mechanism when monetary transfers are infeasible and the agent has private information about a payoff-relevant state of the world. We instead consider a model in which the agent has the ability to quit experimenting, whereas in her model, the principal controls experimentation throughout. Introducing interim constraints on the mechanism adds the history dependence to our mechanism.

Another paper that is close to our own is Henry and Ottaviani (2019), who study a model of regulatory approval when learning takes place through a publicly observable Brownian motion. In their model, both the regulator and the agent possess a common prior about the state. They study the equilibria of the approval process for different configurations of approval and experimentation authority. Our main focus is on the design of optimal mechanisms and the effects of private information rather than the equilibrium outcomes.

Carpenter and Ting (2007) look at a theoretical model of drug approval in which the drug companies are better informed about the state for their drug. They study the resulting equilibria of a discrete time model. They find that the length of experimentation determines the comparative static for the effect of firm size on the number of Type I and Type II errors, which they find support for using data on FDA approval times.

A notable feature of our optimal mechanism is the rigidity in the movement of the approval threshold and the history dependence on only the minimum of past beliefs. This rigidity is somewhat reminiscent of Harris and Holmstrom (1982), who study equilibrium wage contracts in a competitive market. In their model, risk aversion by the agent means wages will be constant until beliefs about the agent's type are high enough that market competition increases wages. Risk aversion drives a similar rigidity in Thomas and Worrall (1988), who study the design of wage contracts. Our model contains no risk aversion, and the rigidity is driven by smoothing over the approval standards. The downward movement of the threshold, towards the agent's preferred level, is also reminiscent of the backloading dynamics as seen in Ray (2002). As beliefs drop, the approval threshold moves towards the agent's preferred level. The dependence on only the minimum of the state variable also arises in Gryglewicz and Kolb (2019), who study the equilibrium strategic pricing game in which the prices of an incumbent firm signal its costs (which may deter entry) and find prices set by firm depend on demand and the minimum over

past demand.

The only other paper we know of in which agency forces drive a dynamically changing threshold is Ely and Szydlowski (2020). They study how a principal can use information disclosure to motivate an agent to work towards completing a project. In contrast to our setting, they study a model in which the principal is the provider of information and the agent controls the threshold for how long he will work.

3 Model

Two players, a (female) regulator R and a (male) agent A, interact over an infinite horizon in continuous time. A project that is up for approval may be of two types: high ($\theta = H$) or low ($\theta = L$). Both players begin the game with a common-prior $\pi_0 = \mathbb{P}(\theta = H)$. Experimentation by A generates information about θ via a publicly-observed Brownian diffusion process X with state-dependent drift[1] starting at X_0; unless stated otherwise, we take $X_0 = 0$. X_t evolves according to

$$dX_t = \mu_\theta dt + \sigma dW_t,$$

where W is a standard one-dimensional Brownian motion on the state space $(\Omega, \mathcal{F}, P)$ and $\mathcal{F} = \{\mathcal{F}_t\}_{0 \leq t \leq \infty}$, where $\mathcal{F}_t = \sigma(X_s, Y_0 : \ 0 \leq s \leq t)$ is the natural filtration of X and, in order to allow for randomized mechanisms, an independent randomization device Y_0 realized at $t = 0$. A history h_t is the realization of Y_0 and the sample path ω of X from time 0 to t. Without loss of generality, we take $\mu_L = -\mu < 0 < \mu = \mu_H$. Unless otherwise specified, we take $X_0 = 0$.

Both players update their beliefs to π_t according to Bayes' rule. It is convenient to instead use the log-likelihood of beliefs $Z_t = \log(\frac{\pi_t}{1-\pi_t})$, which we often simply

[1] Brownian learning has been widely used in the statistics literature on hypothesis testing and the design of clinical trials. The use of continuous time in our model is done purely for mathematical convenience and tractability. The intuition underlying our results is built on optimal stopping arguments that rely on the Brownian motion continuous time setup through the continuity of the belief path and martingale property of beliefs.

call beliefs. Z_t can be written as[2]

$$Z_t = Z_0 + \frac{2\mu}{\sigma^2}(X_t - X_0).$$

We let $z_0 := \log(\frac{\pi_0}{1-\pi_0})$. This transformation of the belief process is useful because both X_t and the initial Z_0 enter linearly into the current Z_t. We denote by $\mathbb{E}^{x,z}$ the expectation over X, Y_0 given $X_0 = x$ and the distribution over θ implied by the beliefs $Z_0 = z$. For notational convenience, we take $\mathbb{E} = \mathbb{E}^{0,z_0}$ and $\mathbb{E}^x = \mathbb{E}^{x,z_x}$ where $z_x = z_0 + \frac{2\mu}{\sigma^2}x$.

R and A share a common discount rate $r \geq 0$ and pay respective flow costs c_R and c_A as long as experimentation continues, with $c_A > c_R \geq 0$. Both players have the option to end the game at *any* time, giving both players a payoff of 0.

R can, at any time, decide to approve or reject the project, thereby ending the game. We let $d_t = 1$ if R approves at time t, and $d_t = 0$ if R rejects at time t. Both players receive 0 if the project is rejected. We normalize both players' payoff from approval when $\theta = H$ to 1 and let a and f be A and R's respective payoffs from approval when $\theta = L$. To avoid trivial cases, we take $f < 0$. We assume $a \in [f, 1]$ so that A is weakly biased in favor of approval and receives a higher utility from approval when $\theta = H$ than when $\theta = L$. The expected utilities for R and A from approval at X_t are $\tilde{u}(X_t) = \frac{e^{Z_t}+f}{1+e^{Z_t}}$ and $\tilde{v}(X_t) = \frac{e^{Z_t}+a}{1+e^{Z_t}}$, respectively, where $Z_t = z_0 + \frac{2\mu}{\sigma^2}(X_t - X_0)$. Both $\tilde{u}(X_t)$ and $\tilde{v}(X_t)$ are martingales in X_t.[3] We call $X_c := \tilde{u}^{-1}(0)$ R's myopic threshold: R prefers approval to rejection if and only if $X_t \geq X_c$.[4]

We give R full commitment power over when to stop and approve or reject the project. We call her rule for deciding when to stop and whether to approve or reject the project a stopping mechanism.

Definition 1. *A **stopping mechanism** is a pair $(\tau, d_\tau) \in \mathbb{T} \times \mathbb{D}$, where $\mathbb{T}$ is the set of stopping times with respect to $\mathcal{F}$ and $\mathbb{D}$ is the set of $\mathcal{F}_\infty$-random variables taking values in $\{0, 1\}$ such that d_τ is $\mathcal{F}_\tau$-measurable.*

[2] A standard derivation (see Shiryaev (2007)) shows that the posterior after history h_t is $\pi_t = \frac{\pi_0 e^{\frac{2\mu}{\sigma^2}(X_t - X_0)}}{1-\pi_0+\pi_0 e^{\frac{2\mu}{\sigma^2}(X_t - X_0)}}$. The formula for Z_t then follows from taking $\log(\frac{\pi_t}{1-\pi_t})$.

[3] This follows by noting that $\tilde{u}$ and $\tilde{v}$ are linear in π_t, which we know is a martingale.

[4] Although it is natural to use Z_t as the state variable, it will be convenient when introducing asymmetric information to write players' utilities as a function of X_t. In the interest of keeping notation consistent throughout, we will write everything in terms of X_t.

Let $u(X_\tau, d_\tau) := \tilde{u}(X_\tau)d_\tau + \frac{c_R}{r}$ and $v(X_\tau, d_\tau) := \tilde{v}(X_\tau)d_\tau + \frac{c_A}{r}$. R and A's expected utility from (τ, d_τ) when $X_0 = x$ and $Z_0 = z$ are given by J and V respectively:

$$J(\tau, d_\tau, x, z) = \mathbb{E}^{x,z}[e^{-r\tau}\tilde{u}(X_\tau)d_\tau - \int_0^\tau e^{-rt}c_R dt] = \mathbb{E}^{x,z}[e^{-r\tau}u(X_\tau, d_\tau)] - \frac{c_R}{r},$$
$$V(\tau, d_\tau, x, z) = \mathbb{E}^{x,z}[e^{-r\tau}\tilde{v}(X_\tau)d_\tau - \int_0^\tau e^{-rt}c_A dt] = \mathbb{E}^{x,z}[e^{-r\tau}v(X_\tau, d_\tau)] - \frac{c_A}{r}.$$

For notational convenience, we will occasionally drop dependence on x in J and V when $x = 0$ and, although $\tilde{u}$ and $\tilde{v}$ depend on Z_0, keep this dependence implicit.[5] Because R has commitment power, it is without loss to focus on mechanisms in which A never takes his outside option.[6] In order to avoid adding unnecessary trivial caveats in the proofs, we will assume throughout the R's optimal mechanism does not stop immediately.

Throughout the paper, our main example is that of drug approval. This situation fits many of the key assumptions of the model: dynamic experimentation, no transfers, delegation of experimentation and observability of experimental results. The FDA cannot make transfers to the company and the length of experimentation is determined by the company; that is, the FDA can not prevent the company from ending experimentation early. The public observability of X_t holds, for example, if R can force A to preregister his experiments and the experimental outcomes are verifiable, as is the case in drug trials. We also view the commitment assumption as reasonable here. Theoretically, this can be justified using a repeated-game logic in which a deviation from the agreed upon mechanism is punished with movement to a bad equilibrium when experimenting on future drugs. Given the large number of drugs the FDA has to consider, maintaining commitment to an agreed upon mechanism should be relatively easy to support.

[5] When taking an expectation $\mathbb{E}^{x,z}[e^{-r\tau}\tilde{v}(X_\tau)d_\tau]$, we will assume that $\tilde{v}$ is defined relative to $Z_0 = z$ and similarly for $\tilde{u}$.

[6] We can always replace A taking the outside option with R rejecting the project. We also note here that R does not formally need the ability to reject the project. If R sets $\tau = \infty$, A will immediately quit. The act of rejection can be taken as short-hand for inducing A to immediately quit by setting $\tau = \infty$. We would like to thank an anonymous referee for pointing this out.

4 Symmetric Information

We start by looking at a simple class of simple benchmark mechanisms, which we call static-threshold mechanisms. These mechanisms approve if X_t rises above a fixed threshold B and reject if X_t falls below a fixed threshold b. We call B the static approval threshold and b the static rejection threshold.

Definition 2. (τ, d_τ) *is a* ***static-threshold mechanism*** *for a pair* $(b, B) \in \mathbb{R}^2$ *with* $b < B$ *if* $\tau = \inf\{t : X_t \notin (b, B)\}$ *and* $d_\tau = \mathbb{1}(X_\tau \geq B)$.

Let $\tau_+(B) := \inf\{t : \ X_t \geq B\}$ and $\tau(b) := \inf\{t : \ X_t \leq b\}$. R and A's expected utility from a static threshold mechanism with (b, B) when $(X_0, Z_0) = (x, z + \frac{2\mu}{\sigma^2}x)$ are $\widetilde{J}$ and $\widetilde{V}$, respectively:[7]

$$\widetilde{J}(B, b, x; z) := J(\tau_+(B) \wedge \tau(b), \mathbb{1}(X_\tau \geq B), x, z + \frac{2\mu}{\sigma^2}x),$$
$$\widetilde{V}(B, b, x; z) := V(\tau_+(B) \wedge \tau(b), \mathbb{1}(X_\tau \geq B), x, z + \frac{2\mu}{\sigma^2}x).$$

We drop dependence on z in $\widetilde{J}$ and $\widetilde{V}$ when $z = z_0$.

It is straightforward to show that both R and A's first-best mechanisms, in which they alone control the length of experimentation and the decision to be made, are static-threshold mechanisms. These mechanisms are stationary: continuation play is independent of the previous history and depends only on the current X_t. Stationarity also implies that, conditional on approval, the probability of Type I error (i.e., approving when $\theta = L$) is independent of the length of experimentation.

The optimality of static-threshold mechanisms turns out to be robust even to mild forms of the agency problem. Consider R's problem in which we limit A to only be able to take his outside option at time zero. R will maximize $J(\tau, d_\tau, z_0)$ subject to a participation constraint $V(\tau, d_\tau, z_0) \geq 0$. Proposition 8 in the Supplementary Material shows the solution to this problem is a static-threshold mechanism. Although R could satisfy A's participation constraint in a myriad of ways, she chooses to do so in a smooth way, by adjusting the approval and rejection thresholds without violating stationarity.

[7] Unless otherwise stated, we will assume that $b < B$ when discussing pairs of thresholds b, B.

4.1 Optimal Mechanism

Including only a time-zero participation constraint is too weak for our problem. To ensure that A does not take his outside option early, R must ensure A's continuation value is weakly positive after all histories h_t until R ends experimentation.

Definition 3. (τ, d_τ) *satisfies the* ***dynamic participation constraint*** *if, after every history h_t, A's continuation value is weakly positive.*

Suppose A were to quit early at some stopping time $\tau' \in \mathbb{T}$. Using this strategy gives A an ex-ante expected utility of $\mathbb{E}[e^{-r(\tau\wedge\tau')}v(X_{\tau\wedge\tau'}, d_\tau \mathbb{1}(\tau < \tau'))] - \frac{c_A}{r}$. We define A's value allowing for deviations to quit early as

$$V^*(\tau, d_\tau, x, z) := \sup_{\tau'\in\mathbb{T}} \quad \mathbb{E}^{x,z}[e^{-r(\tau\wedge\tau')}v(X_{\tau\wedge\tau'}, d_\tau \mathbb{1}(\tau < \tau'))] - \frac{c_A}{r}.$$

As with V and J, we drop dependence on x when $x = 0$. We define a (slightly) relaxed version of R's problem with symmetric information, which we call the SI problem, by only imposing a condition we call $DP(z_0)$:

$$\begin{aligned} [SI]: \quad & \sup_{(\tau,d_\tau)} \quad J(\tau, d_\tau, z_0) \\ \text{subject to } DP(z_0): \quad & V^*(\tau, d_\tau, z_0) \leq V(\tau, d_\tau, z_0). \end{aligned}$$

DP weakens the dynamic participation constraint.[8] For example, DP allows for A's continuation value to be negative on probability-zero events. We can show that every mechanism that satisfies the dynamic participation constraint will satisfy DP.[9] SI is therefore an upper bound on R's utility. We later verify that our solution satisfies the dynamic participation constraint.

Given our benchmark mechanisms, the conjecture that a static-threshold mechanism will be optimal seems natural. We illustrate why this conjecture fails with a simple example below. Simply put, at the point at which R rejects in a static-threshold mechanism, she may be better off lowering the threshold ("cutting A some slack") to incentivize A to continue experimenting. Rejecting the project leaves gains from trade on the table: A would benefit from a lower approval threshold, and, due to the option value of experimentation, R would benefit from continued experimentation.

[8] We drop dependence of $DP(z_0)$ on z_0 where doing so causes no confusion.

[9] See Lemma 23 in the Supplementary Materials.

To simplify the example, we take $c_R = 0$ and $a \geq 0$. Suppose R uses a static approval threshold of $B_1 > X_c$.[10] Because R benefits from the option value of experimentation, she will never reject the project before A decides quit. Let $b_1 = \arg\max_{b \leq 0} \widetilde{V}(B_1, b, 0)$ be the threshold at which A will choose to quit against the approval threshold B_1. To satisfy DP, R must reject when $X_t = b_1$.

