WORKING PAPERS

N° 1677

October 2025

“Dynamic Delegation with Reputation Feedback”

Georgy Lukyanov and Anna Vlasova

Toulouse School of Economics

Dynamic Delegation with Reputation Feedback*

Georgy Lukyanov
Toulouse School of Economics

Anna Vlasova
CREST — École Polytechnique

Abstract

We study dynamic delegation with *reputation feedback*: a long-lived expert advises a sequence of implementers whose effort responds to current reputation, altering outcome informativeness and belief updates. We solve for a recursive, belief-based equilibrium and show that advice is a reputation-dependent cutoff in the expert's signal. A diagnosticity condition—failures at least as informative as successes—implies *reputational conservatism*: the cutoff (weakly) rises with reputation. Comparative statics are transparent: greater private precision or a higher good-state prior lowers the cutoff, whereas patience (value curvature) raises it. Reputation is a submartingale under competent types and a supermartingale under less competent types; we separate boundary hitting into learning (news generated infinitely often) versus no-news absorption. A success-contingent bonus implements any target experimentation rate with a plug-in calibration in a Gaussian benchmark. The framework yields testable predictions and a measurement map for surgery (operate vs. conservative care).

Keywords: Dynamic delegation; expert advice; moral hazard; experimentation; reputational conservatism.
JEL: D82, D83, C73.

1 Introduction

Expert advice shapes high-stakes decisions in many arenas: surgeons decide whether to operate or pursue conservative care; financial analysts issue buy or hold calls that investors act upon; R&D leads greenlight or delay projects; policy consultants argue for reform or the status quo. In all of these settings, outcomes depend not only on whether the idea is good but also on how diligently it is implemented. Implementation, in turn, responds to the expert's standing: trusted experts are heeded and their recommendations are executed with greater effort. This two-way link—reputation raising effort, and effort making outcomes more informative—creates what we call *reputation feedback*. Empirical patterns are consistent with this mechanism: patient adherence rises with clinician trust and communication quality (Haskard Zolnierek and DiMatteo, 2009; Birkhäuer et al., 2017); surgical outcomes vary sharply with surgeon skill and volume (Birkmeyer et al., 2003, 2013); and influential analyst recommendations move implementation scale and outcomes in financial markets (Loh and Stulz, 2011).

We build a tractable dynamic model of expert advice that isolates reputation feedback and turns it into sharp equilibrium predictions. A long-lived expert repeatedly recommends a risky action to

*We thank Costas Cavounidis, Olivier Gossner, Margarita Kirneva, Yukio Koriyama, Chiara Margaria, Harry Pei, Ludvig Sinander, Francesco Squintani, Bruno Strulovici, Stepan Svistunov, Nikhil Vellodi, Maria Ziskelevich, and all the participants of the 2021 SET seminar for useful comments and discussions. All remaining errors are ours. This research is supported by a grant of the French National Research Agency (ANR), "Investissements d'Avenir" (LabEx Ecodec/ANR-11-LABX-0047).

a sequence of short-lived agents (implementers). Success requires both a favorable technological state and implementer effort; recommendations and outcomes are publicly observed and update the expert's perceived competence. Because each agent's optimal effort responds to current reputation, the expert internalizes that today's advice affects both the chance of success and what the outcome will reveal. The equilibrium characterization is recursive and belief-based.

Our main results are threefold. *First*, we show thatfor any public reputation there is a cutoff in the expert's private signal above which she recommends the risky action and below which she advises safety (Theorem 5). This delivers portable comparative statics and clarifies on-path mimicry by less competent experts. *Second*, we characterize *reputation dynamics*: posterior beliefs form a submartingale under a competent expert and a supermartingale under an incompetent one, and the belief process hits trust regions with probability one (Theorem 12). *Third*, we establish *reputational conservatism*: the risky-signal cutoff is increasing in current reputation (Theorem 7). Intuitively, as reputation rises, implementers work harder, so a failure becomes more diagnostic of incompetence while the informativeness of success does not increase; with a convex value of reputation, the downside risk dominates at high reputation. We also provide transparent comparative statics (Theorem 17): better private information and a higher probability that the environment is favorable lower the cutoff and raise experimentation, whereas greater patience amplifies conservatism.

The model yields testable predictions. Holding fundamentals fixed, highly reputed experts recommend the risky action less frequently but achieve higher success rates when they do, because implementers work harder. Transitory reputation shocks (e.g., early successes) reduce the frequency of risky calls but make subsequent failures more revealing and more reputationally costly; negative shocks have the opposite effect. These predictions echo empirical patterns in settings where effort is observable or proxied (e.g., adherence) and where performance differences across experts are documented (Haskard Zolnierek and DiMatteo, 2009; Birkhäuer et al., 2017; Birkmeyer et al., 2003, 2013). They speak to the design of advisory relationships in medicine (operate vs. conservative care), finance (buy vs. hold and trade size), innovation management (greenlight vs. delay), and public policy (reform vs. status quo), where the scale and diligence of implementation are endogenous to the expert's standing.[1]

Two modeling choices drive tractability and clarity. First, we keep the public state one-dimensional (the current reputation about competence) and assume independent period states for the technology. Second, success requires both a good state and implementer effort, which is optimally chosen in response to current reputation and the recommendation. These assumptions allow a recursive solution and clean comparative statics while capturing the core mechanism: reputation feeds effort, and effort feeds learning from outcomes.

Related work spans dynamic career concerns, strategic expert advice, and delegated experimentation. Our contribution is to embed the implementer's effort response within a recursive reputation model and to show how that endogeneity reshapes both policy (a reputation-dependent threshold with conservatism) and learning (martingale dynamics with boundary hitting).

The paper proceeds as follows. Section 2 reviews related work. Section 3 sets up the environment, belief updates, and recursive equilibrium. Section 4 presents the threshold equilibrium, a binary–signal illustration, reputational conservatism, comparative statics, and reputation dynamics with a learning–absorption decomposition. Section 5 treats the surgery application and policy design via success–contingent bonuses, and covers monitoring, committees, endogenous exit, and a continuous–time approximation. Core proofs are in Appendices A–D; algorithms, Gaussian formulas, robustness, committee variants, behavioral twists, and policy calibration appear in the Online

[1] A practical recipe for constructing reputation and effort proxies is in Online Appendix OA2. Computational details for value iteration, cutoff computation, and simulations are in Online Appendix OA3.

Appendix OA2–OA8).

2 Related Literature

This paper connects dynamic career-concerns and reputation, strategic expert advice with reputational motives, and delegated experimentation/bandit models with hidden actions. Our contribution is to embed the implementer's effort response into a recursive reputation model and to characterize how this *reputation–effort complementarity* shapes policies and the law of motion of beliefs.

A first strand is the dynamic reputation and career-concerns tradition, where forward-looking agents distort current behavior to affect future assessments. The classic insight originates with Holmström's (1982) career concerns. Reputations can serve as implicit incentive schemes (Tadelis, 2002), but their informational content evolves endogenously and, under rich observation, can unravel in the long run (Cripps et al., 2004). We take the infinite-horizon viewpoint and show that reputation is a submartingale (supermartingale) for competent (incompetent) experts, with boundary hitting. Our *reputational conservatism*—cutoffs increasing in current reputation—relates to the tension between conservatism and gambling in dynamic career-concerns models (Prendergast and Stole, 1996). In a complementary direction, Mylovanov and Klein (2017) show that sufficiently long horizons can discipline informational biases; our recursive solution clarifies the channel via the diagnosticity of outcomes when clients exert more effort at higher reputation. In relational and repeated settings, reputational incentives can be misaligned and generate inefficient equilibria; Deb et al. (2022) analyze such forces and provide sharp characterizations. The reputation–effort loop we formalize is adviser-side and distinct, but it leads to analogous selection and path-dependence phenomena.

A second strand studies strategic expert advice when advisers care about reputation. In *reputational cheap talk* with continuous signals, experts pool messages to manage reputational risk, yielding coarse disclosure (Ottaviani and Sørensen, 2006). We provide a dynamic counterpart with a state-dependent cutoff that moves with public reputation. Conformity and caution under career concerns have been emphasized in settings where "being wrong alone" is worse than "being wrong with the crowd" (Scharfstein and Stein, 1990; Prendergast, 1993; Morris, 2001). In our model, similar caution emerges through the effort channel: when clients work harder for highly reputed experts, a failure is more revealing, pushing such experts toward conservatism. With short-lived customers, initial honesty may profitably build a reputation that can later be exploited (Ozyurt, 2016). Paradoxically, very competent advisers with spotless records may be fired because their behavior is inferred to be uninformative or biased (Schottmüller, 2019); our law of motion clarifies when occasional visible failures are actually informative (and thus desirable) versus reputation-damaging. Relatedly, "good lies" can build credibility in the long run (Pavesi and Scotti, 2022); in our environment, explicit outcomes discipline such strategies because failures become more diagnostic when reputation is high.

A third strand is delegated experimentation and bandit models with hidden actions. Strategic experimentation highlights informational externalities and dynamic incentives (Bolton and Harris, 1999; Keller et al., 2005); our environment has no cross-agent social-learning externality (agents are short-lived), but it does feature a payoff externality because effort responds to beliefs about the expert, which feeds back into the informativeness of outcomes. Halac et al. (2016) study optimal intertemporal contracts for experimentation with hidden information and hidden effort; we instead characterize reputational incentives when contracts are absent and show when reputation substitutes for explicit incentives. Dynamic signaling through experimentation with changing types is analyzed by Thomas (2019); our cutoff and martingale results are complementary and highlight how the endogenous effort response changes informativeness. Related work on information aggregation and transparency—such as Prat (2005) and more recent contributions like Backus and Little (2020)—suggests that disclosure regimes interact with reputational motives to induce conservatism; our mechanism microfounds this via effort-induced diagnosticity of outcomes.

Finally, markets for advice and compensation design link reputation with pay-for-performance. In financial advice and ratings, incentives and reputation interact in complex ways (Inderst and Ottaviani, 2012; Bolton et al., 2012). Bayesian persuasion (Kamenica and Gentzkow, 2011) provides tools for incentive-aligned disclosure; our one-shot analysis shows that a simple contingent-fee (success bonus) can implement the efficient cutoff by offsetting reputational distortions (Lukyanov et al., 2025). Motivating experimentation often requires tolerating early failures (Manso, 2011); our dynamic analysis shows how reputation can either replicate this tolerance (when effort is low at low reputation) or stall experimentation (when conservatism dominates at high reputation).

Relative to these literatures, our main novelty is to endogenize the implementer's effort in a recursive reputation model and to characterize how this endogeneity—through a stochastic complementarity between reputation and effort—(i) yields a state-dependent cutoff with *reputational conservatism*, (ii) delivers clean martingale dynamics with boundary characterization, and (iii) generates transparent comparative statics in signal precision, priors, and patience. This mechanism is absent in reputational cheap talk without implementation effort and in bandit models without adviser career concerns.

3 Model

3.1 Environment and primitives

Time is discrete, $t = 0, 1, 2, \dots$. A single *expert* (she) interacts in each period with a new short-lived *agent* (he). The expert's *type* $\theta \in \{H, L\}$ (competent vs. less competent) is fixed over time and privately known to the expert. The market holds a *public reputation* (belief) $\pi_t \equiv \mathbb{P}(\theta = H \mid \mathcal{H}_t) \in (0, 1)$, where $\mathcal{H}_t$ is the public history up to the start of period t (defined below). The initial reputation $\pi_0 \in (0, 1)$ is given.

