

Dynamic Screening under Dual Uncertainty

Takakazu Honryo[‡]

October 5, 2026

Working Paper Series No. 69
Faculty of Economics, Doshisha University

Abstract

This paper studies dynamic screening under two layers of uncertainty: the agent privately knows his type, which is informative about the project's quality, yet he may remain uncertain about the project's underlying quality. The decision maker learns from both the agent's willingness to remain under review and evidence arriving over time. Equilibrium screening can proceed in two stages: initial attrition, during which approval may occur without favorable evidence, followed by experimentation, during which approval is withheld until favorable evidence arrives. Deadlines without commitment reduce the decision maker's payoff, whereas deadlines with commitment can improve it by screening out the low type immediately and moving directly to experimentation.

1 Introduction

A decision maker evaluating a project or candidate often faces two sources of uncertainty: the applicant possesses private information, yet may not know the underlying quality with certainty. The approval of new drugs provides a concrete example. A pharmaceutical firm may know more than the regulator about its drug's development and the evidence gathered so far, but promising trial results need not establish the drug's ultimate clinical benefit.¹ Some drugs receive accelerated approval even before their clinical benefit is confirmed,

^{*}Department of Economics, Doshisha University, Email: thonryo@mail.doshisha.ac.jp

[‡]The author is grateful to seminar participants at various venues for their comments, and also to Yusuke Kambe, Takashi Shimizu, and Narumi Teshima for pointing out errors in an earlier version of the paper. This work was supported by JSPS KAKENHI Grant Numbers 25K00623 and 25K05012.

¹The scale of private investment underlying this informational asymmetry is substantial. DiMasi, Grabowski, and Hansen (2016) estimate the average capitalized cost of bringing a new compound to market at approximately \$2.6 billion, an expenditure that the regulator cannot directly verify.

with further trials required to verify that benefit.² Similar problems arise when an investor evaluates a startup, an employer assesses a worker during a probationary period, or a regulator certifies a new technology.

As the review proceeds, these two sources of uncertainty give rise to two learning processes that can pull the decision maker’s beliefs in opposite directions. An applicant’s willingness to remain under review can be good news about his private information, since less confident applicants may withdraw. At the same time, a prolonged failure to produce favorable evidence is bad news about the underlying project. A review process must therefore balance information revealed by who remains against information revealed by what evidence has failed to arrive.

We develop a continuous-time screening model in which an agent submits a project to a decision maker (DM). The model features two layers of uncertainty. First, the agent has private information about the project’s quality: a low type knows that the project is bad, whereas a high type regards it as likely to be good. Second, even the high type is not certain that the project is good, and this residual uncertainty can be resolved only through evidence generated during the review. If the project is genuinely good, favorable evidence arrives according to a Poisson process; if it is bad, no favorable evidence ever arrives. Continuing the review is costly to both parties, and at any time the agent may withdraw, while the DM may terminate the review and make an approval decision.

Our main result is that dynamic screening can exhibit a two-stage equilibrium structure. In the first stage, screening operates through attrition. The low-type agent gradually drops out, while the high type continues. Early in this process, the DM’s belief remains sufficiently favorable for her to approve the project even in the absence of favorable evidence. This possibility of approval sustains the low type’s willingness to remain in the process and thereby allows gradual type screening to occur. The attrition phase thus offers a possible explanation for approval before clinical benefit is confirmed: screening through the exit of low-type applicants can support approval even before conclusive evidence arrives.

At the same time, however, the continued absence of favorable evidence is bad news about the project’s underlying quality. Although low-type exit increases the likelihood that the agent who remains is the high type, the informational effect of signal absence dominates, so that the DM’s posterior declines over time. As the review continues without evidence, the DM therefore becomes increasingly reluctant to approve without a signal. Eventually, the low type is fully exhausted and the equilibrium enters a second stage. From that point onward, approval without favorable evidence no longer occurs: the DM and the remaining high type continue only in the hope that favorable evidence arrives. Screening has thus endogenously shifted from learning about the agent’s private information to learning about the project’s underlying quality.

Depending on parameters, either stage may disappear. When the high type’s residual uncertainty about the project’s quality is sufficiently large, pure experimentation can

²For example, Aduhelm received FDA accelerated approval in 2021 based on its effect on amyloid plaques, with a post-approval trial required to verify clinical benefit. Leqembi likewise received accelerated approval in January 2023 and was converted to traditional approval in July 2023 after the FDA verified its clinical benefit.

arise: the low type exits immediately and screening proceeds solely through stochastic evidence. In this case, type screening is completed at the outset and the review is devoted entirely to learning about project quality. When the high type’s continuation cost is high relative to the rate at which favorable evidence arrives, the experimentation stage cannot be sustained and screening operates solely through attrition. When both stages are sustained, the equilibrium takes the form of sequential screening, with attrition followed by experimentation. When multiple equilibrium forms coexist, the DM has a clear payoff ranking among them: pure experimentation is preferred to sequential screening, which is preferred to pure attrition. Thus, the DM-optimal equilibrium is the most experimentation-intensive form that can be sustained.

The model also yields comparative-statics implications. A higher arrival rate of favorable evidence makes experimentation more attractive and moves the equilibrium toward an earlier transition to experimentation. A higher continuation cost for the high type has the opposite effect by shortening the horizon over which experimentation can be sustained. The high type’s cost therefore affects the DM even though she does not bear it directly: when the agent becomes less willing to continue, the DM loses the informational benefit of waiting for evidence.

A natural regulatory design question that follows from these equilibrium dynamics concerns the role of deadlines. Real-world regulatory institutions routinely operate under explicit time constraints. In the United States, the Prescription Drug User Fee Act (PDUFA) establishes target review goals, typically ten months for standard applications and six months for priority reviews (counted from the 60-day filing date for new molecular entities), which the FDA commits to meeting for a specified share of applications.³ The European Medicines Agency similarly operates under strict procedural timelines.⁴ Notably, these timelines take the form of fixed target dates by which a decision is expected, rather than stochastic mechanisms. These deadlines are motivated by concerns over delayed patient access to beneficial therapies and administrative efficiency, but their informational consequences have received comparatively little formal analysis.

The two-stage structure also clarifies the role of deadlines and commitment. Merely shortening the review horizon is detrimental because it truncates screening that would otherwise occur. Commitment to reject in the absence of favorable evidence has a different effect. When the unconstrained equilibrium contains an attrition stage followed by experimentation, commitment can induce the low type to exit immediately, completing type screening at the outset and moving the process directly to experimentation while preserving the same terminal date. Commitment then strictly benefits the DM. If the equilibrium already consists of pure experimentation, commitment has no value.

When only attrition can be sustained, the effect is ambiguous: commitment screens out the low type immediately but also shortens the high type’s participation by removing the prospect of approval without favorable evidence. When the DM’s delay cost is sufficiently small, the latter effect dominates, making commitment harmful.

³U.S. Food and Drug Administration 2022, PDUFA Reauthorization Performance Goals and Procedures, Fiscal Years 2023-2027.

⁴Under the centralized procedure, standard applications are assessed within 210 active days and accelerated applications within 150 active days. See European Medicines Agency (2019).

The equilibrium analysis also suggests a role for subsidies. When experimentation is constrained by the high type’s continuation cost, reducing that cost can extend the period over which evidence is generated. A subsidy is therefore most useful when the DM would like experimentation to continue substantially longer than the high type would. By contrast, in a pure attrition equilibrium the type-H agent’s participation depends on his cost net of the prospect of approval without evidence, which a subsidy paid to both types leaves unchanged. A subsidy therefore cannot extend the review and only raises the DM’s cost of delay.

Our model can be viewed as a variant of the war of attrition. Representative examples in the literature include Kreps and Wilson (1982) and Ordover and Rubinstein (1986), who study attrition games with asymmetric information in which players employ mixed strategies to resolve dynamic strategic tension. A key difference from the standard war of attrition framework is that payoffs in our model are not discounted over time. Instead, players incur costs in order to continue. This modeling choice has two implications. First, even a firm that ultimately succeeds may incur an ex-post loss. Second, even if the DM can commit to a course of action at the beginning of the game, such commitment may be disadvantageous unless it is implemented probabilistically.

The paper also relates to the strategic experimentation literature with Poisson bandits, initiated by Keller, Rady, and Cripps (2005). Bobtcheff and Levy (2017) study signaling through investment timing; unlike their model, in which information is conveyed solely through the agent’s timing choice, our model features stochastic evidence that arrives independently of the agent’s strategy. Bonatti and Hörner (2017) and Halac and Kremer (2020) study effort incentives under Poisson learning, both featuring moral hazard as the key friction; in our model, by contrast, the central friction is adverse selection, with the agent’s type being private information. Chen, Ishida, and Suen (2021) study reputational concerns in risky experimentation with a Poisson signal structure. A key difference is that in their model, the agent strategically chooses when to abandon a project to signal his ability to the market, whereas in our model, the agent’s dropout decision is driven solely by the prospect of approval rather than reputational concerns. Dong (2025) studies strategic experimentation with asymmetric information between two symmetric players; while sharing the Poisson signal structure, it focuses on mutual encouragement rather than the screening of agent types by a principal.

Our model is also related to the broad literature on dynamic signaling. Daley and Green (2012) analyze a dynamic lemons market in which a seller waits for exogenous news to arrive before trading; in their framework, the exit of low-quality sellers causes the decision maker’s belief to rise over time, as the surviving pool becomes progressively better. Bar-Isaac (2003), Lee and Liu (2013), and Honryo (2018) share this feature: in each of these models, the endogenous exit or mixing of low types drives beliefs upward over time. Our model generates a qualitatively different belief dynamic: because signal absence is bad news about the project’s true quality, the regulator’s belief declines even as low types continue to exit. This simultaneous occurrence of attrition and belief deterioration is a distinctive feature of our equilibrium, and reflects the dominant role of the Poisson signal technology in shaping the regulator’s posterior.

Approval processes have also been studied theoretically by Henry and Ottaviani (2019),

who model the approval process as a persuasion problem in which an informer sequentially conducts research to influence an evaluator’s decision. Our paper differs in focusing on adverse selection rather than the organization of control rights. McClellan (2022) studies how a principal who controls approval should design an approval rule when an agent controls costly experimentation, and Guo (2016) studies the optimal delegation of experimentation to a privately informed agent who is biased toward continuing. In both papers, the principal can commit to a mechanism. In our model, the DM cannot commit, and screening arises in equilibrium from the interaction between the agent’s private information and the arrival of evidence; Section 4 examines how commitment changes this conclusion. Finally, our use of ‘sequential screening’ differs from the mechanism-design literature following Courty and Li (2000): there, the agent’s information arrives sequentially and the principal designs a contract, whereas here screening itself proceeds in two sequential stages in equilibrium, without commitment.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 characterizes the dynamic screening equilibrium and derives comparative statics. Section 4 studies the role of commitment and other policy implications, and Section 5 concludes. Appendix A contains the formal definition of perfect Bayesian equilibrium and Appendix B collects all proofs.

2 Model

Time is continuous and indexed by $t \in [0, \infty] = T$. There are two players: an agent and a decision maker (DM). The agent submits a project (e.g., a newly developed drug) to the DM, who must eventually decide whether to approve or disapprove it. The project has an underlying true quality $\omega \in \{G, B\}$, where G represents a genuinely good quality project and B a bad quality one. The prior probability that the project is of good quality is $p_0 \in (0, 1)$. If the project is approved, the agent receives payoff 1. The DM receives payoff 1 if $\omega = G$ and payoff -1 if $\omega = B$. If the project is disapproved, both players receive 0.

The agent does not directly observe the true quality ω . Instead, the agent can be one of two informational types, indexed by $j \in \{L, H\}$, which is his private information and is not observed by the DM. A type- H agent assigns probability $1 - \epsilon$ to the event that the project is of good quality, $\Pr[\omega = G|H] = 1 - \epsilon$, and probability ϵ to the event that it is of bad quality, $\Pr[\omega = B|H] = \epsilon$. Thus ϵ represents the residual uncertainty of the H -type, or the difficulty in properly assessing the project’s potential. For simplicity, we assume that a type- L agent knows that $\omega = B$; the type- H agent is more optimistic about the project’s quality than the type- L agent. This heterogeneity captures differences in technological maturity. Consistency with the common prior then implies that an agent is H -type with probability $\frac{p_0}{1-\epsilon}$. We assume $p_0 < 1 - \epsilon$, so that both types arise with positive probability.

Before the DM makes a decision, the players may jointly exert effort (e.g., running trials, conducting analyses, or submitting data for the agent and reviewing reports, auditing results, or maintaining oversight for the DM) to generate public evidence about the

project quality. If the true quality is high ($\omega = G$), a favorable signal arrives according to a Poisson process with intensity $\lambda > 0$. If the true quality is low ($\omega = B$), no signal ever arrives. As the signal is perfectly informative, the DM learns that $\omega = G$ with certainty and immediately approves the project upon observing it. Exerting effort incurs a constant flow cost c_j for $j \in \{DM, H, L\}$. For simplicity we set the discount rate to zero. If the game ends at time τ , the total cost incurred by player i is τc_i . At any moment, either player may unilaterally terminate the game. If the DM stops at some point t , she immediately decides whether to approve or disapprove. If the agent stops, the game terminates and the agent's project is rejected.⁵

As long as neither player stops the game, the DM updates her belief continuously based on the absence of a favorable signal. Let p_t denote this continuation belief that $\omega = G$ at time t ; it evolves deterministically according to the no-signal Bayes rule consistent with the agent's strategy, and it summarizes the DM's information along all histories in which the game remains active.

The type- H agent also updates his own belief based on the absence of a signal. Let r_t denote his posterior belief at time t that $\omega = G$, conditional on no signal arrival up to t . Bayesian updating implies

$$r_t = \frac{(1 - \epsilon)e^{-\lambda t}}{\epsilon + (1 - \epsilon)e^{-\lambda t}}, \quad (1)$$

which is fully determined by time.

At time 0, before the review process begins, the agent first decides whether to enter. If the agent does not enter, the project is withdrawn and rejected. If the agent enters, the DM observes the entry decision and updates her belief from p_0 to the post-entry belief p_0^+ . The DM then decides whether to initiate the review process or to stop immediately and approve or disapprove the project. If the DM initiates the review, the dynamic screening process begins.

A type- j agent's mixed strategy is represented by

$$\sigma_j = (d_j, \xi_j),$$

where $d_j \in [0, 1]$ is the probability of **not** entering the review process at time 0, and

$$\xi_j : T \setminus \{0\} \rightarrow [0, \infty]$$

is the quitting hazard after the review has begun. We require ξ_j to be locally integrable on any interval where $\xi_j < \infty$, and interpret $\xi_j(t) = \infty$ as the agent quitting with probability one at t , i.e., an atom in the dropout distribution at t .⁶ We restrict the players' hazard strategies to be piecewise regular: on each interval of a finite partition

⁵Allowing the DM to approve after the agent quits affects only the proofs; none of our results changes.

⁶Formally, the stopping strategy of type j can be represented as a measure on T with both absolutely continuous and atomic components, where the absolutely continuous part has a density given by $\xi_j(t)$ and the atomic part captures sure-quitting events. The formulation here, which allows $\xi_j(t) = \infty$ to represent an atom, is a shorthand for this general representation. We do not need to consider atoms of intermediate size; whenever an atom arises in equilibrium, the agent drops out with probability one conditional on survival until that time.

of the review horizon, each hazard is either zero almost everywhere or strictly positive almost everywhere. Given the sample path of beliefs, the probability of surviving until t without quitting is $\exp(-\int_0^t \xi_j(s) ds)$.

Upon observing the agent's initial entry decision, the DM's belief jumps from p_0 to

$$p_0^+ = \frac{p_0(1 - d_H)}{\frac{p_0}{1-\epsilon}(1 - d_H) + (1 - \frac{p_0}{1-\epsilon})(1 - d_L)}. \quad (2)$$

This is the DM's initial post-entry belief at time 0. After observing entry and updating her belief to p_0^+ , the DM decides whether to initiate the review process or to make an immediate approval decision.

The DM's mixed strategy is represented by a triple

$$\sigma_{DM} = (d_{DM}, \xi_{DM}, \Upsilon),$$

where d_{DM} is the probability that, after observing the agent's entry, the DM stops immediately and either approves or disapproves the project rather than initiating the review process, and

$$\xi_{DM} : T \setminus \{0\} \rightarrow [0, \infty]$$

is a stopping (decision) hazard.⁷ The approval rule

$$\Upsilon : T \setminus \{0\} \rightarrow [0, 1]$$

is a tie-breaking rule that specifies the probability of approval when the DM stops at a belief $p_t = 1/2$. Sequential rationality requires approval if $p_t > 1/2$ and disapproval if $p_t < 1/2$. Since sequential rationality pins down the DM's action for $p_t \neq \frac{1}{2}$, the approval rule is defined only for the tie case $p_t = 1/2$.