Let Mechanism One be a static-threshold mechanism $(b_1, B_1) \in \mathbb{R}^2$ with $b_1 < B_1$. For $b < B$, define $\tau_+(B; b) = \inf\{t < \tau(b) : X_t \geq B\}$ to be the first time X crosses B before reaching b, with $\tau_+(B; b) = \infty$ if $\tau(b) < \tau_+(B)$. Similarly, we define $\tau(b; B) = \inf\{t < \tau_+(B) : X_t \leq b\}$ to be the first time X crosses b before reaching B, with $\tau(b; B) = \infty$ if $\tau(b) > \tau_+(B)$.[11] R's expected utility from Mechanism One is

$$\mathbb{E}[e^{-r\tau_+(B_1;b_1)}\tilde{u}(B_1)].$$

Now consider Mechanism Two, in which R uses the same approval threshold of $B_1 > 0$ until $X_t \notin (b_1, B_1)$. If X_t reaches b_1 first, then, instead of rejecting, R lowers the approval threshold to $B_2 \in (X_c, B_1)$. Because the discounted probability of approval is now higher, A will be willing to continue experimenting until X_t reaches $b_2 = \arg\max_b \widetilde{V}(B_2, b, b_1) < b_1$, so R now rejects at b_2. R's expected payoff from Mechanism Two is

$$\underbrace{\mathbb{E}[e^{-r\tau_+(B_1;b_1)}\tilde{u}(B_1)]}_{\substack{R\text{'s Util. when} \\ \tau_+(B_1)<\tau(b_1)}} + \underbrace{\mathbb{E}[e^{-r\tau(b_1;B_1)}]}_{\substack{\text{Discounted prob. of} \\ \tau(b_1)<\tau_+(B_1)}} \underbrace{\mathbb{E}^{b_1}[e^{-r\tau_+(B_2;b_2)}\tilde{u}(B_2)]}_{\substack{R\text{'s Cont. Value at} \\ \tau(b_1) \text{ when } \tau(b_1)<\tau(B_1)}}.$$

R's payoff when $\tau_+(B_1) < \tau(b_1)$ is the same as that in Mechanism One. However, because R does not reject at $\tau(b_1)$ and $\tilde{u}(B_2) > 0$, she now has a strictly positive continuation value at $\tau(b_1)$, thereby improving on Mechanism One.

Moving out of the class of static-threshold mechanisms, conjecturing the form the optimal policy will take is difficult. The key difficulty comes from the fact that the DP constraint allows the agent to deviate by choosing a $\tau' \in \mathbb{T}$, where $\mathbb{T}$ is infinite-dimensional. For an arbitrary stopping rule (τ, d_τ), finding A's optimal τ'

[10] If $B_1 < X_c$, the static-threshold mechanism would only approve at beliefs which give R negative utility, and, therefore, R would be better off rejecting immediately.

[11] We note that $\mathbb{E}[e^{-r\tau_+(B;b)}] = \mathbb{E}[e^{-r\tau_+(B)}\mathbb{1}(\tau(b) > \tau_+(B))]$, $\mathbb{E}[e^{-r\tau(b;B)}] = \mathbb{E}[e^{-r\tau(b)}\mathbb{1}(\tau(b) < \tau_+(B))]$ and $\mathbb{E}[e^{-r\tau_+(B;b)}] + \mathbb{E}[e^{-r\tau(b;B)}] = \mathbb{E}[e^{-r(\tau(b)\wedge\tau_+(B))}]$.

is infeasible.[12] The standard approach in the dynamic contracting literature (e.g., Sannikov (2008)), to use A's continuation value as an additional state variable, is also intractable. In addition to the continuation value, we must also carry the belief about the state, making the corresponding HJB equation a partial differential equation that is difficult to analyze.

We overcome these difficulties by studying a relaxed version of SI. Although we allow for arbitrary complex and history-dependent mechanisms, the solution depends on the history h_t only through X_t and $M_t := \min_{s \leq t} X_s$, the current minimum of the diffusion path. The optimal mechanism is, roughly speaking, a continuous version of the strategy in our simple example: R decreases the approval threshold whenever A is about to quit, but keeps it fixed as X_t moves back toward the approval threshold.

We now develop some notation to write the optimal mechanism. Suppose R fixes her approval threshold at B. When $Z_0 = z$, it is optimal for A to continue experimenting at $X_t = x$ if his value function $\max_b \widetilde{V}(B, b, x; z)$ is strictly positive—namely, A will choose to quit at $X_t = x$ if and only if $\max_b \widetilde{V}(B, b, x; z) = 0$. We define the threshold at which A will choose to quit as

$$b^*(B; z) = \inf\{x : \max_b \widetilde{V}(B, b, x; z) > 0\}.$$

We note that $b^*(B; z)$ is independent of x. It is straightforward to show that it is optimal for A to quit when $X_t = x$ if and only if $x \leq b^*(B; z)$. Abusing notation slightly, we write $b^*(B; z) = \arg\max_b \widetilde{V}(B, b, x; z)$.[13] We drop dependence on z in b^* when $z = z_0$.

We use b^* to find the maximal static approval threshold that would induce A to optimally choose to quit at $X_t = m$. We define this function as $\underline{B}(m)$, which we show is increasing in m in the Appendix.

$$\underline{B}(m) := \max\{B : b^*(B) = m\}.$$

[12] Even for simple mechanisms, such as a time-varying threshold, solving for A's optimal τ' is difficult and cannot be calculated in closed form.

[13] If $\max_b \widetilde{V}(B, b, x; z) = 0$, then $b^*(B; z) > x$ and $\arg\max_b \widetilde{V}(B, b, x; z) = [x, \infty)$ as all such $b \in [x, \infty)$ lead to A immediately quitting. In such a case, $b^*(B; z) \in \arg\max_b \widetilde{V}(B, b, x; z)$ and so is an optimal quitting threshold.

Let $B(m; B^1)$ be a "capped" version of $\underline{B}$: B remains fixed at B^1 until m is low enough that B begins to follow $\underline{B}(m)$.

$$B(m; B^1) = \begin{cases} B^1 & m > b^*(B^1) \\ \underline{B}(m) & m \leq b^*(B^1). \end{cases}$$

Again, for notational convenience, we occasionally drop dependence on B^1.

Our first main result shows that there exists a B^1 such that the optimal mechanism uses $B(M_t; B^1)$ as its approval threshold. The approval threshold remains at B^1 until X_t decreases low enough that A would find it optimal to quit if the approval threshold were to remain fixed. After reaching this point, which occurs when $X_t = M_t = b^*(B^1)$, R gradually and permanently lowers the approval threshold each time M_t decreases, never raising it again. This adjustment downwards is just enough to keep A indifferent between continuing and quitting when $X_t = M_t \leq b^*(B^1)$; as X_t decreases, A becomes more pessimistic about the state and so, even though he expects the approval threshold to fall when M_t decreases, he is still indifferent between continuing and quitting when $X_t = M_t$ as his increased belief that $\theta = L$ offsets his benefit from the decrease in the approval threshold. The adjustment process for the approval threshold continues until either R approves the project or X_t reaches a fixed lower rejection threshold $\underline{b}$, at which point R rejects the project. This rejection threshold is the point at which continuing to lower the threshold is too costly and R is better off rejecting.

Theorem 1. *There exists* $(\underline{b}, B^1)$ *such that optimal mechanism is*

$$\tau^* = \inf\{t : X_t \notin \big(\underline{b}, B(M_t; B^1)\big)\}, \quad d^*_\tau = \mathbb{1}\big(X_{\tau^*} \geq B(M_{\tau^*}; B^1)\big).$$

Our initial observation that both R and A's first-best mechanisms are static-threshold mechanisms shows the dynamics of our mechanism are driven by the agency problem. This type of history dependence in the approval rule is, to our knowledge, new and shows how strategic interactions can lead to a rich set of dynamics in the design of approval rules.

While the dependence on X_t is natural, the dependence on M_t is perhaps more unusual as it is not payoff relevant. With only a mild form of the agency problem in which R has to only satisfy an ex-ante participation constraint for A, static threshold mechanisms are optimal: R increases A's expected utility in a "smooth"

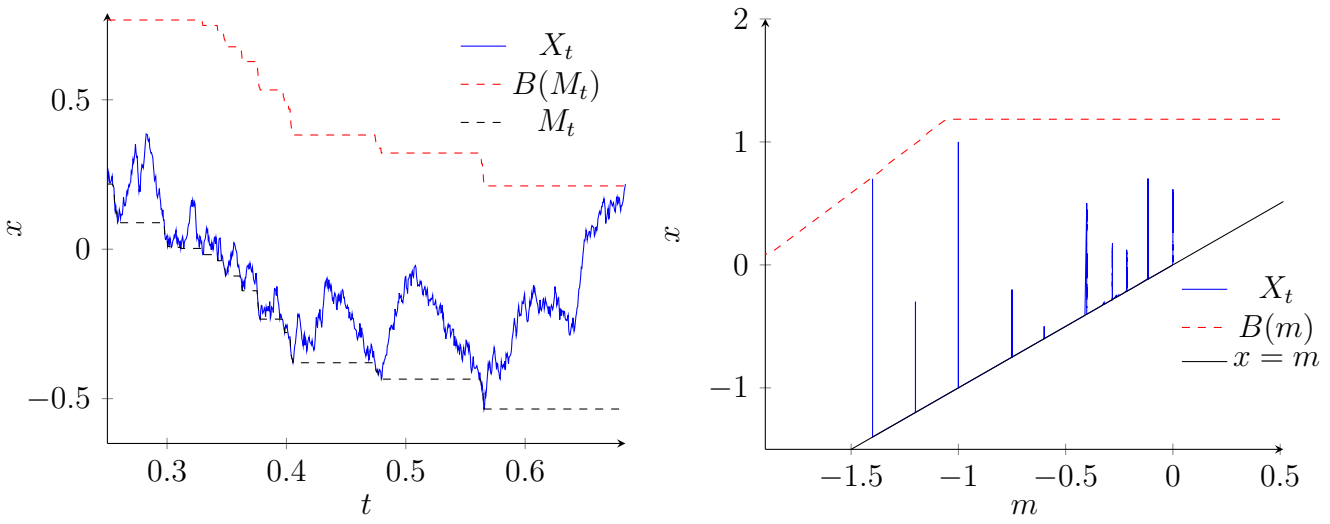

Figure 1: The left graph shows a realized sample path of X_t. In the right graph, the lines extending up from the 45-degree line illustrate another sample path.

manner by lowering the approval threshold but still keeping it fixed. This suggests a smoothing intuition, that R would like to keep the approval threshold fixed as long as possible. However, after a large enough drop in X_t, A will find it optimal to quit unless R modifies the approval threshold in order to increase A's continuation value. There are many ways R could do so, such as lowering the threshold a large amount today before raising it in the future. However, the smoothing intuition suggests that R would prefer to keep the threshold stationary as long as possible. Thus, R will lower the threshold and keep it fixed until X_t again falls low enough that R must increase A's continuation value again; this point occurs when A is most pessimistic about his chances of approval (namely, when $X_t = M_t$). M_t operates as a sufficient statistic for how much R has needed to lower the approval threshold to prevent A from quitting.

Sketch of the Proof

To identify a relaxed version of the SI-problem that is more amenable for our analysis, we make some conjectures regarding the class of τ' that are likely to be binding. As long as R is using a stationary threshold, A's best-response quitting

rule will be a stationary threshold as well. We conjecture a particular class of τ' deviations will be binding, which we call *threshold quitting rules*:

Definition 4. *A uses a* ***threshold quitting rule at*** x *if he quits at time* $\tau(x)$.

A's payoff from (τ, d_τ) when using the quitting rule $\tau(x)$ is $V(\tau \wedge \tau(x), d_\tau(x), x)$, where $d_\tau(x) := d_\tau \mathbb{1}(\tau < \tau(x))$.

We define a relaxed version of SI in which we restrict attention to a finite grid of such quitting rules. The restriction to a finite grid is used for technical reasons, and we look at the limit as this grid becomes arbitrarily fine. Let $b_A^{FB} = \arg\max_b[\max_B \widetilde{V}(B, b, x)]$ be the rejection threshold in A's first best mechanism; we show in the Appendix that if $a \geq 0$, A prefers immediate approval at all values of X_t, in which case we take $b_A^{FB} = -\infty$. It is clear that R can never convince A to experiment below b_A^{FB}.

For some $\underline{X} < 0$, let $\mathcal{T}_N = \{X_n\}_{n=1}^N$ with $X_1 = 0$ and $X_n = X_{n-1} - \delta_N$, where $\delta_N := \frac{|\underline{X}|}{N}$ is the step size of our grid. If $b_A^{FB} > -\infty$, take $\underline{X} = b_A^{FB}$ and if $b_A^{FB} = -\infty$, take $\underline{X}$ going to $-\infty$ as $N \to \infty$ slowly enough that $\delta_N \to 0$. Our relaxed problem replaces the DP constraint with a set of relaxed versions of DP, which we call RDP constraints, one for each $X_n \in \mathcal{T}_N$:

$$RDP(X_n): \ V(\tau \wedge \tau(X_n), d_\tau(X_n), z_0) \leq V(\tau, d_\tau, z_0).$$

Because we have relaxed the constraint set, the solution to this relaxed problem will provide an upper bound on the value of SI.[14]

To solve the relaxed problem, we employ Lagrangian techniques. We construct a Lagrangian with multipliers $\Lambda = (\lambda_1, ..., \lambda_N) \in \mathbb{R}_-^N$:

$$\mathcal{L} = \sup_{(\tau, d_\tau)} \ \mathbb{E}[e^{-r\tau} u(X_\tau, d_\tau) + \sum_{n=1}^{N} \lambda_n \left\{ e^{-r(\tau \wedge \tau(X_n))} v(X_{\tau \wedge \tau(X_n)}, d_\tau(X_n)) - e^{-r\tau} v(X_\tau, d_\tau) \right\}].$$

For an appropriate choice of Λ, the solution to $\mathcal{L}$ will solve our relaxed problem and satisfy complementary slackness conditions. Although selection of the multiplier Λ

[14] It is not obvious that dropping non-threshold constraints is without loss. For many stopping policies that R could use, A's best response will take a more complex form than a threshold quitting rule. For example, if R were to wait until date T and approve if and only if $X_T > B$, then $A's$ optimal quitting rule would take the form $\inf\{t: \ X_t = f(t)\}$ for some increasing f.

can be difficult, the qualitative properties we derive from analyzing the Lagrangian for an arbitrary Λ allows us to pin down the form of the optimal solution.

Let X^1 be the largest $X_n \in \mathcal{T}_N$ for which $RDP(X_n)$ binds in the solution to this relaxed problem. We decompose the problem into before $\tau(X^1)$ and after $\tau(X^1)$. Optimal stopping arguments allow us to establish the optimality of a "local" static-threshold rule: R uses an approval threshold that is constant until $\tau(X^1)$.