In period t, the *state* $\omega_t \in \{0, 1\}$ is drawn i.i.d. with $\mathbb{P}(\omega_t = 1) = \lambda \in (0, 1)$; throughout we write ω for a generic period draw. The expert privately observes a *signal* $s_t \in S \subseteq \mathbb{R}$ about ω_t before issuing a recommendation. For each type θ, the signal distribution conditional on the state admits densities $f_\theta(\cdot \mid \omega)$ with respect to a common reference measure on S, satisfies the monotone likelihood ratio property (MLRP) in s, and the High type H is (strictly) more informative than the Low type L in the Blackwell sense.[2]

If the expert recommends the *risky* action ($a_t = 1$), the agent chooses an *effort* level $e_t \in [0, 1]$ at cost $c(e_t)$, where $c : [0, 1] \to \mathbb{R}_+$ is continuous, strictly increasing and strictly convex, with $c(0) = 0$ and $c'(0) = 0$. If the expert recommends the *safe* action ($a_t = 0$), no effort is exerted ($e_t \equiv 0$). The per-period *outcome* y_t is publicly observed and equals 1 (success) if and only if $a_t = 1$, $\omega_t = 1$, and the agent's effort succeeds; we adopt the reduced-form success technology

$$\mathbb{P}(y_t = 1 \mid a_t = 1, \omega_t, e_t) = e_t \cdot \omega_t, \qquad \mathbb{P}(y_t = 1 \mid a_t = 0) = 0.$$

Thus success requires both a favorable state and implementation effort; failure $y_t = 0$ is observed

[2]Formally, for $\theta \in \{H, L\}$ the likelihood ratio $\ell_\theta(s) \equiv f_\theta(s \mid 1)/f_\theta(s \mid 0)$ is strictly increasing in s, and H Blackwell-dominates L, e.g. there exists a stochastic kernel K with $f_L(\cdot \mid \omega) = \int K(\cdot \mid x)\, f_H(x \mid \omega)\, dx$ for each ω. The binary-signal special case $S = \{0, 1\}$ with accuracies $q_H > 1/2 \geq q_L$ is nested.

otherwise.[3]

The agent's period payoff is $y_t - c(e_t)$. The expert derives a flow payoff $u(\pi_t)$ from her current reputation (e.g. future demand or fees increasing in perceived competence) and discounts the future with factor $\delta \in (0,1)$. We allow an innocuous per-recommendation fee $\phi \geq 0$ if $a_t = 1$; it can be set to zero without affecting the qualitative results.

3.2 Public histories and timing

Public histories are sequences $h_t = ((a_0, y_0), \ldots, (a_{t-1}, y_{t-1})) \in \mathcal{H}_t \equiv (\{0,1\} \times \{0,1\})^t$ with the convention that when $a_\tau = 0$ the associated y_τ is a publicly observed failure ($y_\tau = 0$) and is uninformative about θ. The public belief $\pi_t = \mathbb{P}(\theta = H \mid h_t)$ is computed from Bayes' rule given equilibrium strategies.

The within-period timing is:

1. π_t is public; nature draws ω_t; the expert privately observes $s_t \sim f_\theta(\cdot \mid \omega_t)$.
2. The expert issues a recommendation $a_t \in \{0,1\}$ according to a strategy that may depend on s_t and π_t.
3. The agent, observing (a_t, π_t), forms interim beliefs about θ and ω_t (defined below) and, if $a_t = 1$, chooses effort $e_t \in [0,1]$; if $a_t = 0$, set $e_t = 0$.
4. The outcome $y_t \in \{0,1\}$ is realized and observed; the public belief updates to π_{t+1} by Bayes' rule.

3.3 Strategies and beliefs

A (behavioral) *recommendation strategy* for type θ is a measurable map

$$\alpha_\theta : \ S \times (0,1) \to \Delta(\{0,1\}), \qquad a \mapsto \alpha_\theta(a \mid s, \pi),$$

interpreted as the conditional probability of recommending a after observing (s, π). We restrict attention to *belief-based Markov* strategies that depend on the public history only via π; this is without loss for equilibria in our i.i.d. environment since π is a sufficient statistic for continuation payoffs.[4]

Given (α_H, α_L) and π, define the *advice likelihoods* for each type and state

$$r_\theta(a \mid \omega, \pi) \ \equiv \ \int_S \alpha_\theta(a \mid s, \pi) \, f_\theta(s \mid \omega) \, ds, \qquad \theta \in \{H, L\}, \ \omega \in \{0,1\}.$$

Unconditionally over the state,

$$\mathbb{P}(a \mid \theta, \pi) = \lambda \, r_\theta(a \mid 1, \pi) + (1-\lambda) \, r_\theta(a \mid 0, \pi).$$

The *interim type belief* (posterior on θ) after observing recommendation a is

$$\pi^{\text{rec}}(a; \pi) \ \equiv \ \mathbb{P}(\theta = H \mid a, \pi) \ = \ \frac{\pi \, \mathbb{P}(a \mid H, \pi)}{\pi \, \mathbb{P}(a \mid H, \pi) + (1-\pi) \, \mathbb{P}(a \mid L, \pi)}. \tag{1}$$

[3] We could allow a safe baseline payoff and a small amount of safe-outcome noise. In the baseline we take the safe outcome as uninformative about θ and normalize payoffs so that $y_t \in \{0,1\}$.

[4] A formal argument that π is a payoff-relevant state follows from the i.i.d. state and public observability of (a, y) together with the conditional independence of current signals given (θ, ω_t). See Appendix A for a statement and proof.

Table 1: Notation

Symbol	Meaning
$\theta \in \{H, L\}$	Expert's type: competent H vs. less competent L
$\pi_t \in (0,1)$	Public reputation at start of period t; belief that $\theta = H$
$\omega_t \in \{0,1\}$	Period state (i.i.d. over t), $\mathbb{P}(\omega_t = 1) = \lambda$
$s_t \in S$	Expert's private signal about ω_t
$a_t \in \{0,1\}$	Recommendation: risky 1 vs. safe 0
$e_t \in [0,1]$	Agent's implementation effort (chosen only if $a_t = 1$)
$y_t \in \{0,1\}$	Outcome (public): success 1 vs. failure 0
$c(\cdot)$	Effort cost
$u(\pi)$, δ	Expert's flow utility from reputation; discount factor $\delta \in (0,1)$
$\phi \geq 0$	Per-recommendation fee (relevant only if $a_t = 1$; can be 0)
$r_\theta(a \mid \omega, \pi)$	Advice likelihood induced by type θ given (ω, π)
$\pi^{\text{rec}}(a; \pi)$	Posterior on θ after observing recommendation a
$\lambda(a, \pi)$	Posterior on $\omega = 1$ after observing recommendation a
$e^*(a, \pi)$	Agent's best-response effort
$P_S(1, \pi)$	Success probability conditional on $a = 1$
$\pi^+(\pi),\ \pi^-(\pi)$	Post-outcome reputations after success/failure given $a = 1$
$V(\pi)$	Expert's continuation value at reputation π
$s^*(\pi)$	High-type risky-signal cutoff
$J^+(\pi),\ J^-(\pi)$	Log-likelihood jumps (success/failure) for the type posterior
$\rho(\pi),\ \Lambda(\pi)$	Recommendation and success intensities in the CT limit

The *interim success belief* (posterior on ω) after observing recommendation a is

$$\lambda(a, \pi) \;\equiv\; \mathbb{P}(\omega = 1 \mid a, \pi) \;=\; \frac{\sum_{\theta \in \{H,L\}} \mathbb{P}(\theta \mid \pi)\, \lambda\, r_\theta(a \mid 1, \pi)}{\sum_{\theta \in \{H,L\}} \mathbb{P}(\theta \mid \pi) \big[\lambda\, r_\theta(a \mid 1, \pi) + (1 - \lambda)\, r_\theta(a \mid 0, \pi)\big]}. \tag{2}$$

In particular, $\lambda(a, \pi)$ is well defined whenever $\mathbb{P}(a \mid \pi) > 0$; for off-path recommendations ($\mathbb{P}(a \mid \pi) = 0$) beliefs are specified by standard refinements (we maintain Bayes-consistent selections throughout).

3.4 Agent's best response

Given (a, π), the agent chooses effort $e \in [0,1]$ to maximize expected net benefit

$$e \;\in\; \arg\max_{x \in [0,1]} \; \lambda(a, \pi)\, x - c(x).$$

By strict convexity of c, the solution is unique and characterized by the first-order condition

$$c'\big(e^*(a, \pi)\big) = \lambda(a, \pi), \qquad \text{hence } e^*(a, \pi) = (c')^{-1}\big(\lambda(a, \pi)\big) \in [0, 1], \tag{3}$$

with $e^*(a, \pi)$ increasing in $\lambda(a, \pi)$. Under the common quadratic benchmark $c(x) = \frac{1}{2}x^2$, one has $e^*(a, \pi) = \lambda(a, \pi)$.[5]

[5] General primitives for $c(\cdot)$ (strictly convex C^2 and effort caps) and robustness to outcome/observability frictions (baseline success under $a = 0$, misclassification, partial observability of a) are developed in Online Appendix OA4.

3.5 Outcome probabilities and belief updates

If $a = 1$, the success probability (integrating over ω and effort) given (α_H, α_L) and π is

$$P_S(1,\pi) \;=\; \lambda(1,\pi)\cdot e^*(1,\pi). \tag{4}$$

If $a = 0$, we take outcomes as uninformative and set $P_S(0,\pi) = 0$.

For $a = 1$, the likelihood of outcomes conditional on the type is

$$\begin{aligned}\mathbb{P}(y = 1 \mid \theta, a = 1, \pi) &= e^*(1,\pi)\cdot \mathbb{P}(\omega = 1 \mid \theta, a = 1, \pi)\\ &= e^*(1,\pi)\cdot \frac{\lambda\, r_\theta(1 \mid 1,\pi)}{\lambda\, r_\theta(1 \mid 1,\pi) + (1-\lambda)\, r_\theta(1 \mid 0,\pi)},\\ \mathbb{P}(y = 0 \mid \theta, a = 1, \pi) &= 1 - \mathbb{P}(y = 1 \mid \theta, a = 1, \pi).\end{aligned} \tag{5}$$

Let $\pi^{\mathrm{rec}} = \pi^{\mathrm{rec}}(1;\pi)$ denote the interim type posterior after observing $a = 1$ but before observing y. The *post-outcome* reputations after success and failure, $\pi^+(\pi)$ and $\pi^-(\pi)$, satisfy the Bayes odds-ratio updates

$$\begin{aligned}\frac{\pi^+}{1-\pi^+} &= \frac{\mathbb{P}(y = 1 \mid H, a = 1, \pi)}{\mathbb{P}(y = 1 \mid L, a = 1, \pi)}\cdot \frac{\pi^{\mathrm{rec}}}{1-\pi^{\mathrm{rec}}},\\ \frac{\pi^-}{1-\pi^-} &= \frac{\mathbb{P}(y = 0 \mid H, a = 1, \pi)}{\mathbb{P}(y = 0 \mid L, a = 1, \pi)}\cdot \frac{\pi^{\mathrm{rec}}}{1-\pi^{\mathrm{rec}}}.\end{aligned} \tag{6}$$

If $a = 0$, we set $\pi_{t+1} = \pi_t$ (no new information about θ).[6]

3.6 Expert's problem and value function

Given a public reputation π, the expert anticipates the agent's best response and the induced belief transitions. We consider *Markov* (belief-based) strategies. The expert's continuation value $V(\pi)$ solves the Bellman equation

$$V(\pi) \;=\; \max_{a\in\{0,1\}} \Big\{ u(\pi) + \phi\,\mathbf{1}\{a = 1\} + \delta\cdot \mathbb{E}\big[V(\pi') \mid a, \pi\big] \Big\}, \tag{7}$$

where $\pi' = \pi$ if $a = 0$, and if $a = 1$ then

$$\mathbb{E}\big[V(\pi') \mid a = 1, \pi\big] = P_S(1,\pi)\, V\big(\pi^+(\pi)\big) + \big(1 - P_S(1,\pi)\big)\, V\big(\pi^-(\pi)\big),$$

with $P_S(1,\pi)$ as in (4) and $\pi^\pm(\pi)$ as in (6). A High-type expert's strategy will be shown to admit a cutoff representation in s (Theorem 5 below), while the Low type can mimic by suitable randomization to sustain on-path beliefs.

3.7 Equilibrium

We study *belief-based Markov Perfect Bayesian Equilibria* (MPBE).

Definition 1. An MPBE consists of (i) recommendation strategies (α_H, α_L) measurable in (s,π), (ii) an agent best-response $e^*(a,\pi)$ as in (3), (iii) a value function $V(\pi)$, and (iv) a belief system $(\pi^{\mathrm{rec}}(a;\pi),\ \lambda(a,\pi),\ \pi^\pm(\pi))$ defined by (1)–(6), such that:

[6] Allowing a small safe-outcome signal about θ is straightforward and leaves our main results intact; we focus on the transparent benchmark where safe outcomes carry no θ-information.

1. **Expert optimality.** For each type θ and each $\pi \in (0,1)$, $\alpha_\theta(\cdot \mid s, \pi)$ puts probability one on the set of $a \in \{0,1\}$ that maximize the type-θ Bellman objective (7) given e^* and the belief system.

2. **Agent optimality.** For each (a, π), $e^*(a, \pi)$ solves (3).

3. **Belief consistency.** On-path beliefs are given by Bayes' rule as in (1)–(6). Off-path beliefs are specified to satisfy standard consistency.

3.8 Regularity assumptions

We impose the following standing assumptions, used throughout:

A1 S is an interval of $\mathbb{R}$; for each $\theta \in \{H, L\}$ and $\omega \in \{0,1\}$, $f_\theta(\cdot \mid \omega)$ exists, is continuous in s, and satisfies MLRP; H is (strictly) more informative than L in the Blackwell order.

