

October 2025

“Advising with Threshold Tests:
Complexity, Signaling, and Effort”

Mark Izgarshev and Georgy Lukyanov

Advising with Threshold Tests: Complexity, Signaling, and Effort*

Mark Izgarshev

School number 1239, Vspolny per., 6, Moscow, Russia

Georgy Lukyanov

Toulouse School of Economics, 1, Esplanade de l'Université, Toulouse, France

Abstract

A benevolent advisor observes a project's complexity and posts a pass-fail threshold before the agent chooses effort. The project succeeds only if ability and effort together clear complexity. We compare two informational regimes. In the naive regime, the threshold is treated as non-informative; in the sophisticated regime, the threshold is a signal and the agent updates beliefs. We characterize equilibrium threshold policies and show that the optimal threshold rises with complexity under mild regularity. We then give primitives-based sufficient conditions that guarantee separating, pooling, or semi-separating outcomes. In a benchmark with uniform ability, exponential complexity, and power costs, we provide explicit parameter regions that partition the space by equilibrium type; a standard refinement eliminates most pooling. The results yield transparent comparative statics and welfare comparisons across regimes.

Keywords: threshold tests; signaling; information design; monotone comparative statics; pooling vs. separation.

*We are grateful to Emin Ablyatifov, Jacob Gershman, Arthur Izgarshev, David Li, Anastasia Makhmudova, and Yuliia Tukmakova for helpful comments and suggestions. All remaining errors are our own.

1 Introduction

Most institutions run on thresholds. A medical graduate either clears a licensing bar or does not; an airline first officer either meets a minimum flight-hours rule or sits in the right seat later; a firm’s promotion “bar” is raised or lowered with consequences for who advances and how hard people push. In each case a single cutoff both filters who proceeds and shapes what others learn about the task at hand. Recent policy shifts underscore how consequential such decisions are: the USMLE moved Step 1 to pass/fail¹ to redirect student effort and alter residency-selection incentives (USMLE Program, 2020); in commercial aviation, the post-2010 “1,500-hour rule”² codified in the FAA’s pilot-qualification reforms made the hour threshold itself the pivotal gate into airline cockpits (Federal Aviation Administration, 2013). Thresholds are simple to communicate and enforce—but they also broadcast information about underlying difficulty. When a bar is set high (or lowered), careful observers may infer something about the environment, not just the candidates.

This paper proposes a model to study that dual role. An advisor privately observes project complexity and, before any costly action, commits to a pass–fail threshold that only reveals whether the agent’s latent ability lies above the cutoff. Success requires that ability and effort together clear true complexity, so the test both partitions beliefs and shifts the marginal return to effort.³

We analyze two informational regimes that bracket many applications. In the naive regime, the agent treats the posted threshold as non-informative⁴ about complexity and conditions only on pass/fail. In the sophisticated regime, the threshold is itself a signal: the agent understands the advisor’s policy and updates beliefs about complexity from the cutoff. Our object is the

¹The stated motivations included reducing “score obsession,” rebalancing curricula, and mitigating wellness concerns. The reform also altered residency selection, where Step 1 had functioned as a single-index screen; see the discussion around (USMLE Program, 2020).

²Industry debates emphasize both safety benefits and pipeline costs. For our purposes, once the rule is codified the bar itself becomes an object of inference and incentive; see (Federal Aviation Administration, 2013) for the policy background.

³This dual role echoes education signaling: actions both sort types and affect payoffs. In our case the cutoff simultaneously filters and motivates.

⁴“naive” need not mean irrational. Institutions sometimes explicitly instruct candidates not to infer difficulty from posted rules (“don’t read into the bar”), or norms make such inference taboo. We model this as a benchmark that brackets the fully Bayesian case.

advisor’s threshold policy—the mapping from complexity to cutoffs—chosen to maximize the agent’s expected utility (a benevolent benchmark⁵ that isolates information-and-incentives effects).

We show that optimal thresholds are (weakly) increasing in complexity in both regimes. Formally, we establish increasing differences and apply monotone comparative statics to prove that any optimal selection is nondecreasing in complexity; for interior solutions the derivative is nonnegative (Theorems 1 and 2). Beyond monotonicity, we classify equilibrium threshold policies by informativeness—separating, pooling, or semi-separating—and give primitives-based sufficient conditions for each case that are stated directly in terms of the thickness of the ability distribution near the cutoff, the dispersion of complexity, the curvature of effort costs, and the value of success (Definition 1 and Theorem 3).

A key novelty is to make the *threshold itself* an equilibrium signal about the state of the world (complexity) in a setting where the receiver still chooses costly *post-signal effort*. This tight coupling of endogenous disclosure (via the bar) and subsequent incentives is distinct from standard disclosure or persuasion environments where the receiver’s action is the terminal choice, and from standard threshold/standard-setting models where the bar screens agents but does not reveal the state that determines the productivity of effort. Contemporaneously, Bertola (2025) studies equilibrium failure rates when institutions endogenize pass thresholds. His analysis focuses on how standards map into observed failure frequencies; by contrast, our advisor *observes* complexity and the threshold *signals* that state, which then feeds into the agent’s continuation effort. The two perspectives are complementary: our separating/pooling taxonomy explains when bar movements should be informative about the environment, while his failure-rate lens speaks to observable pass/fail frequencies.

In a benchmark with uniform ability, exponential complexity, and power costs,⁶ we derive explicit, graphable inequalities that compare a feasible separating policy to a conservative pooling benchmark. These yield ready-to-plot separating/pooling regions; with quadratic costs they are fully explicit via

⁵Equivalently, the designer is utilitarian over the agent’s ex-ante payoff. Replacing this with an institutional objective (e.g., weighted success minus effort externalities) preserves our monotonicity results under mild alignment.

⁶Uniform F keeps ability tails uncluttered; exponential G gives memorylessness and tractable posteriors; power costs include quadratic and provide interpretable elasticities of effort.

the Lambert W function, and for general power costs they collapse to a one-dimensional root. We then show how the boundary curves shift with the project’s value and cost curvature. Finally, imposing a standard refinement (the Intuitive Criterion) removes most pooling:⁷ a small threshold deviation can be profitable only for the type that benefits from separation, which forces beliefs that unravel pooling (Theorem 4).

The empirical upshot is straightforward: our sufficient conditions produce transparent diagnostics for when institutions should *separate* (let bars track complexity closely) versus *pool* (hold the line at a constant bar), and our refinement results explain why, in practice, pooling often unravels once stakeholders can infer complexity from observed bar moves (e.g., grading standards, safety minima, or promotion bars that “creep” with market conditions).

Related literature

We connect to (i) signaling/communication, (ii) information design and evidence disclosure, and (iii) monotone comparative statics.

Classic models show how monotone message partitions arise endogenously in education and communication (Spence, 1973; Crawford and Sobel, 1982). Our advisor’s threshold is a one-dimensional signal whose informativeness is determined in equilibrium. Because pooling can often be sustained by beliefs, we rely on the Intuitive Criterion (Cho and Kreps, 1987) and related D1/Divinity logic (Banks and Sobel, 1987) to discipline off-path beliefs in the threshold environment.

The pass–fail rule is a coarse scoring device, akin to the optimal coarse disclosure mechanisms in Rayo and Segal (2010) and to binary experiments in Kamenica and Gentzkow (2011). A closely related strand studies endogenous grading/standards. Boleslavsky and Cotton (2015) analyze how a school sets grading standards to trade off selection and incentives; their bar screens students and shapes effort but does not *signal* the state of the world. Bertola (2025) examines equilibrium failure rates under endogenous pass thresholds, emphasizing how standards translate into observed fail shares. In contrast, our threshold both sorts and *reveals* complexity, and the agent responds

⁷The Intuitive Criterion assigns off-path beliefs to the type that would strictly gain from deviating; here, a small threshold tweak that benefits exactly one T forces posteriors that unravel pooling. See Cho and Kreps (1987) and the related D1/Divinity idea in Banks and Sobel (1987).

with post-signal effort—this two-way interaction between disclosure and incentives is where our comparative statics and taxonomy bite. We also relate to evidence design and selective disclosure: Herresthal (2022) studies *hidden testing* and what evidence is disclosed ex post. Our advisor instead posts an observable bar ex ante; yet the common theme is how a designer chooses a coarse experiment that strategically shapes downstream behavior. Finally, links to certification and grading with endogenous thresholds (Chan et al., 2007; Ostrovsky and Schwarz, 2010) are direct: our results rationalize when coarser (pooling) certification survives refinements and when standards should track underlying difficulty.

We use single-crossing and lattice methods (Milgrom and Shannon, 1994; Topkis, 1998), extended to uncertainty/endogenous information via Athey (2001, 2002) and strict/interval dominance tools (Edlin and Shannon, 1998; Quah and Strulovici, 2009). Log-concavity and related hazard-rate conditions (Bagnoli and Bergstrom, 2005) ensure that threshold and effort best responses move monotonically with complexity and with posterior informativeness, delivering our separating/pooling regions.

Because policy thresholds in education, labor markets, and regulation have measurable behavioral effects, our model provides structure beneath RD-style estimates (Hahn et al., 2001; Imbens and Lemieux, 2008; Cattaneo et al., 2020). It clarifies when endogenously informative cutoffs amplify or attenuate post-threshold behavior and how cost curvature and the distributional shape of complexity govern whether institutions should separate or pool; and it complements the failure-rate perspective emphasized by Bertola (2025) by highlighting when bar movements themselves should be *interpreted* as signals of environmental difficulty.

The paper proceeds as follows. Section 2 lays out the environment and timing. Sections 3 and 4 develop the naive and signaling analyses and prove monotonicity. Section 5 derives the sufficient partitions. Section 6 applies the refinement. Section 7 provides extensions showing robustness to additive technologies, correlation, noisy tests, and testing costs, and quantify the value of observing complexity. Section 8 concludes. All the proofs are in A.

2 Model description

Environment There is a benevolent advisor and a single agent. If the project succeeds, the agent receives a benefit $V > 0$. Success depends on

three primitives: the agent's latent ability $\theta \in \mathbb{R}_+$, his effort $e \in \mathbb{R}_+$, and the project's complexity $T \in \mathbb{R}_+$. Technology is multiplicative:

$$\text{success} \iff \theta e \geq T.$$

Effort is costly. The cost function $C : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfies $C(0) = C'(0) = 0$, $C'(e) > 0$, and $C''(e) \geq 0$ for $e > 0$.

Uncertainty Ability θ is drawn from a commonly known distribution F with density f on \mathbb{R}_+ . Complexity T is drawn from a distribution G with density g on \mathbb{R}_+ . Unless stated otherwise, θ and T are independent.⁸ We will impose shape restrictions only when needed for existence and monotone comparative statics (Sections 3 and 4).

Test Before any effort choice, the advisor can post a pass-fail test⁹ that compares ability to a threshold $\theta^* \geq 0$. The test is costless and does not require effort. It returns an outcome $y \in \{\text{pass}, \text{fail}\}$ with

$$y = \text{pass} \iff \theta \geq \theta^*, \quad y = \text{fail} \iff \theta < \theta^*.$$

The threshold may carry information about T in equilibrium, since the advisor observes T and commits to the test.

Timing

1. Nature draws (θ, T) from $F \times G$. The advisor privately observes T ; the agent observes neither θ nor T .
2. The advisor commits to a threshold $\theta^* \in \Theta \subseteq \mathbb{R}_+$. A *policy* is a measurable mapping $\psi : \mathbb{R}_+ \rightarrow \Theta$, $T \mapsto \psi(T) = \theta^*(T)$.
3. The test is administered and $y \in \{\text{pass}, \text{fail}\}$ is publicly observed.
4. The agent chooses continuation effort $e_y \geq 0$; it can depend on the posted θ^* and on y .

⁸Positive assignment (hard tasks to strong agents) is common in hospitals and professional services. Our correlation extension shows the results persist under MLR-type dependence between θ and T .

⁹Pass/fail is the coarsest scoring rule. Institutions often choose coarse disclosure to manage incentives; see Rayo and Segal (2010). Our endogenous cutoff implements such coarse experiments.

5. Payoffs realize. The agent’s utility is $V \cdot \mathbb{1}\{\theta e_y \geq T\} - C(e_y)$. The advisor is benevolent and maximizes the agent’s expected utility.

Beliefs from the test Conditional on y , the agent knows whether θ lies above or below θ^* . Let the truncated densities be

$$f_{\text{pass}}(\theta \mid \theta^*) = \frac{f(\theta)}{1 - F(\theta^*)} \mathbb{1}\{\theta \geq \theta^*\}, \quad f_{\text{fail}}(\theta \mid \theta^*) = \frac{f(\theta)}{F(\theta^*)} \mathbb{1}\{\theta < \theta^*\}.$$

For complexity, the beliefs used at the effort stage depend on the informational regime defined next.

We analyze two cases:¹⁰

- *naive agent*. The agent treats θ^* as non-informative about T and uses the prior G when choosing e_y . Thus the CDF and density entering the effort problem are $\tilde{G}_{\theta^*} \equiv G$ and $\tilde{g}_{\theta^*} \equiv g$.
- *Sophisticated agent*. The agent is fully Bayesian and understands the advisor’s policy $\psi(\cdot)$. Upon observing an on-path threshold θ^* , he forms the posterior over T ,

$$\mu(t \mid \theta^*) = \frac{g(t) \mathbb{1}\{\psi(t) = \theta^*\}}{\int_{\tau: \psi(\tau) = \theta^*} g(\tau) d\tau},$$

with corresponding $\tilde{G}_{\theta^*}(x) = \int_0^x \mu(t \mid \theta^*) dt$ and density \tilde{g}_{θ^*} . Off-path beliefs are part of equilibrium; when we refine equilibria, we apply the Intuitive Criterion.