We call the stopping time and decision rule (normalized by t) which is induced by (τ, d_τ) after h_t the continuation mechanism.[15] We show that the continuation mechanism is the same at $\tau(X^1)$ for every history up to $\tau(X^1)$. Doing so allows us to prove that, at $\tau(X^1)$, A is indifferent between quitting and continuing to experiment, regardless of the history up to $\tau(X^1)$. The stationarity of the optimal stopping rule prior to $\tau(X^1)$ is key for proving the indifference of A at $\tau(X^1)$.[16]

The mathematical structure of the continuation problem at $\tau(X^1)$ is similar to our problem at $t = 0$, allowing us to apply similar arguments to show that the optimal mechanism at $\tau(X^1)$ is a stationary threshold until the next binding constraint threshold is reached. Repeated application of these arguments allows us to show the optimal mechanism has a sequence of approval thresholds that depend only on the last binding quitting threshold that has been reached, for which M_t is a sufficient statistic. Complementary slackness and optimal stopping arguments allow us to pin down the approval thresholds. Taking the limit of our mechanisms as the grid $\mathcal{T}_N$ becomes arbitrarily fine, we show that a limit mechanism exists and satisfies the dynamic participation constraint.

4.1.1 Features of the Optimal Mechanism

The mechanism has observable implications for the relationship between experimentation length and the probability of error. The decrease in the approval threshold increases the probability of Type I error (i.e., approving a bad project), in contrast to static threshold mechanisms in which the probability of error conditional on approval is constant. Our model predicts a higher probability of Type I error for projects that have taken a long time to be approved relative to projects that

[15] A formal definition of a continuation mechanism is provided in the Appendix.

[16] This property is not true for general (τ, d_τ). If R used a deterministic stopping rule that stopped at some T with probability one, A's $RDP(X^1)$ constraint may bind in expectation at $t = 0$, but, A will in general have a non-zero continuation value at $\tau(X^1)$.

were approved quickly.[17]

Given that beliefs are the payoff-relevant state variable, it is natural to write the mechanism in terms of Z rather than X. Let $M_t^Z := \min_{s \leq t} Z_s$. By Theorem 1 and the fact that Z_t is an affine transformation of X_t, there exists an approval threshold function $B^Z(\cdot)$ and rejection threshold $\underline{b}^Z$ such that the optimal mechanism, when written as a function of Z, is (τ^Z, d_τ^Z) with $\tau^Z = \inf\{t : \ Z_t \notin (\underline{b}^Z, B^Z(M_t^Z))\}$ and $d_\tau^Z = \mathbb{1}(Z_{\tau^Z} \geq B^Z(M_{\tau^Z}^Z))$. Proposition 1 below shows that the optimal approval and rejection thresholds are independent of z_0. This independence from z_0 is standard in single-decision maker problems, but is absent in an agency model if we restrict attention to a choice over static-threshold mechanisms (see Henry and Ottaviani (2019)). Providing flexibility in the design of the approval rule restores independence.

Proposition 1. *The optimal approval and rejection thresholds $B^Z(\cdot)$ and $\underline{b}^Z$ are the same for all z_0.*

The structure of the optimal mechanism is the same under a more general class of players' preferences. We show in the Supplementary Material that our results hold under more general utility functions $\tilde{u}$ and $\tilde{v}$ for R and A. Doing so allows us to extend our results to examples in which the players' utilities from approval at a belief π_τ are not linear in π_τ. In the drug approval example, whether doctors and patients use a particular drug may depend on their belief π_τ about its efficacy. Thus, the belief π_τ will affect how widely the drug is adopted after approval. If R and A care about both the effectiveness of the drug and how widely it is used after approval, then players' payoffs may be non-linear in π_τ.

The mechanism might seem to rely heavily on the commitment assumption. If R could not commit, would she raise the threshold when X_t rises, thereby undermining her past promises to A that incentivized him to keep experimenting? In the Supplementary Material, we look at a version of the model without commitment and show that an equilibrium exists with the same on-path outcome as our optimal mechanism.

[17] Fudenberg et al. (2018) observe a similar relationship between experimentation time and probability of error in a single-decision-maker model when the state space is continuous but not when binary as is the case in our model.

5 Asymmetric Information

We now introduce our other agency friction of private information. That A may possess private information prior to the start of experimentation in our drug-approval example is natural. For example, the company may have acquired information about the drug during the R&D phase or during animal clinical trials, which are not preregistered and are not directly observable by the FDA.

The road map for this section is as follows. After introducing asymmetric information into the model, we describe our main result before providing some additional intuition for its structure. While R was previously concerned only with providing incentives for A to continue experimentation, she must now consider how to distort the mechanism to *decrease* A's utility if he chooses to misreport his private information. Despite this conceptual difference, the structure of the problem is mathematically similar to that in Section 4, allowing us to use many of the same arguments to pin down the structure of the optimal mechanism. We conclude this section with some comparative statics.

We introduce asymmetric information by giving A a private binary signal at $t = 0$ that leads him to update his log-likelihood belief based on the realization of his signal to $z_i \in \{z_\ell, z_h\}$, where $z_\ell < z_h$. We say A is type h if he updates to z_h and is type ℓ if he updates to z_ℓ. The ex-ante probability of z_i is $\mathbb{P}(z_i)$.[18] We restrict attention to $a \in [0, 1]$, so that A wants approval in either state; to rule out a trivial case of the model, we assume $a > 0$ if $z_\ell = -\infty$. For expositional simplicity, we also focus on the case when $c_R = 0$ and z_h is sufficiently high; the case for lower values of z_h is studied in the Supplementary Materials, where we show that our main result in Theorem 2 on the structure of h's optimal mechanism holds for all z_h.[19]

We redefine the stopping mechanism in order to elicit A's private information. Applying the Revelation Principle, R offers a *menu* of stopping mechanisms from which A chooses by reporting his type.

[18] The distributions and accuracy of the binary signal conditional on each θ, along with the prior π_0, will pin down the values of z_h and z_ℓ as well as $\mathbb{P}(z_h)$.

[19] The assumption $c_R = 0$ is reasonable in our drug approval example where the costs of running clinical trials are paid by the company. Although omitted here, a previous version of the paper showed that all results go through when $c_R > 0$ and provided sufficient conditions under which our results on ℓ's optimal mechanism hold for all z_h.

Definition 5. *A **stopping mechanism menu** is a pair of mechanisms* $\{(\tau^i, d_\tau^i)\}_{i=h,\ell}$ *such that* $(\tau^i, d_\tau^i) \in \mathbb{T} \times \mathbb{D}$ *and* R *implements* (τ^i, d_τ^i) *when* A *reports being type* i*.*

When A misreports his type, his beliefs will differ from those of R's, leading to different expectations over the mechanism's outcome. The log-likelihood beliefs of R and a misreporting A will always differ by a constant $\Delta_z := z_h - z_\ell$.

The mechanism must provide incentives for A to truthfully report his type. However, A's report is not the only deviation he can carry out. He could, after misreporting his type, choose to quit prior to R approving or rejecting the project. Type i's value of misreporting his type to be j will depend on when he chooses to quit in j's mechanism. We define a dynamic version of incentive compatibility, which we call DIC, in a similar manner to DP. V^* captures the value of these double deviations.

Definition 6. *A stopping mechanism satisfies **DIC(i)** if for* $j \neq i$

$$V^*(\tau^j, d_\tau^j, z_i) \leq V(\tau^i, d_\tau^i, z_i).$$

Let Γ_i be the set of (τ, d_τ) that satisfy i's dynamic participation constraint. R's mechanism design problem with asymmetric information is given by

$$\sup_{\{(\tau^i, d_\tau^i)\}_{i=\ell,h}} \sum_{i=\ell,h} J(\tau^i, d_\tau^i, z_i) \cdot \mathbb{P}(z_i), \tag{1}$$

$$\text{subject to } (\tau^i, d_\tau^i) \in \Gamma_i,\ DIC(i),\ \forall i \in \{\ell, h\}.$$

Optimal Mechanism

To define the optimal mechanism, we first develop analogous functions to b^* and $\underline{B}(m)$. We must now keep track of what the starting beliefs of A are. Let $b_i^*(B) = b^*(B; z_i)$ be the lower threshold at which type i would quit against an approval threshold of B and $\underline{B}_i(m) = \max\{B :\ b_i^*(B) = m\}$ be the static approval threshold that would induce type i to quit when $X_t = m$.[20]

[20] As we show in the Appendix, when $a \geq 0$, A's utility strictly decreases if the approval threshold increases so there is only a single threshold that induces A to optimally quit at m. Thus, $\underline{B}_i(m) = b_i^{*-1}(m)$.

We define the mechanism for ℓ using two parameters $\eta_\ell = (\underline{b}_\ell, B^1_\ell)$ and for h using four parameters $\eta_h = (\underline{b}_h, b^1_h, B^2_h, B^1_h)$. Given η_i, the approval threshold for type i is $B^i(M;\eta_i)$:

$$B^\ell(m;\eta_\ell) = \begin{cases} B^1_\ell & \text{if } m > b^*_\ell(B^1_\ell), \\ \underline{B}_\ell(m) & \text{if } m \leq b^*_\ell(B^1_\ell), \end{cases}$$

$$B^h(m;\eta_h) = \begin{cases} B^1_h & \text{if } m > b^1_h, \\ \underline{B}_\ell(m) & \text{if } m \in (b^*_\ell(B^2_h), b^1_h], \\ B^2_h & \text{if } m \in (b^*_h(B^2_h), b^*_\ell(B^2_h)], \\ \underline{B}_h(m) & \text{if } m \leq b^*_h(B^2_h). \end{cases}$$

Similar to Section 4, we sometimes drop dependence on η_i.

We now present our main result for this section. Although ℓ's mechanism is essentially the same as his symmetric-information mechanism (SI-mechanism) in Theorem 1, h's mechanism can be *qualitatively* different. When A reports to be type h, the optimal mechanism can give him a fast-track to approval. A *fast-track mechanism* uses a low initial approval threshold B^1_h that remains fixed as long as A remains in the fast-track. A is thrown out of the fast track if the outcomes of the experimentation go poorly; more specifically, if X_t goes below the fixed lower "failure" threshold b^1_h. After being removed from the fast-track, R still allows A to experiment but imposes a more stringent standard as the approval threshold *jumps up* to $\underline{B}_\ell(b^1_h) > B^1_h$, something that never happens in the SI-mechanism. After this jump up, the approval threshold begins to decrease when M_t decreases. Eventually, the approval threshold reaches a level B^2_h at which the mechanism operates as h's SI-mechanism would—namely, the approval threshold remains fixed at B^2_h until X_t reaches the point at which h's participation constraints begin to bind and R again begins to lower the approval threshold.

Theorem 2. *For sufficiently high z_h, there exist η_h, η_ℓ such that the optimal menu $\{(\tau^i, d^i_\tau)\}_{i=\ell,h}$ is given by $\tau^i = \inf\{t : \ X_t \notin (\underline{b}_i, B^i(M_t;\eta_i))\}$ and $d^i_\tau = \mathbb{1}\big(X_{\tau^i} \geq B^i(M_{\tau^i};\eta_i)\big)$.*

To explain the intuition for the jump in a fast-track mechanism, note that R would like to give h a lower approval threshold than ℓ. Lowering h's approval threshold creates incentives for ℓ to misreport his type. To give h a lower approval

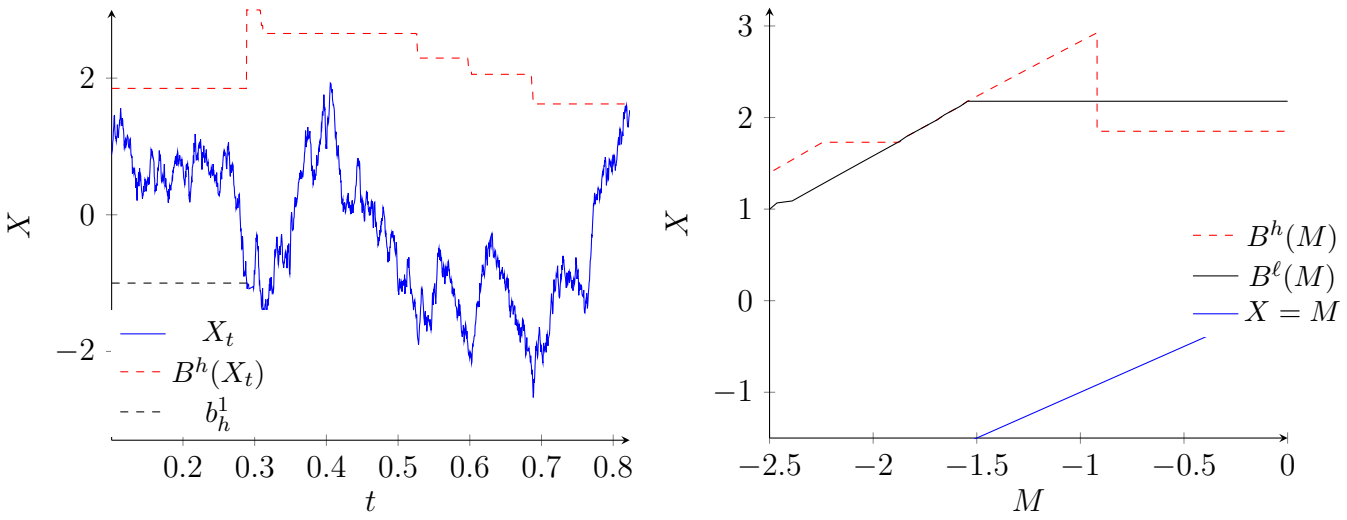

Figure 2: The left graph shows a fast-track mechanism along a realized evidence path over time. The approval threshold for h starts off low but jumps up at $\tau(b_h^1)$. The right graph compares the approval thresholds for h, ℓ.

threshold but still satisfy $DIC(\ell)$, R will need to introduce some distortions into the mechanism.

R could add distortions in many ways. Our results from Section 4 suggest that R optimally provides continuation value to A in a smooth manner, keeping the approval threshold stationary until A is indifferent between continuing and quitting. A similar intuition holds here: R smooths the approval threshold until we reach a history at which, in the continuation mechanism, ℓ is indifferent between continuing and quitting. This first occurs at $\tau(b_h^1)$, the expiry of the fast track. By increasing the approval threshold at $\tau(b_h^1)$, R reduces a misreporting ℓ's continuation value at the cost of a threshold that is above R's preferred level.

The optimal mechanism backloads distortions by increasing the approval threshold only when b_h^1 is first reached. R evaluates the mechanism with initial belief z_h, and a misreporting ℓ evaluates it using z_ℓ, so ℓ will find a decrease in X_t to b_h^1 more likely than R. By backloading the distortions, R is able to use the information generated by X_t to separate types and place distortions after histories that ℓ views as more likely than h does, thereby minimizing R's ex-ante cost of distortions while

still maintaining appropriate incentives for ℓ.[21]

The ability of ℓ to quit early (i.e., a double deviation) limits the amount of punishment R can give; that is, ℓ's continuation value can never be negative. Nevertheless, R could still punish ℓ in many ways. For example, R could reject at b_h^1. However, such a strong punishment is not needed. R could instead raise the approval threshold to $\underline{B}_\ell(b_h^1)$, which would induce ℓ to quit immediately. The approval threshold of $\underline{B}_\ell(b_h^1)$ achieves the same level of punishment for ℓ while still allowing h to continue experimenting.

R need not keep the threshold so high forever. When X_t reaches a new low, ℓ would strictly prefer to quit against the current threshold. The new low by X_t allows R to gradually decrease the approval threshold, leaving ℓ indifferent between quitting and continuing, thereby decreasing the distortions without weakening the punishment from ℓ's perspective. Eventually, when M_t is low enough, R would no longer want to decrease the approval threshold. At this point, which occurs when the approval threshold is at B_h^2, R can cease decreasing the approval threshold until she needs to incentivize h to continue experimentation, at which point she starts decreasing it again.