A2 The cost c is continuously differentiable, strictly increasing and strictly convex on $[0,1]$, with $c(0) = 0$ and $c'(0) = 0$.

A3 The flow utility $u : (0,1) \to \mathbb{R}$ is continuous and weakly increasing; $\delta \in (0,1)$.

A4 The initial belief $\pi_0 \in (0,1)$ and the good-state prior $\lambda \in (0,1)$ are common knowledge.

3.9 Remarks and special cases

The binary-signal case $S = \{0,1\}$ with accuracies $q_H > 1/2 \geq q_L$ provides a transparent benchmark: the High type recommends risk if and only if $s = 1$ when π is sufficiently low, and adopts a stricter rule as π rises; the Low type randomizes to match frequencies on-path. Under quadratic cost $c(e) = \frac{1}{2}e^2$, the agent's effort is $e^*(a, \pi) = \lambda(a, \pi)$ and the success probability for $a = 1$ is $[\lambda(1, \pi)]^2$, which highlights how higher reputation raises effort and thereby increases the diagnostic content of outcomes. The general continuous-signal case is handled by our MLRP/Blackwell assumptions and delivers a cutoff policy in s for the High type.

In the next section we establish the equilibrium cutoff structure (Theorem 5), show that the cutoff is (weakly) increasing in π (*reputational conservatism*), and characterize the resulting reputation dynamics and comparative statics.

4 Core Results

This section establishes the equilibrium structure and the main comparative statics. Throughout we maintain Assumptions A1–A4 from Section 3. Proofs that are longer or purely technical are deferred to the appendices as indicated.[7]

4.1 Reputation–effort feedback

We begin by formalizing the key complementarity between reputation and implementation effort.

Proposition 2. *Fix a belief-based Markov strategy profile (α_H, α_L) and $\pi \in (0,1)$. Let $\pi^{\text{rec}}(a; \pi)$ be the interim type belief after observing a recommendation $a \in \{0,1\}$, and let $e^*(a, \pi)$ be the agent's*

[7] Algorithms and replication details for all figures and numerical objects appear in Online Appendix OA3.

best response. Then $e^*(1,\pi)$ is strictly increasing in $\pi^{\text{rec}}(1;\pi)$, and $\pi^{\text{rec}}(1;\pi)$ is (weakly) increasing in π. Consequently, the conditional success probability after a risky recommendation,

$$P_S(1,\pi) \;=\; \lambda(1,\pi)\cdot e^*(1,\pi),$$

is (weakly) increasing in the public reputation π.

Proof. By (3), $c'(e^*(1,\pi)) = \lambda(1,\pi)$ and $c'' > 0$, hence $e^*(1,\pi)$ is strictly increasing in the interim success belief $\lambda(1,\pi)$. By Bayes' rule, $\lambda(1,\pi)$ is strictly increasing in the interim type belief $\pi^{\text{rec}}(1;\pi)$ (a higher probability of H raises the weight on the more informative recommendation distribution and thus the posterior that $\omega = 1$ under MLRP). Finally, $\pi^{\text{rec}}(1;\pi)$ is (weakly) increasing in π since it is a Bayes posterior in a binary-mixture prior. The conclusion on $P_S(1,\pi)$ follows from (4). □

4.2 Cutoff structure in the private signal

Let the High type's *risky-minus-safe continuation difference* at signal s and reputation π be

$$\Delta_H(s;\pi) \;\equiv\; \Big\{u(\pi) + \phi + \delta\cdot\mathbb{E}\big[V(\pi') \mid a=1, s, \pi\big]\Big\} \;-\; \Big\{u(\pi) + \delta\cdot V(\pi)\Big\}. \tag{8}$$

By construction, recommending $a = 1$ is optimal at (s,π) if and only if $\Delta_H(s;\pi) \geq 0$.

Lemma 3. *Under Assumption A1 and for any fixed $\pi \in (0,1)$, the mapping $s \mapsto \Delta_H(s;\pi)$ is (strictly) increasing. In particular, there exists a (possibly weak) cutoff $s^*(\pi) \in \overline{S}$ such that the High type's optimal recommendation is $a = 1$ if and only if $s \geq s^*(\pi)$.*

Proof. By (8) and (4),

$$\mathbb{E}\big[V(\pi') \mid a=1, s, \pi\big] = P_S(s;\pi)\cdot V\big(\pi^+(\pi)\big) + \big(1 - P_S(s;\pi)\big)\cdot V\big(\pi^-(\pi)\big),$$

where $P_S(s;\pi)$ is the success probability given (s,π) and $a = 1$. Assumption A1 (MLRP and that H is more informative than L) implies that the posterior $\mathbb{P}(\omega = 1 \mid s, a = 1, \pi)$ is strictly increasing in s. By Proposition 2, the agent's effort is strictly increasing in that posterior, hence $P_S(s;\pi)$ is strictly increasing in s. As $V(\pi^+) > V(\pi^-)$ (since V is increasing and $\pi^+ > \pi^-$ by Bayes' rule), it follows that $\mathbb{E}[V(\pi') \mid a = 1, s, \pi]$ is strictly increasing in s, and thus so is $\Delta_H(s;\pi)$. The threshold property follows by monotone selection on $\{s : \Delta_H(s;\pi) \geq 0\}$. □

Theorem 4. *Under (A1)–(A4) and $\delta \in (0,1)$, the expert's value function $V : [0,1] \to \mathbb{R}$ is increasing in π. Moreover, if Condition (6) (failure is at least as diagnostic as success at the cutoff) holds, then V is convex in π. In particular, for any fixed Markov cutoff policy $s(\cdot)$, the associated Bellman operator T^s is a monotone contraction that preserves convexity; hence its unique fixed point V^s is increasing and convex. At the equilibrium cutoff policy $s^*(\cdot)$, $V = V^{s^*}$ inherits these properties.*

Proof (short). Monotonicity: For any bounded V, the map $\pi \mapsto \phi\,\rho(\pi) + \delta\,\mathbb{E}[V(\pi') \mid \pi]$ is increasing because posteriors are increasing in prior odds under MLRP (A1) and V is evaluated at larger posteriors in the first-order stochastic dominance sense. Contraction holds with modulus δ in the sup norm. For convexity, fix a cutoff policy $s(\cdot)$. Under $a = 1$, the posterior takes $\{\pi^+(\pi), \pi^-(\pi)\}$ with probabilities that shift toward the failure branch as reputation (and thus effort) rises; Condition (6) ensures the induced posterior kernel becomes a mean-preserving spread in the convex order as π increases. Hence $\pi \mapsto \mathbb{E}[V(\pi') \mid \pi]$ is convex for convex V, and so is T^sV. The fixed point V^s is therefore increasing and convex; the equilibrium value $V = V^{s^*}$ inherits the properties. Full details appear in Appendix OA1. □

Theorem 5. *For each $\pi \in (0,1)$, in any MPBE there exists a measurable cutoff $s^*(\pi)$ such that the High type recommends $a = 1$ if and only if $s \geq s^*(\pi)$. The Low type can mimic by a (signal-measurable) mixed strategy that matches the induced recommendation frequencies so that Bayes' rule applies on-path.*

Proof. Lemma 3 yields existence of a cutoff for the High type. Given this cutoff and Assumption A1, one constructs the Low type's mixed strategy by equating the induced probabilities of $a = 1$ conditional on ω to ensure well-defined likelihood ratios for on-path belief updates; see Appendix A.4 for existence of an MPBE by a monotone fixed-point argument. □

4.3 Reputational conservatism

We now establish that the cutoff is (weakly) increasing in the public reputation: more highly reputed experts are more conservative.

Condition 6. Let $L^+(\pi) \equiv \frac{r_H(1|1,\pi)}{r_L(1|1,\pi)}$ and $L^-(\pi) \equiv \frac{r_H(1|0,\pi)}{r_L(1|0,\pi)}$ denote the outcome likelihood-ratio jumps after a risky recommendation at public reputation π, evaluated at the equilibrium cutoff $s^*(\pi)$. We assume

$$\log L^+(\pi) \;\leq\; -\log L^-(\pi),$$

with strict inequality on a set of π of positive measure. Equivalently, failures are at least as diagnostic of type as successes (in log-likelihood terms) near the cutoff.

Theorem 7. *Suppose Assumptions A1–A3 hold and V is increasing and convex on $(0,1)$. Then for any $\pi' < \pi''$ in $(0,1)$, one has $s^*(\pi') \leq s^*(\pi'')$.*

Proof outline. Fix π and write $\Delta_H(s;\pi)$ as in (8). The derivative of Δ_H with respect to π at fixed s has two effects: (i) a *diagnosticity effect* via $P_S(s;\pi)$, and (ii) a *baseline effect* via $V(\cdot)$. By Proposition 2, a higher π raises the interim effort and therefore increases the informativeness of outcomes; in particular, the gap $V(\pi^+) - V(\pi^-)$ is multiplied by a larger weight on the bad-news realization when failure occurs (since failure is more likely to be attributed to a bad signal/action when clients worked hard), which reduces the expected continuation value of taking risk, ceteris paribus. Convexity of V implies diminishing returns to further increases in π, so the marginal gain from success is smaller than the marginal loss from failure around high π. Formally, one shows that $\partial_\pi \Delta_H(s;\pi) \leq 0$ for all s; hence the crossing point $\{s : \Delta_H(s;\pi) = 0\}$ weakly increases in π. A complete argument is provided in Appendix B, which establishes the monotone comparative statics using supermodularity of the Bellman operator and Bayes-likelihood monotonicity induced by MLRP. □

Corollary 8. *If $s^*(\pi)$ is increasing in π, then for any two reputations $\pi' < \pi''$ and any signal s with $s^*(\pi') \leq s < s^*(\pi'')$, the High type recommends $a = 1$ at π' but $a = 1$ is not recommended at π''. Thus, as reputation falls, the expert is (weakly) more willing to endorse risk on intermediate signals.*

4.4 Binary-signal worked example

To make Condition 6 and its implications transparent, consider a binary signal $s \in \{h, \ell\}$ with symmetric accuracies. For type $\theta \in \{H, L\}$ and state $\omega \in \{0,1\}$,

$$\mathbb{P}(s = h \mid \omega = 1, \theta) = q_\theta, \qquad \mathbb{P}(s = h \mid \omega = 0, \theta) = 1 - q_\theta,$$

with $q_H > q_L \in (1/2, 1)$ so that H is more informative in the MLRP sense. The advice policy is a (degenerate) cutoff: $a = 1$ if $s = h$ and $a = 0$ if $s = \ell$, unless the risky–safe advantage at h is negative, in which case the expert never recommends risk.

Posterior odds shifts. Let $\text{odds}(\pi) = \pi/(1-\pi)$ and $\text{odds}^{-1}(x) = x/(1+x)$. When $a = 1$ is recommended at $s = h$, the outcome-based likelihood-ratio jumps are

$$L^+ \;=\; \frac{r_H(1 \mid 1)}{r_L(1 \mid 1)} = \frac{q_H}{q_L} \;>\; 1, \qquad L^- \;=\; \frac{r_H(1 \mid 0)}{r_L(1 \mid 0)} = \frac{1-q_H}{1-q_L} \;<\; 1,$$

so that, conditional on $a = 1$, the posterior odds update as

$$\text{odds}(\pi') = \begin{cases} \text{odds}(\pi)\, L^+, & y = 1, \\ \text{odds}(\pi)\, L^-, & y = 0. \end{cases}$$

Thus $\pi^+(\pi) = \text{odds}^{-1}(\text{odds}(\pi)L^+)$ and $\pi^-(\pi) = \text{odds}^{-1}(\text{odds}(\pi)L^-)$.

Diagnosticity asymmetry (Condition 6). In this binary benchmark,

$$\log L^+ + \log L^- \;=\; \log\!\left(\frac{q_H}{q_L} \cdot \frac{1-q_H}{1-q_L}\right) \;\le\; 0 \quad\iff\quad q_H(1-q_H) \;\le\; q_L(1-q_L).$$

Hence Condition 6 reduces to $q_H(1-q_H) \le q_L(1-q_L)$, which holds whenever the high type is sufficiently precise (i.e., q_H close enough to 1 relative to q_L). Intuitively, at the policy cutoff the failure branch ($y = 0$) moves posterior odds farther (down) than the success branch ($y = 1$) moves them up.