Agent’s continuation problem Given θ^* , an outcome $y \in \{\text{pass}, \text{fail}\}$, and regime-specific beliefs about T , the agent solves

$$\max_{e \geq 0} V \cdot \mathbb{E}_{\theta \mid y} [\tilde{G}_{\theta^*}(\theta e)] - C(e),$$

where the expectation is with respect to $f_{\text{pass}}(\cdot \mid \theta^*)$ or $f_{\text{fail}}(\cdot \mid \theta^*)$ depending on y . When interior, first-order conditions are

$$C'(e_{\text{pass}}(\theta^*)) = V \cdot \mathbb{E}_{\theta \sim f_{\text{pass}}(\cdot \mid \theta^*)} [\theta \tilde{g}_{\theta^*}(\theta e_{\text{pass}}(\theta^*))], \quad (1)$$

$$C'(e_{\text{fail}}(\theta^*)) = V \cdot \mathbb{E}_{\theta \sim f_{\text{fail}}(\cdot \mid \theta^*)} [\theta \tilde{g}_{\theta^*}(\theta e_{\text{fail}}(\theta^*))]. \quad (2)$$

¹⁰Think of the naive case as “policy is posted by committee; candidates are told not to infer difficulty,” and the sophisticated case as “everyone knows how the bar moves with conditions.” Both are realistic bracketing assumptions.

Under the regularity imposed below, $e_{\text{pass}}(\theta^*) \geq e_{\text{fail}}(\theta^*)$ (pass truncation raises expected θ) and each effort choice is unique if interior.

Advisor's problem Fix T (observed by the advisor). Anticipating $e_{\text{pass}}(\theta^*)$ and $e_{\text{fail}}(\theta^*)$, the advisor chooses the threshold to maximize the agent's expected utility:

$$U(T, \theta^*) = \int_{\theta^*}^{\infty} \left[V \cdot \mathbf{1}\{\theta e_{\text{pass}}(\theta^*) \geq T\} - C(e_{\text{pass}}(\theta^*)) \right] dF(\theta) \\ + \int_0^{\theta^*} \left[V \cdot \mathbf{1}\{\theta e_{\text{fail}}(\theta^*) \geq T\} - C(e_{\text{fail}}(\theta^*)) \right] dF(\theta). \quad (3)$$

A *threshold policy* is a measurable $\psi : \mathbb{R}_+ \rightarrow \Theta$, $T \mapsto \psi(T)$.

Equilibrium In the naive regime, an equilibrium is a pair $(\psi, \{e_{\text{pass}}, e_{\text{fail}}\})$ such that, for every T , $\psi(T) \in \arg \max_{\theta^*} U(T, \theta^*)$ given $\tilde{G}_{\theta^*} \equiv G$ and efforts solving (1)–(2).

In the sophisticated regime, a *Perfect Bayesian Equilibrium (PBE)* is a triple $(\psi, \{e_{\text{pass}}, e_{\text{fail}}\}, \mu)$ with:

1. *Agent optimality.* For every on-path $\theta^* = \psi(T)$, the efforts $e_{\text{pass}}(\theta^*), e_{\text{fail}}(\theta^*)$ solve (1)–(2) using \tilde{g}_{θ^*} from $\mu(\cdot \mid \theta^*)$.
2. *Advisor optimality.* For every T , $\psi(T) \in \arg \max_{\theta^*} U(T, \theta^*)$ given those efforts and beliefs.
3. *Belief consistency.* For on-path thresholds, $\mu(\cdot \mid \theta^*)$ is Bayes' rule under ψ ; off-path beliefs are specified, and when we refine equilibria they satisfy the Intuitive Criterion.

Regularity We will appeal to the following standing assumptions when proving existence and monotonicity.

Assumption 1. (i) f and g are continuous, strictly positive on compact subsets of \mathbb{R}_+ , and (weakly) log-concave. (ii) C is C^2 , strictly convex for $e > 0$, with C'' continuous and bounded away from 0 on compacts.

Under Assumption 1, the kernels in (1)–(2) are decreasing in e , so each effort is well defined and unique when interior; pass truncation implies $e_{\text{pass}} \geq$

e_{fail} . These properties, combined with the piecewise form of (3), yield increasing differences in (T, θ^*) and the monotonicity results stated in Theorems 1 and 2.

Benchmark parameterization When we specialize, we will use $\theta \sim \text{Unif}[0, 1]$, $T \sim \text{Exp}(\lambda)$, and $C(e) = e^\gamma$ with $\gamma > 1$. This delivers closed-form expressions for the sufficient inequalities that partition the parameter space by equilibrium type (B).

3 Equilibrium analysis: naive agent

The naive regime isolates the *selection* effect of pass/fail without the extra feedback loop from threshold-as-signal. This gives a clean “moving kink” structure in the advisor’s problem that we will reuse verbatim in the signaling case: only the distribution over T changes. The section therefore builds intuition for all that follows—how pass/fail truncation shifts expected ability, how the marginal benefit of raising θ^* behaves, and why monotone comparative statics fall out.

In the naive regime, the agent treats the posted threshold θ^* as non-informative about T and uses the prior G at the effort stage. Hence $\tilde{G}_{\theta^*} \equiv G$ and $\tilde{g}_{\theta^*} \equiv g$ in (1)–(2). The only information created by the test is the pass/fail truncation in ability.

Effort choices Given θ^* and $y \in \{\text{pass}, \text{fail}\}$, the agent solves

$$\max_{e \geq 0} V \cdot \mathbb{E}_{\theta|y}[G(\theta e)] - C(e),$$

so the interior first-order conditions specialize to

$$C'(e_{\text{pass}}(\theta^*)) = V \cdot \mathbb{E}_{\theta \sim f_{\text{pass}}(\cdot|\theta^*)}[\theta g(\theta e_{\text{pass}}(\theta^*))], \quad (4)$$

$$C'(e_{\text{fail}}(\theta^*)) = V \cdot \mathbb{E}_{\theta \sim f_{\text{fail}}(\cdot|\theta^*)}[\theta g(\theta e_{\text{fail}}(\theta^*))]. \quad (5)$$

Pass/fail does two things to the agent’s problem: it shifts the *distribution* of θ (pass FOSD dominates fail) and leaves the *kernel* $e \mapsto \theta g(\theta e)$ decreasing in e under log-concavity. The first drives $e_{\text{pass}} \geq e_{\text{fail}}$; the second guarantees uniqueness (when interior) and underpins all comparative statics by the implicit-function theorem.

Lemma 1. *Suppose Assumption 1 holds. For each threshold θ^* and each outcome $y \in \{\text{pass}, \text{fail}\}$, the agent’s continuation problem admits a maximizer $e_y(\theta^*)$. Whenever the solution is interior, it is unique and satisfies $e_{\text{pass}}(\theta^*) \geq e_{\text{fail}}(\theta^*)$. Moreover, $e_y(\theta^*)$ is (weakly) increasing in V and (weakly) decreasing under first-order stochastic dominance shifts of T toward greater complexity.*

Proof sketch. Strict convexity of C and continuity of the kernels $\theta \mapsto \theta g(\theta e)$ yield at most one interior solution; coercivity of C yields existence. Since f_{pass} first-order stochastically dominates f_{fail} and the right-hand side of (4) is increasing in the distribution of θ under FOSD, we obtain $e_{\text{pass}} \geq e_{\text{fail}}$. Monotonicity in V and under FOSD shifts of T follows by the implicit function theorem. \square

Pass/fail “zooms in” on different regions of ability. Conditioning on pass shifts mass to the right of θ^* , which raises the marginal return to effort in (4); since the right-hand side is decreasing in e , the unique solution must be weakly higher in the pass group. This is the basic selection–incentives complementarity that recurs throughout the paper.

Advisor’s problem Given T and anticipating $e_{\text{pass}}(\theta^*)$, $e_{\text{fail}}(\theta^*)$, the advisor chooses θ^* to maximize $U(T, \theta^*)$ in (3) with $\tilde{G}_{\theta^*} \equiv G$.

Holding fixed the (best-response) efforts $e_{\text{pass}}(\theta^*)$ and $e_{\text{fail}}(\theta^*)$, the advisor’s objective can be read as areas cut off by a right-moving kink at $\max\{\theta^*, T/e_{\text{pass}}(\theta^*)\}$ and $T/e_{\text{fail}}(\theta^*)$. As θ^* increases, those kinks slide right, which is exactly the source of increasing differences in (T, θ^*) .

Theorem 1. *Under Assumption 1, the advisor’s objective $U(T, \theta^*)$ has increasing differences in (T, θ^*) . Hence for every T the set of optimal thresholds is nonempty, and any optimal selection is (weakly) increasing in T . In particular, when the optimizer is interior, $d\theta^*/dT \geq 0$.*

Proof sketch. Write U as a sum of terms of the form $-F(\max\{\theta^*, T/e_{\text{pass}}(\theta^*)\})$ and $-F(T/e_{\text{fail}}(\theta^*))$, with $e_{\text{pass}}, e_{\text{fail}}$ depending only on θ^* . Each term has a kink that moves right with θ^* , generating increasing differences. Topkis then yields monotone optimal selections. \square

Harder environments mechanically call for (weakly) tougher bars. Empirically, any proxy for T that shifts right should be accompanied by nondecreasing θ^* if the institution is close to optimal in this benchmark.

The interior MCS result leaves open what happens near extremes of T . When T is tiny, success saturates and only costs matter near $\theta^* = 0$, creating a no-test corner; when T is huge, success is rare unless either the threshold collapses (expanding fail) or pass effort explodes—so the limiting choice is governed by the continuation-cost tradeoff. The next statement formalizes this split and pins down large- T limits through the cost aggregator $\mathcal{C}(\theta^*)$.

Proposition 1. *Let $U(T, \theta^*)$ be the advisor’s objective in the naive regime. There exists $T_{\text{small}} > 0$ such that $\theta^*(T) = 0$ is optimal for all $T \in [0, T_{\text{small}}]$.*

For every compact interval $[T_\ell, T_u] \subset (0, \infty)$ there exists $K < \infty$ such that, for each $T \in [T_\ell, T_u]$, the function $\theta^ \mapsto U(T, \theta^*)$ attains a maximizer on $[0, K]$. Moreover, for any $T \in (T_\ell, T_u)$ at which $U(T, \cdot)$ is differentiable at the maximizer, the optimizer is interior and satisfies the first-order condition; otherwise the maximizer occurs at a finite kink of $U(T, \cdot)$.*

Define the continuation cost function

$$\mathcal{C}(\theta^*) \equiv (1 - F(\theta^*)) C(e_{\text{pass}}(\theta^*)) + F(\theta^*) C(e_{\text{fail}}(\theta^*)),$$

and extend the policy space to $\bar{\Theta} = [0, \infty]$ by setting $U(T, \infty) := \lim_{\theta \rightarrow \infty} U(T, \theta)$ and $\mathcal{C}(\infty) := \lim_{\theta \rightarrow \infty} \mathcal{C}(\theta)$. Then, uniformly on compact sets of θ^ ,*

$$U(T, \theta^*) = -\mathcal{C}(\theta^*) + o(1) \quad (T \rightarrow \infty).$$

Consequently, any sequence of maximizers $\theta^(T) \in \arg \max_{\theta^* \in [0, \infty]} U(T, \theta^*)$ has limit points contained in*

$$\arg \min \mathcal{C} \subseteq [0, \infty].$$

If $\arg \min \mathcal{C} = \{\hat{\theta}\}$ is a singleton, then $\theta^(T) \rightarrow \hat{\theta}$ as $T \rightarrow \infty$. In particular:*

- (i) *If $\arg \min \mathcal{C} = \{0\}$, then $\theta^*(T) \rightarrow 0$.*
- (ii) *If $\arg \min \mathcal{C} = \{\infty\}$, then $\theta^*(T) \rightarrow \infty$.*
- (iii) *If $\arg \min \mathcal{C} = \{\hat{\theta}\} \subset (0, \infty)$, then $\theta^*(T) \rightarrow \hat{\theta}$.*

Remark. *When $\theta^* \rightarrow \infty$, the pass set is empty and everyone realizes $y = \text{fail}$. The induced effort $e_{\text{fail}}(\theta^*)$ converges to the effort computed under the full prior (since $F(\theta^*) \rightarrow 1$), so $U(T, \infty)$ is the limiting value with an “always fail” test. If one restricts $\Theta = [0, \bar{\theta}]$ with $\bar{\theta} < \infty$, case (ii) above becomes the boundary choice $\theta^*(T) = \bar{\theta}$.*

Corollary 1. *Let $\mathcal{C}(\theta^*) = (1 - F(\theta^*)) C(e_{\text{pass}}(\theta^*)) + F(\theta^*) C(e_{\text{fail}}(\theta^*))$ as in Proposition 1(iii). Suppose $e_{\text{pass}}(\theta^*)$ is (weakly) increasing and $e_{\text{fail}}(\theta^*)$ is (weakly) decreasing in θ^* . Then:*

1. *If there exists $\bar{\theta}$ such that $\theta^* \mapsto (1 - F(\theta^*)) C(e_{\text{pass}}(\theta^*))$ is strictly decreasing on $[\bar{\theta}, \infty)$ and $\theta^* \mapsto F(\theta^*) C(e_{\text{fail}}(\theta^*))$ is nondecreasing there, then \mathcal{C} is strictly decreasing on $[\bar{\theta}, \infty)$ and hence $\arg \min \mathcal{C} = \{\infty\}$. Consequently, $\theta^*(T) \rightarrow \infty$ as $T \rightarrow \infty$.*
2. *If there exists $\bar{\theta} > 0$ such that $\theta^* \mapsto (1 - F(\theta^*)) C(e_{\text{pass}}(\theta^*))$ is nondecreasing on $[0, \bar{\theta}]$ and $\theta^* \mapsto F(\theta^*) C(e_{\text{fail}}(\theta^*))$ is strictly decreasing there, then \mathcal{C} is strictly decreasing on $[0, \bar{\theta}]$ and hence $\arg \min \mathcal{C} = \{0\}$. Consequently, $\theta^*(T) \rightarrow 0$ as $T \rightarrow \infty$.*

In the benchmark $\theta \sim \text{Unif}[0, 1]$, $T \sim \text{Exp}(\lambda)$, and $C(e) = e^\gamma$, these monotonicities hold numerically across the parameter ranges we consider, and the case in (1) (shut-down) is the typical large- T limit.