Our solution verifies the smoothing intuition discussed earlier, albeit with a key difference when compared to the SI-mechanism. With symmetric information, R keeps the threshold fixed until A is about to quit against the current threshold. With asymmetric information, R needs to decrease ℓ's utility of misreporting, which she can do via two ways: raise the approval threshold B_h^1 or induce ℓ to quit *before* he would optimally choose to quit against B_h^1, by setting $b_h^1 > b_\ell^*(B_h^1)$. By raising the failure threshold b_h^1, R is able to lower the approval threshold B_h^1 while still preserving DIC.

5.0.1 Sketch of Proof

In the proof, we derive each (τ^i, d_τ^i) separately. Let V_i be the utility type i gets from truthfully declaring his type in the solution to R's problem. Then (τ^i, d_τ^i) must maximize R's utility from type i among all mechanisms which give type i at least V_i utility and type $j \neq i$ less than V_j utility. The problem for type i's

[21] I would like to thank an anonymous referee for pointing out the benefits of back-loading for generating additional information to separate types.

mechanism is similar to SI, only now we add a modified version of DIC and a promise-keeping constraint $PK_i(V_i')$ to ensure the mechanism delivers an expected utility of at least $V_i' = V^*(\tau^j, d_\tau^j, z_i)$ to i. These V_i' and V_j are chosen by R when designing the optimal mechanism, but for now we can treat them as fixed.[22] R's asymmetric information mechanism for type i solves the problem AM_i given by

$$\begin{aligned}[AM_i]: \sup_{(\tau^i, d_\tau^i)} \; & J(\tau^i, d_\tau^i, z_i) \\ \text{subject to } & (\tau^i, d_\tau^i) \in \Gamma_i, \; PK_i(V_i'): \; V(\tau^i, d_\tau^i, z_i) \geq V_i', \\ & DIC(j, V_j): \; V^*(\tau^i, d_\tau^i, z_j) \leq V_j.\end{aligned}$$

In the Appendix, we show that $DIC(h)$ in 1 is slack when z_h sufficiently high and that, whenever $DIC(h)$ is slack, the solution to R's problem takes the same form as in Theorem 2.[23] When $DIC(h)$ is slack, we can safely drop $PK_h(V_h')$ from AM_h as h's utility from R's choice of (τ^h, d_τ^h) will yield strictly higher utility than $V^*(\tau^\ell, d_\tau^\ell, z_h)$.

To simplify the derivation of (τ^h, d_τ^h), we show that R will never reject above the rejection threshold in h's SI-mechanism. Rejection above this point is dominated by instead using a continuation mechanism with an approval threshold $\underline{B}_h(M_t)$ which will provide the same continuation value for h and ℓ as rejection and will increase R's continuation value.

After any history, using h's SI-mechanism as the continuation mechanism gives an upper bound on R's expected utility from the continuation mechanism of any solution to AM_h. We then take $\underline{m}$ to be the highest level of m at which ℓ would choose to quit immediately at $\tau(m)$ in h's SI-mechanism.[24] Because h is more op-

[22] The exact values of V_i' and V_j chosen by R will depend on the parameters of the problem such as $z_h, z_\ell, \mathbb{P}(z_h)$.

[23] This is the only point at which we use the assumption that either $z_h = \infty$ or $z_h \approx \infty$. Although it is intuitive for $DIC(h)$ to be slack as R and h's incentives are more aligned, there do exist values of z_h, z_ℓ such that $DIC(h)$ is binding. To see why, we note that there are two forces that determine the effect of a change in z_i on R's optimal approval threshold in the symmetric-information case. The first force is that when z_i increases, R requires less evidence before she would like to approve the project. This first force pushes toward giving h a lower approval threshold than ℓ. The second force is that when z_i is lower, R may have to lower the approval threshold to ensure ℓ is willing to experiment. This second force pushes toward giving ℓ a lower approval threshold. This second force can be large enough that $DIC(h)$ binds.

[24] If $z_h = \infty$, so that R always wants to approve h immediately, $\underline{m} = -\infty$.

timistic about his chances of approval, it is reasonable to conjecture that h will be willing to experiment as long as ℓ is. We then propose a relaxed version of AM_h in which we drop h's dynamic participation constraint but fix the continuation mechanism of any choice of (τ, d_τ) at $\tau(\underline{m})$ to be h's SI-mechanism. Furthermore, we relax $DIC(\ell, V_\ell)$ by only considering threshold quitting rules $\tau(X_n)$, with $X_n \in \mathcal{T}_N$. For each X_n, we call the corresponding relaxed version of $DIC(\ell, V_\ell)$ a $RDIC(X_n)$ constraint:[25]

$$RDIC(X_n): \ V(\tau^h \wedge \tau(X_n), d_\tau^h(X_n), z_\ell) \leq V_\ell.$$

The mathematical structure of this relaxed problem is very similar to the relaxed problem in Section 4, with two notable differences: the $RDIC(X_n)$ constraint contains an expectation with respect to z_ℓ rather than z_h, and $RDIC(X_n)$ requires providing sufficient disincentives rather than sufficient incentives as in $RDP(X_n)$. We deal with the first difference by a change of expectation that replaces ℓ's utility function with a modified version of v, allowing us to apply the arguments from Section 4 to find the structure of the optimal stopping rule. The second difference requires some additional arguments to derive the structure of the optimal stopping rule but still possesses enough similarity to relaxed problem in Section 4 that we can use similar arguments to pin down the optimal stopping rule.

In the limit as our grid of $RDIC$ constraints becomes fine, we show that ℓ's continuation value is weakly positive whenever $t < \tau(\underline{m})$. Because h has a higher belief that $\theta = H$ than ℓ does, h's continuation value is strictly positive prior to $\tau(\underline{m})$. When paired with the fact that h's continuation value is always weakly positive after $\tau(\underline{m})$ by Theorem 1, we find that dropping h's dynamic participation constraint in our relaxed problem was without loss.

Solving AM_ℓ turns out to be much simpler. Dropping $DIC(h)$ is equivalent to dropping $DIC(h, V_h)$. Once $DIC(h, V_h)$ is dropped, the only change in AM_ℓ relative to the problem in Section 4 is the inclusion of a promise-keeping constraint for ℓ, which does not qualitatively change the structure of the optimal mechanism. The only important difference is that the addition of the PK_ℓ constraint may lead R to use an approval threshold below X_c, thereby approving at beliefs that give her negative utility. Unlike the model in Section 4, the assumption of commitment by R will be necessary for implementing the solution to R's problem.

[25] We omit dependence on V_ℓ for notational convenience.

It is natural to guess that h's initial threshold B_h^1 will always be weakly lower than ℓ's initial threshold B_ℓ^1. When $B_h^1 < B_\ell^1$, it is not hard to see that h's approval threshold will jump up at b_h^1; that is, h's mechanism is a fast-track one. Although it will often be the case that h's initial threshold is lower, for some parameter values, h's approval threshold will be everywhere higher than ℓ's. To find sufficient conditions for when h's initial threshold will be at least as low as ℓ's, we need to make stronger assumptions on the z_ℓ. In the Supplementary Material, we show that if $z_\ell > \log(-f)$, then h's initial approval threshold is weakly lower than ℓ's.

Comparative Statics

In the absence of private information, increasing the cost c_A unambiguously hurts R, because doing so makes incentivizing experimentation more difficult for her. However, with asymmetric information, this comparative static does not always hold. In the absence of monetary transfers, costly experimentation provides a screening tool. This result can speak to the debate on who should fund drug trials (drug companies or government agencies), providing a reason for requiring the companies to run experiments. A higher cost forces A to have some "skin in the game" and makes eliciting any private information easier. When c_A becomes large, it is possible to screening types at a minimal cost for R.

Proposition 2. *R's value of the optimal mechanism is strictly decreasing in c_A under symmetric information. Under asymmetric information, R's value may increase in c_A.*

We might also wonder whether A having private information about θ is beneficial for R. On one hand, more information is useful for R. On the other hand, private information introduces information rents and can add distortions to the optimal mechanism. Which effect is greater is not obvious ex-ante. Proposition 3 shows that R prefers to have an informed A.

Proposition 3. *The value to R of the optimal mechanism under asymmetric information in which A learns θ perfectly is higher than the value to R of the optimal mechanism under symmetric information with same prior on θ.*

The proof is quite simple. When A is informed, consider the sub-optimal mechanism in which R simply offers the SI-mechanism for both types and lets them

quit whenever they desire. Because ℓ is more pessimistic about the chances of approval, he will quit earlier than an uninformed A. This earlier quitting time by ℓ improves R's utility when compared to the SI-mechanism.

6 Application of the Model: Drug Approval

We now discuss how our model provides insights into the features of the real-world drug approval process. In particular, we focus on two notable features of our results, the decreasing approval threshold and the fast track. These two features distinguish our mechanisms from the optimal mechanism of a single decision-maker, who would choose a static-threshold mechanism.

The insight provided by our decreasing threshold is that lowering the approval threshold can provide incentives for companies to continue experimentation when they might otherwise quit. A single-decision-maker model would use a single trial with a constant approval threshold and no subsequent experimentation after the trial is declared a failure. Yet, we often see drug companies run additional clinical trials after a failed one and receive approval if these later trials are successful. Our results show the benefits of allowing the additional trials in providing incentives for companies to perform additional experimentation.

How such a decreasing threshold can be implemented can be illustrated using the example in Section 4 comparing Mechanism One and Two. Suppose A needs one successful clinical trial for approval. Mechanism One consists of a single trial that lasts until B_1 or b_1 is reached. We call the amount of positive evidence A needs to generate from the beginning of the trial for the trial to be declared a success the *trial standard*. For this first trial, the trial standard is B_1; X_t starts at $X_0 = 0$ and needs to go up, in total, by $B_1 - X_0 = B_1$. Upon reaching b_1, the trial is declared a failure and development of the drug is abandoned.

Mechanism Two (which can be viewed as a rough version of the optimal mechanism) uses the same first trial as Mechanism One, but, instead of rejecting after the first trial fails, allows A to begin a second trial. When the first trial ends at failure for date t, $X_t = b_1$, and needs to reach B_2 for approval, so the standard for the second trial is $B_2 - b_1$. Note that if $B_1 = B_2 - b_1$, lowering the threshold to B_2 when X_t is at b_1 is equivalent to giving A the *same* trial standard for both trials. Keeping the trial standard fixed effectively decreases the approval threshold

upon reaching the end of the first trial at b_1. This discussion shows how allowing multiple trials while keeping the same trial standards leads to a lowering of the approval threshold.

In a recent change to the evidentiary standards for Alzheimer's drugs, the FDA has shown that they understand the need to adjust the bar for approval after failed clinical trials. The failure of numerous clinical trials for drugs to treat Alzheimer's has discouraged investment in their development, as companies have become pessimistic about the chances of successful clinical trials (Cummings and Zhong (2015)). After these repeated failures, the FDA lowered the approval bar by removing one of the criteria for clinical trial success. "U.S. regulators have proposed lowering the bar for clinical trial success... Medicines tested to treat Alzheimer's have had a dismal track record, and the Food and Drug Administration has recognized that goals for clinical trials need to evolve..."[26]; the standards have been subsequently updated since this writing of this article.[27] This change in the bar for approval was viewed as encouraging companies to run clinical trials: the change gave "the field more confidence in being bold about the trials that [the field] design and then carry out" (Reuters) and "signaled [the FDA] is ready to accelerate the development of new drugs for Alzheimer's" (STAT). Although this example concerns the failure of clinical trials for multiple drugs, as opposed to the single drug considered in our model, the economic forces at play are the same as in our model, as approval standards must be adjusted to provide additional incentives for experimentation.[28] Our results explain why such a lowering of the approval standards is optimal and why it happens only after failed clinical trials.

Our fast-track mechanism matches the features of real-world fast-track and expedited-approval programs. Beginning in 1988, the FDA introduced a Fast Track

[26] *Reuters* ("U.S. FDA looks to pave way for earlier-stage Alzheimer's drugs." February 15, 2018). I would like to thank Jorge Lemus for sending me this article.

[27] *STAT* ("FDA's updated standards remove unnecessary barrier to testing Alzheimer's Drugs." March 5, 2018).

[28] A previous version of our paper illustrated this point with an extension to the model in which R interacts sequentially with two agents, each with their own realization of θ. When the θs of the agents are positively correlated and the probability of $\theta = H$ is low enough (as is likely in the Alzheimer's case), after rejecting the first agent, R will give the second one a lower approval threshold than the first agent. The second agent is more pessimistic about his θ than the first agent, and R will need to decrease the approval threshold to provide more incentives for him to experiment.

designation for certain drugs with the purpose of getting promising new drugs that address a serious unmet medical need to the patient earlier. In 1992, the FDA also began offering an Accelerated Approval program with similar features. Companies apply for these programs prior to the start of clinical trials. These programs give drugs lower evidentiary standards for approval and have been succesful in reducing the length of the clincial trial process (Kesselheim and Darrow (2015)). This lower evidentiary standard matches the lower threshold of our model. However, admittance to these expedited approval programs is not permenant. The Fast-Track designation may be removed if, after seeing the result of some clinical trials, the drug is found to "no longer demonstrates a potential to address an unmet medical need" (FDA (2014)). Similarly, the Accelerated Approval program allows drugs to be removed from the program if a "trial required to verify the predicted clinical benefit of the product fails to verify such benefit" (FDA (2014)). In our model, a decrease in X_t to the failure threshold b_h^1 corresponds to the loss of "potential to address an unmet medical need" and failure to "verify predicted clinical benefit." These attributes—the lower approval standards and the removal from the expedited approval program upon poor trial outcomes—fit the qualitative features of our optimal mechanism. Our model explains why lower standards must be paired with the possibility of being thrown out of the program, after which a higher non-fast track standard would be applied.[29]

The goal of these programs is to expedite the approval for drugs that have the potential to address serious or life-threatening medical conditions. Our results suggest that similar programs can be expanded beyond drugs to treat serious conditions. By offering these different tracks, the FDA can shorten clinical trial lengths, thereby benefiting both patients and drug companies by reducing the length of experimentation.

[29]By solving the problems AM_h and AM_ℓ for arbitrary values of V_i, V_j', we can immediately extend our results to include additional costs of misreporting one's type into the model. For example, we can include lying costs (e.g., the employees of a drug company may dislike misrepresenting their beliefs about the efficacy of the drug) and the organizational and time costs of submitting an application for a fast-track. Including such costs will reduce the extent the approval standard will need to jump up after the failed fast-track in order to maintain incentive compatibility.

7 Conclusion

In this paper, we present a model of a hypothesis-testing problem with agency concerns. The optimal mechanism features a history-dependent approval threshold, that can still be solved for in a tractable way and be written as a function of the minimum of the belief process. We find the optimal mechanism when the agent possesses no private information takes the form of a monotonically decreasing approval threshold.

We also apply the model to the case in which the agent has private information, which adds new and distinct features to the optimal mechanism. The optimal solution may take the form of a fast-track mechanism: high types are offered a low starting approval threshold, but if the evidence becomes too unfavorable, the approval threshold jumps up, entering a punishment phase in which it drifts back down slowly. These results show how agency problems may lead to an evolving and history-dependent approval rule.