Cutoff in reputation and conservatism. Let $\Delta_H(s;\pi)$ denote the high type's risky–safe advantage at signal s (defined in Section 4). Because $\Delta_H(h;\pi)$ strictly decreases in π under Condition 6 (decreasing differences), there exists at most one reputation threshold $\bar{\pi} \in (0,1)$ solving

$$\Delta_H(h;\bar{\pi}) = 0.$$

For $\pi < \bar{\pi}$, the high type recommends risk iff $s = h$; for $\pi > \bar{\pi}$, she recommends safe even at $s = h$ (maximal conservatism). In typical parameter regions $\Delta_H(\ell;\pi) < 0$ for all π, so the only relevant margin is the h-signal; if $\Delta_H(\ell;\pi)$ becomes nonnegative at very low π, the risky region expands (weakening conservatism) in the direction predicted by Theorem 7.

Recommendation-only update. If observers also update on the recommendation itself, the recommendation-only likelihood ratio is

$$L^{\text{rec}} \;=\; \frac{\mathbb{P}(a = 1 \mid H)}{\mathbb{P}(a = 1 \mid L)} = \frac{\alpha\, q_H + (1-\alpha)(1-q_H)}{\alpha\, q_L + (1-\alpha)(1-q_L)}\,,$$

so that $\tilde{\pi}(\pi) = \text{odds}^{-1}\big(\text{odds}(\pi)\, L^{\text{rec}}\big)$. The overall continuation term in $\Delta_H(h;\pi)$ then combines $\tilde{\pi}(\pi)$ and the two outcome branches $\pi^{\pm}(\pi)$ as in Section 4. Since implementer effort $e^*(1,\pi)$ rises with π, the success probability $P_S(h;\pi)$ increases with π, which reinforces the asymmetry by making failures more diagnostic when reputation is high.

Interpretation (medicine). Think of $s = h$ as strong clinical/imaging evidence favoring surgery. A highly reputed surgeon (π high) elicits greater adherence/effort from the care team and patient, so an operation that fails is especially revealing (large $|\log L^-|$), whereas a success is less incrementally informative. The policy implication is conservative: at high standing, the bar for recommending surgery rises.

4.5 Comparative statics: precision, prior, and patience

We now promote three comparative statics to main-text propositions. Proofs are short and rely on strict single crossing of $\Delta_H(s;\pi)$ in s; full derivative formulas for the Gaussian benchmark appear in the Online Appendix (OA–C).

Proposition 9. *Fix $\pi \in (0,1)$ and primitives (A1)–(A4). If the high type's signal becomes more informative in the Blackwell sense and/or the low type's becomes weakly less informative (holding α, π fixed), then the risky cutoff $s^*(\pi)$ (weakly) decreases. In the Gaussian benchmark $s \mid (\theta,\omega) \sim \mathcal{N}(\mu_\omega, \sigma_\theta^2)$ with $\sigma_L > \sigma_H$, one has $\partial s^*/\partial(\mu_1 - \mu_0) < 0$, $\partial s^*/\partial \sigma_H > 0$, and $\partial s^*/\partial \sigma_L > 0$ whenever $s^*(\pi) \in (\inf S, \sup S)$.*

Proof. Let $G(s;t) \equiv \Delta_H(s;\pi;t)$ with parameter t indexing informativeness (higher t = more informative H, less informative L). By MLRP/Blackwell dominance, at any fixed s the outcome LLRs satisfy $L^+(t) \uparrow$ and $L^-(t) \downarrow$; hence the continuation term in Δ_H increases pointwise in s, so $\partial G/\partial t > 0$. Strict single crossing gives $\partial G/\partial s > 0$ at $s^*(\pi)$. The implicit-function formula yields

$$\frac{\partial s^*}{\partial t} = -\frac{\partial G/\partial t}{\partial G/\partial s} \; < \; 0.$$

The Gaussian signs follow from the normal tail derivatives of $r_\theta(1 \mid \omega)$ in $(\mu_1 - \mu_0, \sigma_H, \sigma_L)$. □

Proposition 10. *Holding $(\pi$, signals) fixed, $s^*(\pi)$ is (weakly) decreasing in the prior $\alpha \equiv \mathbb{P}(\omega = 1)$.*

Proof. Write $\Delta_H(s;\pi) = \phi + \delta\{\mathbb{E}[V(\pi') \mid a{=}1, s] - \mathbb{E}[V(\pi') \mid a{=}0]\}$. As α increases, $(a{=}1, y{=}1)$ histories become more likely under both types but relatively more under H, so L^+ rises and L^- falls (or falls less). Thus Δ_H increases pointwise in s, i.e., $\partial \Delta_H/\partial \alpha > 0$. With $\partial \Delta_H/\partial s > 0$ at s^*, the implicit-function formula gives $\partial s^*/\partial \alpha < 0$. □

Proposition 11. *Fix π and suppose $\phi \geq 0$ and (A1)–(A4) hold. Then $s^*(\pi)$ is (weakly) increasing in the discount factor δ. Moreover, if V_2 is a mean-preserving spread of V_1 (both increasing), then $s^*_{V_2}(\pi) \geq s^*_{V_1}(\pi)$.*

Proof. At s^*, $\Delta_H(s^*;\pi) = 0 = \phi + \delta\, \Psi(s^*;\pi)$ with $\Psi \equiv \mathbb{E}[V(\pi')] - V(\pi)$. Hence $\Psi(s^*;\pi) = -\phi/\delta \leq 0$. Holding s fixed, $\partial \Delta_H/\partial \delta = \Psi \leq 0$, so by the implicit-function formula and $\partial \Delta_H/\partial s > 0$ we get $\partial s^*/\partial \delta \geq 0$. For curvature, replacing V by a mean-preserving spread lowers Ψ (Jensen gap more negative under the diagnosticity asymmetry), thus lowers Δ_H pointwise in s, which raises the smallest s solving $\Delta_H \geq 0$. □

Remarks. (i) The patience effect is strict when $\phi > 0$ (the flow benefit of recommending risk), because then $\Psi(s^*;\pi) < 0$. (ii) In the Gaussian case, OA–C reports closed-form derivatives of the normal tail terms and the bonus-induced shifts used in policy calibration.

4.6 Reputation dynamics

We next characterize the belief process $\{\pi_t\}_{t \geq 0}$ along the equilibrium path.

Theorem 12. *Fix an MPBE with cutoff policy $s^*(\cdot)$ and let $\{\pi_t\}$ be the induced reputation process. Then:*

1. *If $\theta = H$, $(\pi_t)_{t\geq 0}$ is a submartingale with respect to the public filtration:*

$$\mathbb{E}[\pi_{t+1} \mid \mathcal{H}_t] = \pi_t \quad \textit{and} \quad \mathbb{P}(\pi_{t+1} \geq \pi_t \mid \mathcal{H}_t) > 0$$

   whenever $a_t = 1$ occurs with positive probability.

2. *If $\theta = L$, $(\pi_t)_{t\geq 0}$ is a supermartingale.*

3. *There exist $0 < \underline{\pi} < \overline{\pi} < 1$ (depending on primitives) such that*

$$\mathbb{P}\left(\liminf_{t\to\infty} \pi_t \leq \underline{\pi} \ \textit{or} \ \limsup_{t\to\infty} \pi_t \geq \overline{\pi}\right) = 1.$$

   In particular, with probability one the process hits a high-trust or low-trust region.

Proof sketch. Parts (1)–(2) follow from Bayes' rule and the law of iterated expectations: conditional on the true type, the posterior is a (super/sub)martingale. Part (3) uses that when $a_t = 1$ occurs with positive probability infinitely often, the innovation in the likelihood ratio has nontrivial variance (Appendix C shows a Doob decomposition for the log-odds process and applies optional stopping). When $a_t = 0$ along long stretches, the belief remains constant; however, under our cutoff structure and Assumption A1, either the process eventually visits the risky region infinitely often (yielding learning) or it remains trapped in a no-news region that is separated from the interior by the thresholds, which implies eventual absorption in a low- or high-trust basin. Full details are provided in Appendix C. □

Remark 13. If agents stop consulting the expert when $\pi_t < \underline{\pi}$ (endogenous exit), then π_t is absorbed at $\underline{\pi}$ with positive probability. Conversely, if failures at very high π do not fully overturn beliefs due to prior mass or noisy safe outcomes, then $\overline{\pi} < 1$ can be an attracting region. These cases are covered by the boundary characterization in Appendix C.

To visualize Theorem 12 and the boundary behavior discussed above, Figure 1 plots simulated posterior paths $\{\pi_t\}$ in the Gaussian–quadratic benchmark. The top panel shows trajectories when the true type is H; the bottom panel shows L. Flat segments correspond to periods with $a = 0$ (no experimentation), whereas jumps occur only after a risky recommendation ($a = 1$) and the ensuing outcome. The calibration is chosen for transparency rather than fit; qualitative patterns are robust.

The figure illustrates three predictions of the model. First, under $\theta = H$ the process drifts upward and hits the high-trust region with high probability, while under $\theta = L$ it drifts downward (Theorem 12). Second, learning occurs only when the expert recommends the risky action: flat stretches reflect $a = 0$, and the process advances by likelihood-ratio jumps after $a = 1$ (the visible steps in both panels). Third, jumps down following failures are larger at high reputation than at low reputation because implementers exert more effort when π is high; by contrast, the informational content of success does not increase with effort. This asymmetry is the microfoundation of our reputational conservatism result (Theorem 7) and the boundary-hitting behavior in the dynamics.

4.7 Comparative statics: precision, prior, and patience

We now promote three comparative statics to main-text propositions. Proofs are short and rely on strict single crossing of $\Delta_H(s;\pi)$ in s; full derivative formulas for the Gaussian benchmark appear in the Online Appendix (OA–C).

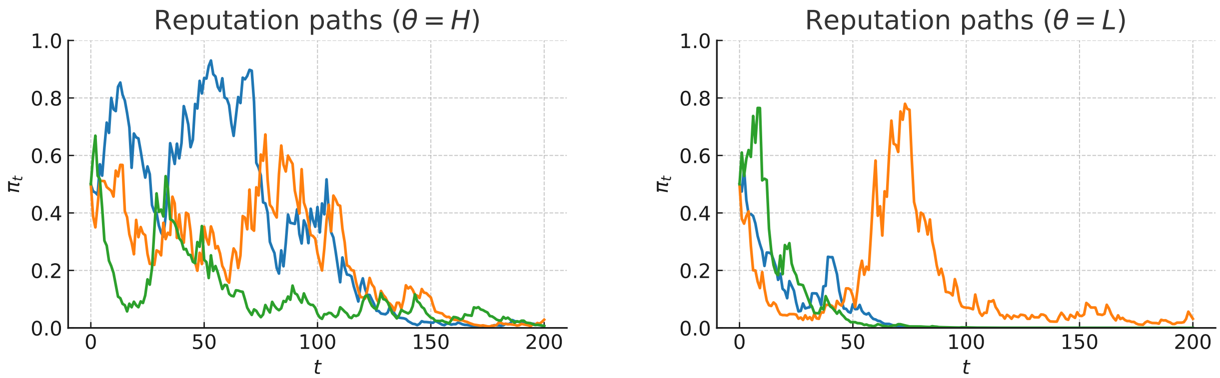

Figure 1: Simulated reputation dynamics in the Gaussian benchmark. Baseline parameters $(\mu_0, \mu_1, \sigma_H, \sigma_L, \lambda, \delta) = (0, 1, 1, 1.7, 0.5, 0.9)$. Each line is one replication starting from $\pi_0 = 0.5$. Flat segments arise when $a = 0$; jumps reflect $(a = 1)$ and the realized outcome.