Conditional on the review being initiated, the DM continues to update her belief from p_0^+ . For $t > 0$, as long as the review remains active and no favorable signal arrives, her continuation belief evolves

$$\frac{dp_t}{dt} = -p_t(1 - p_t) \left\{ \lambda + \frac{r_t - p_t}{r_t(1 - p_t)} (\xi_H(t) - \xi_L(t)) \right\}. \quad (3)$$

The first term, λ , captures the negative informational effect of the failure of favorable evidence to arrive. The second term captures the information conveyed by endogenous dropout. A higher quitting hazard of the H-type lowers the DM's belief, whereas a higher quitting hazard of the L-type raises it.

Since the H-type is the only type who can generate a favorable signal and hence the only type whose continued participation can sustain experimentation, the review process cannot meaningfully continue after the H-type has been exhausted. We therefore define

⁷In this model, once players' strategies are fixed, the no-signal continuation belief is pinned down by calendar time. Hence the DM's decision hazard can be written as a function of time alone, suppressing the redundant dependence on p_t .

⁸It is derived from Bayes' rule; $p_t + dp_t = \frac{p_t(1-\lambda dt)(1-\xi_H dt)}{p_t(1-\lambda dt)(1-\xi_H dt) + \frac{p_t}{r_t}(1-r_t)(1-\xi_H dt) + (1-\frac{p_t}{r_t})(1-\xi_L dt)}$.

the terminal reachable time with respect to the survival of the H -type and the DM. Given the strategy profile (σ_H, σ_{DM}) , a time $t \in T$ is said to be reachable on the equilibrium path if the game survives until t with strictly positive probability. Formally, the set of reachable times $\mathcal{T}(\sigma_H, \sigma_{DM})$ is defined as follows: if $d_H < 1$,

$$\mathcal{T}(\sigma_H, \sigma_{DM}) := \left\{ t \in T \left| \int_0^s \sum_{j \in \{H, DM\}} \xi_j(u) du < \infty \text{ for all } s < t \right. \right\}.$$

If $d_H = 1$, set $\mathcal{T}(\sigma_H, \sigma_{DM}) = \{0\}$. The terminal reachable time is

$$\bar{t} := \sup \mathcal{T}(\sigma_H, \sigma_{DM}).$$

This depends only on the players' strategies and not on the realization of the agent type.

Our solution concept is that of a perfect Bayesian Nash equilibrium (PBE). This requires that at each time, the DM maximizes her continuation value given her belief about the agent's type and the future game play including off-the-equilibrium path, and the agent maximizes his continuation value given the DM's strategy. To ensure dynamic consistency, these continuation values must be characterized as solutions to the associated Hamilton-Jacobi-Bellman equations. This makes the formal definition rather lengthy. For clarity, we collect the complete definition in Appendix A.

3 Sequential Screening

In this section, we characterize the DM-optimal equilibrium. The equilibrium typically exhibits two distinct screening phases – an attrition phase and an experimentation phase – and may endogenously transition from the former to the latter. Depending on parameter specifications, each phase may itself constitute an equilibrium. We refer to the attrition phase and the experimentation phase as the A-phase and the E-phase, respectively.

When both phases arise along the same equilibrium path, an A-phase precedes an E-phase. The transition occurs when the L -type has been completely eliminated from the continuation population, after which further screening operates exclusively through the arrival or absence of favorable evidence. We characterize the two screening phases separately before combining them to describe the full equilibrium.

Throughout the following analysis, we focus on the case in which the H -type's continuation cost is at least as large as the DM's; $c_H \geq c_{DM}$. This case reflects the usual asymmetry between production and evaluation. The agent's effort cost reflects a concentrated sacrifice of specialized labor and cognitive resources required to generate a breakthrough. In contrast, the DM's cost mainly consists of relatively passive auditing and reviewing, as well as ordinary administrative delay costs, rather than active experimentation. Although the opposite case does not substantially alter the equilibrium analysis, this distinction becomes important for the policy implications, where we revisit the role of this assumption.⁹ All formal results below are stated under this maintained condition.

⁹The opposite case is plausible in environments where the DM internalizes large social costs of delay or bears substantial evaluation, legal, or political costs.

The following assumptions are crucial and imposed throughout the following analysis.

Assumption 1: $\lambda > 2c_{DM}$.

Assumption 2: $\epsilon < \frac{1}{2}$.

Assumption 1 ensures that the arrival rate of favorable evidence is sufficiently high relative to the DM's delay cost. As will become clear below, it plays two roles: the absence of favorable evidence exerts a strong downward force on beliefs, and continued review can remain valuable for the DM even when the current belief is at the approval cutoff. Assumption 2 implies that the H -type's initial posterior r_0 exceeds the DM's approval threshold $\frac{1}{2}$. If either assumption is violated, the screening dynamics characterized below do not arise.

Throughout, the DM's expected payoff refers to her ex ante payoff, evaluated before the agent's entry decision.

3.1 Attrition Phase

This subsection explains the A-phase. In this phase, both types of agents remain potentially active, and screening operates through the endogenous dropout of the L -type. Although the gradual exit of the L -type improves the composition of the surviving population, the absence of favorable evidence simultaneously generates bad news about the project's underlying quality. In equilibrium, the latter force dominates, so that the DM's posterior belief declines over time.

Let s denote the exhaustion time of the L -type. Formally, it is the supremum of times at which the L -type survives with positive probability:

$$s = \sup \left\{ t \geq 0 : \int_0^t \xi_L(u) du < \infty \right\}.$$

The A-phase corresponds to the time interval $\mathcal{A} = [0, s)$.

The first property of this phase is that the H -type agent never quits on the equilibrium path, $\xi_H(t) = 0$ on $\mathcal{A}$. At an intuitive level, the H -type has a stronger reason to continue, because he may benefit from the arrival of a favorable signal. This intuition, however, is not by itself sufficient, since the H -type may also face a higher flow cost than the L -type.

The formal argument considers the two possible responses of the L -type when the H -type quits with positive hazard. Suppose, toward a contradiction, that the H -type quits with positive hazard on some nondegenerate interval. Because quitting yields a payoff of zero, value matching requires the H -type's continuation value to be zero throughout that interval.

If the L -type also quits with positive hazard, his continuation value must likewise remain equal to zero throughout the interval. This would not be problematic at an isolated terminal point, but it cannot be sustained over a nondegenerate interval on which continuation remains possible. Intuitively, simultaneous indifference requires the H -type's additional flow cost, $c_H - c_L$, to be exactly offset by his additional instantaneous

benefit from favorable signal arrival, λr_t . Since r_t declines over time, this balance cannot be maintained over a nondegenerate interval.

If instead the L -type does not quit, positive quitting by the H -type worsens the composition of the surviving population and causes the DM's continuation belief p_t to decline. Moreover, keeping the L -type in the game requires the DM to maintain a positive approval hazard. Waiting therefore generates no further separation of the L -type, while the DM's belief continues to deteriorate. The DM would then strictly prefer to approve immediately rather than continue waiting.

Both cases lead to a contradiction. Hence, the H -type cannot quit with positive hazard on any nondegenerate interval, and therefore $\xi_H(t) = 0$ on the equilibrium path.

Since the H -type does not quit on $\mathcal{A}$, the only source of quitting in the $\omega = B$ state is the L -type. The effective quitting hazard observed by the DM in the $\omega = B$ state is therefore given by

$$\Pr[j = L | \omega = B, h_t] \xi_L(t) = \frac{1 - \frac{p_t}{r_t}}{1 - p_t} \xi_L(t),$$

where h_t is the history where no signal has arrived up to time t and the game has continued until t . Let $\xi_B(t)$ denote this effective dropout hazard in the $\omega = B$ state. By using it, the law of motion for the DM's belief (3) is determined as follows:

$$\frac{dp_t}{dt} = -p_t(1 - p_t)(\lambda - \xi_B(t)). \quad (4)$$

The second property of the phase is that the posterior belief never falls below $\frac{1}{2}$ on $\mathcal{A}$. If it were to fall strictly below $\frac{1}{2}$ at any point, the L -type would no longer expect approval from continuing, unless beliefs recovered above $\frac{1}{2}$ at some future time so that the DM would start approving. Since the H -type never quits, such a recovery in beliefs could only occur if the L -type quits along the way. But then the L -type would have an incentive to quit immediately rather than continuing until that belief recovery phase.

The third property is that the DM always maintains a strictly positive hazard of approval on $\mathcal{A}$. The intuition is as follows. Suppose, toward a contradiction, that there exists a pure waiting interval in $\mathcal{A}$ – i.e., during this interval the L -type cannot expect approval. Then the L -type must be expecting approval at the end of the interval and thus does not quit until then once he has entered this interval. However, since the DM would approve at that point anyway, waiting serves no screening purpose – the L -type is not separated out in this interval. Sequential rationality then implies that the DM should approve immediately rather than wait, contradicting the premise.

This provides a key to characterizing the belief path of the DM in the A-phase: the DM's indifference condition. At p_t , the DM's expected payoff from making a decision immediately is $p_t \cdot 1 + (1 - p_t)(-1)$. On the other hand, waiting for a bit longer and making a decision at time $t + dt$ gives her expected payoff of $-c_{DM}dt + p_t \cdot 1 + (1 - p_t)(1 - \xi_B(t)dt)(-1)$. In equilibrium, the DM is indifferent between those options, which yields an arbitrage condition:

$$\xi_B(t) = \frac{c_{DM}}{1 - p_t}. \quad (5)$$

The indifference condition for the DM further implies that the L -type agent always maintains a strictly positive hazard of quitting and is, at every moment, indifferent between quitting and continuing. This implies that the L -type must be exactly compensated for his flow cost by the prospect of immediate approval; along the equilibrium path, the DM's approval hazard is chosen so that the instantaneous benefit from possible approval just offsets the L -type's flow cost, which pins down $\xi_{DM}(t) = c_L$.

Combining (4) and (5) yields the differential equation for p_t in a period of no signal realization on $\mathcal{A}$:

$$\frac{dp_t}{dt} = -p_t(1-p_t) \left(\lambda - \frac{c_{DM}}{1-p_t} \right). \quad (6)$$

The DM's belief is driven not by the mere arrival (or absence) of signals, but by the endogenous incentives it creates for the low type to exit and for the high type to continue. Although experimentation mechanically generates information through signal realizations, it is the endogenous adjustment of the dropout rate that calibrates the DM's incentives and sustains the indifference condition. Conditional on survival and the absence of a signal, the belief declines monotonically over time. Indeed, (6) can be rewritten as $\frac{dp_t}{dt} = -\lambda p_t(\bar{p} - p_t)$ with $\bar{p} \equiv \frac{\lambda - c_{DM}}{\lambda}$, the belief at which the downward pressure from signal absence exactly offsets the upward pressure from L -type exit. Lemma 5 in the Appendix shows that $p_t < \bar{p}$ throughout the A-phase. This stands in contrast to standard dynamic screening models, where the exit of low types causes the DM's belief to rise. Here, because p_t stays below $\bar{p}$, the downward pressure from signal absence dominates the upward pressure from attrition, so that belief deterioration and low-type exit proceed simultaneously throughout the A-phase.

Given a terminal belief p_s at the end of the A-phase, the belief path for $t \leq s$ is obtained by solving (6) backward from the terminal condition p_s . This yields

$$p_t = \frac{\lambda - c_{DM}}{\lambda + \left(\frac{\lambda - c_{DM}}{p_s} - \lambda \right) e^{-(\lambda - c_{DM})(s-t)}}. \quad (7)$$

Thus, once the terminal belief is specified, the entire belief path during the A-phase is uniquely determined.

The A-phase may itself constitute an equilibrium, in which case the disappearance time of the type- L agent coincides with the terminal time: $s = \bar{t}$. We refer to such an equilibrium as a pure attrition equilibrium. In order for this to be an equilibrium, the DM must disapprove with certainty at s ; if she were to approve with a strictly positive probability at s , then from some earlier time onward both types of agents would optimally continue until that moment, which contradicts the above observation. As $p_s \geq \frac{1}{2}$ in this phase, this implies that $p_s = \frac{1}{2}$.

Setting $p_s = \frac{1}{2}$ in the expression (7) gives a closed form of the DM's belief after the initial screening, $\lim_{t \rightarrow 0} p_t$, as

$$\frac{\lambda - c_{DM}}{(\lambda - 2c_{DM}) e^{-(\lambda - c_{DM})s} + \lambda}. \quad (8)$$

Denote this value by $p_0^+(s)$, to make explicit its dependence on the choice of s . This must be consistent with (2) and the agent's initial strategy $d_H = 0$ and $d_L > 0$.

As is also the case for the sequential screening equilibrium discussed later that begins with an A-phase, a key feature of a pure attrition equilibrium is that the DM's equilibrium payoff is determined by the extent of the L -type's initial exit at time zero, d_L . This is because, following the initial screening at time zero, the DM's optimal action is to approve; thus her ex ante equilibrium payoff is $2p_0 - \frac{p_0}{p_0^+}$.¹⁰ Consequently, finding the DM-optimal equilibrium among this class is equivalent to finding the equilibrium with the largest initial exit by the L -type and thus the highest post-initial-screening belief.

Equation (8) thus implies that the equilibrium that maximizes the DM's expected payoff is the one that has the largest s . Two constraints bound s . First, since $p_t \leq r_t$, the terminal condition $p_s = \frac{1}{2}$ requires $r_s \geq \frac{1}{2}$. Hence s cannot exceed the time at which $r_s = \frac{1}{2}$, which we denote by

$$t^* = \frac{1}{\lambda} \ln \left(\frac{1 - \epsilon}{\epsilon} \right).$$

Second, type H must be willing to remain under review until s . During the A-phase, he incurs the flow cost c_H and is approved either upon signal arrival, at rate λr_t , or without evidence, at rate $\xi_{DM}(t) = c_L$. His net flow payoff, $\lambda r_t + c_L - c_H$, is decreasing in t , so he remains until s if and only if $\lambda r_s \geq c_H - c_L$. Let t_A denote the time at which this constraint binds:

$$t_A = \frac{1}{\lambda} \log \left(\frac{(1 - \epsilon)(\lambda - (c_H - c_L))}{\epsilon(c_H - c_L)} \right),$$

with $t_A = 0$ if $(1 - \epsilon)\lambda \leq c_H - c_L$ and $t_A = \infty$ if $c_H \leq c_L$. Let $\bar{s} \equiv \min\{t^*, t_A\}$. The next proposition shows that every $s \in (0, \bar{s}]$ can be supported.

Proposition 1 *For each $s > 0$, a pure attrition equilibrium that satisfies the following conditions exists if and only if $s \leq \bar{s}$ and $p_0 \leq p_0^+(s)$.*

- *The type-H agent enters at time zero, and the DM initiates the review after observing entry, $d_{DM} = d_H = 0$, and the L -type quits at $t = 0$ with probability*

$$d_L = \frac{p_0^+(s) - p_0}{p_0^+(s) \left(1 - \frac{p_0}{1 - \epsilon}\right)}.$$

- *For almost all $t < s$, $\xi_{DM}(t) = c_L$, $\xi_L(t) = \frac{c_{DM}}{1 - p_t/r_t}$, and $\xi_H(t) = 0$.*
- *The DM disapproves at s ; $\Upsilon(s) = 0$.*

The DM-optimal pure attrition equilibrium is the one with $s = \bar{t} = \bar{s}$, which exists if and only if $\bar{s} > 0$ and $p_0 \leq p_0^+(\bar{s})$.

¹⁰Since the DM's posterior conditional on entry is p_0^+ , Bayes' rule implies that the ex ante probability of entry is p_0/p_0^+ . Conditional on entry, approval yields the DM $2p_0^+ - 1$. Her ex ante expected payoff is thus $(p_0/p_0^+)(2p_0^+ - 1) = 2p_0 - \frac{p_0}{p_0^+}$.

Notice that whenever a pure attrition equilibrium exists, the DM-optimal one with $s = \bar{s}$ also exists, because $p_0^+(s)$ is increasing in s . In particular, $\bar{s} = t^*$ if and only if $\lambda \geq 2(c_H - c_L)$. In this case, type H's participation constraint does not bind before the DM's belief reaches $\frac{1}{2}$. The term c_L enters because, even in the absence of favorable evidence, the DM approves with a strictly positive hazard, which gives type H an additional incentive to continue.

Figure 1 illustrates the dynamics of the DM's belief within the DM-optimal pure attrition equilibrium when $\bar{s} \equiv t^*$. The initial screening of agent types causes the DM's belief to jump upward from p_0 to $p_0^+(t^*)$. Subsequently, in the absence of a signal, the belief gradually declines until it reaches $\frac{1}{2}$, the point where the beliefs of the DM and the H -type converge.

Figure 2, on the other hand, depicts the belief dynamics in a typical suboptimal pure attrition equilibrium. Here, the initial screening is less pronounced, starting from a belief level below $p_0^+(t^*)$. The game concludes before t^* , when the H -type's belief would have reached $\frac{1}{2}$; consequently, the extent of gradual screening remains limited.