References

Thomas Balzer and Klaus Janßen. A duality approach to problems of combined stopping and deciding under constraints. *Mathematical Methods of Operations Research*, 55:431–446, 2002.

Dirk Bergemann and Ulrich Hege. The financing of innovation: Learning and stopping. *The RAND Journal of Economics*, 36:719–752, 2005.

Daniel Carpenter and Michael M. Ting. Regulatory errors with endogenous agendas. *American Journal of Political Science*, 51:835–852, 2007.

Morstorf T. Cummings, J.L. and K. Zhong. Alzheimer's disease drug-development pipeline: few candidates, frequent failures. *Alzheimer's Research and Therapy*, 37, 2015.

Jeffrey C. Ely and Martin Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128, 2020.

FDA. Guidance for industry expedited programs for serious conditions – drugs and biologics. 2014.

Drew Fudenberg, Philipp Strack, and Tomasz Strzalecki. Speed, accuracy, and the optimal timing of choices. *American Economic Review*, pages 3651–84, 2018.

George Georgiadis and Balazs Szentes. Optimal monitoring design. *Econometrica*, 2020.

Sebastian Gryglewicz and Aaron Kolb. Strategic pricing in volatile markets. Working Paper, 2019.

Yingni Guo. Dynamic delegation of experimentation. *American Economic Review*, 106:1969–2008, 2016.

Marina Halac, Navin Kartik, and Qingmin Liu. Optimal contracts for experimentation. *The Review of Economic Studies*, 83:1040–1091., 2016.

Marina Halac, Navin Kartik, and Qingmin Liu. Contests for experimentation. *Journal of Political Economy*, 125:1523–1569, 2017.

Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *Review of Economic Studies*, 49:315–333, 1982.

Emeric Henry and Marco Ottaviani. Research and the approval process: the organization of persuasion. *American Economic Review*, 109:911–55, 2019.

Bo Wang Jessica M. Franklin Kesselheim, Aaron S. and Jonathan J. Darrow. Trends in utilization of fda expedited drug development and approval programs, 1987-2014: cohort study. *BMJ*, 315, 2015.

Thomas Kruse and Philipp Strack. Optimal stopping with private information. *Journal of Economic Theory*, 159:702, 2015.

Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems*. Birkhauser Basel, 2006.

Debraj Ray. The time structure of self-enforcing agreements. *Econometrica*, pages 547–582, 2002.

Yuliy Sannikov. A continuous-time version of the principal-agent problem. *Review of Economic Studies*, pages 957–984, 2008.

Albert N Shiryaev. *Optimal stopping rules*, volume 8. Springer Science & Business Media, 2007.

Nancy L Stokey. *The Economics of Inaction: Stochastic Control models with fixed costs.* Princeton University Press, 2009.

Jonathan Thomas and Tim Worrall. Self-enforcing wage contracts. *The Review of Economic Studies*, pages 541–554, 1988.

Abraham Wald. *Sequential Analysis*. John Wiley and Sons, New York, 1947.

Appendices

We now present a formal definition of a continuation mechanism. Without loss, we take $\Omega = C[0,\infty)$, where $C[0,\infty)$ is the space of continuous functions on $[0,\infty)$ and $X_t(\omega) = \omega_t$. Let χ_t be the shift operator defined by $\chi_t\omega = \omega'$, where $\omega'_s = \omega_{s+t}$.

Definition 7. *The* ***continuation mechanism*** *of (τ, d_τ) at h_t is $(\tau[h_t], d_\tau[h_t])$ where, for each ω with history h_t, $\tau[h_t](\chi_t\omega) = \tau(\omega) - t$ and $d_\tau[h_t](\chi_t\omega) = d_\tau(\omega)$.*

A mechanism is stationary in (X, M) when, for each (x, m), its continuation mechanism at each history h_t with $(X_t, M_t) = (x, m)$ is the same.[30] For a mechanism (τ, d_τ) that is stationary in (X, M), when discussing its continuation mechanism at some history $h_{\tau(m)}$, we will simply call it the continuation mechanism of (τ, d_τ) at $\tau(m)$.

It will be notationally useful to allow for $M_0 < X_0$, so we treat M_0 as a parameter and redefine $M_t = \min\{\min_{s\in[0,t]} X_t,\ M_0\}$. We then let $\mathbb{E}_{x,m}$ denote the expectation over X, Y_0 when $X_0 = x$, $M_0 = m$ and $Z_0 = z_0 + \frac{2\mu}{\sigma^2}x$ (i.e., the belief updated from x and z_0); if not specified, we take $M_0 = X_0$. Thus, $\mathbb{E}^x = \mathbb{E}_{x,x}$. Define $\mathbb{E}^x[\cdot|\theta]$ to be the expectation of X, Y_0 given $X_0 = x$ in state θ.

Suppose (τ, d_τ) is stationary in (X, M). For each (x, m) and history h_t such that $(X_t, M_t) = (x, m)$, its continuation mechanism at h_t leads to the same distribution over outcomes as when using (τ, d_τ) from $t = 0$ when $(X_0, M_0) = (x, m)$. For any history h_t with $(X_t, M_t) = (m, m)$, we have $\mathbb{E}^m[e^{-r\tau[h_t]}v(X_{\tau[h_t]}, d_\tau[h_t])] = \mathbb{E}^m[e^{-r\tau}v(X_\tau, d_\tau)]$.

We start by looking at $\widetilde{V}$ and $\widetilde{J}$. It is straightforward to show that both functions are continuous in all arguments. Our next three lemmas give some properties of $\widetilde{V}$ and $\widetilde{J}$.

Lemma 1. *$\widetilde{V}(B, b, x)$ and $\widetilde{J}(B, b, x)$ are single peaked in B on $[x,\infty)$ when $b < x$ and in b on $(-\infty, x]$ when $x < B$.*

[30] Whenever discussing a pair (x, m), we assume $x \geq m$.

Lemma 1 shows how $\widetilde{V}$ and $\widetilde{J}$ change in B and b—for example, $\widetilde{V}(B,b,x)$ is decreasing in B for $B > \arg\max_{B'} \widetilde{V}(B',b,x)$ and is increasing in B for $B < \arg\max_{B'} \widetilde{V}(B',b,x)$.

Lemma 2. *For any $b < x$, $\arg\max_{B \geq x} \widetilde{V}(B,b,x) < \arg\max_{B \geq x} \widetilde{J}(B,b,x)$.*

Lemma 3. *If $z' > z$, $\widetilde{V}(B,b,x;z) \geq 0$ and $x \in (b,B)$, then $\widetilde{V}(B,b,x;z') > \widetilde{V}(B,b,x;z)$.*

For $U \geq 0$, we define a modified version of $\widetilde{J}$, call it $\check{J}$, that replaces the utility of 0 at b with U:

$$\check{J}(B,b,x,U) = \mathbb{E}^x[e^{-r\tau_+(B;b)}u(B,1) + e^{-r\tau(b;B)}(U + \frac{c_R}{r})] - \frac{c_R}{r}.$$

Similar to $\widetilde{V}$ and $\widetilde{J}$, we restrict attention to $b < B$.

Lemma 4. *$\check{J}(B,b,x,U)$ is single peaked in B on $[x,\infty)$ when $b < x$ and in b on $(-\infty,x]$ when $x < B$. $\arg\max_{B \geq x} \check{J}(B,b,x,U)$ is increasing in U.*

A General Stopping Problem

We now introduce a general optimal stopping problem that will be useful in deriving common properties of the solution to R's problems in Sections 4 and 5. Take some $(\xi^1, X^1, \xi^2, ..., \xi^P, X^P) \in \mathbb{R}^{2P}_-$ such that $X^{k+1} < X^k$; throughout the Appendix and Supplementary Materials, when discussing a set $\{X^k\}_{k=1}^P$, we adopt the convention $X^{k+1} < X^k$ for all k, $X^{P+1} = -\infty$ and $X^0 = 0$. For each $d \in \{0,1\}$ and $0 \leq k \leq P$, let $g(x,k,d)$ be a bounded continuous function of x such that $g(X_t,k,d)$ is a martingale in X_t. Define $d_k^x = \arg\max_d g(x,k,d)$ and, for $m \leq 0$, $\kappa(m)$ to be the k such that $m \in (X^{k+1}, X^k]$. We assume that $g(X^{k+1}, k+1, d_k^{X^{k+1}}) - g(X^{k+1}, k, d_k^{X^{k+1}}) \geq -\xi^{k+1}$ for all k.[31] Consider the following stopping problem:

$$\sup_{(\tau,d_\tau)} \mathbb{E}[e^{-r\tau} g(X_\tau, \kappa(M_\tau), d_\tau) + \sum_{k=1}^{P} e^{-r\tau(X^k)} \xi^k \mathbb{1}(\tau \geq \tau(X^k))]. \tag{2}$$

[31] When translating R's problems in Sections 4 and 5 into the form of this general stopping problem, we will directly verify that the corresponding conditions of g, ξ hold.

Proposition 4. *There exists a* (τ^*, d_τ^*) *that solves 2 with* $\tau^* = \inf\{t : \ X_t \notin (\underline{b}, B(M_t))\}$ *for some threshold* $\underline{b}$ *and function* $B(m)$ *that is constant on* $(X^{k+1}, X^k]$ *for all* k. d_τ^* *is a function of only* X_{τ^*} *and* M_{τ^*}.

Proof. It is clear that $d_\tau^* = \arg\max_d \ g(X_\tau, \kappa(M_\tau), d)$, and therefore only depends on (X_τ, M_τ). We focus on deriving τ^*. The continuation value in 2 at $\tau(X^k)$ is[32]

$$K(X^k) = \sup_{(\tau, d_\tau)} \ \mathbb{E}^{X^k}[e^{-r\tau} g(X_\tau, \kappa(M_\tau), d_\tau) + \sum_{j=k+1}^{P} e^{-r\tau(X^j)} \xi^j \mathbb{1}(\tau \geq \tau(X^j))] + \xi^k.$$

Let $G_k(x) = \mathbb{1}(x \leq X^{k+1}) K(X^{k+1}) + \mathbb{1}(x > X^{k+1}) g(x, k, d_k^x)$. The continuation value $F_k(x)$ in 2 at $(X_t, M_t) = (x, m)$ for $k = \kappa(m)$ when $t \neq \tau(X^k)$ is

$$F_k(x) = \sup_{\tau} \ \mathbb{E}^x[e^{-r(\tau \wedge \tau(X^{k+1}))} G_k(X_{\tau \wedge \tau(X^{k+1})})]. \tag{3}$$

At $t = \tau(X^k)$, the instantaneous cost ξ^k is a sunk cost and plays no role in the future decision of when to stop, so $K(X^k) = F_k(X^k) + \xi^k$. It is optimal stop when $(X_t, M_t) = (x, m)$ in 2 if and only if it is optimal to stop when $(X_t, M_t) = (x, m)$ in 3.

We now show that $G_k(x)$ is upper semicontinuous. Continuity at $x \neq X^{k+1}$ is easy to see. Note that $K(X^{k+1}) \geq g(X^{k+1}, k+1, d) + \xi^{k+1}$ for all d. Then $\lim_{x \downarrow X^{k+1}} (G_k(X^{k+1}) - G_k(x)) = K(X^{k+1}) - g(x, k, d_k^x) \geq g(X^{k+1}, k+1, d_k^{X^{k+1}}) - g(X^{k+1}, k, d_k^{X^{k+1}}) + \xi^{k+1} \geq 0$ and $\lim_{x \uparrow X^{k+1}} (G_k(X^{k+1}) - G_k(x)) = 0$. We conclude that G_k is upper semicontinuous.

It easy to see that $F_k(x)$ is continuous. Corollary 2.9 and Remark 2.10 from Peskir and Shiryaev (2006) show that continuity of F_k and upper semicontinuity of G_k imply that an optimal stopping rule for 3 exists[33] and that it is optimal to stop at $t \in [\tau(X^k), \tau(X^{k+1}))$ if and only if $X_t \in \mathcal{D}_k := \{x > X^{k+1} : \ F_k(x) = G_k(x)\}$. Because $t \in [\tau(X^k), \tau(X^{k+1}))$ if and only if $\kappa(M_t) = k$, $\tau^* = \inf\{t : \ X_t \in \mathcal{D}_{\kappa(M_t)}\}$ is an optimal stopping rule.

[32] We adopt the conventions that a sum is zero when it has a lower bound that is higher than its upper bound, so $\sum_{j=P+1}^{P} e^{-r\tau(X^j)} \xi^j \mathbb{1}(\tau > \tau(X^j)) = 0$, and $\xi^0 = 0$.

[33] Formally, because we allow for $\tau = \infty$, we consider, in the terminology of Peskir and Shiryaev (2006), the existence of an optimal Markov time. They also require boundedness of $|G_k|$, which easily follows from the boundedness of $g(x, k, d)$.

Let $b_k = \sup\{x \leq X^k : \ x \in \mathcal{D}_k\}$ and $B_k = \inf\{x \geq X^k : \ x \in \mathcal{D}_k\}$.[34] By continuity of X, the first entrance time into $\mathcal{D}_k$, if it ever occurs, happens when X crosses either b_k or B_k.[35] If $b_k > -\infty$, then $\mathbb{P}(\tau^* > \tau(X^{k+1})) = 0$. Letting $B(m) = B_{\kappa(m)}$ and $\underline{b} = \max_k \ b_k$ (with $\underline{b} = -\infty$ if $b_k = -\infty$ for all k), we have $\tau^* = \inf\{t : \ X_t \notin (\underline{b}, B(M_t))\}$. □

In the Supplementary Material we provide sufficient conditions under which τ^* will be unique,[36] which will be useful for proving the existence of a solution to the relaxed problems we consider in Sections 4 and 5.

B Symmetric Information

We solve a slightly more general version of our relaxed problem in Section 4 in which we add a promise-keeping constraint $PK(V) : V(\tau, d_\tau, z_0) \geq V$ to ensure A's expected utility is at least $V < \sup_{\tau, d_\tau} \ V(\tau, d_\tau, z_0)$:[37]

$$H_N(V) = \sup_{(\tau, d_\tau)} \ J(\tau, d_\tau, z_0)$$
$$\text{subject to } PK(V), \ RDP(X_n) \ \forall X_n \in \mathcal{T}_N.$$

By Theorem 1 of Balzer and Janßen (2002), $H_N(V) = \inf_{\Lambda \in \mathbb{R}^{N+1}_-} \ \sup_{(\tau, d_\tau)} \ \mathcal{L}(\tau, d_\tau, \Lambda)$, where, for $\Lambda = (\lambda_1, ..., \lambda_N, \gamma)$,

$$\mathcal{L}(\tau, d_\tau, \Lambda) = \mathbb{E}[e^{-r\tau}(u(X_\tau, d_\tau) - (\gamma + \sum_{n=1}^{N} \lambda_n) v(X_\tau, d_\tau)) \tag{4}$$
$$+ \sum_{n=1}^{N} e^{-r(\tau \wedge \tau(X_n))} \lambda_n v(X_\tau, d_\tau(X_n))] - \gamma(V + \frac{c_A}{r}).$$

[34] We note that $b_k = -\infty$ if it is never optimal to stop at $X_t \in (X^{k+1}, X^k]$ for $t \in [\tau(X^k), \tau(X^{k+1}))$ and $B_k = \infty$ if it is never optimal to stop at $X_t \geq X^k$ for such t. If $b_k = B_k$, then it is optimal to stop immediately at $\tau(X^k)$.