Proposition 14. *Fix $\pi \in (0, 1)$ and primitives (A1)–(A4). If the high type's signal becomes more informative in the Blackwell sense and/or the low type's becomes weakly less informative (holding α, π fixed), then the risky cutoff $s^*(\pi)$ (weakly) decreases. In the Gaussian benchmark $s \mid (\theta, \omega) \sim \mathcal{N}(\mu_\omega, \sigma_\theta^2)$ with $\sigma_L > \sigma_H$, one has $\partial s^*/\partial(\mu_1 - \mu_0) < 0$, $\partial s^*/\partial\sigma_H > 0$, and $\partial s^*/\partial\sigma_L > 0$ whenever $s^*(\pi) \in (\inf S, \sup S)$.*

Proof. Let $G(s; t) \equiv \Delta_H(s; \pi; t)$ with parameter t indexing informativeness (higher t = more informative H, less informative L). By MLRP/Blackwell dominance, at any fixed s the outcome LLRs satisfy $L^+(t) \uparrow$ and $L^-(t) \downarrow$; hence the continuation term in Δ_H increases pointwise in s, so $\partial G/\partial t > 0$. Strict single crossing gives $\partial G/\partial s > 0$ at $s^*(\pi)$. The implicit-function formula yields

$$\frac{\partial s^*}{\partial t} = -\frac{\partial G/\partial t}{\partial G/\partial s} < 0.$$

The Gaussian signs follow from the normal tail derivatives of $r_\theta(1 \mid \omega)$ in $(\mu_1 - \mu_0, \sigma_H, \sigma_L)$. □

Proposition 15. *Holding $(\pi$, signals$)$ fixed, $s^*(\pi)$ is (weakly) decreasing in the prior $\alpha \equiv \mathbb{P}(\omega = 1)$.*

Proof. Write $\Delta_H(s; \pi) = \phi + \delta\{\mathbb{E}[V(\pi') \mid a{=}1, s] - \mathbb{E}[V(\pi') \mid a{=}0]\}$. As α increases, $(a{=}1, y{=}1)$ histories become more likely under both types but relatively more under H, so L^+ rises and L^- falls (or falls less). Thus Δ_H increases pointwise in s, i.e., $\partial\Delta_H/\partial\alpha > 0$. With $\partial\Delta_H/\partial s > 0$ at s^*, the implicit-function formula gives $\partial s^*/\partial\alpha < 0$. □

Proposition 16. *Fix π and suppose $\phi \geq 0$ and (A1)–(A4) hold. Then $s^*(\pi)$ is (weakly) increasing in the discount factor δ. Moreover, if V_2 is a mean-preserving spread of V_1 (both increasing), then $s^*_{V_2}(\pi) \geq s^*_{V_1}(\pi)$.*

Proof. At s^*, $\Delta_H(s^*; \pi) = 0 = \phi + \delta\,\Psi(s^*; \pi)$ with $\Psi \equiv \mathbb{E}[V(\pi')] - V(\pi)$. Hence $\Psi(s^*; \pi) = -\phi/\delta \leq 0$. Holding s fixed, $\partial\Delta_H/\partial\delta = \Psi \leq 0$, so by the implicit-function formula and $\partial\Delta_H/\partial s > 0$ we get $\partial s^*/\partial\delta \geq 0$. For curvature, replacing V by a mean-preserving spread lowers Ψ (Jensen gap more negative under the diagnosticity asymmetry), thus lowers Δ_H pointwise in s, which raises the smallest s solving $\Delta_H \geq 0$. □

Remarks. (i) The patience effect is strict when $\phi > 0$ (the flow benefit of recommending risk), because then $\Psi(s^*; \pi) < 0$. (ii) In the Gaussian case, OA–C reports closed-form derivatives of the normal tail terms and the bonus-induced shifts used in policy calibration.

4.8 Empirical predictions and measurement

This subsection translates our theory into testable predictions and clarifies how to measure the key objects in data. We focus on settings where outcomes and recommendations are observed and where implementer effort can be proxied (e.g., adherence/compliance in surgery and conservative care Haskard Zolnierek and DiMatteo (2009); Birkhäuer et al. (2017), surgeon skill/volume Birkmeyer et al. (2003, 2013), or trade size/turnover in finance Loh and Stulz (2011)).

4.8.1 Risk taking falls with reputation

By Theorem 7, the risky-signal cutoff is (weakly) increasing in current reputation, so the probability of a risky recommendation declines with reputation. A reduced-form test is

$$\mathbb{P}(a_{it} = 1) = \alpha + \beta \, \mathrm{Rep}_{it-1} + \gamma' X_{it} + \eta_i + \tau_t + \varepsilon_{it},$$

with $\beta < 0$, expert fixed effects η_i, and time (or case-mix) controls X_{it} and τ_t. In surgical data, $a_{it} = 1$ indicates operate vs. conservative care; Rep_{it-1} can be constructed from leave-one-out success rates or Bayesian smoothed scores using only past outcomes (excluding period t). In analyst data, $a_{it} = 1$ is a buy/upgrade vs. hold/downgrade.

4.8.2 Conditional success rises with reputation

Because implementers work harder when reputation is high, the success probability conditional on $a = 1$ increases with reputation (Proposition 2 and Theorem 12). Estimate

$$\mathbb{P}(y_{it} = 1 \mid a_{it} = 1) = \alpha + \theta \, \mathrm{Rep}_{it-1} + \gamma' X_{it} + \eta_i + \tau_t + \varepsilon_{it},$$

with $\theta > 0$. In medicine, success is a composite endpoint (e.g., no 30-day readmission/complication); in finance it can be excess return around recommendation implementation or forecast accuracy; effort proxies include adherence measures, therapy attendance, medication refill behavior, workflow timestamps, or trade size/execution quality.

4.8.3 Failures are more damaging at high reputation

Theorem 12 and the likelihood-ratio calculus imply that the *reputational* drop after a failure is larger when reputation is high, while the reputational gain after success is comparatively stable. Let $\Delta\mathrm{Rep}_{it} \equiv \mathrm{Rep}_{it} - \mathrm{Rep}_{it-1}$ be the ex post revision from an outcome following $a = 1$. Estimate

$$\Delta\mathrm{Rep}_{it} = \alpha + \phi_1 \, \mathbf{1}\{y_{it} = 1\} + \phi_0 \, \mathbf{1}\{y_{it} = 0\} + \psi \, \mathrm{Rep}_{it-1} \times \mathbf{1}\{y_{it} = 0\} + \gamma' X_{it} + \eta_i + \tau_t + \varepsilon_{it},$$

predicting $\psi < 0$. This can be implemented by constructing a transparent, data-driven reputation metric (see "Measurement" below) and running an event-study around outcome realizations.

4.8.4 Comparative statics.

Theorem 17 implies: (i) higher private-signal precision or a higher baseline viability increases experimentation; (ii) greater patience amplifies conservatism at high reputation. Proxy precision by the dispersion/variance of pre-decision diagnostics (imaging quality scores, lab-panel informativeness) or by analyst coverage informativeness; proxy viability by baseline risk scores. Test shifts in the *level* and *slope* of $\mathbb{P}(a = 1)$ in difference-in-differences designs around plausibly exogenous changes in information technology or guidelines.

The paths in Figure 1 illustrate these predictions: upward-drifting, jumpy trajectories under competence and downward-drifting ones under incompetence, with flat stretches when the expert forgoes risk and larger downward jumps after failures at high reputation. These patterns are the visual counterparts of P1–P3 and provide a guide for specification diagnostics in empirical work.

4.9 Comparative statics

We study how primitives shift the cutoff policy and the induced experimentation and learning.

Theorem 17. *Let $s^*(\pi; \cdot)$ denote the equilibrium cutoff as a function of primitives. Then:*

1. ***Signal precision.*** *If the High type's signal becomes more informative in the Blackwell order (holding the Low type fixed), then $s^*(\pi)$ decreases pointwise in π; the expert recommends risk on a (weakly) larger set of signals.*

2. ***Good-state prior*** λ. *The cutoff $s^*(\pi)$ is (weakly) decreasing in λ.*

3. ***Patience*** δ. *If u is sufficiently convex and V inherits convexity, then $s^*(\pi)$ is (weakly) increasing in δ.*

Proof outline. (1) A Blackwell improvement shifts the posterior $\mathbb{P}(\omega = 1 \mid s)$ upward in the monotone likelihood ratio order for each s, hence raises $P_S(s; \pi)$ and the expected value of taking risk; the single-crossing Lemma 3 implies the cutoff weakly falls. (2) A higher prior λ uniformly increases $\mathbb{P}(\omega = 1 \mid s)$ and thus $P_S(s; \pi)$, again lowering the cutoff. (3) With greater patience, the expert puts more weight on the reputational consequences; under convexity of V, the downside from failure grows faster in δ than the upside from success at high π, shifting $\Delta_H(s; \pi)$ downward and raising the cutoff. Appendix D contains full statements and proofs, including conditions under which (3) holds globally and examples (quadratic cost; logistic signals) in which the inequalities are strict. □

4.10 Welfare and experimentation rates

The cutoff policy $s^*(\pi)$ determines an *experimentation rate* at reputation π, namely $\rho(\pi) \equiv \mathbb{P}(s \geq s^*(\pi) \mid \pi, \theta = H)$ on-path. Combining Theorems 7 and 17 yields:

Corollary 18. *Experimentation rates are (weakly) decreasing in π and (weakly) increasing in signal precision and in λ. The expected one-step Kullback–Leibler information gain about θ after a risky recommendation is increasing in π via Proposition 2, but total information accumulation over time can be nonmonotone in π because high π reduces the frequency of experimentation.*

Proof. Immediate from the monotonicity of $s^*(\pi)$ in π and primitives, and from the decomposition of expected information gain into (i) the probability of experimentation and (ii) the informativeness of outcomes conditional on experimentation (which is increasing in π by Proposition 2). □

4.11 Equilibrium selection

The monotone methods used in Appendix A deliver extremal cutoffs when multiple fixed points exist. We select the smallest cutoff policy $s^*(\cdot)$ (the greatest experimentation equilibrium). This selection is natural for comparative statics and yields the sharpest form of reputational conservatism; it also corresponds to the limit of vanishing payoff perturbations (Appendix A.4).

The next section discusses extensions (endogenous exit, partial observability of effort, committees, continuous-time limits) and connects the results to applications. Full proofs are provided in Appendices A–D.

5 Discussion and Extensions

This section develops several extensions that either sharpen the main mechanisms or broaden applicability.[8] In each case we preserve the unified notation from Sections 3–4 and keep the public state one-dimensional (the reputation π). Formal proofs and additional details are provided in the appendices referenced below.

5.1 Flagship application: surgery vs. conservative care

We map the model to surgical decision-making. The expert is a surgeon (or surgical team lead); the risky action ($a = 1$) is operating; the safe action ($a = 0$) is conservative management. The period state $\omega \in \{0, 1\}$ collects anatomical/physiological factors that make surgery succeed when favorable ($\omega = 1$). The surgeon observes a private clinical signal s (history, exam, imaging) and recommends $a \in \{0, 1\}$. The outcome $y \in \{0, 1\}$ is publicly observed (e.g., risk-adjusted success/complication). Implementer effort is adherence by the patient and care team (prehab, preparation, post-op protocols, physical therapy, medication adherence), chosen after observing the recommendation and the surgeon's public reputation π. The recursive equilibrium (*Theorem 5*) yields a reputation-dependent cutoff $s^*(\pi)$; under *Condition 6*, the cutoff is (weakly) increasing in π (*Theorem 7*).

Measurement. Reputation π can be proxied by risk-adjusted historical performance (e.g., rolling outcome indices), surgeon/hospital quality scores, or volume-based measures; s is a summary of clinical predictors (e.g., an ML risk score from EHR and imaging). Effort proxies include adherence indices (medication possession ratio, refill gaps), therapy attendance, completion of pre-op prep and post-op protocols, and peri-operative checklist compliance. The success variable y is a risk-adjusted binary (e.g., no severe complication within 30/90 days).

Identification. Three sources of quasi-random variation are natural. (i) Assignment shocks: on-call/slot-availability rotations quasi-randomize surgeon–patient matches within service lines; this supports IV/event-study designs for how π shifts a and y holding s constant. (ii) Reputation shocks: public report releases or early idiosyncratic successes generate transitory jumps in π, allowing difference-in-differences around discrete events. (iii) Adherence instruments: congestion at therapy clinics, distance or weather shocks on rehab days, or pharmacy disruptions instrument implementer effort without directly shifting s or ω.

Testable predictions. (P1) Holding clinical risk s fixed, higher π reduces the propensity to recommend surgery (negative slope of a on π): reputational conservatism. (P2) Conditional on recommending surgery, higher π raises the hit rate $\mathbb{P}(y = 1 \mid a = 1)$ through higher adherence/effort. (P3) Positive reputation shocks (e.g., early successes) temporarily reduce surgical recommendations and increase the reputational cost of subsequent failures (larger posterior drops after $y = 0$). (P4) Monitoring that makes success less diagnostic (e.g., coarse pass/fail scorecards that move with a) depresses experimentation relative to benchmarks that preserve outcome diagnosticity; by contrast, pre-outcome transparency about effort (checklists, adherence dashboards) raises experimentation by strengthening the expected payoff of $a = 1$.

Design implications. Simple success-contingent bonuses can restore target experimentation by rotating the risky–safe trade-off without requiring complex menus; the calibration maps directly to observed hit-rate shifts. Committee settings (tumor boards) fit the k-of-n extension: higher k lowers pivotality and thus raises s^*, predicting fewer risky recommendations unless countervailed by improved implementation capacity.

[8] Persistent environments and correlated information are treated in Online Appendix OA5.