Figure 1. Optimal pure attrition equilibrium




Figure 2. Suboptimal pure attrition equilibrium

We obtain the following comparative statics.

Corollary 1 *In the DM-optimal pure attrition equilibrium, the DM's expected payoff is*

$$V_{DM}(0) = p_0 \frac{\lambda - 2c_{DM}}{\lambda - c_{DM}} \left[1 - e^{-(\lambda - c_{DM})\bar{s}} \right], \quad \bar{s} = \min\{t^*, t_A\}.$$

It is increasing in λ , decreasing in c_{DM} and ϵ , weakly decreasing in c_H , and weakly increasing in c_L . In particular, it is invariant to c_H and c_L if $\lambda \geq 2(c_H - c_L)$, so that $\bar{s} = t^*$, and it is strictly decreasing in c_H and strictly increasing in c_L if $\lambda < 2(c_H - c_L)$, so that $\bar{s} = t_A$.

Note that when $\lambda > 2c_H$, the pure attrition equilibrium relies on a somewhat delicate specification of off-path beliefs. Near the terminal point, both the type-H agent and the DM would prefer to continue if they expected the other party to do so; the arrival rate of favorable evidence $\frac{1}{2}\lambda$ exceeds their flow costs c_H and c_{DM} . Sustaining termination

therefore requires off-path beliefs that discourage unilateral deviations despite these continuation incentives. In our analysis, we do not impose equilibrium refinements based on the plausibility of off-path beliefs. Even without such refinements, the pure attrition equilibrium is not DM-optimal throughout this parameter region: the equilibria characterized in the following subsections yield the DM a higher payoff.

3.2 Experimentation Phase

We now turn to the E-phase, which may follow an A-phase. In the E-phase screening is driven solely by the arrival of the stochastic evidence. The H -type continues with certainty while his posterior belief remains sufficiently high, waiting for favorable evidence to arrive. In the absence of a signal, his posterior deteriorates over time until it reaches a threshold at which continued experimentation is no longer worthwhile. The H -type then quits and the review process terminates.

At the beginning of this phase, the L -type has been fully exhausted. Thus the E-phase corresponds to the time interval $\mathcal{E} = [s, \bar{t}]$. Hence, conditional on continuation, the DM knows that the agent is of type H , and her posterior coincides with the H -type's posterior, $p_t = r_t$ for all $t \in \mathcal{E}$. Further screening therefore no longer operates through endogenous dropout, but solely through the arrival or absence of favorable evidence. The DM continues the review to give a good project an opportunity to reveal itself. Such continuation can be sustained only if the terminal action depends on whether favorable evidence arrives: if the DM planned to approve even in the absence of a signal, she would strictly prefer to approve immediately and avoid further delay costs. Thus, the E-phase must terminate with rejection in the absence of favorable evidence.

At $t \in \mathcal{E}$, if the project is good, favorable evidence arrives at rate λ , so the probability that it arrives before $\bar{t}$ is $1 - e^{-\lambda(\bar{t}-t)}$. The expected duration before time $\bar{t}$ is $\frac{1}{\lambda}(1 - e^{-\lambda(\bar{t}-t)})$. If the project is bad, favorable evidence never arrives, and the review lasts until $\bar{t}$ and player j ends up paying effort cost $(\bar{t} - t)c_j$. The players' expected payoffs from this policy at time $t \in \mathcal{E}$ are therefore:

$$W_j(t; \bar{t}) = r_t(1 - e^{-\lambda(\bar{t}-t)}) - c_j \left[\frac{r_t}{\lambda}(1 - e^{-\lambda(\bar{t}-t)}) + (1 - r_t)(\bar{t} - t) \right],$$

where $c_j = c_{DM}$ for the DM and $c_j = c_H$ for the H type.

Balancing the informational benefit of giving a good project more time to generate favorable evidence against the additional cost of delay determines the DM's unconstrained optimal stopping time of the E-phase. Maximizing $W_{DM}(0; \bar{t})$ with respect to $\bar{t}$ yields it; however, this stopping time is not implementable when the H type bears a strictly higher delay cost than the DM, because he would quit earlier. Since $c_H \geq c_{DM}$, the feasible terminal time is determined instead by the H type's no-exit constraint. In particular, the phase must terminate no later than the time at which $W_H(0; \bar{t})$ is maximized, i.e., when the instantaneous signal arrival rate, λr_t , equals his delay cost c_H . If $(1 - \epsilon)\lambda > c_H$, this occurs at

$$\bar{t} = \frac{1}{\lambda} \log \left(\frac{(1 - \epsilon)(\lambda - c_H)}{\epsilon c_H} \right). \quad (9)$$

We denote this particular value of $\bar{t}$ by t_E (it is 0 if $(1 - \epsilon)\lambda \leq c_H$). For the E-phase to terminate with rejection at t_E , the DM's posterior must have fallen to or below the approval threshold by that time. This requires $r_{t_E} \leq \frac{1}{2}$, which is equivalent to $\lambda \geq 2c_H$.

The continuation policy is dynamically consistent. Intuitively, once the DM prefers to continue the E-phase rather than stop immediately, she has no incentive to abandon the same plan at a later continuation history. Although the absence of a signal makes the project less attractive over time, the remaining cost of waiting until the planned terminal date also falls. Hence, once continuation is preferred at some date, it remains preferred at all later dates before termination.¹¹

Like the A-phase, the E-phase can form an equilibrium on its own for some parameter values, which we call a pure experimentation equilibrium. For existence, we need the DM to have no incentive to approve once the L type is excluded at time 0, that is, $W_{DM}(0; \bar{t}) \geq 2r_0 - 1$.

Proposition 2 *A pure experimentation equilibrium with $\bar{t} \in [t^*, t_E]$ that satisfies the following conditions exists if and only if $W_{DM}(0; \bar{t}) \geq 2r_0 - 1$ and $\lambda \geq 2c_H$.*

- The type- H agent enters at time zero, and the DM initiates the review after observing entry, $d_{DM} = d_H = 0$, and the L -type quits at $t = 0$ with probability 1, $d_L = 1$.
- For almost all $t < \bar{t}$, $\xi_{DM}(t) = 0$ and $\xi_H(t) = 0$.
- The DM disapproves at $\bar{t}$; $\Upsilon(\bar{t}) = 0$.

The DM-optimal pure experimentation equilibrium satisfies $\bar{t} = t_E$.

Notice that whenever a pure experimentation equilibrium exists, the DM-optimal pure experimentation equilibrium also exists. Indeed, $W_{DM}(0; \bar{t}) < W_{DM}(0; t_E)$ for every $\bar{t} < t_E$; thus if $W_{DM}(0; \bar{t}) \geq 2r_0 - 1$ then $W_{DM}(0; t_E) > 2r_0 - 1$ holds.

The appeal of a pure-experimentation equilibrium is that it eliminates one of the two screening problems at the outset. In general, the DM must screen both the agent's type and the project's quality. Under pure experimentation, the L -type exits immediately at time zero, so the continuation phase is devoted solely to screening the quality of the remaining H -type's project through stochastic evidence. Thus, although pure experimentation is efficient in terms of separating the two screening problems, it tends to be supported in environments that are unfavorable to the DM: ϵ must be sufficiently high, so that even the H -type is sufficiently uncertain about the project's quality.

As will be seen explicitly in the next subsection, the existence condition $W_{DM}(0; \bar{t}) \geq 2r_0 - 1$ also helps explain why the H -type's delay cost matters. In a pure-experimentation equilibrium, the key incentive constraint is the H -type's willingness to stay in the review process while waiting for favorable evidence. As c_H increases, this constraint becomes tighter, making the pure-experimentation equilibrium harder to sustain.

¹¹Mathematically, this is equivalent to the following monotonicity property: if $W_{DM}(t'; \bar{t}) > \max\{2r_{t'} - 1, 0\}$ for some $t' < \bar{t}$, then $W_{DM}(t; \bar{t}) > \max\{2r_t - 1, 0\}$ for all $t \in (t', \bar{t})$.

By contrast, a higher arrival rate of favorable evidence makes pure experimentation easier to sustain. When λ is higher, favorable evidence is more likely to arrive before the terminal date. This increases the DM's value from waiting and also makes it less costly for the H-type to remain in the review process. Thus, the pure-experimentation equilibrium is easier to sustain when evidence arrives more frequently, but harder to sustain when the H-type's delay cost is higher.

Corollary 2 *In a DM-optimal pure experimentation equilibrium, the DM's expected payoff is increasing in λ and decreasing in ϵ , c_{DM} , and c_H .*

While most comparative statics results for the DM's payoff are intuitive, a notable finding is that the payoff decreases with c_H . A higher c_H leads to an earlier terminal time t_E , thereby reducing the screening benefit for the DM. This observation will be relevant to our discussion of subsidies in the next section.

3.3 Sequential Screening Equilibrium

We now turn to the case in which an A-phase and an E-phase appear in sequence. The equilibrium has a natural two-stage interpretation. The first stage is devoted to screening out the L -type through endogenous attrition. Once the L -type has been eliminated, the game moves to a purely collaborative experimentation stage in which the DM and the H-type wait for favorable evidence. In the latter part of the equilibrium path, the DM no longer approves without favorable evidence.

The central theoretical object is the transition time s between the two phases, at which the type- L agent disappears entirely. It must be chosen so that the connection between the two phases does not disturb the within-phase incentive structures characterized in the previous subsections. For this to be the case, the DM's value approaching the transition from the A-phase must coincide with her value at the beginning of the E-phase. This connecting condition implies that, at s , the DM must be indifferent between entering the E-phase and approving;

$$2r_s - 1 = W_{DM}(s; \bar{t}). \quad (10)$$

For any given terminal time $\bar{t}$, the connecting condition (10) uniquely determines the transition time s , whenever such a transition exists (Lemma 4(ii) in the Appendix), and we denote it by $s(\bar{t})$; $2r_{s(\bar{t})} - 1 = W_{DM}(s(\bar{t}); \bar{t})$. The transition time $s(\bar{t})$ in turn determines the extent of initial screening in the A-phase and hence the DM's equilibrium payoff.

To see how the transition time affects initial screening, note that the transition condition requires $p_s = r_s$. Thus, a transition at s selects the point (s, r_s) as the terminal condition for the A-phase belief dynamics. Since r_t is decreasing, an earlier transition corresponds to a higher terminal belief. Consider two transition times $s_1 < s_2$. The A-phase associated with s_2 satisfies $p_t < r_t$ for $t < s_2$, so at $t = s_1$ its belief is strictly below r_{s_1} , the terminal belief of the path associated with s_1 . By uniqueness of the solution to the belief ODE (6), the two belief paths cannot cross. Hence, the path associated with s_1 remains above that associated with s_2 at all earlier dates and, in particular, yields a

higher post-initial-screening belief, $\lim_{t \rightarrow 0} p_t$. An earlier transition therefore requires more initial exit by the L-type and yields a higher payoff to the DM.

It remains to determine the terminal time $\bar{t}$ for the DM-optimal equilibrium. Since the DM prefers the earliest feasible transition, she chooses the terminal time that makes the connecting condition hold at the smallest possible s . Extending the terminal time raises the DM's continuation value $W_{DM}(s; \bar{t})$ at every s as long as $\bar{t} < t_E$, because the instantaneous value of waiting, $\lambda r_{\bar{t}} - c_{DM}$, remains positive up to t_E . A higher continuation value makes the DM willing to enter the E-phase at a higher belief, and hence at an earlier date. Thus $s(\bar{t})$ is strictly decreasing in $\bar{t}$ on $[t^*, t_E]$.¹² Hence, t_E supports the earliest feasible transition and therefore maximizes the DM's payoff among sequential screening equilibria.

The relevant conditions for the existence of the sequential construction are $W_{DM}(0; t_E) < 2r_0 - 1$ and $\lambda \geq 2c_H$. Given $\lambda > 2c_H$, we have $W_{DM}(t^*; t_E) > 2r_{t^*} - 1 = 0$ by Lemma 4(i), and hence the existence of $s(t_E) < t^*$ is ensured. When $\lambda = 2c_H$, we have $t_E = t^*$, and the connecting condition holds at $s = t^* = t_E$; the E-phase is then degenerate, and the construction coincides with the DM-optimal pure attrition equilibrium with $\bar{s} = t^*$. As in a pure-experimentation equilibrium, the condition $\lambda \geq 2c_H$ guarantees that the participation constraint of the H-type is satisfied.

Denote by $p_0^{++}(\bar{t}) = \lim_{t \rightarrow 0} p_t$ the initial post-entry belief induced by this equilibrium construction, with s chosen to be $s(\bar{t})$ and the motion of p_t determined by (7).

Proposition 3 *For each $\bar{t} \in [t^*, t_E]$, a sequential screening equilibrium with terminal time $\bar{t}$ that satisfies the following conditions exists if and only if $W_{DM}(0; \bar{t}) < 2r_0 - 1$, $\lambda \geq 2c_H$, and $p_0 \leq p_0^{++}(\bar{t})$.*

- The type-H agent enters at time zero, and the DM initiates the review after observing entry, $d_{DM} = d_H = 0$, and the L-type quits at $t = 0$ with probability

$$d_L = \frac{p_0^{++}(\bar{t}) - p_0}{p_0^{++}(\bar{t})(1 - \frac{p_0}{1-\epsilon})}.$$

- For almost all $t < s(\bar{t})$, $\xi_{DM}(t) = c_L$, $\xi_L(t) = \frac{c_{DM}}{1 - \frac{p_t}{r_t}}$, and $\xi_H(t) = 0$.
- For almost all $t \in (s(\bar{t}), \bar{t})$, $\xi_{DM}(t) = 0$ and $\xi_H(t) = 0$.
- The DM disapproves at $\bar{t}$; $\Upsilon(\bar{t}) = 0$.

The sequential screening equilibrium that satisfies $(s, \bar{t}) = (t_L, t_E)$, if it exists, payoff-dominates other sequential screening equilibria.

¹²Let $D(s; \bar{t}) = W_{DM}(s; \bar{t}) - (2r_s - 1)$, so that $s(\bar{t})$ is defined by $D(s(\bar{t}); \bar{t}) = 0$. By Lemma 4(ii), $\partial D / \partial s = c_{DM} > 0$ at $s = s(\bar{t})$. Moreover, direct differentiation gives

$$\frac{\partial W_{DM}(s; \bar{t})}{\partial \bar{t}} = [r_s e^{-\lambda(\bar{t}-s)} + 1 - r_s](\lambda r_{\bar{t}} - c_{DM}),$$

which is strictly positive for $\bar{t} < t_E$, since $\lambda r_{\bar{t}} > \lambda r_{t_E} = c_H \geq c_{DM}$. By the implicit function theorem, $s'(\bar{t}) = -(\partial W_{DM}(s; \bar{t}) / \partial \bar{t}) / c_{DM} < 0$. The same expression shows that $W_{DM}(s; \bar{t}) \leq W_{DM}(s; t_E)$ for every s and every $\bar{t} \leq t_E$.




Figure 3. Optimal sequential screening equilibrium




Figure 4. Suboptimal sequential screening equilibrium

Figures 3 and 4 illustrate the belief dynamics in a sequential screening equilibrium. At time zero, the initial exit of the L-type raises the DM's belief from p_0 to the post-screening belief $p_0^{++}(t_E)$. During the subsequent A-phase, the DM's belief p_t gradually declines despite the continued exit of the L-type, reflecting the stronger downward force generated by the absence of favorable evidence. Throughout this phase, $p_t < r_t$, since the continuation pool still contains some L-types. At the transition time s , the L-type is fully exhausted and the two beliefs meet, $p_s = r_s$. Beyond this point, the belief path becomes steeper because the offsetting effect of L-type exit disappears. From then on, screening is driven solely by the absence or arrival of favorable evidence, and the DM's belief coincides with the H-type's posterior, $p_t = r_t$, until the review terminates.

The two figures differ in the amount of screening performed at the initial stage and in the timing of the transition. In the DM-optimal sequential screening equilibrium shown in Figure 3, the transition occurs at the earliest feasible time. This requires the largest initial increase in the DM's posterior and therefore induces the greatest initial exit by the L-type. Figure 4 illustrates a suboptimal sequential screening equilibrium with a later transition. Because the corresponding transition belief r_s is lower, the A-phase can be supported by a lower post-screening belief, implying less initial separation and a lower payoff for the DM.

Corollary 3 *In the DM-optimal sequential screening equilibrium, the DM's expected payoff is increasing in λ and decreasing in ϵ , c_{DM} , and c_H .*

Combining the preceding three propositions yields a simple ranking of the candidate equilibria. Whenever a pure experimentation equilibrium exists, it is optimal for the DM. If no pure experimentation equilibrium exists but a sequential screening equilibrium does, then the sequential screening equilibrium is optimal. Pure attrition is optimal only when neither of these equilibria exists.