Proof. By Proposition 1(iii), $U(T, \theta^*) = -\mathcal{C}(\theta^*) + o(1)$ uniformly on compacts as $T \rightarrow \infty$, so any limit point of maximizers must lie in $\arg \min \mathcal{C}$. In (1), \mathcal{C} is strictly decreasing on $[\bar{\theta}, \infty)$, hence minimized at the endpoint ∞ ; similarly in (2) it is strictly decreasing on $[0, \bar{\theta}]$, hence minimized at 0. The convergence of $\theta^*(T)$ follows because a unique minimizer forces all maximizers of $U(T, \cdot)$ to converge to it as $T \rightarrow \infty$. \square

Case (1) (“shut-down”) says that when the marginal cost of keeping a pass group is eventually dominated by its measure $1 - F(\theta^*)$, the optimal policy is to send everyone to the fail branch at very high T . In many licensing or safety screens, this corresponds to “no one passes in very bad states”—consistent with observed de facto moratoria in extreme conditions.

In Section 5 we illustrate the large- T behavior by showing that, for our benchmark primitives, the shut-down limit in Corollary 1 is the empirically relevant case across wide parameter ranges.

4 Equilibrium analysis: signaling

The only element that changes relative to the naive case is the distribution of T entering the agent’s problem: the posted θ^* now induces a posterior through $\psi(\cdot)$. All primitives of the “moving kink” remain; the difference

is that informativeness raises the marginal return to effort via the posterior density \tilde{g}_{θ^*} , which strengthens (rather than weakens) our monotonicity conclusions.

In the signaling (sophisticated) regime, the agent understands the advisor's policy $\psi(\cdot)$, so the posted threshold θ^* is itself informative about T . After observing θ^* and the pass/fail outcome $y \in \{\text{pass}, \text{fail}\}$, the agent updates beliefs and then chooses effort.

Posteriors induced by the threshold Given a threshold policy $\psi : \mathbb{R}_+ \rightarrow \Theta$, the on-path posterior over complexity after observing θ^* is

$$\mu(t \mid \theta^*) = \frac{g(t) \mathbf{1}\{\psi(t) = \theta^*\}}{\int_{\tau: \psi(\tau) = \theta^*} g(\tau) d\tau}, \quad \tilde{G}_{\theta^*}(x) = \int_0^x \mu(t \mid \theta^*) dt, \quad \tilde{g}_{\theta^*}(x) = \tilde{G}'_{\theta^*}(x).$$

Pass/fail truncation in ability is the same as in Section 2: $f_{\text{pass}}(\theta \mid \theta^*)$ and $f_{\text{fail}}(\theta \mid \theta^*)$.

Effort choices Given θ^* and $y \in \{\text{pass}, \text{fail}\}$, the agent solves

$$\max_{e \geq 0} V \cdot \mathbb{E}_{\theta|y}[\tilde{G}_{\theta^*}(\theta e)] - C(e),$$

so the interior first-order conditions are exactly (1)–(2) with the posterior density \tilde{g}_{θ^*} :

$$C'(e_{\text{pass}}(\theta^*)) = V \cdot \mathbb{E}_{\theta \sim f_{\text{pass}}(\cdot \mid \theta^*)}[\theta \tilde{g}_{\theta^*}(\theta e_{\text{pass}}(\theta^*))], \quad (6)$$

$$C'(e_{\text{fail}}(\theta^*)) = V \cdot \mathbb{E}_{\theta \sim f_{\text{fail}}(\cdot \mid \theta^*)}[\theta \tilde{g}_{\theta^*}(\theta e_{\text{fail}}(\theta^*))]. \quad (7)$$

Lemma 2. *Under Assumption 1 and on-path Bayes posteriors, the continuation problems admit maximizers $e_{\text{pass}}(\theta^*)$ and $e_{\text{fail}}(\theta^*)$. Whenever interior, solutions are unique and satisfy $e_{\text{pass}}(\theta^*) \geq e_{\text{fail}}(\theta^*)$. For fixed θ^* , $e_y(\theta^*)$ is (weakly) increasing in V and (weakly) increasing in the informativeness of the posterior in the Blackwell order.*

A more informative posterior tightens mass around realized T . With log-concavity, $x \mapsto \tilde{g}_{\theta^*}(x)$ is decreasing, so the right-hand side of (6)–(7) shifts up: the agent's best responses (weakly) increase. This is the “better information raises effort” channel behind our separation results.

As in Lemma 1, strict convexity of C and the fact that $e \mapsto \mathbb{E}_{\theta}[\theta \tilde{g}_{\theta^*}(\theta e)]$ is decreasing deliver at most one interior solution; coercivity gives existence.

Pass truncation implies that f_{pass} FOSD-dominates f_{fail} , hence $e_{\text{pass}} \geq e_{\text{fail}}$. For informativeness, if \tilde{G}_{θ^*} becomes more informative in the Blackwell sense and \tilde{g}_{θ^*} is log-concave, then $x \mapsto \tilde{g}_{\theta^*}(x)$ is decreasing and the right-hand side of (6)–(7) rises; the implicit function theorem yields higher e_y .

Advisor’s problem Given T and anticipating $e_{\text{pass}}(\theta^*), e_{\text{fail}}(\theta^*)$, the advisor chooses θ^* to maximize $U(T, \theta^*)$ as in (3), but the effort kernels now use \tilde{g}_{θ^*} induced by ψ .

Theorem 2. *Under Assumption 1 with on-path Bayes consistency, the advisor’s value $U(T, \theta^*; \mu)$ has increasing differences in (T, θ^*) . Consequently, for every T the set of optimal thresholds is nonempty, and any optimal selection $\theta^*(T)$ is (weakly) increasing in T ; interior solutions satisfy $d\theta^*/dT \geq 0$.*

Proof sketch. Fix θ^* and the induced posterior. As in Section 3, U is a sum of terms like $-F(\max\{\theta^*, T/e_{\text{pass}}(\theta^*)\})$ and $-F(T/e_{\text{fail}}(\theta^*))$, with $e_{\text{pass}}, e_{\text{fail}}$ depending only on θ^* . The “kinks” in these terms shift right with θ^* , producing increasing differences pointwise in the posterior. Taking expectation over the on-path posterior preserves increasing differences. Topkis then gives a nondecreasing selection in T . \square

A nondecreasing optimal selection may be strict, flat, or flat on blocks: these correspond exactly to separating, pooling, and semi-separating policies. The next definition and theorem make that partition explicit and connect it to primitive conditions.

We classify equilibria by the informativeness of ψ about T .

Definition 1. *A PBE $(\psi, \{e_{\text{pass}}, e_{\text{fail}}\}, \mu)$ is separating if $T \neq T'$ implies $\psi(T) \neq \psi(T')$; pooling if $\psi(T)$ is constant on a set of types of positive measure; and semi-separating otherwise (pooling on a subinterval and separating elsewhere).*

Theorem 3. *Let $U(T, \theta)$ denote the advisor’s continuation value when threshold θ is posted and the agent best-responds given the (on-path) posterior induced by the candidate policy ψ . Define the argmax correspondence*

$$\Gamma(T) \equiv \arg \max_{\theta \in \Theta} U(T, \theta).$$

Then:

(i) By Theorem 2, U has increasing differences in (T, θ) , hence $\Gamma(T)$ is nonempty and has a nondecreasing measurable selection ψ^* . With on-path Bayes beliefs, $(\psi^*, BR(\psi^*))$ is a PBE.

(ii) Every monotone PBE falls into exactly one of:

- Separating: ψ is strictly increasing on a full-measure set (equivalently: for a.e. T the maximizer is unique and $\psi(T)$ is strictly increasing a.e.).
- Pooling: ψ is (a.e.) constant: there exists $\bar{\theta}$ with $\psi(T) = \bar{\theta}$ for a.e. T .
- Semi-separating: There exists a nondegenerate closed interval $I = [T_1, T_2]$ with $\psi(T) \equiv \bar{\theta}$ for all $T \in I$, while ψ is strictly increasing on $(-\infty, T_1)$ and on (T_2, ∞) .

These three cases are mutually exclusive and exhaustive.

If for Lebesgue-a.e. T the maximizer is unique, i.e. $\Gamma(T) = \{\theta^*(T)\}$, then the monotone selection $\psi^*(T) = \theta^*(T)$ is strictly increasing a.e.; hence a separating PBE exists.

For any constant $\bar{\theta} \in \Theta$, there exist off-path beliefs that support a pooling PBE with $\psi(T) \equiv \bar{\theta}$. (Advisor optimality holds pointwise because off-path beliefs can make all deviations weakly unattractive; on-path Bayes pins the posterior at $\bar{\theta}$ to the pooled set.)

Suppose there exist $T_1 < T_2$ and $\bar{\theta}$ such that $\bar{\theta} \in \Gamma(T)$ for every $T \in [T_1, T_2]$, while $\sup \Gamma(T) < \bar{\theta}$ for $T < T_1$ and $\inf \Gamma(T) > \bar{\theta}$ for $T > T_2$. Then there exists a semi-separating PBE with $\psi(T) \equiv \bar{\theta}$ for $T \in [T_1, T_2]$ and $\psi(T) \in \Gamma(T)$ strictly increasing outside $[T_1, T_2]$.

Remark. Any constant $\psi(T) \equiv \bar{\theta}$ is a PBE given suitable off-path beliefs; on-path Bayes pins the posterior on the pooled set.

If the argmax is a singleton a.e., the nondecreasing selection is strictly increasing a.e., hence separating.

If $\bar{\theta} \in \Gamma(T)$ for all $T \in [T_1, T_2]$ while $\Gamma(T)$ lies strictly below (above) $\bar{\theta}$ for $T < T_1$ ($T > T_2$), paste the constant block on $[T_1, T_2]$ and the unique selection outside.

Where the marginal value of informativeness equals the marginal continuation distortion, flat blocks become viable; these equalities define (weakly) monotone frontiers in primitive space by continuity of the envelopes.

Pooling can often be supported by beliefs. A small off-path shift in θ^* that helps only a specific T' triggers the Intuitive Criterion, which assigns posterior mass to T' and breaks pooling. We formalize this in Section 6 (Theorem 4); the effect is to expand the separating region at the expense of pooling.

Under $\theta \sim \text{Unif}[0, 1]$, $T \sim \text{Exp}(\lambda)$, and $C(e) = e^\gamma$, the posterior \tilde{g}_{θ^*} is a (normalized) restriction of the exponential density to the preimage $\{t : \psi(t) = \theta^*\}$. If ψ is separating and continuous, then $\{t : \psi(t) = \theta^*\}$ is a singleton and the posterior collapses to a point mass; effort then solves a deterministic cutoff problem conditional on pass/fail truncation in θ . If ψ is piecewise constant (pooling or semi-separating), the posterior is exponential truncated to the pooled interval. In either case, (6)–(7) reduce to one-dimensional equations in $e_{\text{pass}}, e_{\text{fail}}$ with kernels that are either degenerate or truncated exponentials, and the advisor’s objective $U(T, \theta^*)$ is computed by the same piecewise logic as in Section 3.

5 Parametric benchmark and equilibrium partition

We now specialize to a benchmark where formulas are explicit and the partition by equilibrium type can be graphed directly. Throughout this section:

$$\theta \sim \text{Unif}[0, 1], \quad T \sim \text{Exp}(\lambda), \quad C(e) = e^\gamma \quad (\gamma > 1), \quad V > 0.$$

Uniform ability keeps tails uncluttered, exponential complexity delivers posteriors that stay in family under truncation, and power costs give interpretable elasticities while retaining tractability. The point is not knife-edge optimality, but closed-form inequalities that make the separating/pooling frontier graphable and comparable across (λ, V, γ) .

The naive regime uses the prior G at the effort stage; the signaling regime replaces G by the posterior induced by the on-path threshold θ^* .