5.2 Endogenous exit and demand for advice

In many markets, consultation ceases when reputation is sufficiently low. We model this via an absorbing lower region: fix $\underline{\pi} \in (0,1)$ and assume that if the public reputation after any period satisfies $\pi_{t+1} < \underline{\pi}$, no further agents arrive and the game terminates with continuation value 0.

The Bellman equation (7) then becomes

$$V(\pi) = \max_{a \in \{0,1\}} \Big\{ u(\pi) + \phi \, \mathbf{1}\{a=1\} + \delta \Big[P_S(1,\pi) \, V(\pi^+(\pi)) + (1 - P_S(1,\pi)) \, \widehat{V}(\pi^-(\pi)) \Big] \Big\},$$

where $\widehat{V}(x) = V(x)$ if $x \geq \underline{\pi}$ and $\widehat{V}(x) = 0$ if $x < \underline{\pi}$. The advice cutoffs are defined as in Theorem 5.

Proposition 19. *Suppose Assumptions A1–A3 hold and $\underline{\pi} \in (0,1)$ is an absorbing lower boundary. Then:*

1. *A (possibly weak) cutoff policy exists for the High type at each $\pi \geq \underline{\pi}$, with the Low type mimicking on-path as before.*

2. *The cutoff $s^*(\pi)$ is (weakly) increasing in π on $[\underline{\pi}, 1)$ and has a (weak) kink at the locus where $\pi^-(\pi) = \underline{\pi}$.*

3. *Near the lower boundary, the sign of the change in $s^*(\pi)$ relative to the no-exit benchmark depends on the local comparison of the success jump $V(\pi^+) - V(\underline{\pi})$ and the status-quo value $V(\pi) - V(\underline{\pi})$: if $V(\pi^+) - V(\underline{\pi})$ is large relative to $V(\pi) - V(\underline{\pi})$, the expert* gambles for resurrection *(a lower cutoff than without exit); otherwise the policy is more conservative.*

The result highlights that exit shapes policy through a discrete downside at failure. In applications where a single failure at low π credibly triggers exit, the expert is tempted to take chances if the upside is meaningful, consistent with turnarounds by struggling advisers. Appendix E provides a full proof and examples under quadratic cost and logistic signal families.

5.2.1 Failure-driven exit risk

The next result formalizes that failures are more likely to trigger exit when current reputation is higher.

Corollary 20. *Fix an exit boundary $\underline{\pi} \in (0,1)$ and consider a period in which the expert recommends risk and the outcome is a failure $(a = 1, y = 0)$. For any signal realization s at which both reputations $\pi' < \pi''$ would recommend risk, the failure-updated posterior satisfies*

$$\pi^-(\pi'', s) \; < \; \pi^-(\pi', s).$$

Consequently, for any measurable set of signals on which both reputations recommend risk, the probability of crossing the boundary,

$$\mathbb{P}\{\pi^-(\pi, s) < \underline{\pi} \mid a = 1, \; y = 0\},$$

is (weakly) increasing in π, and strictly increasing whenever the risky region is nondegenerate and Assumption A1 holds on a set of positive measure.

Proof. Write odds as $O(x) = x/(1-x)$. After a risky recommendation and failure,

$$O(\pi^-(\pi, s)) \; = \; O(\pi^{\text{rec}}(1; \pi, s)) \cdot e^{J^-(\pi, s)}.$$

By Lemma 27, the recommendation-only posterior $\pi^{\text{rec}}(1;\pi,s)$ is increasing in π, hence $O(\pi^{\text{rec}}(1;\pi,s))$ is increasing in π. By Lemma 31, $J^-(\pi,s)$ is strictly decreasing in the implementer's effort $e^*(1,\pi;s)$, and $e^*(1,\pi;s)$ is increasing in π; thus $J^-(\pi,s)$ is (strictly) decreasing in π, so $e^{J^-(\pi,s)}$ is (strictly) decreasing in π. The product of an increasing term and a decreasing term is (strictly) decreasing here because e^{J^-} dominates in the log-odds update: $\partial_\pi \log O(\pi^-) = \partial_\pi \log O(\pi^{\text{rec}}(1)) + \partial_\pi J^- < 0$ whenever Assumption A1 implies $p_H(\cdot) - p_L(\cdot) > 0$ and e^* responds to π.[9] Hence $O(\pi^-(\pi,s))$ and therefore $\pi^-(\pi,s)$ are strictly decreasing in π for any fixed s in the overlapping risky region, proving the pointwise claim.

For the probability statement, fix $\pi' < \pi''$ and any measurable set S of signals on which both reputations recommend $a = 1$. The crossing event $\{\pi^-(\pi,s) < \underline{\pi}\}$ is an *upward-closed* set in s (failures become more damaging at higher s since e^* increases in s), and by the pointwise monotonicity above its indicator is weakly larger at π'' than at π' on S. Integrating against the (absolutely continuous) conditional failure densities on S yields a weakly larger probability at π'', with strict inequality whenever the risky region has positive measure and J^- is strictly decreasing. □

5.3 Partial observability of effort

The main model assumes that effort e is unobserved; only the binary outcome y is public. In many settings, implementation intensity is partially observable (e.g., audit trails, compliance logs). Let $\tilde{e} \in [0,1]$ be a public signal of effort drawn from a kernel $g(\tilde{e} \mid e)$ satisfying the Blackwell order in an informativeness parameter $\kappa \in [0,\infty)$ (higher κ means more precise monitoring; $\kappa = 0$ recovers the baseline with no information about e).

The agent's best response remains (3), since monitoring does not alter his current payoff. However, belief updates now condition on $(a, y, \tilde{e})$ rather than (a, y).

Proposition 21. *Under Assumptions A1–A3 and for any $\pi \in (0,1)$, the advice cutoff $s^*(\pi;\kappa)$ is (weakly) decreasing in the informativeness κ of effort monitoring. In particular, more transparent implementation makes failures less diagnostic of type (holding fixed a and the signal distribution), which raises the expected continuation value of recommending risk and lowers the cutoff.*

Proof. Fix $\pi \in (0,1)$ and consider the risky branch ($a = 1$). Let $\tilde{e}_\kappa$ denote the public effort monitor with informativeness parameter κ, ordered in the Blackwell sense (so $\kappa' > \kappa$ means $\tilde{e}_{\kappa'}$ is a strict refinement/less garbled version of $\tilde{e}_\kappa$). The implementer's effort is chosen before the outcome and is independent of θ conditional on (a,π,s); the signal $\tilde{e}_\kappa$ is therefore also independent of θ given (a,π,s).

Let $\pi^{-,\kappa}$ be the failure-updated posterior about θ under monitor $\tilde{e}_\kappa$. As the monitoring σ-field is refined from κ to κ', Bayes posteriors form a martingale:

$$\mathbb{E}[\pi^{-,\kappa'} \mid \tilde{e}_\kappa] \;=\; \pi^{-,\kappa} \quad \text{a.s.}$$

By Assumption A3, V is increasing and convex, hence by Jensen's inequality,

$$\mathbb{E}\big[V(\pi^{-,\kappa'}) \mid \tilde{e}_\kappa\big] \;\geq\; V\big(\mathbb{E}[\pi^{-,\kappa'} \mid \tilde{e}_\kappa]\big) \;=\; V(\pi^{-,\kappa}) \quad \text{a.s.}$$

Taking expectations yields $\mathbb{E}[V(\pi^{-,\kappa'})] \geq \mathbb{E}[V(\pi^{-,\kappa})]$. The success branch is unaffected by κ because J^+ does not depend on effort. Therefore the risky–safe continuation difference $\Delta_H(s;\pi,\kappa)$ is (weakly) larger at κ' than at κ for every s. By strict single crossing of $\Delta_H(\cdot;\pi,\kappa)$ in s (Appendix A), the cutoff is (weakly) lower under the more informative monitor: $s^*(\pi;\kappa') \leq s^*(\pi;\kappa)$, with strict inequality when the refinement is strict and the risky region has positive probability. □

[9] Equivalently, $\partial_\pi J^-(\pi,s) = -\frac{(p_H - p_L)\,\partial_\pi e^*}{(1-e^* p_H)(1-e^* p_L)} < 0$ by Lemma 31 and $\partial_\pi e^* > 0$.

Intuitively, when the market can separate "low effort" from "bad idea," reputational downside risk from a failure shrinks; this weakens the reputational force behind conservatism. Appendix F formalizes the Blackwell comparison and the induced likelihood-ratio ordering of $(y, \tilde{e})$.[10]

5.4 Multiple implementers and committees

Two extensions illustrate how aggregation on the recommendation side and scale on the implementation side shape experimentation.

In a committee with n experts where the risky action is taken if $\sum_i a^i \geq k$, an individual expert's outcome-based gain from recommending risk equals the single-expert gain scaled by the *pivot probability* $\zeta_k(\pi)$—the chance her recommendation flips the committee decision. Appendix G shows

$$\zeta_k(\pi) = \lambda \binom{n-1}{k-1} \rho_1(\pi)^{k-1} \big(1 - \rho_1(\pi)\big)^{n-k} + (1-\lambda) \binom{n-1}{k-1} \rho_0(\pi)^{k-1} \big(1 - \rho_0(\pi)\big)^{n-k},$$

with $\rho_\omega(\pi)$ the risky-recommendation frequency of a representative *other* expert conditional on ω. The cutoff structure carries over: in a symmetric equilibrium the High type uses a threshold $s^*(\pi; k)$ (Appendix G). Because only the outcome channel is scaled by $\zeta_k(\pi)$—the signaling content of one's own recommendation is unchanged—a higher threshold k lowers pivotality and thus raises the cutoff under mild conditions ($\rho_\omega(\pi) \leq k/n$), reducing experimentation; see Lemma 54 and Proposition 57.

Suppose that following a risky recommendation there are $m \geq 1$ implementers (e.g., a project team or multiple trial sites), each choosing effort $e_i \in [0, 1]$ at cost $c(e_i)$, independently across i. The project succeeds if at least one implementer succeeds, yielding the reduced-form success probability

$$\mathbb{P}(y = 1 \mid a = 1, \omega, \mathbf{e}) = 1 - \prod_{i=1}^{m} (1 - \omega\, e_i).$$

Under symmetry and common knowledge of (a, π), each implementer solves $c'(e_i) = \lambda(1, \pi)$, so $e_i = e^*(1, \pi)$ and

$$P_S^{(m)}(1, \pi) = 1 - \Big(1 - \lambda(1, \pi)\, e^*(1, \pi)\Big)^m.$$

Proposition 22. *Fix π. Then $P_S^{(m)}(1, \pi)$ is strictly increasing in m, and the High type's cutoff $s_m^*(\pi)$ is (weakly) decreasing in m. Thus larger implementation scale (more parallel effort) lowers the bar for recommending risk.*

The result quantifies a natural complementarity: reputational incentives are stronger when a risky recommendation triggers broader implementation capacity, because the chance of success—and hence the expected reputational gain—is higher. Appendix G provides the details and variants with success thresholds (e.g., k-out-of-m success) and correlated efforts.[11]

5.5 Policy design: success-contingent bonuses

We show that a simple success bonus implements any target experimentation rate at a given reputation. Let $\beta = (\beta_1, -\beta_0)$ denote transfers after $(a, y) = (1, 1)$ and $(1, 0)$, with limited liability $\beta_1 \geq 0$, $\beta_0 \geq 0$. The expert's risky–safe advantage becomes

$$\Delta_H^\beta(s; \pi) \;=\; \Delta_H(s; \pi) \;+\; \alpha\, \beta_1 \;-\; (1 - \alpha)\, \beta_0,$$

[10] Variants with pre- vs. post-outcome monitors, multi-level monitoring, action-dependent noise (Blackwell comparisons), and verifiable disclosure are developed in Online Appendix OA7.

[11] General monotone aggregation rules, exchangeable/heterogeneous committees, and comparative statics beyond k-of-n are in Online Appendix OA6.

so the induced cutoff $s^*_\beta(\pi)$ solves $\Delta^\beta_H(s^*_\beta(\pi);\pi) = 0$, and the induced risky rate is $\rho(\pi;\beta) \equiv \sum_{\omega\in\{0,1\}} \mathbb{P}(\omega)\, r_H(1 \mid \omega,\pi;\beta)$.