Theorem 1 1. Suppose that $\lambda \geq 2c_H$. If

$$c_{DM} > \frac{\epsilon\lambda(\lambda - 2c_H)}{(1 - \epsilon)\lambda - c_H + \epsilon(\lambda - c_H) \log \frac{(1 - \epsilon)(\lambda - c_H)}{\epsilon c_H}}, \quad (11)$$

and $p_0 \leq p_0^{++}(t_E)$ then the DM-optimal equilibrium is the sequential screening equilibrium with $(s, \bar{t}) = (t_L, t_E)$. If the inequality (11) is reversed, the DM-optimal equilibrium is the pure experimentation equilibrium with $(s, \bar{t}) = (0, t_E)$.

2. If $\lambda \in [2(c_H - c_L), 2c_H)$ and $p_0 \leq p_0^+(t^*)$, the DM-optimal equilibrium is the pure attrition equilibrium with $\bar{t} = t^*$. If $\lambda \in ((c_H - c_L)/r_0, 2(c_H - c_L))$ and $p_0 \leq p_0^+(t_A)$, the DM-optimal equilibrium is the pure attrition equilibrium with $\bar{t} = t_A$.

3. In other cases, no equilibrium in which the review is initiated exists. The DM does not initiate the review and instead decides immediately on the basis of her prior, approving if $p_0 > \frac{1}{2}$ and disapproving if $p_0 < \frac{1}{2}$.

Condition (11) is equivalent to $W_{DM}(0; t_E) < 2r_0 - 1$. The theorem clarifies how λ and c_H affect the cutoff value of c_{DM} between pure experimentation and sequential screening. A higher λ shifts this cutoff in favor of pure experimentation, whereas a higher c_H shifts it in favor of sequential screening. The exact boundary depends on the H-type's continuation constraint through c_H . It is useful, however, to isolate a simpler force that rules out pure experimentation independently of c_H . A sufficient condition is given by $\epsilon < \frac{c_{DM}}{\lambda}$.¹³ When this inequality holds, the H-type agent's posterior is sufficiently high at the initial stage. The DM then has too strong an incentive to approve immediately rather than wait for favorable evidence. Therefore, a pure experimentation equilibrium cannot be sustained.

Finally, the equilibrium characterization in this section carries an implication for regulatory authorities: imposing a deadline without committing to a specific approval strategy may be counterproductive. Indeed, setting a deadline $\hat{t} < t_E$ exogenously precludes any equilibria where $\bar{t} > \hat{t}$, thereby constraining the efficiency of the screening process and ultimately diminishing the DM's payoff.

Corollary 4 *Suppose that the review must end by an exogenous deadline $\hat{t}$, at which the DM makes her decision without commitment. Then the DM's expected payoff in the DM-optimal equilibrium is weakly lower than in the model without a deadline.*

4 Commitment and Policy Implications

This section examines two interventions in the screening process. We first show how commitment affects the DM's payoff and then examine how a subsidy to the agent changes the duration and form of screening. Finally, we examine the uncertainty-free limit, which motivates a form of commitment we call terminal recognition.

¹³Let $k = c_{DM}/\lambda$. Maximizing $W_{DM}(0; \bar{t}) = (1 - \epsilon)(1 - k)(1 - e^{-\lambda \bar{t}}) - k\epsilon\lambda \bar{t}$ over $\bar{t} \geq 0$ without the type-H agent's constraint, the maximum is attained where $\lambda r_{\bar{t}} = c_{DM}$, i.e., at $e^{\lambda \bar{t}} = \frac{(1-\epsilon)(1-k)}{\epsilon k}$, provided that this expression exceeds one. Substituting gives

$$\max_{\bar{t} \geq 0} W_{DM}(0; \bar{t}) - (2r_0 - 1) = (\epsilon - k) - \epsilon k \log \frac{(1 - \epsilon)(1 - k)}{\epsilon k},$$

which is strictly negative whenever $\epsilon \leq k$, since the logarithm is then positive. If instead $\frac{(1-\epsilon)(1-k)}{\epsilon k} \leq 1$, then $W_{DM}(0; \bar{t}) \leq 0 < 2r_0 - 1$ for all $\bar{t}$. In either case, $W_{DM}(0; t_E) < 2r_0 - 1$, so no pure experimentation equilibrium exists.

4.1 Value of Commitment

The analysis in the previous section shows that shortening the time horizon within the dynamic screening equilibrium does not improve the DM's payoff. Although an authority may hope that imposing a tighter review timeline would discipline the screening process and lead to a better outcome, such an intervention is futile in the absence of a corresponding commitment on how the terminal decision will be made. To derive policy implications, this section therefore considers a setting in which the DM can commit ex ante to a strict deadline together with a commitment not to approve without signal arrival, and evaluates whether such commitment improves her payoff relative to the equilibrium.

Commitment to the no-approval rule induces the type-L agent, who can never generate favorable evidence, to exit immediately. The type-H agent, by contrast, remains active in the hope of generating such evidence until the announced deadline. The policy therefore eliminates the attrition phase and implements a pure experimentation outcome. Its value depends on whether the gain from immediately screening out the type-L agent outweighs the loss from forgoing approval without evidence during the attrition phase.

Mathematically, the DM's commitment problem can be formulated as follows:

$$\begin{aligned} \max_{\tau} W_{DM}(0; \tau), \\ s.t. \ W_H(t; \tau) \geq 0 \text{ for all } t \leq \tau. \end{aligned}$$

In this formulation, the constraint is the type- H agent's participation constraint, which guarantees that the H -type continues until τ .

An immediate observation is that a pure experimentation equilibrium implements the commitment solution. Therefore, whenever such an equilibrium exists, commitment has no value to the DM. Moreover, whenever a sequential screening equilibrium becomes optimal, commitment strictly increases the DM's payoff. The reason is that the commitment solution and the sequential screening equilibrium attain the same terminal time, t_E . In the experimentation phase of a sequential screening equilibrium, the approval hazard has already fallen to zero, so the continuation problem from that point onward coincides with the corresponding commitment problem. Commitment therefore raises the DM's payoff by eliminating the preceding attrition phase while preserving the same terminal time.

Finally, consider the region where pure attrition is optimal. Here the effect of commitment is ambiguous, because approval without favorable evidence plays two roles in the pure attrition equilibrium. On the one hand, it keeps the type-L agent in the review process, which is costly to the DM.

On the other hand, it keeps the type-H agent in the review until $\bar{s}$: favorable evidence alone would sustain his participation only until t_E , and $t_E < \bar{s}$ because $\lambda < 2c_H$ and $c_L > 0$. Commitment removes both effects at once. The type-L agent exits immediately, but the type-H agent quits at t_E .

The comparison is transparent because the DM is indifferent throughout the attrition phase. Her equilibrium payoff therefore equals the payoff from waiting for favorable evidence until $\bar{s}$, holding the agents' equilibrium strategies fixed. This benchmark differs

from commitment precisely because the agents' strategies would not stay fixed under commitment. Commitment is harmful if and only if the gain from the longer experimentation horizon exceeds the DM's delay cost of the type-L agent's residual presence. When c_{DM} is small, the latter cost is negligible and commitment is harmful. When c_{DM} is close to $\lambda/2$, the attrition phase yields the DM almost nothing, and commitment is beneficial provided that experimentation alone can sustain the type-H agent's participation for a nondegenerate period, that is, $(1 - \epsilon)\lambda > c_H$. (Proposition 4(iii) also requires $p_0 \leq \frac{1}{2}$, which guarantees that the pure attrition equilibrium exists for all such c_{DM} .)

This distinction is relevant to drug regulation. The FDA's accelerated approval pathway permits approval before a drug's clinical benefit has been conclusively verified, subject to subsequent confirmatory trials. Our analysis suggests why a policy that categorically rules out approval under uncertainty may be too restrictive: the prospect of approval without conclusive evidence can be what keeps promising applicants engaged long enough for evidence to arrive.

Overall, commitment not to approve is valuable as an implementation device for the pure experimentation outcome. The value of commitment arises only when the H-type's incentive to continue experimentation is sufficiently strong. When this incentive is weak, the equilibrium without commitment may instead deliver a higher payoff to the DM.

Proposition 4 *The value of commitment depends on the form of the DM-optimal equilibrium as follows.*

- (i) *If the DM-optimal equilibrium is a pure experimentation equilibrium, commitment does not change the DM's payoff.*
- (ii) *If the DM-optimal equilibrium is a sequential screening equilibrium, commitment strictly increases the DM's payoff.*
- (iii) *If the DM-optimal equilibrium is a pure attrition equilibrium, commitment strictly decreases the DM's payoff if and only if*

$$\frac{p_0}{1 - \epsilon} [W_{DM}(0; \bar{s}) - W_{DM}(0; t_E)] > c_{DM}\Lambda_L, \quad (12)$$

where Λ_L is the ex ante expected time the type-L agent spends under review in the DM-optimal pure attrition equilibrium. In particular, commitment strictly decreases the DM's payoff when c_{DM} is sufficiently small. By contrast, if $(1 - \epsilon)\lambda > c_H$ and $p_0 \leq \frac{1}{2}$, commitment strictly increases the DM's payoff when c_{DM} is sufficiently close to $\lambda/2$.

Improving upon a pure attrition equilibrium requires a more sophisticated form of commitment. In particular, the DM must be able to use a stochastic mechanism that permits approval even in the absence of favorable evidence. For example, suppose that the DM can commit to a flow approval rate of c_L along the no-evidence path. This makes the type-L agent indifferent between remaining active and exiting; under the usual tie-breaking rule in favor of exit, the type-L agent exits immediately and is screened out from the outset. Equivalently, a flow approval rate slightly below c_L makes immediate exit strictly optimal for the type-L agent.

Such a mechanism, however, may be difficult to implement in practice. It requires both stochastic approval and the possibility of approving a project without verifiable favorable

evidence. These features are difficult to reconcile with verifiability and transparency, which are essential for ensuring that the DM follows the prescribed rule and for preventing ex post manipulation. In particular, when approval is granted without supporting evidence, it may be difficult for outside parties to distinguish an outcome generated by the stipulated randomization device from favoritism or an opportunistic deviation by the DM. We therefore view this mechanism primarily as a theoretical benchmark.

4.2 Subsidies

The preceding analysis also has implications for policy interventions other than commitment. In particular, a subsidy that reduces the agent's effort cost can improve screening by extending the terminal date within a given equilibrium regime. The key observation is that, when the DM's cost of delay is lower than the high type's effort cost, the DM would prefer screening to continue beyond the point at which the high type is willing to remain. A subsidy relaxes this constraint by inducing the high type to stay longer. If instead the DM's cost of delay is higher, the DM would already prefer to stop earlier than the H-type agent. In that case, there is no reason to provide the H-type with additional incentives to prolong his participation. In the following analysis, let $\varphi \geq 0$ denote a per-unit-time subsidy paid to the agent while the agent remains active.

Within a pure attrition equilibrium a subsidy cannot extend the review horizon. Because the DM cannot observe the agent's type, the subsidy is paid to both types while they remain active. The type-L agent's indifference condition then lowers the DM's approval hazard to $\xi_{DM}(t) = c_L - \varphi$, so the type-H agent's net flow payoff from continuing, $\lambda r_t + (c_L - \varphi) - (c_H - \varphi) = \lambda r_t + c_L - c_H$, is unaffected by φ .¹⁴ Hence the terminal date of the attrition phase is unchanged. At the same time, the DM's effective delay cost rises to $c_{DM} + \varphi$, which raises the dropout hazard required for her indifference, $\xi_B(t) = (c_{DM} + \varphi)/(1 - p_t)$, and thereby lowers the post-screening belief p_0^+ consistent with the same terminal belief. By Corollary 1, the DM's equilibrium payoff strictly decreases. A subsidy is therefore counterproductive in a pure attrition equilibrium.

To highlight the within-equilibrium effect of a subsidy, therefore, it is natural to focus first on a pure experimentation equilibrium. In this case, the subsidy does not change the nature of the equilibrium; instead, its main effect is to alter how long experimentation can be sustained. Intuitively, a subsidy relaxes the type-H agent's incentive constraint by reducing the effective cost of remaining active, and may therefore allow the DM to extend the experimentation phase. At the same time, the subsidy itself is costly to the DM. This creates a simple trade-off between a longer experimentation horizon and the direct cost of subsidization.

Under the subsidy, the type H agent's effective flow cost is $c_H - \varphi$, while the DM's flow cost is $c_{DM} + \varphi$. For each subsidy level φ , let $t_E(\varphi)$ denote the terminal date of experimentation induced by the type-H agent's incentive constraint. The DM's expected

¹⁴This requires $\varphi < c_L$, so that the approval hazard remains positive, and $\lambda > 2(c_{DM} + \varphi)$, so that Assumption 1 continues to hold with the DM's effective delay cost.

payoff is expressed as:

$$W_{DM}(0; t_E(\varphi), \varphi) = r_0(1 - e^{-\lambda t_E(\varphi)}) - (c_{DM} + \varphi) \left[\frac{r_0}{\lambda}(1 - e^{-\lambda t_E(\varphi)}) + (1 - r_0)t_E(\varphi) \right].$$

Then the subsidy is effective if

$$\frac{dW_{DM}(0; t_E(\varphi), \varphi)}{d\varphi} = \frac{\partial W_{DM}(0; t_E(\varphi), \varphi)}{\partial t_E} \frac{dt_E}{d\varphi} + \frac{\partial W_{DM}(0; t_E(\varphi), \varphi)}{\partial \varphi} > 0,$$

evaluated at $\varphi = 0$, where the first term is the positive effect of extending the terminal time, while the second one is the negative effect of the subsidy payment. Intuitively, such a policy is more likely to be effective when the disagreement between the two players over how long experimentation should continue is particularly large: the DM would prefer experimentation to continue for longer, whereas the type- H agent has a stronger incentive to quit. A subsidy can reduce this conflict by bringing the agent's incentives closer to the DM's preferred experimentation horizon. Also, it is more likely to be effective when the initial experimentation period is short, since the subsidy is paid only until t_E . In particular, when λ is close to $2c_H$ and r_0 is close to $\frac{1}{2}$, the terminal time t_E is close to t^* , making the direct subsidy cost small.

Evaluating the two terms at $\varphi = 0$ and simplifying yields

$$\frac{dW_{DM}(0; t_E(\varphi), \varphi)}{d\varphi} \Big|_{\varphi=0} = \frac{(1 - r_0)\lambda(c_H - c_{DM})}{c_H(\lambda - c_H)^2} - \left\{ \frac{r_0\lambda - c_H}{\lambda(\lambda - c_H)} + \frac{1 - r_0}{\lambda} \log \frac{r_0(\lambda - c_H)}{(1 - r_0)c_H} \right\}.$$

The above intuition can be made precise by considering the boundary case in which $c_{DM} = 0$, $\lambda = 2c_H$, and $r_0 = \frac{1}{2}$. At this boundary, we have $t_E(0) = t^*$ and the marginal effect of an increase in subsidy from $\varphi = 0$, $\frac{dW_{DM}(0; t_E(\varphi), \varphi)}{d\varphi} \Big|_{\varphi=0}$, is given by $\frac{1}{c_H} > 0$. Hence, a small subsidy is strictly beneficial at the boundary. This boundary point itself lies outside Assumption 2 and the region supporting pure experimentation. However, since the marginal effect is continuous in the parameters, it remains positive in an open neighborhood of the boundary point, and this neighborhood intersects the region supporting a pure experimentation equilibrium: for any $\lambda > 2c_H$ and $\epsilon < \frac{1}{2}$, the inequality (11) is reversed for all sufficiently small $c_{DM} > 0$. Therefore, there exists a nonempty open set of parameter values for which a small subsidy strictly increases the DM's payoff.

By contrast, when r_0 is sufficiently close to one, such a subsidy reduces the DM's payoff. As the uncertainty about the type- H agent's project quality vanishes, the screening benefit from extending the experimentation period converges to zero, whereas the marginal subsidy cost converges to $1/\lambda$, the expected time until a signal arrives.

The preceding analysis assumes, for simplicity, that the subsidy is paid throughout the entire experimentation phase. This is not essential, however. The type- H agent's participation constraint is satisfied strictly before the terminal date and binds only at the end of the experimentation phase. The subsidy can therefore be back-loaded and paid only near the terminal date, where it is needed to extend the experimentation horizon. If favorable evidence arrives earlier, the project is approved and no subsidy is paid.

The advantage of such back-loading is even more pronounced in a sequential screening equilibrium. During the attrition phase, the type- L agent remains active together with the

type-H agent, so a subsidy paid at this stage must also be paid to the type-L agent. Moreover, subsidizing participation during the attrition phase may delay the type-L agent's exit and thereby interfere with the screening role of attrition. By postponing the subsidy until after the type-L agent has exited, the DM can target the transfer to the experimentation phase, where it relaxes the type-H agent's binding participation constraint without subsidizing the type-L agent or distorting the preceding attrition phase.