[35] In order for X to reach any $x > B_k$ from $X_t = X^k$, it must travel through B_k. Similarly for X to reach $x < b_k$ from $X_t = X^k$, it must travel through b_k.

[36] Whenever discussing uniqueness of an optimal stopping rule, we ignore differences on probability zero events; the stopping rule used on these events does not impact ex-ante payoffs and, thus, can be set to coincide with the stopping rule elsewhere without loss.

[37] The solution when $V = \sup_{\tau, d_\tau} \ V(\tau, d_\tau, z_0)$ is the solution to a single-decision maker problem with A's preferences, which is a static-threshold mechanism.

Moreover, the inf is achieved at some $\widehat{\Lambda} = (\widehat{\lambda}_1, ..., \widehat{\gamma})$. We say that (τ, d_τ) and $\widehat{\Lambda}$ satisfy the complementary slackness conditions for $RDP(X_n)$ if (τ, d_τ) satisfies $RDP(X_n)$ and $\widehat{\lambda}_n < 0$ only if $RDP(X_n)$ binds under (τ, d_τ), with an analogous definition for $PK(V)$. Let $\mathcal{L}^*(\Lambda) = \sup_{(\tau,d_\tau)} \mathcal{L}(\tau, d_\tau, \Lambda)$. Take $\widehat{\Lambda} \in \arg\min_{\Lambda \in \mathbb{R}_-^{N+1}} \mathcal{L}^*(\Lambda)$ and $(\tau, d_\tau) \in \arg\max_{(\tau', d'_\tau)} \mathcal{L}(\tau', d'_\tau, \widehat{\Lambda})$. If (τ, d_τ) and $\widehat{\Lambda}$ satisfy complementary slackness conditions for all constraints, then (τ, d_τ) solves $H_N(V)$ by Theorem 1 of Balzer and Janßen (2002).[38] We use this result to establish the existence of a solution to $H_N(V)$ and show that it is stationary in (X, M).

Proposition 5. *There exists a $(\tau_N^*, d_{N,\tau}^*)$ that solves $H_N(V)$. Let $\mathcal{B}_N = \{X_n : RDP(X_n) \text{ binds under } (\tau_N^*, d_{N,\tau}^*)\} = \{X^1, ..., X^P\}$. There exist a threshold $\underline{b}_N$ and function $B_N(m)$ that is constant on $(X^{k+1}, X^k]$ for each $0 \leq k \leq P$ such that $\tau_N^* = \inf\{t : X_t \notin (\underline{b}_N, B_N(M_t))\}$ and $d_{N,\tau}^* = \mathbb{1}(X_{\tau_N^*} \geq B_N(M_{\tau_N^*}))$. If $B_N(0) > 0$, then $B_N(m) > m$ for all $m < 0$.*

The fact that $(\tau_N^*, d_{N,\tau}^*)$ is stationary in (X, M) allows us to write A's continuation value[39] at $\tau(X_n)$ as $\rho(X_n) := \mathbb{E}^{X_n}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)] - \frac{c_A}{r}$.

Lemma 5. *$\rho(X_n) \geq 0$, with equality if and only if $RDP(X_n)$ binds.*

Proof. $RDP(X_n)$ implies

$$
\begin{aligned}
&\mathbb{E}[e^{-r\tau_N^*}\mathbb{1}(\tau_N^* < \tau(X_n))v(X_{\tau_N^*}, d_{N,\tau}^*) + e^{-r\tau(X_n)}\mathbb{1}(\tau_N^* \geq \tau(X_n))\frac{c_A}{r}] \qquad (5)\\
&= \mathbb{E}[e^{-r(\tau_N^* \wedge \tau(X_n))} v(X_{\tau_N^* \wedge \tau(X_n)}, d_{N,\tau}^*(X_n))]\\
&\leq \mathbb{E}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)]\\
&= \mathbb{E}\big[e^{-r\tau_N^*}\mathbb{1}(\tau_N^* < \tau(X_n))v(X_{\tau_N^*}, d_{N,\tau}^*) + e^{-r\tau(X_n)}\mathbb{1}(\tau_N^* \geq \tau(X_n))\mathbb{E}^{X_n}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)]\big]\\
&= \mathbb{E}[e^{-r\tau_N^*}\mathbb{1}(\tau_N^* < \tau(X_n))v(X_{\tau_N^*}, d_{N,\tau}^*) + e^{-r\tau(X_n)}\mathbb{1}(\tau_N^* \geq \tau(X_n))(\rho(X_n) + \frac{c_A}{r})].
\end{aligned}
$$

with equality if and only if $RDP(X_n)$ binds. Using the first and last lines, we have $\rho(X_n) \geq 0$, with equality if and only if $RDP(X_n)$ binds. □

[38] This theorem requires a Slater condition that there exist a mecahnism in which all constraints are slack. In the Supplementary Material we show that this holds in our problem and discuss two other technical conditions Balzer and Janßen (2002) require.

[39] Where no confusion is caused, we will leave implicit the mechanism under which we are evaluating A's continuation value, here $(\tau_N^*, d_{N,\tau}^*)$.

Next, we provide a characterization of $\mathcal{B}_N$.

Lemma 6. $\mathcal{B}_N = \{X_n \in \mathcal{T}_N : \ X_n \leq X^1\}$.

Proof. $RDP(X_n)$ binds whenever $\mathbb{P}(\tau(X_n) \geq \tau_N^*) = 1$. Thus, it binds for all $X_n < \underline{b}_N$ and, if $B_N(0) = 0$, for all X_n (so that $X^1 = 0$). For the sake of contradiction, suppose $B_N(0) > 0$, $\underline{b}_N \leq X^k - \delta_N$ and $X^{k+1} < X^k - \delta_N$. First consider the case in which $X^{k+1} \geq \underline{b}_N$. By the stationarity of $(\tau_N^*, d_{N,\tau}^*)$ in (X, M) and the fact that $B_N(m) > X^k$ is constant on $(X^{k+1}, X^k]$, for all $m \in (X^{k+1}, X^k]$ we have $\mathbb{E}_{X^k,m}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)] = \mathbb{E}^{X^k}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)]$. By Lemma 5, $\rho(X^k) = \rho(X^{k+1}) = 0 < \rho(X^k - \delta_N)$, so $\mathbb{E}^x[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)] = \frac{c_A}{r}$ for $x \in \{X^k, X^{k+1}\}$. If taking the expectation over X when $(X_0, M_0) = (X^k - \delta_N, X^k - \delta_N)$, whenever $\tau_+(X^k) < \tau(X^{k+1})$ we have $M_{\tau_+(X^k)} \in (X^{k+1}, X^k)$. Thus,

$$\begin{aligned}
\rho(X^k - \delta_N) &= \mathbb{E}^{X^k - \delta_N}\big[e^{-r\tau_+(X^k; X^{k+1})} \mathbb{E}_{X^k, M_{\tau_+(X^k)}}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)] \\
&\quad + e^{-r\tau(X^{k+1}; X^k)} \mathbb{E}^{X^{k+1}}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{N,\tau}^*)]\big] - \frac{c_A}{r} \\
&= \mathbb{E}^{X^k - \delta_N}\big[e^{-r(\tau_+(X^k) \wedge \tau(X^{k+1}))} \frac{c_A}{r}\big] - \frac{c_A}{r} < 0,
\end{aligned}$$

contradicting $\rho(X^k - \delta_N) > 0$. If $X^{k+1} < \underline{b}_N$, we can replace X^{k+1} with $\underline{b}_N$ and the argument is identical. □

Take $\underline{B}_N(m) := \max\{B : \ \widetilde{V}(B, m - \delta_N, m) = 0\}$. We next state properties of $\underline{B}_N$ before using it to characterize $B_N(m)$.

Lemma 7. *$\underline{B}_N(m)$ is continuous and increasing in m on $[b_A^{FB} + \delta_N, \infty)$, with $\lim_{N\to\infty} \underline{B}_N(m) = \underline{B}(m)$ for $m \geq b_A^{FB}$. $\underline{B}(m)$ is continuous.*

Lemma 8. *For each $1 \leq k \leq P$ and $m \in (X^{k+1}, X^k]$ such that $X^{k+1} \geq \underline{b}_N$, $B_N(m) = \underline{B}_N(X^k)$. If $X^1 < 0$, then $B_N(0) < \underline{B}_N(X^1 + \delta_N)$.*

Proof. Suppose $X^{k+1} \geq \underline{b}_N$. By $\rho(X^{k+1}) = 0$, $\mathbb{E}^{X^{k+1}}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{\tau_N^*}^*)] = \frac{c_A}{r}$, so

$$\begin{aligned}
\rho(X^k) &= \mathbb{E}^{X^k}\big[e^{-r\tau_+(B_N(X^k); X^{k+1})} v(B_N(X^k), 1) + e^{-r\tau(X^{k+1}; B_N(X^k))} \mathbb{E}^{X^{k+1}}[e^{-r\tau_N^*} v(X_{\tau_N^*}, d_{\tau_N^*}^*)]\big] \\
&= \mathbb{E}^{X^k}\big[e^{-r\tau_+(B_N(X^k); X^{k+1})} v(B_N(X^k), 1) + e^{-r\tau(X^{k+1}; B_N(X^k))} \frac{c_A}{r}\big] - \frac{c_A}{r} \\
&= \widetilde{V}(B_N(X^k), X^{k+1}, X^k). \qquad (6)
\end{aligned}$$

Because $RDP(X^k)$ binds, we have $0 = \rho(X^k) = \widetilde{V}(B_N(X^k), X^{k+1}, X^k)$. By the single-peakedness of $\widetilde{V}$ in B, at most two possible values of B, say $B^1 < B^2$, satisfy $\widetilde{V}(B, X^{k+1}, X^k) = 0$. If only one solution exists, we are done. If two exist, single-peakedness implies $B^1 < \arg\max_B \widetilde{V}(B, X^{k+1}, X^k) < B^2$. By Lemma 6, $X^{k+1} = X^k - \delta_N$, so $B_2 = \underline{B}_N(X^k)$.

Let $U \geq 0$ be R's continuation value at $\tau(X^{k+1})$.[40] Holding $B_N(m)$ fixed for $m \leq X^{k+1}$, R's preferred approval threshold at $\tau(X^k)$ is $\arg\max_B \check{J}(B, X^{k+1}, X^k, U)$. By Lemmas 2 and 4, we have

$$\begin{aligned} \arg\max_B \check{J}(B, X^{k+1}, X^k, U) &\geq \arg\max_B \check{J}(B, X^{k+1}, X^k, 0) \qquad (7) \\ &= \arg\max_B \widetilde{J}(B, X^{k+1}, X^k) \geq \arg\max_B \widetilde{V}(B, X^{k+1}, X^k). \end{aligned}$$

By the single-peakedness of $\widetilde{V}$ and $\check{J}$, B^1 is Pareto dominated by $\arg\max_B \widetilde{V}(B, X^{k+1}, X^k)$ and so cannot be optimal. Thus, $B_N(X^k) = B^2 = \underline{B}_N(X^k)$.

By the same arguments as in 6, $\rho(X^1 + \delta_N) = \widetilde{V}(B_N(0), X^1, X^1 + \delta_N)$. Then $\widetilde{V}(B, X^1, X^1 + \delta_N) < 0$ for all $B > \underline{B}_N(X^1 + \delta_N)$ since $\widetilde{V}(B, X^1, X^1 + \delta_N)$ is equal to 0 at $B = \underline{B}_N(X^1 + \delta_N)$ and is decreasing in B for $B > \underline{B}_N(X^1 + \delta_N)$. Because $\rho(X^1 + \delta_N) > 0$, it must be that $B_N(0) < \underline{B}_N(X^1 + \delta_N)$. □

We now use Lemma 7 to pin down the limit approval threshold.

Lemma 9. $\lim_{N\to\infty} B_N(M_t) = B(M_t; B^1)$ *for* $B^1 := \lim_{N\to\infty} B_N(0)$.

Proof. Given Lemmas 7 and 8, it suffices to show $B_N(X^1) \leq B_N(0)$. Suppose $B_N(X^1) > B_N(0) \geq 0$. Returning to the Lagrangian in 4 and fixing the continuation value $K(X^1)$ at $\tau(X^1)$, the fact that $B_N(0)$ is the optimal threshold implies

$$B_N(0) = \arg\max_B \mathbb{E}[e^{-r\tau_+(B;X^1)}(u(B,1) - \widehat{\gamma}v(B,1)) + e^{-r\tau(X^1;B)}K(X^1)].$$

Now we consider the choice of $B_N(X^1)$ in our Lagrangian. By standard dynamic programming arguments, the choice at $t \in [\tau(X^k), \tau(X^{k+1}))$ of the optimal threshold to use between $\tau(X^1)$ and $\tau(X^2)$ is the same for all values of X_t. Taking $(X_t, M_t) = (X^1, 0)$ and fixing the continuation value $F_1(X^1)$ at s such that

[40] If $U < 0$, then R would be strictly better off rejecting at $\tau(X^{k+1})$; because A's continuation value at $\tau(X^{k+1})$ is zero, rejecting at $\tau(X^{k+1})$ would provide the same continuation value to A.

$(X_s, M_s) = (X^1, m)$, we have

$$B_N(X^1) = \arg\max_B \ \mathbb{E}_{0,m}[e^{-r\tau_+(B;X^1)}(u(B,1) - (\widehat{\gamma} + \widehat{\lambda}^1)v(B,1)) + e^{-r\tau(X^1;B)}F_1(X^1)].$$

Using $K(X^1) = F_1(X^1) + \widehat{\lambda}^1 \frac{c_A}{r}$, optimality of $B_N(0)$ and $B_N(X^1)$ implies

$$\begin{aligned}
&\mathbb{E}[e^{-r\tau_+(B_N(0);X^1)}](u(B_N(0),1) - \widehat{\gamma}v(B_N(0),1)) + \mathbb{E}[e^{-r\tau(X^1;B_N(0))}](F_1(X^1) + \widehat{\lambda}^1\frac{c_A}{r}) \\
&\geq \mathbb{E}[e^{-r\tau_+(B_N(X^1);X^1)}](u(B_N(X^1),1) - \widehat{\gamma}v(B_N(X^1),1)) + \mathbb{E}[e^{-r\tau(X^1;B_N(X^1))}](F_1(X^1) + \widehat{\lambda}^1\frac{c_A}{r}), \\
&\mathbb{E}_{0,m}[e^{-r\tau_+(B_N(X^1);X^1)}](u(B_N(X^1),1) - (\widehat{\gamma} + \widehat{\lambda}^1)v(B_N(X^1),1)) + \mathbb{E}_{0,m}[e^{-r\tau(X^1;B_N(X^1))}]F_1(X^1) \\
&\geq \mathbb{E}_{0,m}[e^{-r\tau_+(B_N(0);X^1)}](u(B_N(0),1) - (\widehat{\gamma} + \widehat{\lambda}^1)v(B_N(0),1)) + \mathbb{E}_{0,m}[e^{-r\tau(X^1;B_N(0))}]F_1(X^1)
\end{aligned}$$

Note that $\mathbb{E}_{0,m}[e^{-r\tau_+(B;X^1)}] = \mathbb{E}[e^{-r\tau_+(B;X^1)}]$ for all B and similarly for $\tau(X^1;B)$. Adding the above inequalities together and simplifying, we get

$$\begin{aligned}
&\mathbb{E}[e^{-r\tau_+(B_N(X^1);X^1)}]v(B_N(X^1),1) + \mathbb{E}[e^{-r\tau(X^1;B_N(X^1))}]\frac{c_A}{r} \\
&\geq \mathbb{E}[e^{-r\tau_+(B_N(0);X^1)}]v(B_N(0),1) + \mathbb{E}[e^{-r\tau(X^1;B_N(0))}]\frac{c_A}{r},
\end{aligned}$$

which implies that $\widetilde{V}(B_N(X^1), X^1, 0) \geq \widetilde{V}(B_N(0), X^1, 0)$.