Theorem 23. *Fix $\pi \in (0,1)$, assume (A1)–(A4) and strictly positive signal densities at the cutoff. Under limited liability ($\beta_0 = 0$), the map $\beta_1 \mapsto \rho(\pi;\beta_1)$ is continuous and strictly increasing with*

$$\lim_{\beta_1\downarrow 0} \rho(\pi;\beta_1) = \rho(\pi;0), \qquad \lim_{\beta_1\uparrow\infty} \rho(\pi;\beta_1) = 1.$$

Hence, for any target $\rho^\star \in \big(\rho(\pi;0),\,1\big)$ there exists a unique $\beta_1(\rho^\star)$ such that $\rho(\pi;\beta_1(\rho^\star)) = \rho^\star$. With affine transfers $(\beta_1,-\beta_0)$, uniqueness extends to any $\rho^\star \in (0,1)$ via the net wedge $\alpha\beta_1 - (1-\alpha)\beta_0$.

Proof. By strict single crossing, $s \mapsto \Delta_H(s;\pi)$ is strictly increasing; adding $\alpha\beta_1$ shifts Δ_H up and yields a unique $s^*_{\beta_1}(\pi)$, strictly decreasing and continuous in β_1. Since $\rho(\pi;\cdot)$ is a continuous, strictly decreasing function of the cutoff (with positive densities at s^*), it is continuous and strictly increasing in β_1. As $\beta_1 \downarrow 0$, $s^*_{\beta_1}(\pi) \to s^*(\pi)$ and $\rho(\pi;\beta_1) \to \rho(\pi;0)$; as $\beta_1 \uparrow \infty$, $s^*_{\beta_1}(\pi) \to -\infty$ and $\rho(\pi;\beta_1) \to 1$. The affine case follows by replacing $\alpha\beta_1$ with $\alpha\beta_1 - (1-\alpha)\beta_0$. Full details are in Online Appendix OA8.

□

Proposition 24. *Let $S(\rho)$ denote per-experiment surplus from $a = 1$ (concave, differentiable). The implementer chooses $\beta_1 \geq 0$ to maximize*

$$U(\beta_1) \;=\; S(\rho(\pi;\beta_1)) \;-\; \alpha\,\rho(\pi;\beta_1)\,\beta_1,$$

with interior FOC

$$S'(\rho^\star) \;=\; \alpha\,\beta_1^\star \;+\; \alpha\,\rho^\star\,\frac{1}{\rho'(\beta_1^\star)}\,, \qquad \rho^\star = \rho(\pi;\beta_1^\star), \quad \rho'(\beta_1^\star) > 0.$$

If a budget B imposes $\alpha\,\rho(\pi;\beta_1)\,\beta_1 \leq B$, the same FOC holds with a positive multiplier added to the right-hand side at the optimum.

Proof. Envelope differentiation: $U'(\beta_1) = S'(\rho)\rho'(\beta_1) - \alpha[\rho + \beta_1\rho'(\beta_1)]$. Setting $U'(\beta_1^\star) = 0$ and dividing by $\rho'(\beta_1^\star) > 0$ yields the FOC. Concavity of S and monotonicity of ρ ensure global optimality. See OA–D for regularity and the budgeted case. □

One-line Gaussian mapping. With $s \mid (\theta,\omega) \sim \mathcal{N}(\mu_\omega,\sigma_\theta^2)$ and strictly positive densities at $s^*_{\beta_1}(\pi)$,

$$\rho'(\beta_1) \;=\; \frac{\alpha\Big[(1-\alpha)\,f_H(s^*_{\beta_1}(\pi)\mid 0) + \alpha\,f_H(s^*_{\beta_1}(\pi)\mid 1)\Big]}{\Delta'_H(s^*_{\beta_1}(\pi);\pi)} \;>\; 0, \tag{9}$$

and $\frac{ds^*_{\beta_1}}{d\beta_1} = -\alpha/\Delta'_H(s^*_{\beta_1}(\pi);\pi)$. These expressions give a plug-in calibration for $\beta_1^\star$ via Proposition 24. Derivations and numerical recipes are reported in OA–D.

5.6 Continuous-time approximation

Consider a vanishing-period approximation with period length $\Delta > 0$ and scale primitives so that (i) signals arrive with intensity $O(\Delta^{-1})$ and the High type's recommendation region induces a recommendation intensity $\rho(\pi)$; (ii) conditional on $a = 1$, successes arrive as a Poisson process with intensity $\Lambda(\pi) \equiv \lambda(1,\pi)\,e^*(1,\pi)$; and (iii) belief updates follow log-odds jumps of size $J^+(\pi)$ at success and $J^-(\pi)$ at failure as in (6). Let $L_t \equiv \log\frac{\pi_t}{1-\pi_t}$.

Proposition 25. *Under regularity conditions on $\rho, \Lambda, J^{\pm}$ (boundedness, Lipschitz continuity), as $\Delta \to 0$ the log-odds process L_t under the cutoff policy converges weakly to a jump-diffusion process with local drift and variance*

$$\mu(\pi) = \rho(\pi)\left[\Lambda(\pi)\, J^{+}(\pi) + (1 - \Lambda(\pi))\, J^{-}(\pi)\right],$$
$$\sigma^2(\pi) = \rho(\pi)\left[\Lambda(\pi)\,(J^{+}(\pi))^2 + (1 - \Lambda(\pi))\,(J^{-}(\pi))^2\right],$$

where $\pi = \frac{e^L}{1+e^L}$. In particular, $\mu(\pi)$ is (weakly) increasing in π by Proposition 2, capturing the reinforcement of reputation in continuous time.

The limit provides a tractable approximation for inference and policy experiments and connects our discrete-time structure to continuous-time reputation models. Appendix H states the functional central limit theorem used and verifies conditions in a canonical logistic-signal specification.

5.7 Heterogeneous agents and costs

If the arriving agent's cost function c_i varies across periods (e.g., drawn i.i.d. from a known distribution), the best response becomes $e_i^*(1, \pi) = (c_i')^{-1}(\lambda(1, \pi))$. The expected effort conditional on $(a = 1, \pi)$ is then $E[e_i^*(1, \pi)]$, which is still increasing in $\lambda(1, \pi)$. All results in Sections 4 carry through with e^* replaced by its expectation; in particular, Proposition 2 and Theorem 7 continue to hold.

5.8 Policy and design: contingent fees in dynamic settings

Our companion paper Lukyanov et al. (2025) analyzes a one-shot version with contingent compensation. In the dynamic environment here, a per-period success bonus $B \geq 0$ (paid only when $a = 1$ and $y = 1$) augments the continuation objective by $\delta\, B$ at success.

Proposition 26. *With a per-period success bonus B, the High type's cutoff $s^*(\pi; B)$ is (weakly) decreasing in B for all π. Moreover, for small B the marginal effect satisfies*

$$\frac{\partial s^*(\pi; B)}{\partial B}\Big|_{B=0} < 0,$$

and the effect is larger in magnitude at higher π (where reputational conservatism is strongest).

Thus, modest explicit incentives can partially undo reputational conservatism and restore experimentation at high reputation.[12]

These extensions illustrate that the main forces—reputation–effort feedback, state-dependent conservatism, and path-dependent learning—are robust. Transparency in implementation and larger implementation scale reduce conservatism; endogenous exit creates kinks and can induce gambling near the lower boundary; and the continuous-time approximation yields a convenient representation for empirical or quantitative work.

[12]Planner's welfare and the first-order condition linking $\partial\rho/\partial\beta$ to per-experiment surplus and the budget shadow price, as well as affine $(\beta_1, -\beta_0)$ contracts (including corners and implementation sets), are presented in Online Appendix OA8. See also the minimal-implementation formulas in (31).

6 Conclusion

This paper develops a dynamic theory of expert advice with reputation–effort feedback. A long-lived expert repeatedly recommends a risky action to short-lived implementers who optimally choose effort in response to the expert's current reputation. Embedding this complementarity in a recursive formulation delivers a tractable belief-based Markov equilibrium. We prove that the competent expert's policy is a cutoff in her private signal that depends on public reputation, that the cutoff is (weakly) increasing in reputation (*reputational conservatism*), and that reputation evolves as a submartingale for competent experts (supermartingale for incompetent ones), hitting boundary regions with probability one. Comparative statics clarify how signal precision, the good-state prior, and patience shift both the cutoff and the law of motion of beliefs.

Two forces drive these results. First, higher reputation elicits greater implementation effort, which raises success probabilities and—crucially—amplifies the informational content of *failures* without increasing the informational content of successes. Second, with an increasing and convex value of reputation, the downside from a visible failure grows faster than the upside from success at high reputation. These forces jointly tilt established experts toward caution and outsiders toward "gambling for resurrection," organizing a range of behaviors observed in advisory markets.

The framework is intentionally parsimonious—one-dimensional public state π, i.i.d. technological state ω, and MLRP signals—yet portable. It applies to financial analysts (buy vs. hold), surgeons (operate vs. conservative care), policy consultants (reform vs. status quo), and R&D leadership (greenlight vs. delay). In each case, reputational status shapes the intensity of implementation, which in turn shapes what outcomes reveal about ability. The equilibrium characterization yields transparent, testable predictions:

1. Holding fundamentals fixed, experts with higher current reputation recommend the risky action less frequently; when they do, success rates are higher because clients work harder.

2. Transitory shocks that raise perceived competence (e.g., an early success) reduce experimentation frequency but increase the informational content of each failure; the opposite holds for negative shocks.

3. Improved signal precision or a higher prior probability of viability increases experimentation across the reputation spectrum; greater patience amplifies conservatism at high reputation.

From a design perspective, reputation partly substitutes for explicit incentives: even without contracts, high reputation alone can sustain implementation effort. But reputation can also *over*-discipline, stalling experimentation when caution is excessive. Simple instruments can counterbalance this conservatism. We show that monitoring of implementation (transparent effort) and larger implementation scale reduce the cutoff by shrinking the reputational downside of failure. In dynamic environments, modest success-contingent transfers play the same role, rotating the risky–safe tradeoff most where reputational forces are strongest.

The analysis suggests empirical and quantitative agendas. On the empirical side, the model points to reduced-form tests using exogenous shifts in perceived competence (e.g., plausibly random early outcomes) and administrative measures of effort or compliance as proxies for the mechanism. On the structural side, the Gaussian–quadratic benchmark furnishes a workhorse specification for estimating π-dependent cutoffs, experimentation rates, and belief dynamics from panel advice data with outcomes.

The paper also clarifies the boundary of tractability and indicates natural extensions. Allowing persistent technological states or cross-period spillovers (ω_t Markov) would introduce an exploration–exploitation tradeoff atop reputational concerns, connecting our setup to dynamic bandits with dual learning about technology and ability. Richer communication (multi-message or continuous recommendations) and heterogeneous implementers (costs or observability) can be accommodated within the recursive approach at the cost of additional state variables. Finally, endogenous market frictions—entry, exit, and competition among experts—would convert the reputation–effort feedback into an industry-level selection mechanism.

We view the main contribution as conceptual and methodological: a clean recursive formulation that isolates the interaction between reputation and implementer effort, yields sharp comparative statics and dynamics, and remains close to data. We hope the framework can serve as a baseline for studying advice, disclosure, and implementation in organizations and markets where outcomes are jointly produced by ideas and execution.

A Proof of Theorem 5

This appendix proves that the competent expert ($\theta = H$) has a cutoff policy: for each public reputation $\pi \in (0,1)$ there is a threshold $s^*(\pi) \in S$ such that she recommends the risky action iff her private signal $s \geq s^*(\pi)$. The proof proceeds in three steps: (i) single crossing of the marginal value in s; (ii) existence and uniqueness of the zero; (iii) measurability and equilibrium consistency.