4.3 Uncertainty-Free Limit and Terminal Recognition

We may consider the uncertainty-free limit in which $\epsilon = 0$. In this case, the type-H agent knows with certainty that $\omega = G$, whereas the type-L agent knows that $\omega = B$. Thus, conditional on the agent's type, there is no residual uncertainty about the project's quality. The DM, however, still does not observe the agent's type and must therefore distinguish a good project held by a type-H agent from a bad project held by a type-L agent. The uncertainty that remains is entirely due to adverse selection.

Favorable evidence continues to play an informational role because its arrival reveals that the project is good. In its absence, the DM's posterior is affected by two opposing forces. The failure of favorable evidence to arrive puts downward pressure on her belief, whereas the gradual exit of type-L agents puts upward pressure on it. In the uncertainty-free limit, these forces can exactly offset each other indefinitely. The equilibrium quitting and approval hazards keep the DM's posterior constant at a stationary level, while the type-H agent's posterior remains equal to one.

Consequently, an equilibrium consisting solely of an A-phase can persist indefinitely. Unlike the case with residual uncertainty, the type-H agent's belief does not deteriorate over time. His participation therefore imposes no finite terminal date, and no transition to an E-phase is required. This stationary infinite-horizon equilibrium is the distinctive feature of the uncertainty-free limit and, as shown below, yields the highest payoff to the DM.

Proposition 5 *Consider the uncertainty-free limit case: $\epsilon = 0$. Let*

$$\bar{p} = \frac{\lambda - c_{DM}}{\lambda}.$$

For each $s \in (0, \infty]$, a pure attrition equilibrium that satisfies the following conditions exists if and only if $\lambda \geq c_H - c_L$ and $p_0 \leq p_0^+(s)$.

· The type-H agent enters at time zero, and the DM initiates the review after observing entry, $d_{DM} = d_H = 0$, and the L-type quits at $t = 0$ with probability

$$d_L = \frac{p_0^+(s) - p_0}{p_0^+(s)(1 - p_0)}.$$

· For almost all $t < s$,

$$\xi_{DM}(t) = c_L, \quad \xi_L(t) = \frac{c_{DM}}{1 - p_t}, \quad \text{and} \quad \xi_H(t) = 0.$$

· The DM disapproves at s ; $\Upsilon(s) = 0$.

When $s = \infty$, the equilibrium is stationary: $p_t = \bar{p}$ for all $t > 0$, and there is no terminal disapproval. This stationary equilibrium is the DM-optimal pure attrition equilibrium.

Notice that whenever a finite-horizon pure attrition equilibrium exists, the stationary infinite-horizon equilibrium also exists. Indeed, $p_0^+(s) < \bar{p}$ for every finite s . The stationary equilibrium can be obtained as the limit of the finite-horizon equilibria as $s \rightarrow \infty$. Moreover, it yields the DM a higher payoff than any finite-horizon pure attrition equilibrium and is therefore the DM-optimal pure attrition equilibrium.

Along this stationary equilibrium path, the absence of a signal puts downward pressure on the DM's posterior belief that an active agent is type H, while the gradual exit of type-L agents puts upward pressure on it. These two effects exactly offset each other, leaving the posterior belief constant over time. The stationary belief $\bar{p}$ is determined by setting $\frac{dp_t}{dt} = 0$ in (6).

If there is a finite terminal time in this setting, the stationary equilibrium construction is no longer possible and the belief path is forced to drift down to the DM's indifference threshold before the terminal time. This allows us to strengthen our previous results: in the uncertainty-free limit, the DM's expected payoff is strictly lower in the finite-horizon model than in the infinite-horizon counterpart.

Corollary 5 *In the uncertainty-free limit case, suppose that $\lambda \geq c_H - c_L$ and $p_0 < \bar{p}$. Then the DM's expected payoff is strictly lower under a finite horizon than under an infinite horizon.*

Along the stationary equilibrium path in the uncertainty-free limit, the DM's posterior remains at $\bar{p} > 1/2$, so she always prefers approval to rejection, even in the absence of favorable evidence. This suggests a form of commitment that preserves type screening without requiring rejection at the terminal date: terminal recognition. Under this rule, an agent who remains active until a specified date is approved even if no favorable evidence has arrived.

The type-L agent knows that his project cannot generate favorable evidence. If the deadline is sufficiently distant, the cost of waiting for guaranteed approval exceeds its benefit, inducing him to exit immediately. The type-H agent, by contrast, may obtain approval earlier when favorable evidence arrives and is guaranteed approval at the deadline if it does not. The rule can therefore separate the two types at the outset while preserving the high type's incentive to participate. Although this idea is suggested by the uncertainty-free limit, it also applies when ϵ is small: the type-L agent's incentive is unchanged, while only a small fraction of type-H projects approved at the deadline are bad.

Terminal recognition is, of course, a theoretical benchmark: as discussed in Section 4.1, a rule permitting approval without favorable evidence may be difficult to justify or implement in practice. Abstracting from these concerns, consider its effect in the uncertainty-free limit. To induce the type-L agent to exit immediately, the deadline t_D

must make his payoff from remaining active non-positive; thus $t_D \geq 1/c_L$. The DM therefore chooses the earliest deadline consistent with immediate exit, $t_D = 1/c_L$. At this deadline, the type-L agent is indifferent between exiting and remaining active; immediate exit can be supported by the tie-breaking rule. Provided that $\lambda > c_H$, the type-H agent is willing to participate even though approval without favorable evidence is unavailable before the deadline. Terminal recognition thus induces immediate exit by the type-L agent while ensuring eventual approval of the type-H agent. The DM's ex ante payoff is then $p_0[1 - \frac{c_{DM}}{\lambda}(1 - e^{-\lambda/c_L})]$, which exceeds her payoff in the stationary equilibrium, $p_0 \frac{\lambda - 2c_{DM}}{\lambda - c_{DM}} = p_0[1 - \frac{c_{DM}}{\lambda - c_{DM}}]$: terminal recognition removes the type-L agent at the outset and caps the type-H agent's review time at t_D .

The presumption that the type-H agent participates relies on $\lambda > c_H$.¹⁵ Outside this parameter region, terminal recognition need not preserve his participation. When c_L is small relative to c_H , excluding the type-L agent requires a late deadline, $t_D = 1/c_L$. The type-H agent may then prefer to exit immediately rather than bear the expected cost of remaining active until a signal arrives or the deadline is reached. In that case, terminal recognition fails to achieve its intended screening outcome. This limitation may be relevant in settings where higher-quality projects face greater costs of continued evaluation, while lower-quality projects can remain in the process through relatively inexpensive procedural compliance. It offers one possible reason why terminal recognition is rarely used in practice.

5 Conclusion

This paper studies a dynamic screening environment in which information is revealed through stochastic evidence and endogenous dropout. The belief dynamics of this environment are qualitatively distinct from those in standard dynamic screening models: low-type exit and belief deterioration occur simultaneously in equilibrium, driven by the dominant informational content of signal absence under the Poisson technology.

The equilibrium characterization and the deadline analysis have concrete policy implications for regulatory design. In pharmaceutical approval, our analysis suggests that pressure to shorten review timelines, as embodied in PDUFA target dates (U.S. Food and Drug Administration, 2022), may be counterproductive when it operates without an accompanying commitment on the approval strategy. A regulator who shortens the review horizon without committing to a strict rejection rule at the deadline truncates the very process through which low-quality applications are endogenously weeded out, reducing the informativeness of the review and the regulator's expected payoff. By contrast, a credible commitment to reject absent evidence can restore and even improve upon the screening outcome, provided that signals arrive sufficiently fast relative to participation costs. The same logic extends beyond pharmaceutical regulation: safety certification of

¹⁵More precisely, it requires $\lambda > c_H(1 - e^{-\frac{\lambda}{c_L}})$. This requirement is weaker than $\lambda > c_H$: the latter is sufficient, but not necessary. This is because terminal recognition guarantees the type-H agent approval at the deadline even if no favorable evidence has arrived, thereby increasing his continuation value and relaxing his participation constraint.

artificial intelligence systems, venture capital staging, and probationary employment evaluation all involve analogous trade-offs between the length of the evaluation horizon and the purity of the resulting screening.

A natural limitation of the present analysis is that the DM is modeled as a payoff-maximizing entity unconstrained by internal organizational pressures. An important empirical regularity documented by Carpenter, Zucker, and Avorn (2008) and Carpenter, Chattopadhyay, Moffitt, and Nall (2012) points to a distinct mechanism: PDUFA deadlines induce a piling of drug approvals just before the deadline, and these approvals are associated with significantly higher rates of post-market safety problems. While our model explains how deadlines can undermine screening quality through the truncation of endogenous low-type exit, it does not account for this piling phenomenon, which suggests an additional channel whereby deadline pressure distorts decision-making on the regulator's side. Incorporating such agency frictions, whereby the regulator's own incentives lead to premature approval rather than optimal stopping under deadline pressure, into a dynamic screening framework represents a natural and empirically motivated direction for future research.

Appendix A: Perfect Bayesian Equilibrium

Given the DM's tie-breaking rule $\Upsilon(t)$ at $p_t = \frac{1}{2}$, let $U_{DM}(p)$ and $U_A(p)$ denote the DM's and the agent's expected payoffs, respectively, evaluated at the moment the DM takes an action, net of each player's sunk costs;

$$U_{DM}(p_t) = \begin{cases} 2p_t - 1 & \text{if } p_t \geq \frac{1}{2}, \\ 0 & \text{if } p_t < \frac{1}{2} \end{cases}, \quad \text{and } U_A(p_t) = \begin{cases} 1 & \text{if } p_t > \frac{1}{2}, \\ \Upsilon(t) & \text{if } p_t = \frac{1}{2}, \\ 0 & \text{if } p_t < \frac{1}{2} \end{cases}.$$

Let τ_s denote the stopping time at which a signal arrives, with the convention that $\tau_s = \infty$ if no signal arrives. Given the players' strategies σ_j and σ_{DM} and the realized agent type j , let τ_A and τ_{DM} denote the stopping times induced by σ_j and σ_{DM} , respectively, and define the stopping time of the game by

$$\tau = \min\{\tau_s, \tau_A, \tau_{DM}\},$$

which is a random variable.

For each player $i \in \{DM, H, L\}$, let $V_i(t)$ denote player i 's continuation value at t .

A tuple $(\sigma_L, \sigma_H, \sigma_{DM}, \{p_t\}_{t \in T}, V_H, V_L, V_{DM})$ is a Perfect Bayesian Equilibrium if and only if the following conditions are satisfied.

1. Dynamic optimality:

Given $(\sigma_L, \sigma_H, \sigma_{DM}, \{p_t\}_{t \in T})$, the continuation values of players are continuous on $[0, \bar{t}]$, where $\bar{t}$ is as defined in the main text, and satisfy the following conditions: For

almost all $t \in (0, \bar{t}]$, the following value-consistency conditions hold for each player;¹⁶

$$\begin{aligned} V_{DM}(t) &= \mathbb{E}_\tau [G_{DM}(\tau) - (\tau - t) c_{DM} | \tau \geq t], \\ V_H(t) &= \mathbb{E}_\tau [G_H(\tau) - (\tau - t) c_H | \tau \geq t], \\ \text{and } V_L(t) &= \mathbb{E}_\tau [G_L(\tau) - (\tau - t) c_L | \tau \geq t], \end{aligned}$$

where

$$G_{DM}(\tau) = \begin{cases} U_{DM}(1) & \text{if } \tau = \tau_s, \\ 0 & \text{if } \tau = \tau_A, \\ U_{DM}(p_\tau) & \text{if } \tau = \tau_{DM}, \end{cases}$$

and similarly for $G_H(\tau)$ and $G_L(\tau)$, which are players' terminal payoffs at stopping time τ , as determined by the event that ends the game and the corresponding terminal belief.

For almost all $t \in (0, \bar{t})$, the players' continuation values satisfy the Hamilton–Jacobi–Bellman (HJB) equations¹⁷:

$$\frac{dV_{DM}(t)}{dt} = c_{DM} - p_t \lambda (1 - V_{DM}(t)) + \left[\frac{p_t}{r_t} \xi_H(t) + (1 - p_t) \xi_B(t) \right] V_{DM}(t), \quad (13)$$

where $\xi_B(t)$ is set to be zero and $p_t = r_t$ for all $t > s$. The term $\xi_{DM}(t)[U_{DM}(p_t) - V_{DM}(t)]$ associated with the DM's own stopping is omitted from (13), since it vanishes under the value-matching condition whenever $\xi_{DM}(t) > 0$.

$$\frac{dV_H(t)}{dt} = c_H - \lambda r_t (1 - V_H(t)) + \xi_{DM}(t)[V_H(t) - U_A(p_t)], \quad (14)$$

$$\frac{dV_L(t)}{dt} = c_L + \xi_{DM}(t)(V_L(t) - U_A(p_t)). \quad (15)$$

For almost all $t < \bar{t}$, the following value-matching conditions hold for each player;

$$\begin{aligned} V_{DM}(t) &= V_H(t) = V_L(t) = 1 \text{ for all } t \text{ such that } p_t = 1, \\ V_{DM}(t) &\geq U_{DM}(p_t) \text{ with equality if } \xi_{DM}(t) > 0, \\ V_j(t) &\geq 0 \text{ with equality if } \xi_j(t) > 0 \text{ for } j \in \{H, L\}. \end{aligned}$$

Similarly;

$$\begin{aligned} d_j &> 0 \text{ only if } V_j(0) = 0 \text{ for } j \in \{H, L\}, \\ \text{and } d_{DM} &> 0 \text{ only if } V_{DM}(0) \leq \max\{\lim_{t \rightarrow 0} 2p_t - 1, 0\}. \end{aligned}$$

¹⁶Thus, in this formulation, past effort costs are sunk.

¹⁷These are derived from the recursive forms:

$$V_{DM}(t) = -c_{DM}dt + \lambda p_t dt + [1 - (p_t \lambda + (p_t/r_t) \xi_H(t) + (1 - p_t) \xi_B(t))dt] V_{DM}(t + dt, p_{t+dt}),$$

$$V_H(t) = -c_H dt + \lambda r_t dt + \xi_{DM}(t) U_A(p_t) dt + (1 - \lambda r_t dt - \xi_{DM}(t) dt) V_H(t + dt),$$

and

$$V_L(t) = -c_L dt + \xi_{DM}(t) U_A(p_t) dt + (1 - \xi_{DM}(t) dt) V_L(t + dt).$$

2. Bayes consistency:

$$\lim_{t \rightarrow 0} p_t = \frac{p_0(1 - d_H)}{\frac{p_0}{(1-\epsilon)}(1 - d_H) + (1 - \frac{p_0}{(1-\epsilon)})(1 - d_L)}$$

and for all $t \in (0, \bar{t})$, (3) holds.

Beliefs and strategies off the equilibrium path are arbitrary as long as they do not affect on-path optimality.

Appendix B: Proofs

In the following proofs, statements like “on (t_1, t_2) ” are to be read in the almost-everywhere sense.

Lemma 1 *In an equilibrium, $\xi_H(t) = 0$ for almost all t on $\mathcal{A}$.*

Proof. Suppose, to the contrary, that the claim fails. Then there exists an open interval $(t_1, t_2) \subset \mathcal{A}$ such that $\xi_H(t) > 0$ for $t \in (t_1, t_2)$. By value matching, $V_H(t) = 0$ throughout this interval and thus $\frac{dV_H(t)}{dt} = 0$ on (t_1, t_2) . The H-type’s HJB equation (14) therefore implies

$$c_H = \lambda r_t + \xi_{DM}(t) U_A(p_t). \quad (16)$$

Because r_t is strictly decreasing, this equality requires $\xi_{DM}(t) > 0$.

Now suppose that $V_L(t) = 0$ on some nondegenerate subinterval of (t_1, t_2) . Using (14) and (15), we obtain

$$\frac{d}{dt}(V_H(t) - V_L(t)) = c_H - c_L - \lambda r_t = 0,$$

which is impossible because r_t is not constant on any nondegenerate interval. Thus $V_L(t) > 0$ and $\xi_L(t) = 0$ on (t_1, t_2) .

Because $\xi_H(t) > 0$ while $\xi_L(t) = 0$, p_t is strictly decreasing on (t_1, t_2) . We first show that $p_t > \frac{1}{2}$ on (t_1, t_2) . If $p_t < \frac{1}{2}$ on a subset of positive measure, then $U_A(p_t) = 0$ there, and (16) reduces to $c_H = \lambda r_t$, which is impossible because r_t is strictly decreasing. Hence $p_t \geq \frac{1}{2}$ almost everywhere on (t_1, t_2) , and since p_t is continuous and strictly decreasing, $p_t > \frac{1}{2}$ on (t_1, t_2) . In particular, $r_t \geq p_t > \frac{1}{2}$.