Let $U \geq 0$ be R's continuation value under $(\tau_N^*, d_{N,\tau}^*)$ at $\tau(X^1)$. R's ex-ante expected utility is $\check{J}(B_N(0), X^1, 0, U)$. Suppose $B_N(0) < \arg\max_B \ \widetilde{V}(B, X^1, 0)$. By the same arguments as in 7, $\arg\max_B \ \check{J}(B, X^1, 0, U) \geq \arg\max_B \ \widetilde{V}(B, X^1, 0)$, so R could increase $B_N(0)$ and make both players better off, a contradiction. Thus, $B_N(0) \geq \arg\max_B \ \widetilde{V}(B, X^1, 0)$. By single-peakedness of $\widetilde{V}$, $B_N(0) < B_N(X^1)$ implies that $\widetilde{V}(B_N(X^1), X^1, 0) < \widetilde{V}(B_N(0), X^1, 0)$, a contradiction. Therefore, $B_N(X^1) \leq B_N(0)$. □

The next Lemma shows A's continuation value is strictly positive at every history h_t with X_t above the last $RDP(X_n)$ constraint to be reached.

Lemma 10. *For $m \in (\max\{X^{k+1}, \underline{b}_N\}, X^k]$ and $x \in (X^k, B_N(X^k)]$, A's continuation value after any history h_t with $(X_t, M_t) = (x, m)$ is strictly positive.*

Proof. Take any history h_t with $(X_t, M_t) = (x, m)$ for any $x \in (X^k, B_N(X^k)]$ and $m \in (\max\{X^{k+1}, \underline{b}_N\}, X^k]$. Suppose $X^{k+1} \geq \underline{b}_N$. A's continuation value at h_t is

$\mathbb{E}_{x,m}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)] - \frac{c_A}{r}$, which is the same for all $m \in (X^{k+1}, X^k]$. Because $\rho(X^k) = \rho(X^{k+1}) = 0$, we have

$$\begin{aligned}
0 = \rho(X^k) &= \mathbb{E}^{X^k}\left[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)\right] - \frac{c_A}{r} \\
&= \mathbb{E}^{X^k}[e^{-r\tau_+(x;X^k)}\mathbb{E}_{x,M_{\tau_+(x)}}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)] + e^{-r\tau(X^{k+1};x)}\mathbb{E}^{X^{k+1}}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)]] - \frac{c_A}{r} \\
&= \mathbb{E}^{X^k}\left[e^{-r\tau_+(x;X^{k+1})}\mathbb{E}_{x,X^k}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)] + e^{-r\tau(X^{k+1};x)}\frac{c_A}{r}\right] - \frac{c_A}{r},
\end{aligned}$$

which implies $\mathbb{E}_{x,X^k}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)] - \frac{c_A}{r} > 0$. Because $\mathbb{E}_{x,m}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)] = \mathbb{E}_{x,X^k}[e^{-r\tau_N^*}v(X_{\tau_N^*}, d_{N,\tau}^*)]$, we conclude that A's continuation value at h_t is positive. The proof when $\underline{b}_N > X^{k+1}$ is analogous. □

We can now formally prove Theorem 1 .

Proof of Theorem 1. Let $\underline{b} := \lim_{N\to\infty} \underline{b}_N$ and define $\tau^* = \inf\{t : X_t \notin (\underline{b}, B(M_t; B^1))\}$ and $d_\tau^* = \mathbb{1}(X_{\tau^*} \geq B(M_{\tau^*}; B^1))$, which is the limit of $(\tau_N^*, d_{N,\tau}^*)$. Because each $(\tau_N^*, d_{N,\tau}^*)$ yields an upper-bound on our full problem, so does (τ^*, d_τ^*).[41] We only need to verify that (τ^*, d_τ^*) satisfies the dynamic participation constraint. By Lemma 10, in the limit as $N \to \infty$, we see that A's continuation value is weakly positive after every history h_t with $X_t \geq \lim_{N\to\infty} X_N = b_A^{FB}$. If $\underline{b} \geq b_A^{FB}$, then (τ^*, d_τ^*) satisfies the dynamic participation constraint and solves our full problem.

We now show $\underline{b} \geq b_A^{FB}$. If not, then $b_A^{FB} > -\infty$. We have $b_A^{FB} + \delta_N = X_N \in \mathcal{B}_N$ and so, when $\underline{b}_N < b_A^{FB}$, $\widetilde{V}(B_N(b_A^{FB} + \delta_N), \underline{b}_N, b_A^{FB} + \delta_N) = 0$. By definition of b_A^{FB}, it is optimal for A to quit at $\tau(b_A^{FB})$ for any approval threshold; so, for any B and $x \in (b_A^{FB}, B)$, $\widetilde{V}(B, b, x)$ is strictly decreasing in b for $b < b_A^{FB}$. Because $\underline{b}_N$ is bounded away from b_A^{FB}, for large N we have $\widetilde{V}(B, \underline{b}_N, b_A^{FB} + \delta_N) < 0$ for all B, a contradiction. Therefore, we must have $\underline{b} \geq b_A^{FB}$. □

We now show that $B(\underline{b}; B^1) \geq X_c$ when $V = 0$. If $B^1 < X_c$, then $\tilde{u}(B(m; B^1)) < 0$ for all $m < 0$, so R would be better off rejecting immediately. The approval threshold will never decrease below X_c, as R would be better off rejecting at $\tau(b^*(X_c))$. Because A's continuation value is equal to 0 at $\tau(m)$ for all $m \leq b^*(B^1)$, rejecting at $\tau(b^*(X_c))$ delivers the same expected utility for A. Therefore, $\underline{b} \geq b^*(X_c)$, which implies that $B(m; B^1) \geq X_c$ for all $m \geq \underline{b}$.

[41] Continuity of players' payoffs with respect to this sequence of mechanisms is shown in the Supplementary Materials.

C Asymmetric Information

We start with several auxiliary lemmas. Let $\widetilde{V}_i(B,b,x) = \widetilde{V}(B,b,x;z_i)$. Define $\underline{B}_{N,i}(m) := \{B : \widetilde{V}_i(B, m-\delta_N, m) = 0\}$, which is unique by Lemma 11 below. By Lemma 7, $\lim_{N\to\infty} \underline{B}_{N,i} = \underline{B}_i(m)$.

Lemma 11. *For any $a \geq 0$, $b < x$ and $i \in \{h,\ell\}$, $\widetilde{V}_i(B,b,x)$ is strictly decreasing in B on $[x,\infty)$.*

Lemma 12. $\underline{B}_\ell(m) < \underline{B}_h(m)$.

Lemma 13. *For some $\underline{b} < 0$ and increasing function $B(m) > \max\{\underline{B}_\ell(m), \underline{b}\}$, let $\tau' = \inf\{t : X_t \notin (\underline{b}, B(M_t))\}$ and $d'_\tau = \mathbb{1}(X_{\tau'} \geq B(M_{\tau'}))$. Then $V^*(\tau', d'_\tau, z_\ell) = 0$.*

Lemma 14. *Take R's problem in 1 and drop $DIC(h)$. If z_h is sufficiently high, the this relaxed problem satisfies $DIC(h)$.*

For the rest of the proof of Theorem 2, we will drop $DIC(h)$ and solve AM_ℓ and AM_h for an arbitrary value of z_h. Dropping $DIC(h)$ means we can drop $PK_h(V'_h)$ from AM_h and $DIC(h,V_h)$ from AM_ℓ.

Optimal Mechanism for ℓ

When $DIC(h,V_h)$ is dropped, AM_ℓ corresponds to the SI-problem when $z_0 = z_\ell$ with a PK constraint. The structure of (τ^ℓ, d^ℓ_τ) follows from Theorem 1.

Optimal Mechanism for h

Unless otherwise specified, expectations below are taken with respect to $Z_0 = z_h$. We let $z_i(x) = z_i + \frac{2\mu}{\sigma^2}x$, $\tilde{v}_i(x) = \frac{e^{z_i(x)}+a}{1+e^{z_i(x)}}$, $v_i(x,d) = \tilde{v}_i(x)d + \frac{c_A}{r}$ and $\underline{b}^S$ and $B^S(m)$ be the rejection and approval thresholds, respectively, in the SI-problem when $z_0 = z_h$ and $V = 0$. If $z_h = \infty$, we take $\underline{b}^S = -\infty$ and $B^S(m) = -\infty$ for all m.

Lemma 15. *It is never optimal in AM_h to reject at any history h_t with $X_t > \underline{b}^S$.*

Proof. As argued in the proof of Theorem 1, $B^S(\underline{b}^S) \geq X_c$ (with X_c defined relative to $z_0 = z_h$). Take any (τ, d_τ) that satisfies all constraints in AM_h and rejects at $X_\tau > \underline{b}^S$ with positive probability. Let $\tau' = \inf\{t : X_t \notin (\underline{b}^S, \underline{B}_h(M_t))\}$ and

$d'_\tau = \mathbb{1}(X_{\tau'} \geq \underline{B}_h(M_{\tau'}))$. Suppose R used (τ', d'_τ) as the continuation mechanism of (τ, d_τ) rather than rejecting at $X_\tau > \underline{b}^S$. Because $\underline{B}_h(m)$ is increasing in m and $\underline{B}_h(M_{\tau'}) \geq \underline{B}_h(\underline{b}^S) \geq X_c$, R's continuation value at h_s from the continuation mechanism (τ', d'_τ) is strictly positive, as she is approving with positive probability in the future at $X_{\tau'} > X_c$ and never approves at $X_{\tau'} < X_c$. It is easily verified that h's continuation value under (τ', d'_τ) is 0, the same as rejection, and (τ', d'_τ) satisfies h's dynamic participation constraint. Because $\underline{B}_\ell(m) < \underline{B}_h(m)$, by Lemma 13 ℓ will optimally choose to quit immediately under (τ', d'_τ). If (τ, d_τ) rejects at h_s with $X_s > \underline{b}^S$, both types of A receive the same continuation value at h_s from (τ', d'_τ) as under rejection, and, hence, using (τ', d'_τ) as the continuation mechanism at h_s does not change either type's continuation value prior to s. Thus, using (τ', d'_τ) represents an improvement over rejecting. □

Let $\bar{\tau} = \inf\{t : \ X_t \notin (\underline{b}^S, B^S(M_t))\}$ and $\bar{d}_\tau = \mathbb{1}(X_{\bar{\tau}} \geq B^S(M_{\bar{\tau}}))$; then, $(\bar{\tau}, \bar{d}_\tau)$ solves our SI-problem when $z_0 = z_h$. Define $\underline{m}$ as the $m \leq 0$ such that $B^S(m) = \underline{B}_\ell(m)$; if $B^S(m) > \underline{B}_\ell(m)$ $\forall m \leq 0$, take $\underline{m} = 0$ and if $z_h = \infty$, take $\underline{m} = -\infty$. Because $\underline{B}_\ell(m) < \underline{B}_h(m)$, $B^S(m') > \underline{B}_\ell(m')$ for all $m' < \underline{m}$. By Lemma 13, for any mechanism (τ, d_τ) that uses $(\bar{\tau}, \bar{d}_\tau)$ as its continuation mechanism at $\tau(\underline{m})$, if ℓ has not already quit by $\tau(\underline{m})$, then ℓ will choose to quit immediately at $\tau(\underline{m})$.

In any solution (τ', d'_τ) to AM_h, the continuation mechanism of (τ', d'_τ) at $h_{\tau(\underline{m})}$ will deliver to each type i some continuation value $W_i(h_{\tau(\underline{m})})$. The continuation mechanism will solve the problem of maximizing R's continuation value at $h_{\tau(\underline{m})}$ subject h's dynamic participation constraint and constraints ensuring that h receives a continuation value of at least $W_h(h_{\tau(\underline{m})})$ and ℓ receives a continuation value of at most $W_\ell(h_{\tau(\underline{m})})$. Lowering $W_h(h_{\tau(\underline{m})})$ to 0 and raising $W_\ell(h_{\tau(\underline{m})})$ to ∞ increases R's value of this problem and this new problem is equivalent to maximizing R's value subject to h's dynamic participation constraint, which is equal to our SI-problem when $(X_0, Z_0) = (\underline{m}, z_h(\underline{m}))$. By Proposition 1, the solution to the SI-problem is $(\bar{\tau}, \bar{d}_\tau)$.[42] Therefore, $(\bar{\tau}, \bar{d}_\tau)$ generates an upper-bound on R's continuation value at $\tau(\underline{m})$ under any solution to AM_h.

In the SI-problem when $z_0 = z_h$ and $V = 0$, by the same arguments as in our example in Section 4, R will never reject while the approval threshold is above

[42] We discuss this point in more detail after the proof of Proposition 1 in the Supplementary Materials, in which we argue that taking $(X_0, Z_0) = (x, z_0 + \frac{2\mu}{\sigma^2}x)$, the optimal thresholds (as a function of (X, M)) are independent of x.

X_c. She will also never set an initial approval threshold below X_c, so the approval threshold must decrease for low enough m; for such m, $B^S(m) = \underline{B}_h(m)$. Thus, $\underline{m} > \underline{b}^S$ if $\underline{b}^S > -\infty$. By Lemma 15, we can set $d_\tau = 1$ for all $\tau < \tau(\underline{m})$ without loss.