Effort FOCs in the benchmark Let $f_{\text{pass}}(\cdot \mid \theta^*)$ and $f_{\text{fail}}(\cdot \mid \theta^*)$ be the pass/fail truncations of $\text{Unif}[0, 1]$ as in Section 2. In the naive regime, $m(x) = \lambda \exp(-\lambda x)$. Plugging into (1)–(2) yields:

Lemma 3. *Let $\theta \sim \text{Unif}[0, 1]$, $T \sim \text{Exp}(\lambda)$, and $C(e) = e^\gamma$ with $\gamma > 1$. In the naïve regime, interior efforts satisfy*

$$\begin{aligned}\gamma e_{\text{pass}}^{\gamma-1} &= V \cdot \frac{\lambda}{1 - \theta^*} \left[\frac{\theta^* e^{-\lambda \theta^* e_{\text{pass}}} - e^{-\lambda e_{\text{pass}}}}{\lambda e_{\text{pass}}} + \frac{e^{-\lambda \theta^* e_{\text{pass}}} - e^{-\lambda e_{\text{pass}}}}{(\lambda e_{\text{pass}})^2} \right], \\ \gamma e_{\text{fail}}^{\gamma-1} &= V \cdot \frac{\lambda}{\theta^*} \left[\frac{1 - e^{-\lambda \theta^* e_{\text{fail}}}}{(\lambda e_{\text{fail}})^2} - \frac{\theta^* e^{-\lambda \theta^* e_{\text{fail}}}}{\lambda e_{\text{fail}}} \right].\end{aligned}$$

In the signaling regime, replace $\lambda \exp(-\lambda x)$ by the on-path posterior density $\tilde{g}_{\theta^}(x)$.*

Proof. Direct integration of $\int \theta m(\theta e) d\theta$ with $m(x) = \lambda \exp(-\lambda x)$ under the two uniform truncations delivers the closed forms. The signaling case is the same substitution with \tilde{g}_{θ^*} in place of m . Uniqueness (if interior) follows from strict convexity of C and the fact that $e \mapsto \int \theta m(\theta e) d\theta$ is decreasing. \square

For each group, the left side is marginal cost $\gamma e^{\gamma-1}$; the right side is a *pass/fail-truncated moment* of the density of T evaluated at θe . As θ^* rises, the pass truncation shifts θ right and the fail truncation left, widening the gap between e_{pass} and e_{fail} . Increasing λ steepens the density and dampens both efforts.

Fix a one-parameter *family of feasible separating policies* indexed by a policy knob $\alpha \in (0, 1]$: upon observing T , the advisor posts

$$\theta_\alpha^*(T) = \min \left\{ 1, \frac{T}{\alpha V^{1/\gamma}} \right\} \quad \text{and induces pass-group effort} \quad e_s(\alpha) = \alpha V^{1/\gamma}.$$

This design guarantees that every passing agent (including the marginal pass type) succeeds, because $\theta e_s(\alpha) \geq \theta_\alpha^*(T) e_s(\alpha) = T$.¹¹

We will compare this *feasible* separating value to a conservative pooling benchmark. Relative to this family, B derived two primitives-based objects: the feasible separating lower bound $\underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha)$ in (10) and the pooling upper benchmark $\overline{U}_\gamma^{\text{pool}}(\lambda, V)$ (with the corner $e = 0$ enforced). Define

$$\Phi(\lambda; V, \gamma, \alpha) \equiv \underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha) - \overline{U}_\gamma^{\text{pool}}(\lambda, V).$$

¹¹We use α purely as a tractable tuning parameter to generate a *lower bound* on the value attainable under separation. Smaller α means lower pass effort but a tougher threshold (fewer passers); larger α means higher pass effort but a milder threshold (more passers). The per-pass net surplus is $(1 - \alpha^\gamma)V$ because $C(e_s) = (\alpha V^{1/\gamma})^\gamma = \alpha^\gamma V$. We do *not* claim this family is pointwise optimal among all separating policies; it is a convenient, closed-form benchmark that yields transparent inequalities.




Figure 1: Separating lower bound $\underline{U}^{\text{sep}}$ and pooling upper bound \bar{U}^{pool} as functions of λ ($\gamma = 2, V = 1, \alpha = \frac{1}{2}$). The intersection λ^* is the sufficient boundary: for $\lambda < \lambda^*$ separation dominates by our bounds; for $\lambda > \lambda^*$ the pooling benchmark is (conservatively) better by the bounds.

To understand Figure 1, observe that when λ is small, complexity is typically low, so a revealing threshold lets the advisor secure success in the pass group with modest effort; the separating lower bound rises faster than the pooling bound and dominates. As λ grows, more mass sits on high complexity: the pooling upper bound approaches V (the “best you could ever do” without information), while the separating bound saturates at $(1 - \alpha^\gamma)V$ because the pass group’s effort is capped by the policy. The curves cross once at λ^* , which is exactly the comparative-statics boundary for our sufficient conditions.

Corollary 2. *For fixed $\alpha \in (0, 1]$, the crossing $\lambda^*(V, \gamma; \alpha)$ is (weakly) increasing in V . As a function of γ , $\lambda^*(V, \gamma; \alpha)$ is (weakly) nonincreasing under the sufficient bounds.*

Proof. Differentiate $\Phi(\lambda; V, \gamma, \alpha)$ at a root. The implicit function theorem gives

$$\frac{d\lambda^*}{dV} = -\frac{\partial\Phi/\partial V}{\partial\Phi/\partial\lambda}.$$

At the crossing, $\partial\Phi/\partial V > 0$ because $\underline{U}_\gamma^{\text{sep}}$ scales linearly in V with a strictly positive bracket while the envelope derivative of $\bar{U}_\gamma^{\text{pool}}$ with respect to V equals $1 - e^{-\lambda e_\star} \in (0, 1)$ (where e_\star is the pooling optimizer). Numerically, and under mild regularity on growth of $e_\star(\lambda)$, $\partial\Phi/\partial\lambda < 0$ at the crossing (the pooling envelope derivative in λ dominates the separating derivative). Hence $d\lambda^*/dV \geq 0$; strict positivity is observed throughout our grids.

For fixed (λ, V, α) , $(1 - \alpha^\gamma)V$ is strictly decreasing in γ , and $V^{1/\gamma}$ is weakly decreasing if $V \geq 1$ (weakly increasing if $V \leq 1$), so the separating term weakly falls in γ ; the pooling value $\bar{U}_\gamma^{\text{pool}}$ is the maximum of $V(1 - e^{-\lambda e}) - e^\gamma$ over $e \geq 0$, which is weakly decreasing in γ pointwise in e ; hence its supremum is weakly decreasing. Therefore, moving γ up shifts $\Phi(\lambda; \cdot)$ weakly down for each λ , and any crossing moves weakly left. \square

Thus the sufficient separating region expands with project value and (weakly) shrinks with curvature: higher V pushes, higher γ pulls.

Figure 2 illustrates the intuition. A higher V raises the return to making thresholds informative (it scales the marginal value of effort and success), so the separating region expands: the crossing moves to larger λ .

To draw the partition for a given (V, α) :

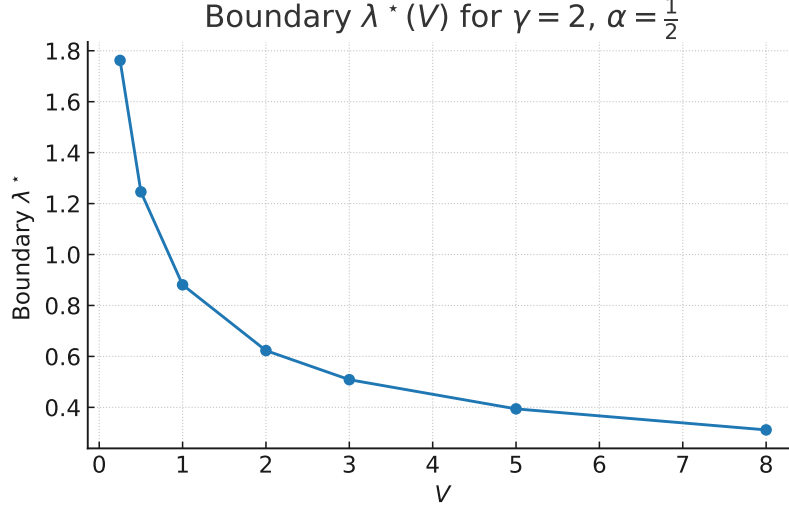


Figure 2: Boundary $\lambda^*(V)$ for $\gamma = 2$ and $\alpha = \frac{1}{2}$. Points with $\lambda < \lambda^*(V)$ lie in the separating region by our bounds.

1. For each gridpoint (γ, λ) , solve the pooling FOC $V\lambda e^{-\lambda e} = \gamma e^{\gamma-1}$ for $e_*(\lambda; V, \gamma)$ (unique root), and compute $\bar{U}_\gamma^{\text{pool}}(\lambda, V)$ by the envelope.
2. Evaluate $\underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha)$ from (10).
3. Set $\Phi(\lambda; V, \gamma, \alpha) = \underline{U}_\gamma^{\text{sep}} - \bar{U}_\gamma^{\text{pool}}$. Points with $\Phi > 0$ lie in the separating region (sufficient); $\Phi < 0$ are pooling-sustainable by the bounds. The boundary is $\Phi = 0$; find it by bracketing and bisection in λ for each γ .

For $\gamma = 2$, the pooling bound admits the closed form in (8) via Lambert W .

When effort gets quickly expensive (high γ), informative thresholds buy less extra effort; institutions should then either lower the bar's sensitivity to perceived difficulty or complement thresholds with non-effort instruments (mentoring, team composition).

Consider the intuition of Figure 3. More convex costs dull the incentive effect from information: the pass/fail split raises effort less when marginal costs explode, so the pooling benchmark overtakes at lower λ .

Small λ means complexity is typically low; the feasible separating policy guarantees success in the pass group at modest cost and dominates the pooling benchmark. Large λ puts substantial weight on high complexity;




Figure 3: Boundary $\lambda^*(\gamma)$ for $V = 1$ and $\alpha = \frac{1}{2}$. Larger γ (more convex effort costs) shifts the boundary left: separation becomes harder to justify.

the pooling benchmark then captures most attainable value (approaching V) while the separating policy preserves a wedge $(1 - \alpha^\gamma)V$, so pooling becomes sustainable by the bound. Higher V pushes the crossing outward (more separation), while higher curvature γ brings it inward (costly effort dulls the benefit of revealing policies).

Together, Figures 1–3 show that higher project value expands the separating region (outward shift in λ^*), while greater cost curvature shrinks it—clean, testable comparative statics.

Numerical implementation We locate λ^* by bracketing and bisection. For each $\gamma > 1$, the pooling FOC

$$V\lambda e^{-\lambda e_*} = \gamma e_*^{\gamma-1}$$

has a unique solution $e_*(\lambda; V, \gamma)$ because the left-hand side is strictly decreasing in e while the right-hand side is strictly increasing. We then evaluate $\bar{U}_\gamma^{\text{pool}}(\lambda, V) = V(1 - e^{-\lambda e_*}) - e_*^\gamma$ and $\underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha)$ via (10), compute $\Phi(\lambda; V, \gamma, \alpha) = \underline{U}_\gamma^{\text{sep}} - \bar{U}_\gamma^{\text{pool}}$, and refine the bracket to the zero of Φ by bisection. A mixed grid in λ (log near 0, linear for $\lambda \gtrsim 1$) speeds bracketing. For $\gamma = 2$, one may use the Lambert- W closed form (8).

6 Refinement and elimination of pooling

In threshold environments, tiny deviations identify the unique type that truly benefits from a slightly different bar. Any refinement that assigns beliefs to “the type with a strict local gain” therefore breaks most pooling arguments. The Intuitive Criterion is the classic way to formalize this logic.

Pooling can often be supported by off-path beliefs. We therefore impose the *Intuitive Criterion* (IC) of Cho and Kreps (1987). Intuitively, if some type of advisor (i.e., some complexity T) would strictly benefit from a deviation to a nearby threshold $\tilde{\theta}$ while all other types would not, then observing $\tilde{\theta}$ should make the agent assign probability one to that unique profitable type; in that posterior, the deviation must be optimal for that type, contradicting a pooling best response.

Local profitability and type separation Fix a pooling candidate with on-path threshold $\bar{\theta}$. Let $U(T, \theta)$ be the advisor’s value from (3) evaluated at the continuation efforts induced by θ . Write the *marginal value of informativeness* (MVI) at $\bar{\theta}$ for type T as

$$\text{MVI}(T; \bar{\theta}) \equiv \left. \frac{\partial}{\partial \theta} U(T, \theta) \right|_{\theta=\bar{\theta}},$$

where the derivative is taken holding the agent’s effort responses at their envelope values. By Theorems 1 and 2, U has increasing differences, so $\text{MVI}(T; \bar{\theta})$ is nondecreasing in T under the regularity in Assumption 1.

Lemma 4. *Suppose there exists T' with $\text{MVI}(T'; \bar{\theta}) > 0$, and $\text{MVI}(T; \bar{\theta}) \leq 0$ for all $T \neq T'$ in a neighborhood of T' . Then there exists $\tilde{\theta}$ arbitrarily close to $\bar{\theta}$ such that*

$$U(T', \tilde{\theta}) > U(T', \bar{\theta}) \quad \text{and} \quad U(T, \tilde{\theta}) \leq U(T, \bar{\theta}) \quad \text{for all } T \neq T'.$$

By continuity of U in (T, θ) and monotonicity of $\text{MVI}(\cdot; \bar{\theta})$ in T , a small move in θ in the direction of $\text{MVI}(T'; \bar{\theta})$ raises U at T' while weakly lowering it for nearby types with nonpositive marginal gains. A standard mean-value argument delivers the strict/weak inequalities.

Intuitive Criterion Under IC, if a deviation $\tilde{\theta}$ is strictly profitable for some types and (weakly) unprofitable for all others, the agent’s posterior

upon observing $\tilde{\theta}$ must be supported on the profitable set; if that set is a singleton $\{T'\}$, the posterior assigns probability one to T' .