Since $\xi_{DM}(t) > 0$, value matching gives $V_{DM}(t) = U_{DM}(p_t) = 2p_t - 1$ on (t_1, t_2) , so that $1 - V_{DM}(t) = 2(1 - p_t)$. Substituting this into (13) with $\xi_B(t) = 0$ yields

$$\frac{dV_{DM}(t)}{dt} = c_{DM} - 2\lambda p_t(1 - p_t) + \frac{p_t}{r_t} \xi_H(t) (2p_t - 1).$$

On the other hand, by (3) with $\xi_L(t) = 0$,

$$\frac{dV_{DM}(t)}{dt} = 2 \frac{dp_t}{dt} = -2\lambda p_t(1 - p_t) - \frac{2p_t(r_t - p_t)}{r_t} \xi_H(t).$$

Equating the two expressions, the terms involving λ cancel, and we obtain

$$\frac{p_t}{r_t} \xi_H(t) (1 - 2r_t) = c_{DM}.$$

The left-hand side is strictly negative because $r_t > \frac{1}{2}$ and $\xi_H(t) > 0$, whereas the right-hand side is strictly positive, a contradiction. We therefore conclude that $\xi_H(t) = 0$ on $\mathcal{A}$. ■

Lemma 2 *In an equilibrium, $p_t \geq \frac{1}{2}$ on $\mathcal{A}$.*

Proof. Suppose, toward a contradiction, that there exists some $t' \in \mathcal{A}$ such that $p_{t'} < \frac{1}{2}$.

We first show that there must exist some $t'' > t'$ such that $p_{t''} = \frac{1}{2}$. If this were not the case, then $p_t < \frac{1}{2}$ for all $t \geq t'$. In this case, $U_A(p_t) = 0$ for all $t \geq t'$, and the continuation value of the L -type agent at time t' is given by the expected payoff from waiting until the stopping time τ ,

$$V_L(t') = \mathbb{E}_\tau [U_A(p_\tau) - (\tau - t') c_L | \tau \geq t'] = -\mathbb{E}_\tau [(\tau - t') c_L | \tau \geq t'] < 0.$$

This contradicts the value matching condition $V_L(t') \geq 0$. Hence, there must exist some $k > t'$ such that $p_k = \frac{1}{2}$. Let k be the smallest such value; thus $p_t < \frac{1}{2}$ and $U_A(p_t) = 0$ for all $t \in [t', k)$.

Since p_t is continuous, there exists an interval $[t_1, t_2]$ contained in $(t', k]$ with $t_1 < t_2$ such that $\frac{dp_t}{dt} > 0$ for all $t \in [t_1, t_2]$. This implies that $\xi_L(t) > 0$ for all $t \in [t_1, t_2]$, which implies that $V_L(t) = 0$ on $[t_1, t_2]$.

From (15) and $U_A(p_t) = 0$, it follows that $\frac{dV_L(t)}{dt} = c_L > 0$ for $t \in [t_1, t_2]$. This contradicts $V_L(t) = 0$ for all $t \in [t_1, t_2]$. We therefore conclude that $p_t \geq \frac{1}{2}$ on $\mathcal{A}$. ■

Lemma 3 *In an equilibrium, $\xi_{DM}(t) = c_L$ and $\xi_B(t) = \frac{c_{DM}}{1-p_t}$ for almost all t on $\mathcal{A}$.*

Proof. We first show that $\xi_{DM}(t) > 0$ for almost all t on $\mathcal{A}$. Suppose, toward a contradiction, that there exists an open interval $(t_1, t_2) \subset \mathcal{A}$ such that $\xi_{DM}(t) = 0$ on (t_1, t_2) .

From value matching, if $\xi_L(t) > 0$ for some $t \in (t_1, t_2)$, then $V_L(t) = 0$. On the other hand, from the HJB for the L -type agent, we have $\frac{dV_L(t)}{dt} = c_L > 0$ for all $t \in (t_1, t_2)$. It then follows that, for any $t \in (t_1, t_2)$,

$$V_L(t) = c_L(t - t_1) + V_L(t_1) > V_L(t_1) \geq 0,$$

which contradicts $V_L(t) = 0$. Therefore, $\xi_L(t) = 0$ on (t_1, t_2) .

Next, let $t_3 = \inf\{t : \int_{t_1}^t \xi_{DM}(u) du > 0\} \geq t_2$. From the above argument, $\xi_L(t) = 0$ on (t_1, t_3) . Because $\xi_j(t) = 0$ on (t_1, t_3) for $j \in \{DM, H, L\}$, in (t_1, t_3) the game stops only when the signal arrives.

Then, by the martingale property of Bayesian beliefs, waiting over (t_1, t_3) yields the DM $2p_{t_1} - 1$ in expectation net of the delay cost incurred, so that $V_{DM}(t_1) < U_{DM}(p_{t_1})$, which is a contradiction. Therefore we may conclude that $\xi_{DM}(t) > 0$ on $\mathcal{A}$. This further implies value matching for the DM, so that $V_{DM}(t) = 2p_t - 1$ on $\mathcal{A}$. Substituting this expression into the HJB equation for the DM (13) yields $\xi_B(t) = \frac{c_{DM}}{1-p_t}$ on $\mathcal{A}$.

By a symmetric argument, $\xi_B(t) > 0$ for almost all t implies $\xi_{DM}(t) = c_L$ on $\mathcal{A}$. ■

Lemma 4 Fix $\bar{t} \leq t_E$, and define $D(s) = W_{DM}(s; \bar{t}) - (2r_s - 1)$ for $s < \bar{t}$.

- (i) $W_{DM}(s; \bar{t}) > 0$ for all $s < \bar{t}$.
- (ii) At any s with $D(s) = 0$, $D'(s) = c_{DM} > 0$. Hence D has at most one zero; if $D(\hat{s}) = 0$, then $D(s) < 0$ for all $s < \hat{s}$ and $D(s) > 0$ for all $s \in (\hat{s}, \bar{t})$.
- (iii) If $W_{DM}(s'; \bar{t}) > \max\{2r_{s'} - 1, 0\}$ for some $s' < \bar{t}$, then $W_{DM}(s; \bar{t}) > \max\{2r_s - 1, 0\}$ for all $s \in (s', \bar{t})$.

Proof. (i) The continuation value satisfies

$$\frac{\partial W_{DM}(s; \bar{t})}{\partial s} = c_{DM} - \lambda r_s (1 - W_{DM}(s; \bar{t})) \quad (17)$$

with terminal condition $W_{DM}(\bar{t}; \bar{t}) = 0$. Solving backward gives

$$W_{DM}(s; \bar{t}) = \int_s^{\bar{t}} \exp\left(-\int_s^u \lambda r_v dv\right) (\lambda r_u - c_{DM}) du.$$

Since r_t is decreasing and $\bar{t} \leq t_E$, we have $r_u \geq r_{\bar{t}} \geq r_{t_E} = c_H/\lambda \geq c_{DM}/\lambda$ for all $u \in [s, \bar{t}]$. The integrand is therefore strictly positive on a set of positive measure and nonnegative elsewhere, so $W_{DM}(s; \bar{t}) > 0$.

(ii) Using (17) and $\frac{dr_s}{ds} = -\lambda r_s(1 - r_s)$, at any s with $D(s) = 0$ we have $1 - W_{DM}(s; \bar{t}) = 2(1 - r_s)$, and hence

$$D'(s) = c_{DM} - 2\lambda r_s(1 - r_s) + 2\lambda r_s(1 - r_s) = c_{DM} > 0. \quad (18)$$

Let $D(\hat{s}) = 0$. By (18), $D > 0$ on some interval $(\hat{s}, \hat{s} + \delta)$. Suppose $D(u) \leq 0$ for some $u \in (\hat{s}, \bar{t})$, and let $u^* = \inf\{u > \hat{s} : D(u) \leq 0\}$. By continuity, $D(u^*) = 0$ and $D > 0$ on $(\hat{s}, u^*)$, which contradicts $D'(u^*) > 0$. Hence $D > 0$ on $(\hat{s}, \bar{t})$. A symmetric argument shows $D < 0$ for all $s < \hat{s}$.

(iii) By (i), it suffices to show $D(s) > 0$ for all $s \in (s', \bar{t})$. Since $D(s') > 0$, (ii) implies that D has no zero after s' : any zero $\hat{s} > s'$ would require $D(s') < 0$. By continuity, $D(s) > 0$ for all $s \in (s', \bar{t})$. ■

Lemma 5 In an equilibrium, $p_t < \bar{p} \equiv \frac{\lambda - c_{DM}}{\lambda}$ and p_t is strictly decreasing on A .

Proof. By Lemma 3, p_t follows (6) on A , which can be rewritten as

$$\frac{dp_t}{dt} = -\lambda p_t (\bar{p} - p_t). \quad (19)$$

Since $p_t \geq \frac{1}{2} > 0$ on A by Lemma 2, (19) implies that $\frac{dp_t}{dt} < 0$ if and only if $p_t < \bar{p}$. It therefore suffices to show that $p_t < \bar{p}$ on A .

Let $p_s = \lim_{t \uparrow s} p_t$. The belief path on A is given by (7), in which

$$\frac{\lambda - c_{DM}}{p_s} - \lambda = \frac{\lambda(\bar{p} - p_s)}{p_s}.$$

Hence, if $p_s < \bar{p}$, the denominator of (7) exceeds λ , and $p_t < \frac{\lambda - c_{DM}}{\lambda} = \bar{p}$ for all $t < s$. We now show $p_s < \bar{p}$. The A-phase ends in one of two ways.

Case 1: The review terminates with disapproval at s . Then $p_s = \frac{1}{2}$, and Assumption 1 ($\lambda > 2c_{DM}$) is equivalent to $\bar{p} > \frac{1}{2}$. Hence $p_s < \bar{p}$.

Case 2: The A-phase is followed by an E-phase that terminates at $\bar{t}$. Then $p_s = r_s$ and the connecting condition (10) holds, so that $D(s) = 0$, where $D(u) = W_{DM}(u; \bar{t}) - (2r_u - 1)$. Let $k = c_{DM}/\lambda$ and $q = r_{\bar{t}}$. Two bounds on q hold. First, $q \leq \frac{1}{2}$, since the DM disapproves at $\bar{t}$. Second, $q \geq c_H/\lambda \geq k$, since $\bar{t} \leq t_E$ and $c_H \geq c_{DM}$. By (1), $r_u e^{-\lambda(\bar{t}-u)} = (1 - r_u) \frac{q}{1-q}$, so

$$W_{DM}(u; \bar{t}) = (1 - k) \frac{r_u - q}{1 - q} - k(1 - r_u)\lambda(\bar{t} - u). \quad (20)$$

If $r_0 \leq \bar{p}$, then $r_s < r_0 \leq \bar{p}$ because r_t is strictly decreasing and $s > 0$.

Suppose instead that $r_0 > \bar{p}$. Since r_t is continuous and strictly decreasing, and $q \leq \frac{1}{2} < \bar{p}$, there exists a unique $\sigma \in (0, \bar{t})$ such that $r_\sigma = \bar{p} = 1 - k$. Substituting into (20) yields

$$D(\sigma) = \frac{k(k - q)}{1 - q} - k^2\lambda(\bar{t} - \sigma) < 0,$$

where the inequality follows from $k \leq q$, $k > 0$, and $\sigma < \bar{t}$. By Lemma 4(ii), D is negative before its unique zero and positive after it. Since $D(\sigma) < 0$ and $D(s) = 0$, it follows that $s > \sigma$, and hence $r_s < r_\sigma = \bar{p}$.

In both cases $p_s < \bar{p}$, which completes the proof. ■

Lemma 6 *Let $\pi_H = p_0/(1 - \epsilon)$. In any pure attrition or sequential screening equilibrium with terminal time $\bar{t}$, the DM's ex ante payoff equals*

$$\pi_H W_{DM}(0; \bar{t}) - c_{DM} \Lambda_L, \quad \Lambda_L = (1 - \pi_H)(1 - d_L) \int_0^{\bar{t}} \exp\left(-\int_0^t \xi_L(v) dv\right) dt,$$

where Λ_L is the ex ante expected time the type-L agent spends under review. Moreover, $\Lambda_L > 0$.

Proof. On $(0, \bar{t})$, $\xi_H(t) = 0$: on $\mathcal{A}$ by Lemma 1, and on $\mathcal{E}$ by construction. Hence the DM's continuation value V_{DM} satisfies (13) with $\xi_H = 0$ on $(0, \bar{t})$. This is a linear ODE in V_{DM} whose coefficients do not involve ξ_{DM} . Its terminal condition is $V_{DM}(\bar{t}) = 0$, because the DM disapproves at $\bar{t}$ and $p_{\bar{t}} \leq \frac{1}{2}$.

Consider the plan of never stopping before $\bar{t}$ and disapproving at $\bar{t}$, holding the agents' equilibrium strategies fixed. Its value satisfies the same linear ODE with the same terminal condition, and hence coincides with V_{DM} on $(0, \bar{t})$. Under this plan, the project is approved

if and only if $\omega = G$ and favorable evidence arrives before $\bar{t}$; the type-H agent remains under review until the signal arrives or $\bar{t}$ is reached, and the type-L agent until he exits or $\bar{t}$ is reached. Taking expectations over the agent's type and entry decision, the DM's ex ante payoff is therefore

$$\pi_H \left\{ (1 - \epsilon)(1 - e^{-\lambda \bar{t}}) - c_{DM} \left[\frac{1 - \epsilon}{\lambda} (1 - e^{-\lambda \bar{t}}) + \epsilon \bar{t} \right] \right\} - c_{DM} \Lambda_L = \pi_H W_{DM}(0; \bar{t}) - c_{DM} \Lambda_L.$$

Finally, since $p_t < r_t$ on $\mathcal{A}$, the post-entry belief is strictly below $r_0 = 1 - \epsilon$. By (2) with $d_H = 0$, this implies $d_L < 1$. Since $\pi_H < 1$ and the integrand is positive near $t = 0$, $\Lambda_L > 0$. ■

Proof of Proposition 1 and Corollary 1. To prove that the prescribed strategies are supported as an equilibrium, suppose that $s \leq \bar{s}$ and $p_0 \leq p_0^+(s)$. Let;

$$\begin{aligned} V_L(t) &= 0, \quad V_{DM}(t) = 2p_t - 1, \\ \text{and } V_H(t) &= r_t \left(1 - \frac{c_H}{\lambda + c_L} \right) (1 - e^{-(\lambda + c_L)(s-t)}) + (1 - r_t) \left(1 - \frac{c_H}{c_L} \right) (1 - e^{-c_L(s-t)}), \end{aligned}$$

for $t \in (0, s)$, where r_t and p_t follow (1) and (6), respectively. Then, it is verified that the value-matching conditions, $V_{DM}(t) = U_{DM}(p_t)$, $V_L(t) = 0$, and $V_H(t) > 0$ for all t on $\mathcal{A}$, are satisfied.

Conversely, if $s > t^*$, then $r_s < \frac{1}{2} = p_s$, which contradicts $p_t \leq r_t$. Suppose next that $s > t_A$, so that $\lambda r_s < c_H - c_L$. Under the prescribed strategies, type H's continuation value is

$$V_H(t) = \int_t^s \exp \left(- \int_t^u (\lambda r_v + c_L) dv \right) (\lambda r_u + c_L - c_H) du,$$

where the integrand is decreasing in u because r_u is decreasing. By continuity of r_t , there exists $\delta > 0$ such that $\lambda r_u + c_L - c_H < 0$ for all $u \in (s - \delta, s)$, and hence $V_H(t) < 0$ for all $t \in (s - \delta, s)$. This violates the value-matching condition $V_H(t) \geq 0$: type H would strictly prefer to quit, contradicting $\xi_H(t) = 0$.

Suppose next that $p_0 > p_0^+(s)$. Since $p_s = \frac{1}{2}$ and the belief path on the A-phase is uniquely determined by (6), the post-entry belief must equal $p_0^+(s)$. With $d_H = 0$, Bayes consistency (2) then requires

$$d_L = \frac{p_0^+(s) - p_0}{p_0^+(s) \left(1 - \frac{p_0}{1 - \epsilon} \right)} < 0,$$

which is infeasible. Hence, if either condition is violated, the equilibrium construction described in the proposition is impossible.

To see that the equilibrium with $s = \bar{s}$ payoff-dominates other equilibria, observe that the DM's equilibrium payoff is $\frac{p_0}{p_0^+(s)} (2p_0^+(s) - 1)$, as it is optimal for her to immediately approve after the initial screening, which occurs with probability $\frac{p_0}{p_0^+(s)}$. This value is increasing in $p_0^+(s)$. As $p_0^+(s)$ is also increasing in s and s cannot exceed $\bar{s}$, the equilibrium with $s = \bar{s}$ attains the highest payoff.