Let $\overline{U} = J(\bar{\tau}, \bar{d}_\tau, \underline{m}, z_h(\underline{m}))$, which is an upper-bound on R's continuation value in any solution to AM_h at $\tau(\underline{m})$. We can find an upper-bound on R's ex-ante value of AM_h by solving a relaxed program in which we drop h's dynamic participation constraint while fixing, for any choice of mechanism (τ, d_τ), the continuation mechanism of (τ, d_τ) at $\tau(\underline{m})$ to be $(\bar{\tau}, \bar{d}_\tau)$, set $d_\tau = 1$ for $\tau < \tau(\underline{m})$ and replace DIC with a finite number of $RDIC$ constraints:

$$\begin{aligned} H_h = \sup_{\tau} \ & \mathbb{E}[e^{-r(\tau\wedge\tau(\underline{m}))}\big(u(X_\tau,1)\mathbb{1}(\tau<\tau(\underline{m})) + \overline{U}\mathbb{1}(\tau\geq\tau(\underline{m}))\big)] \\ & \text{subject to } \forall X_n \in \mathcal{T}_{\mathcal{N}}: \ X_n > \underline{m}, \\ & RDIC(X_n): \mathbb{E}^{0,z_\ell}[e^{-r(\tau\wedge\tau(X_n))}v_\ell(X_{\tau\wedge\tau(X_n)}, \mathbb{1}(\tau<\tau(X_n)))] \leq V_\ell + \frac{c_A}{r}. \end{aligned} \tag{8}$$

If $\underline{m} = 0$, we are done. Let us therefore assume $\underline{m} < 0$. We now change the expectation in the $RDIC$ constraints from being taken with respect to z_ℓ to z_h. Using $z_\ell = z_h - \Delta_z$, we have

$$\begin{aligned} & \mathbb{E}^{0,z_\ell}[e^{-r(\tau\wedge\tau(X_n))}v(X_{\tau\wedge\tau(X_n)}, d_\tau(X_n))] \\ & = \frac{e^{z_\ell}}{1+e^{z_\ell}}\mathbb{E}[e^{-r(\tau\wedge\tau(X_n))}(d_\tau(X_n)+\frac{c_A}{r})|H] + \frac{1}{1+e^{z_\ell}}\mathbb{E}[e^{-r(\tau\wedge\tau(X_n))}(ad_\tau(X_n)+\frac{c_A}{r})|L] \\ & = \frac{1+e^{z_h}}{1+e^{z_\ell}}\Bigg(\frac{e^{z_h}}{1+e^{z_h}}\mathbb{E}[e^{-r(\tau\wedge\tau(X_n))}e^{-\Delta_z}(d_\tau(X_n)+\frac{c_A}{r})|H] \\ & + \frac{1}{1+e^{z_h}}\mathbb{E}[e^{-r(\tau\wedge\tau(X_n))}(ad_\tau(X_n)+\frac{c_A}{r})|L]\Bigg) \\ & = \mathbb{E}^{0,z_h}[e^{-r(\tau\wedge\tau(X_n))}\hat{v}_\ell(X_{\tau\wedge\tau(X_n)}, d_\tau(X_n))], \end{aligned} \tag{9}$$

where $\hat{v}_\ell(x,d) := \frac{1+e^{z_h}}{1+e^{z_\ell}} \cdot \frac{e^{z_h(x)-\Delta_z}(d+\frac{c_A}{r})+ad+\frac{c_A}{r}}{1+e^{z_h(x)}}$.

Let $N' = |\{X_n \in \mathcal{T}_N : \ X_n > \underline{m}\}|$ and, for $\Lambda = (\lambda_1, ..., \lambda_{N'})$, define $\mathcal{L}(\tau, \Lambda)$ as

$$\begin{aligned} \mathcal{L}(\tau,\Lambda) = \mathbb{E}[e^{-r(\tau\wedge\tau(\underline{m}))}\big(u(X_\tau,1)\mathbb{1}(\tau<\tau(\underline{m})) + \overline{U}\mathbb{1}(\tau\geq\tau(\underline{m}))\big) \\ + \sum_{n=1}^{N'}\lambda_n e^{-r(\tau\wedge\tau(X_n))}\hat{v}_\ell(X_\tau, \mathbb{1}(\tau<\tau(X_n)))] - \sum_{n=1}^{N'}\lambda_n(V_\ell + \frac{c_A}{r}) \end{aligned} \tag{10}$$

We now redefine $\mathcal{L}^*(\Lambda) = \sup_\tau \mathcal{L}(\tau, \Lambda)$. By Theorem 1 of Balzer and Janßen (2002), $H_h = \mathcal{L}^*(\widehat{\Lambda})$ for some $\widehat{\Lambda} \in \arg\min_{\Lambda \in \mathbb{R}^{N'}_-} \mathcal{L}^*(\Lambda)$. Our next result uses this Lagrangian approach to show the existence of a solution to H_h.

Proposition 6. *There exists a τ_N^* that solves H_h. Let $\mathcal{B}_N = \{X^1, ..., X^P\}$ be the set of X_n such that $RDIC(X_n)$ binds under τ_N^*. There exists a function $B_N(m)$ that is constant on $(X^{k+1}, X^k]$ for each $0 \leq k \leq P$ such that $\tau_N^* = \inf\{t : X_t \geq B_N(M_t)\}$.*

Let $d^*_{N,\tau}(X_n) = \mathbb{1}(\tau_N^* < \tau(X_n))$. We can use the stationarity of τ_N^* to pin down ℓ's continuation value at $\tau(X_n)$ under $(\tau_N^*, 1)$ when ℓ plans to quit at $\tau(X_n - \delta_N)$; this continuation value is given by $\rho_\ell(X_n)$ with

$$\rho_\ell(X_n) = \mathbb{E}^{X_n, z_\ell(X_n)}[e^{-r\tau_+(B_N(X_n); X_n - \delta_N)} v_\ell(B_N(X_n), 1) + e^{-r\tau(X_n - \delta_N; B_N(X_n))} \frac{c_A}{r}] - \frac{c_A}{r}.$$

Lemma 16. *$\rho_\ell(X^k) \leq 0$, with equality if and only if $RDIC(X^k - \delta_N)$ binds. $\rho_\ell(X^k + \delta_N) \geq 0$, with equality if and only if $RDIC(X^k + \delta_N)$ binds.*

Proof. $RDIC(X^k - \delta_N)$ implies

$$\begin{aligned}
V_\ell + \frac{c_A}{r} &\geq \mathbb{E}^{0, z_\ell}[e^{-r(\tau_N^* \wedge \tau(X^k - \delta_N))} v_\ell(X_{\tau_N^* \wedge \tau(X^k - \delta_N)}, d^*_{N,\tau}(X^k - \delta_N))] \\
&= \mathbb{E}^{0, z_\ell}[e^{-r\tau_N^*} \mathbb{1}(\tau_N^* < \tau(X^k)) v_\ell(X_{\tau_N^*}, 1) + \big(e^{-r\tau(X^k)} \mathbb{1}(\tau_N^* \geq \tau(X^k)) \\
&\quad \cdot \mathbb{E}^{X^k, z_\ell(X^k)}[e^{-r\tau_+(B_N(X^k); X^k - \delta_N)} v_\ell(B_N(X^k), 1) + e^{-r\tau(X^k - \delta_N; B_N(X^k))} \frac{c_A}{r}]\big)] \\
&= \mathbb{E}^{0, z_\ell}[e^{-r\tau_N^*} \mathbb{1}(\tau_N^* < \tau(X^k)) v_\ell(X_{\tau_N^*}, 1) + e^{-r\tau(X^k)} \mathbb{1}(\tau_N^* \geq \tau(X^k))(\rho_\ell(X^k) + \frac{c_A}{r})],
\end{aligned}$$

with equality if and only if $RDIC(X^k - \delta_N)$ binds. Similarly, $RDIC(X^k)$ implies

$$V_\ell + \frac{c_A}{r} = \mathbb{E}^{0, z_\ell}[e^{-r\tau_N^*} \mathbb{1}(\tau_N^* < \tau(X^k)) v_\ell(X_{\tau_N^*}, 1) + e^{-r\tau(X^k)} \mathbb{1}(\tau_N^* \geq \tau(X^k)) \frac{c_A}{r}].$$

Putting these equations together and simplifying, we get $\rho_\ell(X^k) \leq 0$, with equality if and only if $RDIC(X^k - \delta_N)$ binds.

$RDIC(X^k + \delta_N)$ implies

$$V_\ell + \frac{c_A}{r} \geq \mathbb{E}^{0, z_\ell}[e^{-r\tau_N^*} \mathbb{1}(\tau_N^* < \tau(X^k + \delta_N)) v_\ell(X_{\tau_N^*}, 1) + e^{-r\tau(X^k + \delta_N)} \mathbb{1}(\tau_N^* \geq \tau(X^k + \delta_N)) \frac{c_A}{r}],$$

with equality if and only if $RDIC(X^k + \delta_N)$ binds. $RDIC(X^k)$ implies

$$V_\ell + \frac{c_A}{r} = \mathbb{E}^{0,z_\ell}[e^{-r\tau_N^*}\mathbb{1}(\tau_N^* < \tau(X^k + \delta_N))v_\ell(X_{\tau_N^*}, 1) + e^{-r\tau(X^k+\delta_N)}\mathbb{1}(\tau_N^* \geq \tau(X^k + \delta_N))(\rho_\ell(X^k + \delta_N) + \frac{c_A}{r})].$$

Putting these last two equations together and simplifying, we get $\rho(X^k + \delta_N) \geq 0$ with equality if and only if $RDIC(X^k + \delta_N)$ binds. □

Let $\overline{X}_\ell^N = \max\{X_n \in \mathcal{B}_N\}$ and $\underline{X}_\ell^N = \min\{X_n \in \mathcal{B}_N\}$. By the choice of X_n in our constraint set, $\underline{X}_\ell^N > \underline{m}$. Our next lemmas characterize the set of binding constraints and give properties of the approval threshold. The proof is similar to the characterization of $\mathcal{B}_N$ in the proof of Theorem 1.

Lemma 17. $\mathcal{B}_N = \{X_n \in \mathcal{T}_N : \underline{X}_\ell^N \leq X_n \leq \overline{X}_\ell^N\}$.

Lemma 18. $B_N(X^k) = \underline{B}_{N,\ell}(X^k)$ *for all* $X^k > \underline{X}_\ell^N$.

Proof. The Lemma follows immediately from $\rho(X^k) = \widetilde{V}_\ell(B_N(X^k), X^k - \delta_N, X^k) = 0$ by Lemmas 16 and 17. □

Lemma 19. $B_N(\underline{X}_\ell^N) = B^S(\underline{m})$ *and* $\lim_{N\to\infty} \underline{X}_\ell^N = \underline{m}$.

Proof. We first look at $B^S(0)$ in the SI-problem. $B^S(\underline{m}) = \underline{B}_\ell(\underline{m}) < \underline{B}_h(\underline{m})$, so $\underline{m} \geq b_h^*(B^S(0))$, the first point at which h is indifferent between continuing and quitting. Thus, $B^S(m)$ must be constant for all $m > \underline{m}$, which implies $B^S(0) = B^S(\underline{m})$. Remember that, in the SI-problem, R's continuation value under $(\bar{\tau}, \bar{d}_\tau)$ at $\tau(\underline{m})$ is $\overline{U}$. R's continuation value under $(\bar{\tau}, \bar{d}_\tau)$ at $\tau(\underline{X}_\ell^N)$ is $\check{J}(B^S(0), \underline{m}, \underline{X}_\ell^N, \overline{U})$. Because h's continuation value is strictly positive at all $t < \tau(\underline{m})$, R could change the approval threshold in the SI-problem slightly for $t < \tau(\underline{m})$ (and revert back to $B^S(M_t)$ at $t \geq \tau(\underline{m})$) without violating the dynamic participation constraint. Because such a modification is not optimal, $B^S(0) = \arg\max_B \check{J}(B, \underline{m}, \underline{X}_\ell^N, \overline{U})$.

We now return to the problem H_h. At $\tau(\underline{X}_\ell^N)$, all $\hat{v}_\ell$ terms have dropped out of our Lagrangian in 10 and the continuation value at $\tau(\underline{X}_\ell^N)$ is $\check{J}(B_N(\underline{X}_\ell^N), \underline{m}, \overline{X}_\ell^N, \overline{U})$. Optimality of B_N implies $B_N(\underline{X}_\ell^N) = \arg\max_B \check{J}(B, \underline{m}, \underline{X}_\ell^N, \overline{U}) = B^S(\underline{m})$.

Let $\underline{x}_\ell = \lim_{N\to\infty} \underline{X}_\ell^N$ and suppose $\underline{x}_\ell > \underline{m}$. For large N, $\underline{X}_\ell^N - \delta_N > \underline{m}$ and $\rho_\ell(\underline{X}_\ell^N) = \widetilde{V}_\ell(B_N(\underline{X}_\ell^N), \underline{X}_\ell^N - \delta_N, \underline{X}_\ell^N) < 0$. Thus, $B_N(\underline{X}_\ell^N) > \underline{B}_{N,\ell}(\underline{X}_\ell^N)$ and

$$B^S(\underline{m}) = B_N(\underline{X}_\ell^N) > \underline{B}_{N,\ell}(\underline{X}_\ell^N) > \underline{B}_\ell(\underline{m}),$$

where the last inequality follows from $\lim_{N\to\infty} \underline{B}_{N,\ell}(\underline{X}_\ell^N) = \underline{B}_\ell(\underline{x}_\ell) > \underline{B}_\ell(\underline{m})$, a contradiction of the $B^S(\underline{m}) = \underline{B}_\ell(\underline{m})$. Therefore, $\lim_{N\to\infty} \underline{X}_\ell^N = \underline{m}$. □

Let $b_h^1 := \lim_{N\to\infty} \overline{X}_\ell^N$ and define a function B^h as $B^h(m) = \lim_{N\to\infty} B_N(m)$ for $m > \underline{m}$ and $B^h(m) = B^S(m)$ for $m \leq \underline{m}$. Because the approval threshold $B_N(m)$ is constant above $\overline{X}_\ell^N$, $B^h(m)$ is constant for all $m > b_h^1$. Because $\underline{B}_\ell(m) = \lim_{N\to\infty} \underline{B}_{N,\ell}(m)$ and $B_N(X^k) = \underline{B}_{N,\ell}(X^k)$, $B^h(m) = \underline{B}_\ell(m)$ for all $m \in (\underline{m}, b_h^1]$. Taking $\underline{b}_h = \underline{b}^S$, $B_h^1 = B^h(0)$, and $B_h^2 = B^S(\underline{m})$, the limit of the solutions to our relaxed problems as $N \to \infty$ takes the form $\tau^h = \inf\{t : X_t \notin (\underline{b}_h, B^h(M_t; \eta_h))\}$ and $d_\tau^h = \mathbb{1}(X_{\tau^h} \geq B^h(M_{\tau^h}; \eta_h))$ with $\eta_h = (\underline{b}_h, b_h^1, B_h^2, B_h^1)$. We now show that (τ^h, d_τ^h) solves AM_h.

Proof. Because the solution to the relaxed problem for each N delivers an upper-bound on R's expected utility in AM_h and satisfies all $RDIC$ constraints, it is easy see that (τ^h, d_τ^h) delivers an upper-bound as well and ℓ's expected utility from (τ^h, d_τ^h) when quitting at $\tau(\underline{m})$ is weakly less than V_ℓ.[43] We only need to show that (τ^h, d_τ^h) satisfies the constraints in AM_h. The same arguments as in Lemma 10 imply ℓ's continuation value when planning to quit at $\tau(\underline{m})$ is weakly positive at $t < \tau(\underline{m})$. Therefore, ℓ cannot strictly increase his expected utility by quitting before $\tau(\underline{m})$. By Lemma 13, ℓ would prefer to immediately quit at $\tau(\underline{m})$. Therefore, (τ^h, d_τ^h) satisfies DIC. It is easy to see that h's continuation value is positive at $t < \tau(\underline{m})$. Because the continuation mechanism at $\tau(\underline{m})$ is the SI-mechanism with $X_0 = \underline{m}$ and $Z_0 = z_h + \frac{2\mu}{\sigma^2}\underline{m}$, Theorem 1 implies h's continuation value is always weakly positive at $t \geq \tau(\underline{m})$ as well and so the dynamic participation constraint for h holds. □

[43] Continuity of players' payoffs with respect to this sequence of mechanisms follows from the same arguments in the Supplementary Materials used for Theorem 1.