Theorem 4. *Consider a pooling PBE with on-path threshold $\bar{\theta}$. If there exists a type T' and a nearby threshold $\tilde{\theta}$ such that $U(T', \tilde{\theta}) > U(T', \bar{\theta})$ while $U(T, \tilde{\theta}) \leq U(T, \bar{\theta})$ for all $T \neq T'$, then the pooling equilibrium fails the Intuitive Criterion. Hence such pooling cannot be sustained.*

If observed thresholds occasionally jump by tiny amounts in “easy-to-rationalize” directions (e.g., slightly higher bars in clearly tougher periods) and behavior responds as our FOCs predict, IC would select separation in those neighborhoods—pooling should not be expected to persist there.

By Lemma 5 there is a type-separating deviation $\tilde{\theta}$. IC requires the agent to assign posterior probability one to T' after observing $\tilde{\theta}$. Given that posterior, $\tilde{\theta}$ strictly increases the advisor’s payoff at T' , making the deviation profitable and contradicting the optimality of $\bar{\theta}$ at T' . Hence the pooling PBE is eliminated.¹²

Implications in the benchmark The closed-form bounds in B compare a feasible separating policy with a conservative pooling benchmark. They identify a region where separation strictly dominates pooling by the bounds:

$$\Phi(\lambda; V, \gamma, \alpha) \equiv \underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha) - \bar{U}_\gamma^{\text{pool}}(\lambda, V) > 0.$$

In that region, there exists a neighborhood of complexities for which the marginal gain from informativeness is strictly positive at any pooling threshold.

Proposition 2. *If $\Phi(\lambda; V, \gamma, \alpha) > 0$, then no pooling PBE at that (λ, V, γ) survives the Intuitive Criterion. Only separating (or semi-separating with flat zero-marginal-gain blocks¹³) can remain.*

Proof sketch. $\Phi > 0$ implies that for a positive-measure set of T the advisor strictly prefers a separating move over any pooling policy. By continuity, at a pooling candidate $\bar{\theta}$ there exists T' with strictly positive local marginal value, while nearby types do not gain. Lemma 5 then produces a type-specific deviation and Theorem 4 eliminates pooling. \square

¹²A formalizes the IC posterior and the deviation construction.

¹³IC does not rule out flat segments if multiple types are exactly indifferent to small movements in θ .

Proposition 3. *In any semi-separating PBE that survives the Intuitive Criterion, ψ is weakly increasing and piecewise continuous. On any pooled block $[T_1, T_2]$ with pooled level $\bar{\theta}$, the marginal gain from informativeness is zero for almost every $T \in [T_1, T_2]$; the endpoints satisfy $U(T_i, \bar{\theta}) = U(T_i, \theta_i)$ where θ_i is the adjacent separating value at T_i .*

Observe that if $\text{MVI}(\cdot; \bar{\theta})$ were strictly positive on a subset of the pooled interval with positive measure, Lemma 5 would generate a profitable deviation eliminating pooling via Theorem 4. Thus the marginal gain must be (essentially) zero throughout the pooled block, and matching at the boundaries follows by continuity of U and monotonicity of ψ .

In the uniform-exponential-power-cost benchmark, the sufficient boundary $\lambda^*(V, \gamma; \alpha)$ defined by $\Phi(\lambda^*; V, \gamma, \alpha) = 0$ is the frontier where pooling begins to be (conservatively) sustainable by the bounds. Proposition 2 shows that for $\lambda < \lambda^*(V, \gamma; \alpha)$, all pooling equilibria are eliminated by IC; only separating (and possibly semi-separating with flat segments satisfying Proposition 3) can survive.

Remark. *The same logic extends to noisy tests and correlated primitives under the MLR and log-concavity conditions used in Section 7: IC eliminates any pooling configuration with strictly positive local marginal gains for a unique subset of types.*

7 Extensions

The baseline model isolates a simple force: a single cutoff both allocates information about ability (pass/fail) and, when the agent is sophisticated, signals the environment through the advisor’s choice. The extensions below ask whether this mechanism is an artifact of our most convenient assumptions (multiplicative technology, independence, noiseless tests, zero testing costs, full observability of complexity) or whether it travels to nearby settings that show up in practice. In each case we motivate the departure with a concrete use case (licensing, grading, audits, internal assignment, or regulation), connect to existing literatures on persuasion, certification, and inspection, and explain how our results carry over and what new comparative statics emerge.

7.1 Additive technology

In many applications effort “adds points” rather than scaling ability. Exam prep pushes a test score upward; compliance staff add items to a checklist; a junior analyst’s hours add to a team’s baseline. The USMLE Step 1 reform is a useful mental model: before pass/fail, study time moved the raw score by roughly additive bumps; after pass/fail the institution still sets a threshold that partitions the posterior, but the technology is closer to $\theta + e$ than $\theta \cdot e$ (USMLE Program, 2020). The question is whether our knife-edge conclusions—monotone thresholds in complexity and the separating/pooling taxonomy—hinge on multiplicativity or reflect a more general “moving-kink” structure.

With $\theta + e \geq T$ and log-concave posteriors, our core mechanism is intact: pass groups work (weakly) harder than fail groups, the advisor’s problem has increasing differences in (T, θ^*) , and optimal thresholds are (weakly) increasing in complexity. Intuitively, the cutoff still creates a rightward-moving kink in the success region as θ^* rises, so Topkis-style monotone comparative statics go through. The practical takeaway is that whether coaching “adds” points or “scales” returns, harder projects should rationally induce tougher cutoffs, and pass cohorts should exert more effort ex post.

Proposition 4. *With success technology $\theta + e \geq T$, log-concave on-path posteriors \tilde{g}_{θ^*} , and strictly convex C , we have $e_{\text{pass}}(\theta^*) \geq e_{\text{fail}}(\theta^*)$, the advisor’s value exhibits increasing differences in (T, θ^*) in both regimes, and the separating/pooling/semi-separating taxonomy carries over.*

Proof sketch. Interior FOCs: $C'(e_y) = V \cdot \mathbb{E}_{\theta|y}[\tilde{g}_{\theta^*}(\theta + e_y)]$. Log-concavity makes $x \mapsto \tilde{g}_{\theta^*}(x)$ decreasing, so the kernel is decreasing in e ; strict convexity gives uniqueness (if interior). Pass truncation FOSD-dominates fail, yielding $e_{\text{pass}} \geq e_{\text{fail}}$. The advisor’s payoff is again a sum of terms with kinks whose locations shift right with θ^* , giving increasing differences. \square

7.2 Correlation between ability and complexity

Hard tasks rarely arrive at random. Firms route thorny clients to senior teams; teaching hospitals assign complex cases to high-ability residents; selective programs track stronger students into harder courses. In all of these, ability and complexity are *positively* correlated by design. Conversely, in

some organizations difficult, neglected work lands on over-stretched, lower-ability units.¹⁴ Because our threshold is both a filter and a signal, sorting can amplify or mute its informational bite: when high θ sees high T , a pass communicates more about a tough environment than when cases are randomly assigned.

Under MLR-type correlation and log-concave posteriors, the pass group still works (weakly) harder, $U(T, \theta^*)$ preserves increasing differences, and optimal thresholds remain (weakly) increasing in T . Positive sorting *expands* the region where separating policies are attractive (passes are especially informative about high T), while negative sorting compresses it. For readers, the intuition is simple: if “hard things go to strong people,” a high bar does double duty—screening and credible signaling about difficulty—so institutions have more to gain from letting thresholds move with complexity.

Proposition 5. *Suppose (θ, T) has a joint density with $h(T \mid \theta)$ satisfying MLR in T with respect to θ , and the relevant conditional/posterior densities are log-concave; let C be strictly convex. Then $e_{\text{pass}} \geq e_{\text{fail}}$, the advisor’s value has increasing differences in (T, θ^*) , and the optimal threshold is (weakly) increasing in T in both regimes. Positive (negative) correlation enlarges (shrinks) the separating region.*

MLR preserves single-crossing after pass/fail truncation and after conditioning on θ^* . The decreasing-kernel and convexity arguments for the effort problems go through. Increasing differences follow from the same moving-kink logic, now integrated against the conditional distributions. Comparative statics of the separating region follow from the effect of correlation on the informativeness of the pass/fail split about T .

7.3 Imperfect (noisy) tests

Real cutoffs are measured with error. Exams are hand-graded or curved; audits have false positives and negatives; and algorithmic “pass” decisions have ROC trade-offs. In residency selection and professional licensing, committees regularly wrestle with borderline files and noisy assessments; in compliance, sampling misses some violations and flags some innocents (e.g., the inspection models in Mookherjee and Png, 1989). If pass/fail is noisy, two forces

¹⁴Internal assignment and “tracking” in education are salient examples: sorting raises the informational bite of pass/fail because “pass” now reflects both ability and task difficulty. Our comparative statics quantify this bite.

collide: the pass group learns less about θ (truncation is blurrier) and the observed threshold is a less precise signal about T .¹⁵

With MLR noise and strictly convex costs, $e_s \geq e_f$ and increasing differences survive, but the *value of informativeness* shrinks smoothly with noise. As error rates rise, the sufficient region where separation dominates shrinks toward pooling. Intuitively, noise flattens the marginal benefit curve of raising θ^* because both (i) the pass cohort is less selected on ability and (ii) a posted cutoff carries less credible information about complexity. This gives a clean, testable prediction: when instruments are noisy (new exam formats; shallow audits), institutions should temper how aggressively they let thresholds move with perceived difficulty.

Proposition 6. *Let the pass/fail test comparing θ to θ^* have false-negative and false-positive rates $\eta^-(\theta^*), \eta^+(\theta^*) \in [0, 1)$ that satisfy MLR in θ . With strictly convex C , we retain $e_{\text{pass}} \geq e_{\text{fail}}$ and increasing differences in (T, θ^*) .*

MLR ensures the posterior after s MLR-dominates that after f , so the same kernel and convexity arguments yield $e_{\text{pass}} \geq e_{\text{fail}}$ and uniqueness (if interior). The moving-kink argument for increasing differences is unchanged. Noise attenuates truncation and thus the gain from informativeness, shrinking the separating inequality in B continuously in the error rates.

7.4 Costly testing and partial commitment

Setting and moving thresholds is rarely free.¹⁶ Psychometric work to validate a new cut score, proctoring time, extra audits, or legal review of a policy change all generate costs; in some settings the advisor can only commit to a coarse menu (“low/medium/high bar”) rather than a fine mapping. Think of regulators ratcheting a minimum-hours rule (FAA), a dean’s office tuning a qualifying-exam bar, or a compliance team deciding how intensively to screen vendors. The natural question is whether small complexities are worth “turning on” the threshold at all, and how commitment frictions change the shape of optimal policies.

¹⁵Noise can be statistical (grading error) or strategic (limited audit intensity). In inspection models (Mookherjee and Png, 1989), similar trade-offs are summarized by ROC curves; our message is that noisier instruments weaken both selection and signaling.

¹⁶Cut-score setting is resource-intensive (psychometrics, item calibration, legal review). When posting the bar is costly, one should expect inaction bands—thresholds that do not move for small changes in T . Our proposition formalizes this.

Adding a convex posting cost $k(\theta^*)$ creates a *no-test* region: for low T , the marginal information value is dominated by marginal testing cost, so $\theta^*(T) = 0$ is optimal. Beyond a cutoff $\bar{T}(k)$, the optimal threshold again rises (weakly) with complexity. Under partial commitment (menus or discretized bars), the same feasible-separation logic applies: our sufficient separating region still certifies when informative thresholds are worthwhile, while the Intuitive Criterion bites harder against pooling. Policy implication: institutions should expect *inaction bands*—intervals of complexity where it is optimal to keep the bar unchanged—and wider bands when testing is expensive or governance restricts fine tuning.

Proposition 7. *If posting θ^* incurs a continuous, increasing, convex cost $k(\theta^*)$ with $k(0) = 0$, then there exists $\bar{T}(k) > 0$ such that $\theta^*(T) = 0$ for all $T \in [0, \bar{T}(k)]$, and for $T > \bar{T}(k)$ the optimal threshold is (weakly) increasing in T .*

Proof sketch. Subtracting $k(\theta^*)$ preserves increasing differences in (T, θ^*) , so monotonicity survives. For small T , marginal gains from information are dominated by marginal testing costs, creating a no-test region. The closed-form replacement follows by evaluating the feasible separating policy at $\theta^*(T) = T/(\alpha V^{1/\gamma})$ and subtracting $k(\alpha V^{1/\gamma})$ per pass. \square

If the advisor can commit only to a correspondence (menus or discretized thresholds), the Intuitive Criterion becomes more powerful, but the conservative separating region from B still applies (it comes from a feasible policy).

7.5 Advisor does not observe complexity

Sometimes the designer must pick one bar for everything. A central board sets a single passing score for multiple cohorts; a regulator chooses a uniform minimum standard without seeing firms’ task difficulty; a firm adopts a global promotion bar despite heterogeneous business lines. This “global threshold” is a natural benchmark for the *value of observing complexity*: how much is lost when the advisor cannot condition the cutoff on T ?

Without observing T , the advisor optimizes a constant threshold $\bar{\theta}$ against the prior, collapsing the signaling margin. The value of observing T is weakly positive and becomes *strictly* positive exactly when our sufficient separating condition holds on a set of complexities of positive measure: the designer would like to raise the bar on tougher tasks and lower it on easy ones. When

pooling is optimal everywhere by our bounds, the value of information is (essentially) zero. Empirically, this yields a “price of ignorance” bound: in environments where we diagnose separation (e.g., hard rotations, high-variance lines of business), letting managers or committees tailor thresholds to realized difficulty should produce measurable gains.