To establish the claims in Corollary 1, substitute (8) into $2p_0 - p_0/p_0^+(\bar{s})$ to obtain

$$V_{DM}(0) = p_0 \frac{\lambda - 2c_{DM}}{\lambda - c_{DM}} [1 - e^{-(\lambda - c_{DM})\bar{s}}].$$

The first factor is positive by Assumption 1, increasing in λ , and decreasing in c_{DM} , and the bracket is increasing in $(\lambda - c_{DM})\bar{s}$. Both t^* and t_A are independent of c_{DM} and decreasing in ϵ . Moreover, $\lambda t^* = \log \frac{1-\epsilon}{\epsilon}$ and $\lambda t_A = \log \frac{1-\epsilon}{\epsilon} + \log \frac{\lambda - (c_H - c_L)}{c_H - c_L}$ are nondecreasing in λ , so that $(\lambda - c_{DM})\bar{s} = (1 - c_{DM}/\lambda) \lambda \bar{s}$ is increasing in λ . Finally, t^* does not depend on c_H or c_L , whereas t_A is decreasing in $c_H - c_L$. The claims follow. ■

Proof of Proposition 2 and Corollary 2. Suppose that $W_{DM}(0; \bar{t}) \geq 2r_0 - 1$ and $\lambda \geq 2c_H$. Consider the following strategy profile; $d_L = 1$, $d_H = 0$, $\xi_{DM}(t) = \xi_H(t) = 0$ for $t < \bar{t}$ and $\xi_{DM}(\bar{t}) = \infty$.

We have $W_{DM}(t; \bar{t}) > \max\{2r_t - 1, 0\}$ for all $t \in (0, \bar{t})$. If $W_{DM}(0; \bar{t}) > 2r_0 - 1$, this follows from Lemma 4(iii). If $W_{DM}(0; \bar{t}) = 2r_0 - 1$, Lemma 4(ii) implies $W_{DM}(t; \bar{t}) > 2r_t - 1$ for all $t \in (0, \bar{t})$, and Lemma 4(i) implies $W_{DM}(t; \bar{t}) > 0$.

Then, by simply letting $V_H(t) = W_H(t; \bar{t})$, $V_{DM}(t) = W_{DM}(t; \bar{t})$, and $V_L(t) = 0$ for all $t \leq \bar{t}$, we can verify that the equilibrium conditions are satisfied.

We now prove that the DM-optimal pure experimentation equilibrium has $\bar{t} = t_E$. At $t = 0$,

$$W_{DM}(0; \bar{t}) = (1 - \epsilon)(1 - e^{-\lambda \bar{t}}) - c_{DM} \left[\frac{1 - \epsilon}{\lambda} (1 - e^{-\lambda \bar{t}}) + \epsilon \bar{t} \right],$$

and hence

$$\frac{\partial W_{DM}(0; \bar{t})}{\partial \bar{t}} = (1 - \epsilon)(\lambda - c_{DM})e^{-\lambda \bar{t}} - \epsilon c_{DM} = [\epsilon + (1 - \epsilon)e^{-\lambda \bar{t}}](\lambda r_t - c_{DM}).$$

The DM's unconstrained optimal stopping time t_{DM} satisfies $\lambda r_{t_{DM}} = c_{DM}$. Since $c_H \geq c_{DM}$, we have $t_E \leq t_{DM}$. It follows that

$$\frac{\partial W_{DM}(0; \bar{t})}{\partial \bar{t}} > 0 \text{ for all } \bar{t} < t_E.$$

Thus, among all pure experimentation equilibria, the DM prefers the largest sustainable terminal time. Therefore, provided that $W_{DM}(0; t_E) \geq 2r_0 - 1$, the DM-optimal pure experimentation equilibrium has $\bar{t} = t_E$.

This also shows that if $W_{DM}(0; t_E) < 2r_0 - 1$, then $W_{DM}(0; \bar{t}) < 2r_0 - 1$ for all choices of $\bar{t}$; thus there is no pure experimentation equilibrium.

To establish the claims in Corollary 2, observe that, by using (9), the DM's expected payoff can be written as

$$W_{DM}(0; t_E) = \frac{(\lambda - c_{DM})((1 - \epsilon)\lambda - c_H)}{\lambda(\lambda - c_H)} - \frac{\epsilon c_{DM}}{\lambda} \log \left(\frac{(1 - \epsilon)(\lambda - c_H)}{\epsilon c_H} \right),$$

from which the claims follow. The DM's ex ante payoff is $\pi_H W_{DM}(0; t_E)$ with $\pi_H = p_0/(1 - \epsilon)$. Since π_H does not depend on λ , c_{DM} , or c_H , the claims regarding these

parameters follow from the expression above. For ϵ , let $k = c_{DM}/\lambda$, $h = c_H/\lambda$, and $u = \epsilon/(1 - \epsilon)$, which is increasing in ϵ . Using $e^{-\lambda t_E} = \frac{h}{1-h}u$ and $\lambda t_E = \log \frac{1-h}{hu}$, we obtain

$$\pi_H W_{DM}(0; t_E) = p_0 \left[(1 - k) \left(1 - \frac{h}{1-h}u \right) - ku \log \frac{1-h}{hu} \right],$$

so that

$$\frac{\partial}{\partial u} \pi_H W_{DM}(0; t_E) = p_0 \left[k - \frac{(1-k)h}{1-h} - k \log \frac{1-h}{hu} \right] < 0,$$

because $(1-k)h/(1-h) \geq h \geq k$ by $k \leq h$, and $\log \frac{1-h}{hu} = \lambda t_E > 0$. ■

Proof of Proposition 3 and Corollary 3. If $\bar{t} = t^*$, then $W_{DM}(t^*; t^*) = 0 = 2r_{t^*} - 1$, so the connecting condition holds at $s(\bar{t}) = t^*$; the E-phase is degenerate, and the construction reduces to the pure attrition equilibrium of Proposition 1 with $s = t^*$. Suppose therefore that $\bar{t} > t^*$. From Lemma 4(i), $W_{DM}(t^*; \bar{t}) > 0 = 2r_{t^*} - 1$. Then $W_{DM}(0; \bar{t}) < 2r_0 - 1$ and the continuity of $W_{DM}(s; \bar{t})$ in s ensures that there is $s' \in (0, t^*)$ such that $W_{DM}(s'; \bar{t}) = 2r_{s'} - 1$. From Lemma 4(ii), this s' is uniquely determined, and we denote it by $s(\bar{t})$. Therefore, under $W_{DM}(0; \bar{t}) < 2r_0 - 1$ and $\lambda \geq 2c_H$, the value $p_0^{++}(\bar{t})$ is defined.

To prove that the prescribed strategies are supported as an equilibrium, suppose that $W_{DM}(0; \bar{t}) < 2r_0 - 1$, $\lambda \geq 2c_H$, and $p_0 \leq p_0^{++}(\bar{t})$. For $t \in \mathcal{A}$, let;

$$\begin{aligned} V_L(t) &= 0, \quad V_{DM}(t) = 2p_t - 1, \\ \text{and } V_H(t) &= r_t \left(1 - \frac{c_H}{\lambda + c_L} \right) (1 - e^{-(\lambda + c_L)(s(\bar{t}) - t)}) + (1 - r_t) \left(1 - \frac{c_H}{c_L} \right) (1 - e^{-c_L(s(\bar{t}) - t)}) \\ &\quad + [r_t e^{-(\lambda + c_L)(s(\bar{t}) - t)} + (1 - r_t) e^{-c_L(s(\bar{t}) - t)}] W_H(s(\bar{t}); \bar{t}). \end{aligned}$$

For $t \in \mathcal{E}$,

$$V_L(t) = 0, \quad V_H(t) = W_H(t; \bar{t}) \text{ and } V_{DM}(t) = W_{DM}(t; \bar{t}).$$

As $V_L(t) = 0$ and $V_H(t) > 0$ for all $t < \bar{t}$, the agent's strategies are optimal. By construction, $V_{DM}(t)$ is continuous on $[0, \bar{t}]$ and $V_{DM}(s(\bar{t})) = W_{DM}(s(\bar{t}); \bar{t}) = 2r_{s(\bar{t})} - 1$ and $V_{DM}(t) = W_{DM}(t; \bar{t}) > \max\{2r_t - 1, 0\}$ for all $t \in \mathcal{E}$. Thus the DM's strategy is also optimal.

Next, we show that, among all sequential screening equilibria, the equilibrium characterized by $(s, \bar{t}) = (t_L, t_E)$ yields the highest payoff to the DM. We first establish that, among sequential screening equilibria, the DM prefers the one that has the earlier transition time. For a given transition time s , the A-phase belief path is pinned down by the terminal condition $p_s = r_s$ and the law of motion (6).

Consider two transition times $s_1 < s_2$, each of which corresponds to an equilibrium, and denote the corresponding A-phase belief paths by p^1 and p^2 . Since the equilibrium associated with s_2 remains in the A-phase at s_1 ,

$$p_{s_1}^2 < r_{s_1} = p_{s_1}^1.$$

By uniqueness of the solution to the belief ODE, the two belief paths cannot cross. Hence, $p_t^1 > p_t^2$ for all $t < s_1$. In particular, $\lim_{t \downarrow 0} p_t^1 > \lim_{t \downarrow 0} p_t^2$. Since the DM's equilibrium payoff is $2p_0 - \frac{p_0}{\lim_{t \downarrow 0} p_t}$, the DM strictly prefers a smaller transition time s .

It remains to show that the equilibrium that supports the earliest transition time is the one with $\bar{t} = t_E$. Define

$$F(s) = W_{DM}(s; t_E) - (2r_s - 1).$$

From Lemma 4, $F(t_L) = 0$ and $F(s) < 0$ for all $s < t_L$.

Now consider any feasible terminal time $\bar{t}$. By the expression for $\partial W_{DM}(s; \bar{t}) / \partial \bar{t}$ in footnote 12, $W_{DM}(s; \bar{t}) \leq W_{DM}(s; t_E)$ for every s and every feasible $\bar{t}$.

Suppose, contrary to the claim, that some sequential screening equilibrium has a transition time $s < t_L$. Its connecting condition would imply $W_{DM}(s; \bar{t}) = 2r_s - 1$.

Together with $W_{DM}(s; \bar{t}) \leq W_{DM}(s; t_E)$, this gives $W_{DM}(s; t_E) \geq 2r_s - 1$, or equivalently, $F(s) \geq 0$. However, this contradicts $F(s) < 0$ for all $s < t_L$. Therefore, no sequential screening equilibrium has a transition time earlier than t_L . Hence, the DM-optimal sequential screening equilibrium is characterized by $(s, t) = (t_L, t_E)$, as long as it exists.

To see that Corollary 3 holds, we first consider the effect of ϵ .

Let $\Delta_E = t_E - t_L$. Since $r_{t_E} = c_H / \lambda$, (1) implies

$$r_{t_L} = \frac{c_H e^{\lambda \Delta_E}}{\lambda - c_H + c_H e^{\lambda \Delta_E}}.$$

When expressed in terms of Δ_E , the connecting condition $W_{DM}(t_L; t_E) = 2r_{t_L} - 1$ is written as

$$\left(1 - \frac{c_{DM}}{\lambda}\right) \frac{r_{t_L} - c_H / \lambda}{1 - c_H / \lambda} - c_{DM}(1 - r_{t_L})\Delta_E = 2r_{t_L} - 1,$$

which is independent of ϵ . Thus Δ_E is independent of ϵ . Hence, it follows that

$$\frac{\partial r_{t_L}}{\partial \epsilon} = 0 \text{ and } \frac{\partial t_L}{\partial \epsilon} = \frac{\partial t_E}{\partial \epsilon} = -\frac{1}{\lambda \epsilon (1 - \epsilon)} < 0. \quad (21)$$

From (7),

$$p_0^{++} = \frac{\lambda - c_{DM}}{\lambda + \left(\frac{\lambda - c_{DM}}{r_{t_L}} - \lambda\right) e^{-(\lambda - c_{DM})t_L}}.$$

Since

$$r_{t_L} < \frac{\lambda - c_{DM}}{\lambda},$$

by Lemma 5, this expression is increasing in t_L . Therefore, from (21), p_0^{++} , and hence the DM's expected payoff is decreasing in ϵ .

We next consider the effect of λ . We see that the transition time is decreasing in λ . For this purpose, let

$$r_L = r_{t_L}, \quad h = \frac{c_H}{\lambda}, \quad \text{and} \quad k \equiv \frac{c_{DM}}{\lambda}.$$

Recall that the H-type's posterior satisfies

$$\frac{r_t}{1 - r_t} = \frac{1 - \epsilon}{\epsilon} e^{-\lambda t}.$$

Hence,

$$\frac{r_L/(1 - r_L)}{r_{t_E}/(1 - r_{t_E})} = e^{\lambda(t_E - t_L)}.$$

Since $r_{t_E} = c_H/\lambda = h$, it follows that

$$\lambda(t_E - t_L) = \log \left(\frac{r_L(1 - h)}{h(1 - r_L)} \right).$$

Using this expression, we can write the DM's continuation payoff at t_L as

$$W_{DM}(t_L; t_E) = (1 - k) \frac{r_L - h}{1 - h} - k(1 - r_L) \log \left(\frac{r_L(1 - h)}{h(1 - r_L)} \right).$$

Therefore, the connecting condition $W_{DM}(t_L; t_E) = 2r_{t_L} - 1$ can be written as

$$\Gamma(r_L, h, k) = 0,$$

where

$$\Gamma(r, h, k) = (1 - k) \frac{r - h}{1 - h} - k(1 - r) \log \left(\frac{r(1 - h)}{h(1 - r)} \right) - (2r - 1).$$

At the connecting point, the argument used in the proof of Lemma 4(ii) gives $\frac{\partial F(s)}{\partial s}|_{s=t_L} = c_{DM} > 0$. Since $\frac{dr_t}{dt} = -\lambda r_t(1 - r_t)$, it follows that

$$\Gamma_r(r_L, h, k) = -\frac{k}{r_L(1 - r_L)} < 0.$$

Moreover,

$$\Gamma_k(r, h, k) = -\frac{r - h}{1 - h} - (1 - r) \log \left(\frac{r(1 - h)}{h(1 - r)} \right) < 0,$$

because $r_L > h$, and

$$\Gamma_h(r, h, k) = -\frac{(1 - r)(1 - k/h)}{(1 - h)^2} \leq 0,$$

where the last inequality follows from $c_{DM} \leq c_H$, and hence $k \leq h$.

Because

$$\frac{dk}{d\lambda} = -\frac{k}{\lambda} \text{ and } \frac{dh}{d\lambda} = -\frac{h}{\lambda},$$

differentiation of $\Gamma(r_L, h, k) = 0$ gives

$$\Gamma_r \frac{dr_L}{d\lambda} + \Gamma_h \frac{dh}{d\lambda} + \Gamma_k \frac{dk}{d\lambda} = 0.$$

Consequently,

$$\frac{dr_L}{d\lambda} = \frac{h\Gamma_h + k\Gamma_k}{\lambda\Gamma_r} > 0.$$

Finally, the H-type's posterior formula implies

$$t_L = \frac{1}{\lambda} \log \left(\frac{(1-\epsilon)(1-r_L)}{\epsilon r_L} \right). \quad (22)$$

Differentiating with respect to λ yields

$$\frac{dt_L}{d\lambda} = -\frac{1}{\lambda^2} \log \left(\frac{(1-\epsilon)(1-r_L)}{\epsilon r_L} \right) - \frac{1}{\lambda r_L(1-r_L)} \frac{dr_L}{d\lambda}.$$

Since $t_L > 0$ and $dr_L/d\lambda > 0$, both terms on the right-hand side are negative. Therefore,

$$\frac{dt_L}{d\lambda} < 0.$$

Substituting (22) into (7), and writing $x = r_L$ and $y = \lambda t_L = \log \frac{(1-\epsilon)(1-x)}{\epsilon x} > 0$, we obtain

$$p_0^{++} = \frac{1-k}{1+T}, \quad T = \frac{1-k-x}{x} e^{-(1-k)y} = \frac{1-k-x}{x} \left(\frac{\epsilon x}{(1-\epsilon)(1-x)} \right)^{1-k},$$

where $1-k-x > 0$ by Lemma 5. Since

$$\frac{\partial \log T}{\partial x} = -\frac{1}{1-k-x} - \frac{k}{x} + \frac{1-k}{1-x} < 0,$$

p_0^{++} is increasing in $x = r_L$. Moreover, holding x fixed and letting $a = 1-k-x$ and $z = e^{-(1-k)y} \in (0, 1)$,

$$\frac{\partial \log p_0^{++}}{\partial k} = -\frac{1}{1-k} + \frac{T}{1+T} \left(\frac{1}{a} - y \right) = -\frac{1}{1-k} + \frac{z(1-ay)}{x+az} < 0,$$

since $z(1-ay)(1-k) < z(x+a) \leq x+az$. Hence p_0^{++} is decreasing in $k = c_{DM}/\lambda$. Since $dr_L/d\lambda > 0$ and $dk/d\lambda < 0$, the DM's expected payoff is increasing in λ .

For c_{DM} and c_H , implicit differentiation of $\Gamma(r_L, h, k) = 0$ gives $dr_L/dk = -\Gamma_k/\Gamma_r < 0$ and $dr_L/dh = -\Gamma_h/\Gamma_r \leq 0$, with strict inequality when $c_H > c_{DM}$. Since p_0^{++} is increasing in r_L and decreasing in k , the DM's expected payoff is decreasing in c_{DM} and in c_H . ■

Proof of Theorem 1. Let $\pi_H = p_0/(1-\epsilon)$. By the characterization in Section 3, any equilibrium in which the review is initiated is a pure attrition, pure experimentation, or sequential screening equilibrium. We first rank these equilibria and the outcome without review.

Step 1 (pure experimentation versus the other forms). Suppose that $\lambda \geq 2c_H$, so that $\bar{s} = t^* \leq t_E$. Every pure attrition equilibrium has a terminal time $s \leq \bar{s} \leq t_E$, and every

sequential screening equilibrium has a terminal time $\bar{t} \leq t_E$. By Lemma 6, the DM's payoff in any such equilibrium with terminal time $\bar{t}$ is

$$\pi_H W_{DM}(0; \bar{t}) - c_{DM} \Lambda_L < \pi_H W_{DM}(0; \bar{t}) \leq \pi_H W_{DM}(0; t_E),$$

where the last inequality holds because $\partial W_{DM}(0; \bar{t}) / \partial \bar{t} > 0$ for $\bar{t} < t_E$ (see the proof of Proposition 2). The right-hand side is the DM's payoff in the pure experimentation equilibrium with $\bar{t} = t_E$.

Step 2 (sequential screening versus pure attrition). Suppose that $\lambda > 2c_H$ and that the sequential screening equilibrium with $(s, \bar{t}) = (t_L, t_E)$ exists. Since $W_{DM}(t^*; t_E) > 0 = 2r_{t^*} - 1$ by Lemma 4(i), part (ii) of the same lemma implies $t_L < t^*$, and hence $r_{t_L} > \frac{1}{2}$. Let p^S denote the belief path of this equilibrium on its A-phase, and let p^A denote the belief path of a pure attrition equilibrium with terminal time $s \leq \bar{s} = t^*$. If $s > t_L$, then $p_{t_L}^A < r_{t_L} = p_{t_L}^S$. If $s \leq t_L$, then $p_s^A = \frac{1}{2} < r_{t_L} \leq p_s^S$, since p^S is decreasing on its A-phase (Lemma 5). In either case, by uniqueness of the solution to (6), the two paths cannot cross, so that $p_0^{++}(t_E) > p_0^+(s)$. Since the DM's payoff in both equilibria is $2p_0 - p_0 / \lim_{t \downarrow 0} p_t$, the sequential screening equilibrium yields a strictly higher payoff. When $\lambda = 2c_H$, the two equilibria coincide (Section 3.3).

Step 3 (comparison with no review). Without review, the DM obtains $\max\{2p_0 - 1, 0\}$. In a pure attrition or sequential screening equilibrium, her payoff is $2p_0 - p_0 / \hat{p}$ with $\hat{p} = \lim_{t \downarrow 0} p_t$. Since $\hat{p} \geq p_0$ and $\hat{p} > \frac{1}{2}$, this is at least $2p_0 - 1$ and strictly positive. In a pure experimentation equilibrium, her payoff is $\pi_H W_{DM}(0; t_E) \geq \pi_H(2r_0 - 1) = 2p_0 - \pi_H \geq 2p_0 - 1$, and it is strictly positive by Lemma 4(i). Hence any equilibrium with review weakly dominates the outcome without review.

Part 1. Suppose that $\lambda \geq 2c_H$. As noted after the theorem, (11) is equivalent to $W_{DM}(0; t_E) < 2r_0 - 1$. If (11) is reversed, the pure experimentation equilibrium with $\bar{t} = t_E$ exists by Proposition 2 and is DM-optimal by Steps 1 and 3. If (11) holds, no pure experimentation equilibrium exists (proof of Proposition 2). If, in addition, $p_0 \leq p_0^{++}(t_E)$, the sequential screening equilibrium with $(s, \bar{t}) = (t_L, t_E)$ exists by Proposition 3 and is DM-optimal by Proposition 3 and Steps 2 and 3.

Part 2. Suppose that $\lambda < 2c_H$. Pure experimentation and sequential screening equilibria require $\lambda \geq 2c_H$ (Propositions 2 and 3), so only pure attrition equilibria can exist. By Proposition 1, the DM-optimal one has $s = \bar{s}$ and exists if and only if $\bar{s} > 0$ and $p_0 \leq p_0^+(\bar{s})$. Here $\bar{s} = t^*$ if $\lambda \geq 2(c_H - c_L)$ and $\bar{s} = t_A$ otherwise, and $\bar{s} > 0$ if and only if $(1 - \epsilon)\lambda > c_H - c_L$, that is, $\lambda > (c_H - c_L)/r_0$. Optimality follows from Step 3.

Part 3. If the conditions in Part 2 fail, no equilibrium with review exists. Under the conditions of Part 1, suppose that (11) holds and $p_0 > p_0^{++}(t_E)$. Then no pure experimentation equilibrium exists. A sequential screening equilibrium with terminal time $\bar{t}$ requires $p_0 \leq p_0^+(\bar{t}) \leq p_0^{++}(t_E)$, where the second inequality holds because $s(\bar{t}) \geq t_L$ and an earlier transition yields a higher post-entry belief (proof of Proposition 3). A pure attrition equilibrium requires $p_0 \leq p_0^+(s) < p_0^{++}(t_E)$ by Step 2. Hence no equilibrium with review exists, and the DM decides on the basis of her prior. ■

Proof of Corollary 4. Consider the game in which the review must end by $\hat{t}$, at which the DM approves if $p_{\hat{t}} > \frac{1}{2}$ and disapproves if $p_{\hat{t}} < \frac{1}{2}$. The arguments of Lemmas 1–5 do

not depend on the length of the horizon and apply on $(0, \hat{t})$. In particular, an A-phase cannot end at $\hat{t}$ with approval with positive probability, for the same reason that the DM must disapprove at the end of a pure attrition equilibrium, and an E-phase must end with disapproval, which requires $r_{\bar{t}} \leq \frac{1}{2}$. Hence every equilibrium of the game with the deadline in which the review is initiated is an equilibrium characterized in Section 3 with terminal time $\bar{t} \leq \hat{t}$. If no such equilibrium exists, the DM obtains the payoff without review, which by Step 3 of the proof of Theorem 1 is no higher than her payoff in the DM-optimal equilibrium without a deadline. Moreover, the loss is strict whenever $\hat{t}$ is smaller than the terminal time of the DM-optimal equilibrium without a deadline and some equilibrium with review exists under the deadline, by the strict monotonicity of the DM's payoff in the terminal time established in Propositions 1–3 and the ranking in the proof of Theorem 1. ■

Proof of Proposition 4. Let $\pi_H = p_0/(1 - \epsilon)$. Under a commitment to reject at τ in the absence of favorable evidence, the type-L agent can never be approved and, since $c_L > 0$, exits at time zero. The type-H agent remains active until τ if and only if $\lambda r_t \geq c_H$ for all $t \leq \tau$, that is, $\tau \leq t_E$. If $(1 - \epsilon)\lambda \leq c_H$, so that $t_E = 0$, he does not participate. Since $\partial W_{DM}(0; \tau)/\partial \tau > 0$ for $\tau < t_E$ (see the proof of Proposition 2), the DM's payoff under commitment is

$$U^C = \pi_H W_{DM}(0; t_E),$$

with the convention $W_{DM}(0; 0) = 0$.

(i) In the DM-optimal pure experimentation equilibrium, the type-L agent exits at time zero and the review ends at t_E , so the DM's payoff is $\pi_H W_{DM}(0; t_E) = U^C$.

(ii) By Lemma 6, the DM's payoff in the DM-optimal sequential screening equilibrium is $\pi_H W_{DM}(0; t_E) - c_{DM} \Lambda_L < U^C$, since $\Lambda_L > 0$.

(iii)

In the DM-optimal pure attrition equilibrium, the terminal time is $\bar{s}$. By Lemma 6, the DM's ex ante equilibrium payoff is

$$U^A = \pi_H W_{DM}(0; \bar{s}) - c_{DM} \Lambda_L,$$

where Λ_L is the ex ante expected time the type-L agent spends under review. Hence $U^A - U^C = \pi_H [W_{DM}(0; \bar{s}) - W_{DM}(0; t_E)] - c_{DM} \Lambda_L$, which proves the equivalence with (12).

Since pure attrition is DM-optimal only if $\lambda < 2c_H$ (Theorem 1), we have $t_E < t^*$; since $c_L > 0$, we also have $t_E < t_A$. Hence $t_E < \bar{s}$. Neither $\bar{s}$ nor t_E depends on c_{DM} , $W_{DM}(0; t)$ is continuous in c_{DM} , and $0 \leq \Lambda_L \leq (1 - \pi_H)\bar{s}$, where $\bar{s} \leq t^* < \infty$. As $c_{DM} \rightarrow 0$, the right-hand side of (12) converges to zero, whereas the left-hand side converges to $\pi_H(1 - \epsilon)(e^{-\lambda t_E} - e^{-\lambda \bar{s}}) > 0$. Hence (12) holds for all sufficiently small c_{DM} .

Suppose next that $(1 - \epsilon)\lambda > c_H$ and $p_0 \leq \frac{1}{2}$. Then $t_E > 0$, and since $(1 - \epsilon)\lambda > c_H > c_H - c_L$, also $t_A > 0$, so that $\bar{s} > 0$. By (8) and Assumption 1, $p_0^+(s) > \frac{1}{2}$ for every $s > 0$; hence $p_0^+(\bar{s}) > \frac{1}{2} \geq p_0$, and the DM-optimal pure attrition equilibrium exists for every $c_{DM} < \lambda/2$. By Corollary 1,

$$U^A = p_0 \frac{\lambda - 2c_{DM}}{\lambda - c_{DM}} [1 - e^{-(\lambda - c_{DM})\bar{s}}] \rightarrow 0 \quad \text{as } c_{DM} \rightarrow \lambda/2.$$

On the other hand, Lemma 4(i) implies $W_{DM}(0; t_E) > 0$ for every $c_{DM} \leq \lambda/2$, since $\lambda/2 < c_H$. By continuity, $U^C = \pi_H W_{DM}(0; t_E)$ is bounded away from zero as $c_{DM} \rightarrow \lambda/2$. Hence $U^C > U^A$ for c_{DM} sufficiently close to $\lambda/2$. ■

Proof of Proposition 5 and Corollary 5. When $\epsilon = 0$, $r_t = 1$ for all t , so that $t^* = \infty$, and type H's net flow payoff from continuing in the A-phase, $\lambda + c_L - c_H$, is constant. For finite s , the construction in the proof of Proposition 1 applies with $r_t = 1$: the belief path is given by (7) with $p_s = \frac{1}{2}$, $\xi_L(t) = c_{DM}/(1 - p_t)$, type H is willing to remain until s if and only if $\lambda + c_L - c_H \geq 0$, and $d_L \geq 0$ if and only if $p_0 \leq p_0^+(s)$.

For $s = \infty$, let $p_t = \bar{p}$ for all $t > 0$, which solves (6). Then $\xi_L(t) = c_{DM}/(1 - \bar{p}) = \lambda$. Substituting $V_{DM} = 2\bar{p} - 1$ and $\xi_B = c_{DM}/(1 - \bar{p})$ into (13) gives $dV_{DM}/dt = 2\bar{p}[c_{DM} - \lambda(1 - \bar{p})] = 0$; $V_L = 0$ satisfies (15) with $\xi_{DM} = c_L$; and $V_H = (\lambda + c_L - c_H)/(\lambda + c_L) \geq 0$ is the stationary solution of (14). By (8), $p_0^+(s) \uparrow \bar{p}$ as $s \rightarrow \infty$, so Bayes consistency requires $p_0 \leq \bar{p} = p_0^+(\infty)$. Since the DM's payoff $2p_0 - p_0/p_0^+(s)$ is increasing in $p_0^+(s)$ and $p_0^+(s) < \bar{p}$ for every finite s , the stationary equilibrium is the DM-optimal pure attrition equilibrium.

For Corollary 5, let the horizon be $T < \infty$. Since $r_t = 1 > \frac{1}{2}$, an E-phase cannot end with disapproval, so no pure experimentation or sequential screening equilibrium exists. By the argument in the proof of Corollary 4, every equilibrium with review is a pure attrition equilibrium with $s \leq T$, and its payoff is at most $2p_0 - p_0/p_0^+(T) < 2p_0 - p_0/\bar{p}$. If no equilibrium with review exists, the DM obtains $\max\{2p_0 - 1, 0\}$, which is strictly below $2p_0 - p_0/\bar{p}$ because $\frac{1}{2} < \bar{p}$ and $p_0 < \bar{p}$. ■

References

- [1] Bar-Isaac, Heski, “Reputation and Survival: Learning in a Dynamic Signalling Model,” *Review of Economic Studies*, 2003, 70 (2), 231–251.
- [2] Bobtcheff, Catherine and Raphael Levy, “More Haste, Less Speed? Signaling through Investment Timing,” *American Economic Journal: Microeconomics*, 2017, 9 (3), 148–186.
- [3] Bonatti, Alessandro and Johannes Hörner, “Career Concerns with Exponential Learning,” *Theoretical Economics*, 2017, 12 (1), 425–475.
- [4] Carpenter, Daniel P, Evan James Zucker, and Jerry Avorn, “Drug-Review Deadlines and Safety Problems,” *New England Journal of Medicine*, 2008, 358(13), 1354–1361.
- [5] Carpenter, Daniel P, Jacqueline Chattopadhyay, Susan Moffitt, and Clayton Nall, “The Complications of Controlling Agency Time Discretion: FDA Review Deadlines and Postmarket Drug Safety,” *American Journal of Political Science*, 2012, 56 (1), 98–114.
- [6] Chen, Chia-Hui, Junichi Ishida, and Wing Suen, “Reputation Concerns in Risky Experimentation,” *Journal of the European Economic Association*, 2021, 19 (4), 1981–2021.

- [7] Courty, Pascal and Hao Li, “Sequential Screening,” *Review of Economic Studies*, 2000, 67 (4), 697–717.
- [8] Daley, Brendan and Brett Green, “Waiting for News in the Market for Lemons,” *Econometrica*, 2012, 80 (4), 1433–1504.
- [9] DiMasi, Joseph A, Henry G. Grabowski, and Ronald W. Hansen, “Innovation in the Pharmaceutical Industry: New Estimates of R&D Costs,” *Journal of Health Economics*, 2016, 47, 20–33.
- [10] Dong, Miaomiao, “Strategic Experimentation with Asymmetric Information,” 2025, SSRN Working Paper.
- [11] European Medicines Agency, “The Evaluation of Medicines, Step-by-Step,” 2019.
- [12] Guo, Yingni, “Dynamic Delegation of Experimentation,” *American Economic Review*, 2016, 106 (8), 1969–2008.
- [13] Halac, Marina and Ilan Kremer, “Experimenting with Career Concerns,” *American Economic Journal: Microeconomics*, 2020, 12 (1), 260–288.
- [14] Henry, Emeric and Marco Ottaviani, “Research and the Approval Process: The Organization of Persuasion,” *American Economic Review*, 2019, 109 (3), 911–955.
- [15] Honryo, Takakazu, “Dynamic Persuasion,” *Journal of Economic Theory*, 2018, 178, 36–58.
- [16] Keller, Godfrey, Sven Rady, and Martin Cripps, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, 73 (1), 39–68.
- [17] Kreps, David, and Robert Wilson, “Reputation and Imperfect Information,” *Journal of Economic Theory*, 1982, 27 (2), 253–279.
- [18] Lee, Jihong and Qingmin Liu, “Gambling Reputation: Repeated Bargaining with Outside Options,” *Econometrica*, 2013, 81 (4), 1601–1672.
- [19] McClellan, Andrew, “Experimentation and Approval Mechanisms,” *Econometrica*, 2022, 90 (5), 2215–2247.
- [20] Ordoover, Janusz A and Ariel Rubinstein, “A Sequential Concession Game with Asymmetric Information,” *Quarterly Journal of Economics*, 1986, 101 (4), 879–888.
- [21] U.S. Food and Drug Administration, “PDUFA Reauthorization Performance Goals and Procedures Fiscal Years 2023 Through 2027,” 2022. <https://www.fda.gov/media/151712/download>.

Declaration of generative AI and AI-assisted technologies:

Statement: During the preparation of this work the author used ChatGPT and Claude in order to improve the readability and language of the manuscript. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.