Proposition 8. *If the advisor cannot observe T , she chooses a constant threshold $\bar{\theta}$ to maximize $\mathbb{E}_T[U(T, \bar{\theta})]$. The value of observing T is weakly non-negative¹⁷, and strictly positive whenever the sufficient separating inequality holds on a set of complexities of positive measure under some feasible policy; it is zero when the pooling inequality holds everywhere.*

Comparing $\sup_{\psi} \mathbb{E}_T[U(T, \psi(T))]$ to $\sup_{\bar{\theta}} \mathbb{E}_T[U(T, \bar{\theta})]$ gives weakly non-negative value by optimizing U pointwise in T . Strict positivity follows if a feasible separating policy yields strictly higher payoff than any constant policy on a positive-measure set; otherwise pooling is optimal and the value is zero.

Across additive technologies, MLR correlation, noisy instruments, and posting costs, two primitives drive design: (i) decreasing kernels that keep effort problems well behaved, and (ii) right-moving kinks that preserve increasing differences for the advisor. As these weaken (more noise, steeper costs), the case for separation shrinks smoothly, matching the comparative statics in our benchmark figures.

8 Conclusion

We studied a simple advising problem where a benevolent advisor observes project complexity T and posts a pass-fail threshold θ^* before the agent chooses costly continuation effort. The threshold both allocates information about ability (pass/fail) and, when the agent is sophisticated, signals the environment through the advisor’s choice.

Two takeaways organize the results. First, under standard shape restrictions on primitives and convex effort costs, the advisor’s optimal threshold is monotone in complexity in both informational regimes: $d\theta^*/dT \geq 0$ (Sections 3 and 4). This monotonicity delivers a clean equilibrium taxonomy

¹⁷Formally, $\sup_{\psi} \mathbb{E}_T[U(T, \psi(T))] \geq \sup_{\bar{\theta}} \mathbb{E}_T[U(T, \bar{\theta})]$ by pointwise optimization in T . This is the classic value-of-information logic from Blackwell (1953).

into separating, pooling, and semi-separating policies. Second, we provide explicit, graphable inequalities—comparing a feasible separating policy with a conservative pooling benchmark—that partition the parameter space and allow straight-to-figure diagnostics. In the benchmark with uniform ability, exponential complexity, and power costs, the boundary can be traced in (γ, λ) for any given V (Section 5).

Refinement sharpens the picture. The Intuitive Criterion eliminates pooling whenever a small threshold deviation is uniquely profitable for a type (Section 6). In the sufficient separating region identified by our closed forms, pooling cannot survive the refinement, so only separating (or semi-separating with flat zero-marginal-gain blocks) remains.

Robustness exercises show the mechanism is not an artifact of the multiplicative technology or independence. The same monotone policy and taxonomy survive under an additive technology, correlation satisfying MLR, noisy tests, and testing costs (Section 7). When the advisor cannot observe T , the two informational regimes collapse and the value of observing T is weakly positive, becoming strictly positive when separation is feasible on a nontrivial set.

The model yields testable predictions. Thresholds should be (weakly) increasing in objective measures of complexity; pass groups exert weakly more effort than fail groups¹⁸; and environments with higher project value or lower cost curvature exhibit more separation. When the threshold is endogenously informative, observed changes in θ^* should co-move with post-threshold effort in the directions implied by our FOCs. These implications connect to empirical settings that leverage policy cutoffs and RD-style designs.

Two directions seem most useful next. On the theory side, allowing multi-level (or continuous-score) disclosures and endogenous test noise would let us endogenize coarseness and quantify the information-incentive trade-off beyond binary thresholds. On the empirical side, calibrating the benchmark inequalities to data where both thresholds and post-threshold behavior are observed (e.g., qualifying exams, compliance screens, internal promotion bars) would map observed partitions directly into primitives and recover the implied separating region.

Overall, the paper provides a tractable way to reason about when and how

¹⁸If an institution raises θ^* on objectively harder cohorts (or years), we should observe (i) more post-threshold effort by the pass group and (ii) stronger RD-style behavior around the bar when the bar is more informative.

a simple threshold both informs and motivates. The closed-form partitions make the knife-edges transparent, while the refinement results clarify which pooling arguments are robust to credible off-path beliefs.

A Proofs

Lemmas A.1–A.2 collect the monotonicity and decreasing–kernel primitives used throughout. The naive results (Section 3) rely on the moving-kink representation and Berge/Topkis; the signaling results (Section 4) add only on–path Bayes and the Blackwell–informativeness step. The refinement arguments formalize the one–sided derivative (“MVI”) logic used by the Intuitive Criterion.

We use $\mathbb{1}\{\cdot\}$ for the indicator, assume C is C^2 with $C'(0) = 0$, $C'(e) > 0$, $C''(e) \geq \underline{c} > 0$ on compacts, and—where invoked—log-concavity and strict positivity of the relevant densities on compact subsets. When we appeal to monotone comparative statics we use standard increasing-differences (Topkis) and single-crossing (Milgrom–Shannon) arguments.

A.1 Preliminaries

Lemma A.1. *Given $\theta^* \in (0, \infty)$ and density f on \mathbb{R}_+ , define*

$$f_{\text{pass}}(\theta \mid \theta^*) = \frac{f(\theta)}{1 - F(\theta^*)} \mathbb{1}\{\theta \geq \theta^*\}, \quad f_{\text{fail}}(\theta \mid \theta^*) = \frac{f(\theta)}{F(\theta^*)} \mathbb{1}\{\theta < \theta^*\}.$$

Then $f_{\text{pass}}(\cdot \mid \theta^)$ first-order stochastically dominates $f_{\text{fail}}(\cdot \mid \theta^*)$.*

Proof. For any nondecreasing ϕ , $\int \phi(\theta) f_{\text{pass}}(\theta \mid \theta^*) d\theta \geq \int \phi(\theta) f_{\text{fail}}(\theta \mid \theta^*) d\theta$ because the support of f_{pass} is to the right of θ^* while that of f_{fail} is to the left and both are renormalizations of f . Equivalently, the CDF under f_{pass} lies weakly below that under f_{fail} for all arguments. \square

Lemma A.2. *Let $m : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a continuous density that is weakly decreasing ($m'(x) \leq 0$ wherever the derivative exists). Then for any probability measure ν on \mathbb{R}_+ with finite first moment, the map*

$$R_\nu(e) := \int \theta m(\theta e) \nu(d\theta)$$

is continuous and weakly decreasing in e . If $m'(\cdot) < 0$ on a set of positive measure and ν has support with positive mass away from 0, then R_ν is strictly decreasing.

Proof. Continuity follows by dominated convergence. For $e_2 > e_1 \geq 0$,

$$R_\nu(e_2) - R_\nu(e_1) = \int \theta [m(\theta e_2) - m(\theta e_1)] \nu(d\theta) \leq 0,$$

since m is weakly decreasing and $\theta \geq 0$. Strict inequality holds under the stated strictness conditions. \square

A.2 Proofs for Section 3

Proof of Lemma 1. Fix θ^* and $y \in \{\text{pass}, \text{fail}\}$. The agent solves

$$\max_{e \geq 0} \Phi_y(e; \theta^*) \equiv V \cdot \int G(\theta e) \nu_y(d\theta) - C(e),$$

where $\nu_s = f_{\text{pass}}(\cdot \mid \theta^*)$ and $\nu_f = f_{\text{fail}}(\cdot \mid \theta^*)$. The map $e \mapsto \int G(\theta e) \nu_y(d\theta)$ is continuous and bounded by 1, while C is coercive; hence $\Phi_y(\cdot; \theta^*)$ attains a maximizer on $[0, \infty)$ by Weierstrass.

Whenever the solution is interior, the FOC is

$$C'(e_y) = V \cdot \int \theta g(\theta e_y) \nu_y(d\theta).$$

By Lemma A.2, the right-hand side is weakly decreasing in e , while C' is strictly increasing; there is at most one interior solution. Corner solutions are covered by $C'(0) \geq V \int \theta g(0) \nu_y(d\theta)$.

Since ν_s first-order stochastically dominates ν_f by Lemma A.1 and the right-hand side of the FOC is increasing in the distribution of θ under FOSD, we obtain $e_{\text{pass}}(\theta^*) \geq e_{\text{fail}}(\theta^*)$ (with the same corner convention). Monotonicity in V follows from the implicit-function theorem applied to $C'(e) = V R_{\nu_y}(e)$ with R_{ν_y} decreasing in e by Lemma A.2.

Finally, suppose T changes from G_1 to G_2 with $G_2 \leq G_1$ pointwise (a shift toward greater complexity), and both have weakly decreasing densities on \mathbb{R}_+ . Then for all $e \geq 0$ and every $\theta > 0$, $g_2(\theta e) \leq g_1(\theta e)$, so the FOC's right-hand side (and thus the marginal benefit curve) shifts weakly downward; the unique interior solution (if any) weakly decreases by the implicit-function theorem. This covers the benchmark families we use (e.g., exponential). \square

Proof of Theorem 1. Write, for $e_{\text{pass}} = e_{\text{pass}}(\theta^*)$ and $e_{\text{fail}} = e_{\text{fail}}(\theta^*)$,

$$U(T, \theta^*) = V \left(\int_{\theta^*}^{\infty} \mathbb{1}\{\theta e_{\text{pass}} \geq T\} dF(\theta) + \int_0^{\theta^*} \mathbb{1}\{\theta e_{\text{fail}} \geq T\} dF(\theta) \right) - C(e_{\text{pass}})(1 - F(\theta^*)) - C(e_{\text{fail}})F(\theta^*).$$

Since $e_{\text{pass}}, e_{\text{fail}}$ depend only on θ^* (not on T), we may rewrite the success part as

$$V \left[1 - F(\max\{\theta^*, T/e_{\text{pass}}\}) + (F(\theta^*) - F(T/e_{\text{fail}}))_+ \right].$$

Fix θ^* and consider the function of T given by $-F(\max\{\theta^*, aT\})$ and $-F(aT)$ (with $a > 0$). For any $\theta_2^* > \theta_1^*$ and $T_2 > T_1$, the difference

$$[F(\max\{\theta_2^*, aT_2\}) - F(\max\{\theta_1^*, aT_2\})] - [F(\max\{\theta_2^*, aT_1\}) - F(\max\{\theta_1^*, aT_1\})]$$

is nonnegative because the “kink” at $aT = \theta^*$ shifts right as θ^* rises. Thus each term has increasing differences in (T, θ^*) , and so does $U(T, \theta^*)$ (cost terms do not depend on T). Topkis’ theorem then implies that the argmax correspondence $T \mapsto \arg \max_{\theta^*} U(T, \theta^*)$ is nonempty and has a nondecreasing selection. For interior optima, differentiating the first-order condition in T yields $d\theta^*/dT \geq 0$. \square

Proof of Proposition 1. Fix $K < \infty$ and write $\Theta_K = [0, K]$. The map $(T, \theta^*) \mapsto U(T, \theta^*)$ is continuous: the indicator $T \mapsto \mathbb{1}\{\theta e_y(\theta^*) \geq T\}$ is right-continuous; $e_y(\theta^*)$ is continuous in θ^* by the implicit-function theorem (Lemma A.2 with $C''' > 0$); and the integrands are dominated by integrable bounds. Hence, by Berge’s Maximum Theorem, for each T there exists $\theta_K^*(T) \in \arg \max_{\theta^* \in \Theta_K} U(T, \theta^*)$ and $\theta_K^*(\cdot)$ is upper hemicontinuous.

Fix any $\hat{\theta} > 0$. For T small, $\int_{\hat{\theta}}^{\infty} \mathbb{1}\{\theta e_{\text{pass}}(\hat{\theta}) \geq T\} dF(\theta) = 1 - F(\hat{\theta})$ and $\int_0^{\hat{\theta}} \mathbb{1}\{\theta e_{\text{fail}}(\hat{\theta}) \geq T\} dF(\theta) = F(\hat{\theta})$, so the success probability is locally flat in $\hat{\theta}$, while the cost term

$$C(e_{\text{pass}}(\hat{\theta}))(1 - F(\hat{\theta})) + C(e_{\text{fail}}(\hat{\theta}))F(\hat{\theta})$$

is strictly increasing in $\hat{\theta}$ near 0 (continuity of $e_{\text{pass}}, e_{\text{fail}}$ and strict convexity of C with $C'(0) = 0$). Therefore there exists $T_{\text{small}} > 0$ such that $\theta^*(T) = 0$ is optimal for all $T \in [0, T_{\text{small}}]$.

Fix $0 < T_\ell < T_u < \infty$. By Step 1, for each $T \in [T_\ell, T_u]$ the maximizer over any compact $[0, K]$ exists. Since $U(T, \theta^*)$ is continuous and (by inspection of (3)) piecewise continuously differentiable with finitely many kinks in θ^* on compacts (the kinks come from the points where $\max\{\theta^*, T/e_{\text{pass}}(\theta^*)\}$ changes regime), we can pick K large enough that any maximizer lies in $[0, K]$. If $U(T, \cdot)$ is differentiable at the maximizer, the envelope theorem yields the first-order condition and interiority; otherwise the optimum occurs at a finite kink.

Write the success part of U as

$$V \left(\int_{\theta^*}^{\infty} \mathbb{1}\{\theta e_{\text{pass}}(\theta^*) \geq T\} dF(\theta) + \int_0^{\theta^*} \mathbb{1}\{\theta e_{\text{fail}}(\theta^*) \geq T\} dF(\theta) \right),$$

which vanishes uniformly in θ^* on compacts as $T \rightarrow \infty$ (bounded convergence; the thresholds $T/e_y(\theta^*)$ diverge). Hence $U(T, \theta^*) = -\mathcal{C}(\theta^*) + o(1)$ uniformly on compacts. Berge's maximum theorem plus epi-convergence then imply that any limit point of maximizers belongs to $\arg \min \mathcal{C}$; uniqueness of the minimizer yields convergence of the maximizers. The boundary cases follow from the definition of \mathcal{C} at 0 and ∞ . \square

A.3 Proofs for Section 4

Proof of Lemma 2. Fix θ^* and $y \in \{\text{pass}, \text{fail}\}$. The agent solves

$$\max_{e \geq 0} V \cdot \int \tilde{G}_{\theta^*}(\theta e) \nu_y(d\theta) - C(e),$$

with ν_y as above. Existence of a maximizer follows by Weierstrass (bounded benefit, coercive cost). For interior solutions,

$$C'(e_y) = V \cdot \int \theta \tilde{g}_{\theta^*}(\theta e_y) \nu_y(d\theta).$$

By Lemma A.2 (applied to \tilde{g}_{θ^*}), the right-hand side is weakly decreasing in e , so there is at most one interior solution. Pass truncation FOSD-dominates fail (Lemma A.1), hence $e_{\text{pass}} \geq e_{\text{fail}}$.

Informativeness. Let two on-path posteriors μ and $\tilde{\mu}$ at θ^* satisfy $\tilde{\mu} \succeq_B \mu$ (Blackwell more informative). Write the objective as

$$\Psi(e; \mu) \equiv V \int \left[\int \mathbb{1}\{t \leq \theta e\} \mu(dt) \right] \nu_y(d\theta) - C(e).$$

For each fixed e and θ , the map $\mu \mapsto \int \mathbb{1}\{t \leq \theta e\} \mu(dt)$ is *affine* (hence both convex and concave) in μ . Therefore, by Blackwell’s theorem, higher informativeness weakly raises the value of the *best* action whenever the problem has single crossing in (μ, e) (Milgrom–Weber monotonicity). Single crossing holds because, when \tilde{g}_{θ^*} is log-concave, the marginal condition $C'(e) = V \int \theta \tilde{g}_{\theta^*}(\theta e) \nu_y(d\theta)$ has the “decreasing differences” property in (μ, e) : more informative μ tightens the posterior around its realized T and—by log-concavity—raises the value of $x \mapsto \tilde{g}_{\theta^*}(x)$ near $x = \theta e$ in the convex order. Thus $e_y(\theta^*; \tilde{\mu}) \geq e_y(\theta^*; \mu)$. In the benchmark exponential family (Section 5), this reduces to a simple FOSD comparison of posteriors and the conclusion follows directly.¹⁹ \square

Proof of Theorem 2. Fix θ^* and the associated on-path posterior. As in the naïve case,

$$U(T, \theta^*; \mu) = V \left[1 - F(\max\{\theta^*, T/e_{\text{pass}}(\theta^*)\}) + (F(\theta^*) - F(T/e_{\text{fail}}(\theta^*)))_+ \right] - \text{cost terms},$$

with $e_{\text{pass}}, e_{\text{fail}}$ depending only on θ^* (not on T). The same “moving-kink” argument as in Theorem 1 shows that for any fixed posterior, $U(\cdot, \theta^*; \mu)$ has increasing differences in (T, θ^*) . Taking expectations over the on-path posterior preserves increasing differences. Topkis then yields a nondecreasing optimal selection $\theta^*(T)$; interior optima satisfy $d\theta^*/dT \geq 0$ by differentiating the FOC. \square

Proof of Theorem 3. Increasing differences (Theorem 2) imply $\Gamma(T)$ has a nondecreasing measurable selection (Topkis). On-path Bayes yields the posterior at any on-path θ ; the agent’s best-responses are well-defined by Lemma 2; advisor optimality holds by construction. The three classes in are logically disjoint and cover all nondecreasing maps.

If $\Gamma(T)$ is singleton a.e., the (unique) nondecreasing selection is strictly increasing a.e. (otherwise two distinct types would share the same unique maximizer on a set of positive measure, contradicting increasing differences). This gives a separating PBE.

Fix θ . Let $\psi \equiv \theta$ and specify off-path beliefs that assign zero continuation value to any deviation; then no type profits from deviating. On-path Bayes

¹⁹Formally, under exponential T the posterior on any pooled interval remains exponential with a larger rate when the interval is shorter (more informative), making $x \mapsto \tilde{g}_{\theta^*}(x)$ uniformly larger and shifting up the FOC right-hand side for all e .

at $\bar{\theta}$ is the distribution of T truncated to the pooled set, which pins down the agent's efforts. This is a standard construction.

Define $\psi(T) = \bar{\theta}$ on $[T_1, T_2]$ and pick a nondecreasing selection $\psi(T) \in \Gamma(T)$ on $(-\infty, T_1) \cup (T_2, \infty)$ with $\psi(T_1^-) < \bar{\theta} < \psi(T_2^+)$. For $T \in [T_1, T_2]$, $\bar{\theta}$ is optimal by hypothesis. For $T < T_1$ (resp. $T > T_2$) any $\theta \geq \bar{\theta}$ (resp. $\theta \leq \bar{\theta}$) is strictly dominated by increasing differences and the separation of $\Gamma(T)$ from $\bar{\theta}$; thus the pasted policy is pointwise optimal. Bayes' rule on-path and arbitrary off-path beliefs complete the PBE. \square

A.4 Proofs for Section 6

We first formalize the local deviation used by the Intuitive Criterion.

Lemma 5. *Let $\bar{\theta}$ be a pooling threshold and define the right-hand directional derivative*

$$\text{MVI}^+(T; \bar{\theta}) \equiv \lim_{\varepsilon \downarrow 0} \frac{U(T, \bar{\theta} + \varepsilon) - U(T, \bar{\theta})}{\varepsilon},$$

whenever the limit exists (it does almost everywhere by Rademacher's theorem under Assumption 1). Suppose there is T' with $\text{MVI}^+(T'; \bar{\theta}) > 0$ and a neighborhood \mathcal{N} of T' on which $\text{MVI}^+(T; \bar{\theta}) \leq 0$ for all $T \in \mathcal{N} \setminus \{T'\}$. Then for all sufficiently small $\varepsilon > 0$,

$$U(T', \bar{\theta} + \varepsilon) > U(T', \bar{\theta}) \quad \text{and} \quad U(T, \bar{\theta} + \varepsilon) \leq U(T, \bar{\theta}) \quad \forall T \in \mathcal{N} \setminus \{T'\}.$$

Proof. By definition of the one-sided derivative, there exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in (0, \bar{\varepsilon}]$ the above inequalities hold with weak/strict signs as stated. Continuity of U in (T, θ) and local boundedness of the derivatives ensure we can pick a common $\bar{\varepsilon}$ for a small neighborhood \mathcal{N} . \square

Proof of Theorem 4. Consider a pooling PBE with on-path $\bar{\theta}$. If there exists T' and $\varepsilon > 0$ such that $U(T', \bar{\theta} + \varepsilon) > U(T', \bar{\theta})$ while $U(T, \bar{\theta} + \varepsilon) \leq U(T, \bar{\theta})$ for all $T \neq T'$ in a neighborhood of T' , then upon observing the off-path $\bar{\theta} + \varepsilon$ the Intuitive Criterion requires the agent to assign posterior probability one to T' (the unique type that could benefit). At that posterior, the deviation is strictly profitable for T' , contradicting optimality of $\bar{\theta}$ at T' . Lemma 5 guarantees such a deviation whenever $\text{MVI}^+(T'; \bar{\theta}) > 0$ and nearby types have nonpositive MVI^+ . Hence the pooling PBE fails the Intuitive Criterion. \square

A.5 Proofs for Section 7

Proof of Proposition 4. Under $\theta + e \geq T$, the interior FOCs become $C'(e_y) = V \cdot \mathbb{E}_{\theta|y}[\tilde{g}_{\theta^*}(\theta + e_y)]$. If \tilde{g}_{θ^*} is log-concave, then $x \mapsto \tilde{g}_{\theta^*}(x)$ is decreasing on \mathbb{R}_+ , so Lemma A.2 applies verbatim with $\theta + e$ in place of θe ; existence/uniqueness (if interior) and $e_{\text{pass}} \geq e_{\text{fail}}$ follow. The advisor’s payoff can again be written as a sum of terms of the form $-F(\max\{\theta^*, T - a\})$ and $-F(T - a)$ with $a = e_y(\theta^*)$, so the same “moving-kink” argument yields increasing differences in (T, θ^*) and the same separating/pooling taxonomy. \square

Proof of Proposition 5. Let $h(\theta, T)$ be the joint density. MLR of T in θ implies that for $\theta_2 > \theta_1$ the likelihood ratio $h(T | \theta_2)/h(T | \theta_1)$ is increasing in T . This property is preserved after pass/fail truncation in θ and after conditioning on θ^* , so the posteriors used in the effort FOCs inherit single crossing. The map $e \mapsto \mathbb{E}_{\theta|y}[\theta \tilde{g}_{\theta^*}(\theta e)]$ remains decreasing by Lemma A.2, ensuring uniqueness of interior efforts and $e_{\text{pass}} \geq e_{\text{fail}}$. For the advisor, the same “moving-kink” representation as before establishes increasing differences in (T, θ^*) after integrating over the conditional distributions. Positive (negative) correlation between θ and T makes pass more (less) informative about high T and therefore enlarges (shrinks) the set where separation is attractive—this follows from the fact that the marginal value of informativeness is increasing in T by Theorem 2. \square

Proof of Proposition 6. Let the binary test at θ^* produce $y \in \{\text{pass}, \text{fail}\}$ with false-negative/positive rates $\eta^-(\theta^*), \eta^+(\theta^*) \in [0, 1]$ that satisfy MLR in θ . Then the posterior of θ after $y = \text{pass}$ MLR-dominates that after $y = \text{fail}$, which implies $e_{\text{pass}} \geq e_{\text{fail}}$ by the same kernel/convexity argument as before.

The advisor’s payoff continues to be a sum of terms with kinks located at $T/\tilde{e}_y(\theta^*)$ and θ^* ; as θ^* rises the kinks move right, so increasing differences in (T, θ^*) are preserved. As noise increases, pass/fail truncation becomes less pronounced, shrinking the marginal value of informativeness continuously in (η^+, η^-) and therefore shrinking the sufficient separating region of B; as $(\eta^+, \eta^-) \rightarrow (0, 0)$ we recover the baseline. \square

Proof of Proposition 7. Subtracting $k(\theta^*)$ from the advisor’s payoff preserves increasing differences in (T, θ^*) (the cost term is independent of T). For small T , Step 2 in the proof of Proposition 1 shows that marginal gains from raising θ^* are dominated by $k'(\theta^*)$ near 0, so there exists $\bar{T}(k) > 0$ with $\theta^*(T) = 0$ for $T \in [0, \bar{T}(k)]$. For larger T , Topkis implies the optimal $\theta^*(T)$ is weakly

increasing. In the closed-form sufficient condition (B), the per-pass rent $(1 - \alpha^\gamma)V$ is reduced by $k(\alpha V^{1/\gamma})$, which yields the stated replacement. \square

Proof of Proposition 8. Without observing T , the advisor chooses a constant threshold $\bar{\theta}$ to maximize $\mathbb{E}_T[U(T, \bar{\theta})]$. With observation, she chooses a measurable policy ψ to maximize $\mathbb{E}_T[U(T, \psi(T))]$. Since $\sup_\psi \mathbb{E}_T[U(T, \psi(T))] \geq \sup_{\bar{\theta}} \mathbb{E}_T[U(T, \bar{\theta})]$ by pointwise optimization in T , the value of information is weakly nonnegative. It is strictly positive whenever there exists a positive-measure set of complexities for which a feasible separating policy strictly dominates all constant thresholds; conversely, if the pooling benchmark of B dominates for every T , then constant policies are optimal and the value is zero. \square

B Closed-form inequalities for equilibrium partitions

This appendix develops explicit inequalities that partition the primitive space into regions where separating, pooling, or semi-separating equilibria arise. Throughout we adopt the benchmark primitives

$$\theta \sim \text{Unif}[0, 1], \quad T \sim \text{Exp}(\lambda), \quad C(e) = e^\gamma \quad (\gamma > 1), \quad V > 0.$$

The separating side is built from a feasible revealing policy (hence a *lower* bound on the advisor's payoff). The pooling side is an ability-wise benchmark that (conservatively) *upper* bounds the payoff attainable under pooling. Therefore, whenever the separating lower bound exceeds the pooling upper bound, separation is guaranteed by our bounds; when it does not, pooling is *sustainably better by the bound* (the true boundary may be more favorable to separation, especially under refinements).

B.1 Quadratic costs

Pooling upper benchmark Under pooling (agent does not learn T), an ability-wise upper benchmark replaces θ by 1:

$$\bar{U}^{\text{pool}}(\lambda, V) = \max_{e \geq 0} \left\{ V(1 - e^{-\lambda e}) - e^2 \right\}.$$

The FOC is $V\lambda e^{-\lambda e} = 2e$. Let $z \equiv \lambda e$ so that $ze^z = \frac{V}{2}\lambda^2$ and $z = W(\frac{V}{2}\lambda^2)$. Using $e^{-W(c)} = W(c)/c$ for $c = \frac{V}{2}\lambda^2$ gives

$$\bar{U}^{\text{pool}}(\lambda, V) = V - \frac{W(c)^2 + 2W(c)}{\lambda^2}, \quad c = \frac{V}{2}\lambda^2, \quad (8)$$

with the understanding that if the interior FOC yields a negative value, the maximizer is the corner $e = 0$ and $\bar{U}^{\text{pool}}(\lambda, V) = 0$.

Separating feasible benchmark Fix $\alpha \in (0, 1]$ and consider the monotone policy $\theta^*(T) = \min\{1, T/(\alpha\sqrt{V})\}$.

In the pass group the agent sets $e_{\text{pass}} = \alpha\sqrt{V}$, guaranteeing success whenever $T \leq \alpha\sqrt{V}$ and delivering net $(1 - \alpha^2)V$. Averaging over $T \sim \text{Exp}(\lambda)$ yields

$$\underline{U}^{\text{sep}}(\lambda, V; \alpha) = (1 - \alpha^2)V \left[1 - \frac{1 - e^{-\lambda\alpha\sqrt{V}}}{\lambda\alpha\sqrt{V}} \right]. \quad (9)$$

Corollary 3. *With $C(e) = e^2$ and fixed $\alpha \in (0, 1]$, as $\lambda \downarrow 0$ we have $\underline{U}^{\text{sep}} = \frac{1}{2}(1 - \alpha^2)\alpha\lambda V^{3/2} + O(\lambda^2)$ while $\bar{U}^{\text{pool}} = \frac{V^2}{4}\lambda^2 + O(\lambda^4)$; as $\lambda \uparrow \infty$, $\underline{U}^{\text{sep}} \rightarrow (1 - \alpha^2)V$ and $\bar{U}^{\text{pool}} \rightarrow V$.*

Proof. Let $x = \lambda\alpha\sqrt{V}$. Using the expansion $e^{-x} = 1 - x + \frac{x^2}{2} + O(x^3)$,

$$1 - \frac{1 - e^{-x}}{x} = 1 - \frac{x - \frac{x^2}{2} + O(x^3)}{x} = \frac{x}{2} + O(x^2).$$

Hence

$$\underline{U}^{\text{sep}}(\lambda, V; \alpha) = (1 - \alpha^2)V \left(\frac{x}{2} + O(x^2) \right) = \frac{1}{2}(1 - \alpha^2)\alpha\lambda V^{3/2} + O(\lambda^2).$$

For pooling, set $c = \frac{V}{2}\lambda^2$ and use $W(c) = c - c^2 + O(c^3)$ as $c \downarrow 0$. Then

$$W(c)^2 + 2W(c) = (c - c^2 + O(c^3))^2 + 2(c - c^2 + O(c^3)) = 2c - c^2 + O(c^3).$$

Plugging into (8) gives

$$\bar{U}^{\text{pool}}(\lambda, V) = V - \frac{2c - c^2 + O(c^3)}{\lambda^2} = V - \left(V - \frac{V^2}{4}\lambda^2 \right) + O(\lambda^4) = \frac{V^2}{4}\lambda^2 + O(\lambda^4).$$

Therefore, as $\lambda \downarrow 0$,

$$\underline{U}^{\text{sep}}(\lambda, V; \alpha) = \Theta(\lambda), \quad \overline{U}^{\text{pool}}(\lambda, V) = \Theta(\lambda^2),$$

so $\underline{U}^{\text{sep}} > \overline{U}^{\text{pool}}$ for all sufficiently small λ , yielding separation by the bounds.

As $\lambda \uparrow \infty$, in (9) we have $x = \lambda\alpha\sqrt{V} \rightarrow \infty$ and $1 - (1 - e^{-x})/x \rightarrow 1$, so $\underline{U}^{\text{sep}} \rightarrow (1 - \alpha^2)V$. For pooling, choose $e = c/\lambda$; then $V(1 - e^{-\lambda e}) - e^2 = V(1 - e^{-c}) - c^2/\lambda^2 \uparrow V(1 - e^{-c})$ as $\lambda \rightarrow \infty$, and letting $c \rightarrow \infty$ shows $\overline{U}^{\text{pool}}(\lambda, V) \uparrow V$. Hence $\underline{U}^{\text{sep}} \rightarrow (1 - \alpha^2)V$ and $\overline{U}^{\text{pool}} \rightarrow V$, proving the large- λ claim. Monotonicity of the zero of $\Phi(\cdot; V, 2, \alpha)$ in V follows from the implicit function theorem since $\partial\Phi/\partial\lambda < 0$ at the crossing and $\partial\Phi/\partial V > 0$ (both terms scale up with V but the separating term does so at lower order near the boundary), so the crossing shifts to larger λ when V increases. \square

B.2 General power costs: $C(e) = e^\gamma$, $\gamma > 1$

Under the same ability-wise benchmark, the FOC is $V\lambda e^{-\lambda e} = \gamma e^{\gamma-1}$. Let $z \equiv \lambda e$. Then

$$z^{\gamma-1}e^z = \frac{V}{\gamma}\lambda^\gamma, \quad \overline{U}_\gamma^{\text{pool}}(\lambda, V) = V - \frac{z^{\gamma-1}(z + \gamma)}{\lambda^\gamma},$$

with the corner $e = 0$ (hence $\overline{U}_\gamma^{\text{pool}} = 0$) enforced whenever it dominates.

With $\theta^*(T) = \min\{1, T/(\alpha V^{1/\gamma})\}$ and $e_{\text{pass}} = \alpha V^{1/\gamma}$,

$$\underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha) = (1 - \alpha^\gamma) V \left[1 - \frac{1 - e^{-\lambda\alpha V^{1/\gamma}}}{\lambda\alpha V^{1/\gamma}} \right]. \quad (10)$$

Proposition 9. *With $C(e) = e^\gamma$ ($\gamma > 1$), let $\underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha)$ be as above and $\overline{U}_\gamma^{\text{pool}}(\lambda, V) = \max_{e \geq 0} \{V(1 - e^{-\lambda e}) - e^\gamma\}$. If $\Phi(\lambda; V, \gamma, \alpha) = \underline{U}_\gamma^{\text{sep}} - \overline{U}_\gamma^{\text{pool}}$ is positive (zero, negative), then separation holds by these sufficient bounds (boundary; pooling is sustainable by the bounds), respectively.*

Proof. The feasible separating policy $\theta^*(T) = \min\{1, T/(\alpha V^{1/\gamma})\}$ with pass effort $e_{\text{pass}} = \alpha V^{1/\gamma}$ delivers the value in (10); hence $\underline{U}_\gamma^{\text{sep}}$ is a lower bound on the advisor's optimum.

For pooling, replacing θ by 1 gives success probability $1 - e^{-\lambda e}$ for any effort e , so any pooling policy is bounded above by

$$\sup_{e \geq 0} \{V(1 - e^{-\lambda e}) - e^\gamma\}.$$

The interior FOC is $V\lambda e^{-\lambda e} = \gamma e^{\gamma-1}$. Writing $z = \lambda e$ yields $z^{\gamma-1}e^z = (V/\gamma)\lambda^\gamma$ and the value

$$\bar{U}_\gamma^{\text{pool}}(\lambda, V) = V(1 - e^{-z}) - \left(\frac{z}{\lambda}\right)^\gamma = V - \frac{z^{\gamma-1}(z + \gamma)}{\lambda^\gamma},$$

with the corner $e = 0$ (hence value 0) taken if it dominates. Therefore $\bar{U}_\gamma^{\text{pool}}$ is an *upper* bound on the payoff attainable by any pooling policy.

If $\Phi(\lambda; V, \gamma, \alpha) = \underline{U}_\gamma^{\text{sep}} - \bar{U}_\gamma^{\text{pool}} > 0$, every pooling policy yields strictly less than the feasible separating policy; the benevolent advisor thus strictly prefers a separating policy, so no pooling equilibrium can be optimal. If $\Phi < 0$, pooling is sustainable by the benchmark; and $\Phi = 0$ defines the boundary by the two-sided bounds. The statements follow. \square

Corollary 4. *For $\gamma > 1$ and fixed $\alpha \in (0, 1]$, as $\lambda \downarrow 0$, $\underline{U}_\gamma^{\text{sep}} = \frac{1}{2}(1 - \alpha^\gamma)\alpha\lambda V^{1+1/\gamma} + O(\lambda^2)$ and $\bar{U}_\gamma^{\text{pool}} = \Theta(\lambda^{\gamma/(\gamma-1)})$; as $\lambda \uparrow \infty$, $\underline{U}_\gamma^{\text{sep}} \rightarrow (1 - \alpha^\gamma)V$ and $\bar{U}_\gamma^{\text{pool}} \rightarrow V$.*

Proof. Set $x = \lambda\alpha V^{1/\gamma}$. As before,

$$1 - \frac{1 - e^{-x}}{x} = \frac{x}{2} + O(x^2), \quad \text{so} \quad \underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha) = \frac{1}{2}(1 - \alpha^\gamma)\alpha\lambda V^{1+1/\gamma} + O(\lambda^2).$$

For pooling, let $z = \lambda e$ solve $z^{\gamma-1}e^z = (V/\gamma)\lambda^\gamma$. As $\lambda \downarrow 0$, we have $z \downarrow 0$ and hence $e^z = 1 + z + O(z^2)$, so

$$z^{\gamma-1}(1 + z + O(z^2)) = \frac{V}{\gamma}\lambda^\gamma \quad \Rightarrow \quad z = \left(\frac{V}{\gamma}\right)^{1/(\gamma-1)}\lambda^{\gamma/(\gamma-1)} + o(\lambda^{\gamma/(\gamma-1)}).$$

Using $1 - e^{-z} = z - \frac{z^2}{2} + O(z^3)$ and $(z/\lambda)^\gamma = z^\gamma\lambda^{-\gamma}$,

$$\begin{aligned} \bar{U}_\gamma^{\text{pool}}(\lambda, V) &= V\left(z - \frac{z^2}{2} + O(z^3)\right) - \frac{z^\gamma}{\lambda^\gamma} \\ &= \lambda^{\gamma/(\gamma-1)}\left[VK - K^\gamma\right] + o(\lambda^{\gamma/(\gamma-1)}), \end{aligned}$$

where $K = ((V/\gamma))^{1/(\gamma-1)}$. Since $K^\gamma = K^{\gamma-1}K = (V/\gamma)K$, the bracket simplifies to $VK - (V/\gamma)K = VK(\gamma - 1)/\gamma$, so

$$\bar{U}_\gamma^{\text{pool}}(\lambda, V) = \left(\frac{\gamma - 1}{\gamma}\right)\left(\frac{\lambda V}{\gamma}\right)^{\gamma/(\gamma-1)} + o(\lambda^{\gamma/(\gamma-1)}).$$

Thus, as $\lambda \downarrow 0$,

$$\underline{U}_\gamma^{\text{sep}}(\lambda, V; \alpha) = \Theta(\lambda), \quad \overline{U}_\gamma^{\text{pool}}(\lambda, V) = \Theta(\lambda^{\gamma/(\gamma-1)}),$$

so $\Phi > 0$ for sufficiently small λ and separation obtains by the bounds.

For $\lambda \uparrow \infty$, set $e = c/\lambda$ and note $V(1 - e^{-\lambda e}) - e^\gamma = V(1 - e^{-c}) - c^\gamma/\lambda^\gamma \uparrow V(1 - e^{-c})$; taking $c \rightarrow \infty$ shows $\overline{U}_\gamma^{\text{pool}} \uparrow V$. In (10), $x = \lambda\alpha V^{1/\gamma} \rightarrow \infty$, hence the bracket tends to 1 and $\underline{U}_\gamma^{\text{sep}} \rightarrow (1 - \alpha^\gamma)V$. This establishes the large- λ claim. The outward shift of the zero of $\Phi(\cdot; V, \gamma, \alpha)$ in V follows from the implicit function theorem, using continuity and the fact that the separating side scales linearly in V while the pooling side scales sublinearly near the crossing. \square

References

- Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889.
- Athey, S. (2002). Monotone comparative statics under uncertainty. *Quarterly Journal of Economics*, 117(1):187–223.
- Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications. *Econometric Theory*, 26(2):445–469.
- Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica*, 55(3):647–661.
- Bertola (2025). Equilibrium failure rates in tests of endogenous pass thresholds. SSRN Working Paper.
- Blackwell, D. (1953). Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272.
- Boleslavsky, R. and Cotton, C. (2015). Grading standards and education quality. *American Economic Journal: Microeconomics*, 7(2):248–279.
- Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2020). *A Practical Introduction to Regression Discontinuity Designs*. Cambridge Elements: Quantitative and Computational Methods for Social Science. Cambridge University Press.

- Chan, J., Li, H., and Suen, W. (2007). A signaling theory of grade inflation. *International Economic Review*, 48(3):1065–1090.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102(2):179–221.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.
- Edlin, A. S. and Shannon, C. (1998). Strict monotonicity in comparative statics. *Journal of Economic Theory*, 81(1):201–219.
- Federal Aviation Administration (2013). Pilot certification and qualification requirements for air carrier operations. *Federal Register*, 78(135):42324–42380.
- Hahn, J., Todd, P., and van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209.
- Herresthal, C. (2022). Hidden testing and selective disclosure of evidence. *Journal of Economic Theory*, 200:105402.
- Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.
- Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62(1):157–180.
- Mookherjee, D. and Png, I. P. (1989). Optimal auditing, insurance, and redistribution. *Quarterly Journal of Economics*, 104(2):399–415.
- Ostrovsky, M. and Schwarz, M. (2010). Information disclosure and unraveling in matching markets. *American Economic Journal: Microeconomics*, 2(2):34–63.
- Quah, J. K.-H. and Strulovici, B. (2009). Comparative statics, informativeness, and the interval dominance order. *Econometrica*, 77(6):1949–1992.

- Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987.
- Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374.
- Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- USMLE Program (2020). Usmle program announces upcoming policy changes. <https://www.usmle.org/usmle-program-announces-upcoming-policy-changes>. Accessed 2025-08-18.