A.1 Preliminaries

Fix $\pi \in (0,1)$. Let $\Delta_H(s;\pi)$ denote the High type's risky-minus-safe continuation gain when others' strategies are fixed at the symmetric profile described in Section 3. With our timing and observables,

$$\Delta_H(s;\pi) = \phi + \delta \left\{ \underbrace{\mathbb{E}\big[V(\pi') \mid a=1,\pi,s\big]}_{\text{recommendation \& outcome}} - \underbrace{\mathbb{E}\big[V(\pi') \mid a=0,\pi,s\big]}_{\text{recommendation only}} \right\},$$

where V is the continuation value and π' the next-period reputation. Define the recommendation-only posterior $\pi^{\text{rec}}(a;\pi,s)$ (from observing a) and, conditional on $a=1$, the outcome posteriors $\pi^+(\pi,s)$ and $\pi^-(\pi,s)$ after success/failure. Let $P_S(s;\pi)$ be the success probability under $a=1$. Then

$$\begin{aligned}\Delta_H(s;\pi) = \phi + \delta\Big(& P_S(s;\pi)\, V(\pi^+(\pi,s)) \\ & + \big(1 - P_S(s;\pi)\big)\, V(\pi^-(\pi,s)) - V(\pi^{\text{rec}}(0;\pi,s)) \Big).\end{aligned} \tag{10}$$

A.2 Likelihood ratios and posterior monotonicity

Let $f_\theta(\cdot \mid \omega)$ denote the density of s under type $\theta \in \{H, L\}$ and state $\omega \in \{0,1\}$. Assumption A1 (MLRP and Blackwell domination) implies the *type* log-likelihood ratio

$$\Lambda_\theta(s) \equiv \log \frac{f_H(s \mid \omega)}{f_L(s \mid \omega)}$$

is (strictly) increasing in s for each ω; likewise, the *state* likelihood ratio

$$\Lambda_\omega(s) \equiv \log \frac{f_\theta(s \mid \omega = 1)}{f_\theta(s \mid \omega = 0)}$$

is (strictly) increasing in s for each θ. Bayes' rule therefore gives:

Lemma 27. *For every $\pi \in (0,1)$: (i) the recommendation-only posterior $\pi^{\text{rec}}(1;\pi,s)$ is (strictly) increasing in s and $\pi^{\text{rec}}(0;\pi,s)$ is (strictly) decreasing in s; (ii) the posterior success belief $\lambda(1,\pi;s)$ after $a=1$ is (strictly) increasing in s; hence the agent's best-response effort $e^*(1,\pi;s)$ solving $c'(e)=\lambda(1,\pi;s)$ is (strictly) increasing in s.*

Proof. Part (i) follows since $\mathbb{P}(a=1 \mid s,\theta)$ is (weakly) higher under H than L for larger s in any best response (MLRP ensures the monotone likelihood ratio for the event $a=1$ in s), so the posterior odds about θ after observing a are increasing in s. Part (ii) is the standard MLRP implication for the ω-posterior given s. Strict monotonicity obtains whenever the signal has strictly increasing likelihood ratios. The mapping e^* is increasing because c' is strictly increasing by Assumption A2. □

The log-likelihood jumps in π after outcomes when $a=1$ are

$$J^+(\pi,s) = \log \frac{p_H(\pi,s)}{p_L(\pi,s)}, \tag{11}$$

$$J^-(\pi,s) = \log \frac{1-e^*(1,\pi;s)\,p_H(\pi,s)}{1-e^*(1,\pi;s)\,p_L(\pi,s)}, \tag{12}$$

where $p_\theta(\pi,s) \equiv \mathbb{P}(\omega=1 \mid a=1,\theta,\pi,s)$. Note J^+ is independent of effort, while J^- is strictly decreasing in e^* and hence in s.

A.3 Single crossing and existence

Lemma 28. *Under Assumptions A1–A3, $s \mapsto \Delta_H(s;\pi)$ is continuous and strictly increasing on S.*

Proof. Continuity follows from continuity of the primitives and the dominated convergence theorem. To prove monotonicity, examine (8). First, $s \mapsto \pi^{\text{rec}}(0;\pi,s)$ is decreasing by Lemma 27(i) while V is increasing (A3), so $s \mapsto -V(\pi^{\text{rec}}(0;\pi,s))$ is increasing. Second, $s \mapsto P_S(s;\pi)$ is increasing because both $\lambda(1,\pi;s)$ and $e^*(1,\pi;s)$ increase with s (Lemma 27(ii) and A2). Third, conditional on $a=1$, $s \mapsto \pi^+(\pi,s)$ increases with s and $s \mapsto \pi^-(\pi,s)$ decreases with s because J^+ is independent of e^* while J^- becomes more negative as e^* rises; since V is increasing and convex (A3), the map

$$s \;\longmapsto\; P_S(s;\pi)\,V(\pi^+(\pi,s)) + \big(1-P_S(s;\pi)\big)\,V(\pi^-(\pi,s))$$

is increasing. Adding the three components yields the claim; strictness follows from strict MLRP and strict convexity of c (hence strict increase of e^*). □

Lemma 29. *For each π, there exist $s_-, s_+ \in S$ with $s_- < s_+$ such that $\Delta_H(s_-;\pi) < 0$ and $\Delta_H(s_+;\pi) > 0$.*

Proof. As $s \to \inf S$, $\lambda(1,\pi;s)$ and e^* become small, so $P_S(s;\pi)$ is arbitrarily small, while the recommendation-only posterior after $a=0$ approaches its upper support; hence the Jensen gain from switching to risk is dominated by the safe continuation and $\Delta_H < 0$ for s low enough. As $s \to \sup S$, both $\lambda(1,\pi;s)$ and e^* approach their upper supports so that $P_S(s;\pi)$ and $V(\pi^+)$ make the experimentation option dominant; thus $\Delta_H > 0$ for s high enough. The argument uses that V is increasing and bounded on $(0,1)$. □

Proof of Theorem 5. By Lemma 28, $s \mapsto \Delta_H(s;\pi)$ is strictly increasing and continuous; by Lemma 29 it changes sign. Hence there exists a unique $s^*(\pi) \in S$ with $\Delta_H(s^*(\pi);\pi)=0$, and the High type's best reply is the cutoff policy $a=1$ iff $s \geq s^*(\pi)$. Measurability and existence of a symmetric MPBE follow by standard fixed-point arguments (Appendix A.4). □

A.4 Existence of MPBE

Under Assumptions A1–A4 and the cutoff structure of Theorem 5, the set of symmetric strategy profiles with cutoff policies is a nonempty, compact, convex lattice under the pointwise order; the induced best–response correspondence is nonempty, convex–valued, and monotone (single–crossing of the marginal value in s implies monotone best responses). By Tarski's fixed–point (or, equivalently, by Kakutani applied to a compact convex subset of L^∞), a symmetric MPBE exists. The continuity of posteriors and V ensures the usual measurable–selection requirements.

B Proof of Theorem 7

We prove that the cutoff $s^*(\pi)$ is (weakly) increasing in the public reputation π. The proof uses Topkis' monotone comparative statics: it suffices to show that $\Delta_H(s;\pi)$ has *decreasing differences* in (s,π) (i.e., the cross-partial $\partial_{s\pi}\Delta_H \le 0$ in the sense of monotone differences).

B.1 Decomposing the marginal value

Write $\Delta_H(s;\pi) = \phi + \delta\{\Gamma_1(s;\pi) - \Gamma_0(s;\pi)\}$ with

$$\Gamma_0(s;\pi) \equiv V\big(\pi^{\text{rec}}(0;\pi,s)\big),$$
$$\Gamma_1(s;\pi) \equiv P_S(s;\pi)\,V\big(\pi^+(\pi,s)\big) + \big(1-P_S(s;\pi)\big)V\big(\pi^-(\pi,s)\big).$$

The dependence on π enters through (i) the recommendation posteriors π^{rec} (because the prior odds enter Bayes' rule), and (ii) P_S and the outcome posteriors $\pi^\pm$ via the agent's effort best reply $e^*(1,\pi;s)$ and the jumps $J^\pm$.

B.2 Decreasing differences

Lemma 30. *The map $(s,\pi)\mapsto\Gamma_0(s;\pi)$ has decreasing differences.*

Proof. By Bayes' rule, $\pi^{\text{rec}}(0;\pi,s) = \sigma(\text{logit}(\pi) + J_0^{\text{rec}}(s))$, where $\sigma(x) = \frac{1}{1+e^{-x}}$ and $J_0^{\text{rec}}(s)$ is the (type) log-likelihood ratio of observing $a=0$. Lemma 27(i) implies $J_0^{\text{rec}}(s)$ is decreasing in s. The map $(x,z)\mapsto\sigma(x+z)$ has decreasing differences because σ is increasing and concave. Composing with the increasing V preserves the property (Topkis). □

B.3 LLR asymmetry and effort

We record a basic property used repeatedly in the decreasing–differences proof.

Lemma 31. *When $a=1$, the success jump $J^+(\pi,s)$ is independent of effort, while the failure jump $J^-(\pi,s)$ is strictly decreasing in the implementer's effort $e^*(1,\pi;s)$. Consequently, as s or π rise (both increase e^*), failures become (weakly) more damaging to reputation, whereas the informativeness of success is unchanged.*

Proof. By (11), $J^+(\pi,s) = \log\big(p_H(\pi,s)/p_L(\pi,s)\big)$ does not involve e^*. By the same display,

$$J^-(\pi,s) = \log\frac{1-e^*(1,\pi;s)\,p_H(\pi,s)}{1-e^*(1,\pi;s)\,p_L(\pi,s)},$$

whose derivative with respect to $e^*(1,\pi;s)$ equals $-\big(p_H(\pi,s)-p_L(\pi,s)\big)/\big((1-e^*p_H)(1-e^*p_L)\big)<0$ by Assumption A1. Since $e^*(1,\pi;s)$ increases in s and π (Lemma 27), the comparative statics follow. □

Lemma 32. *The map $(s,\pi) \mapsto \Gamma_1(s;\pi)$ has decreasing differences.*

Proof. Three effects matter.

(i) *Effort amplification.* By Lemma 27(ii), $s \mapsto \lambda(1,\pi;s)$ is increasing. For fixed s, the agent's success belief $\lambda(1,\pi;s)$ is also increasing in π because the recommendation posterior $\pi^{\text{rec}}(1;\pi,s)$ is increasing in π and raises the weight on the more informative type H; therefore the best response $e^*(1,\pi;s)$ is increasing in both s and π. As a result, $P_S(s;\pi) = \lambda(1,\pi;s)\, e^*(1,\pi;s)$ is supermodular in (s,π).

(ii) *LLR asymmetry.* The success jump J^+ is independent of effort (hence of π), while the failure jump J^- is strictly decreasing in $e^*(1,\pi;s)$ (see (12)); thus s and π jointly make failures more damaging but do not change the informativeness of success. This creates decreasing differences in the pair (s,π) for the outcome-updated value $V(\pi^{\pm}(\pi,s))$: higher π reduces the marginal gain from raising s because the downside risk of a failure (whose probability rises with s) becomes more severe.

(iii) *Jensen's inequality.* Since V is increasing and convex, the map $q \mapsto q\,V(\pi^+) + (1-q)\,V(\pi^-)$ has decreasing differences in (q,π) when $\pi^{\pm}$ move apart in π mainly through the more negative J^-. Combining (i)–(iii) delivers decreasing differences of Γ_1. □

Proof of Theorem 7. By Lemmas 30 and 32, $\Delta_H(s;\pi) = \phi + \delta\{\Gamma_1(s;\pi) - \Gamma_0(s;\pi)\}$ has decreasing differences in (s,π). Lemma 28 gives strict single crossing in s. Topkis' monotone comparative statics theorem then implies the smallest (and here unique) solution $s^*(\pi)$ to $\Delta_H(s;\pi) = 0$ is (weakly) increasing in π. □

C Proof of Theorem 12

We prove that public reputation $\{\pi_t\}$ is a submartingale when the true type is H and a supermartingale when the true type is L, and that it reaches trust regions with probability one. The key step is an exact expression for the conditional drift of π_t in terms of likelihood ratios.

C.1 One-step update and conditional drift

Let L_{t+1} be the one-step likelihood ratio between histories under H and L (recommendation and, if applicable, outcome) given $\mathcal{F}_t$:

$$L_{t+1} \;\equiv\; \frac{\mathbb{P}(\text{obs}_{t+1} \mid \theta = H, \mathcal{F}_t)}{\mathbb{P}(\text{obs}_{t+1} \mid \theta = L, \mathcal{F}_t)} \;\in (0,\infty).$$

Bayes' rule yields the posterior odds recursion

$$\frac{\pi_{t+1}}{1-\pi_{t+1}} = \frac{\pi_t}{1-\pi_t}\, L_{t+1}, \qquad \pi_{t+1} = g(\pi_t, L_{t+1}), \quad g(\pi, L) \equiv \frac{\pi L}{\pi L + 1 - \pi}.$$

Hence

$$\pi_{t+1} - \pi_t \;=\; \pi_t(1-\pi_t)\,\frac{L_{t+1} - 1}{\pi_t L_{t+1} + 1 - \pi_t}. \tag{13}$$

Lemma 33. *For every t, under the probability measure with $\theta = H$,*

$$\mathbb{E}[\pi_{t+1} - \pi_t \mid \mathcal{F}_t, \theta = H] \;\geq\; 0,$$

with equality iff $L_{t+1} = 1$ almost surely (i.e., the period-t observation is uninformative about type). Symmetrically, under $\theta = L$ the conditional drift is ≤ 0, with equality iff $L_{t+1} = 1$ almost surely.

Proof. Let $a \equiv \pi_t \in (0,1)$ and $h(x) \equiv (x-1)/(ax+1-a)$. Then (13) gives $\pi_{t+1}-\pi_t = a(1-a)\,h(L_{t+1})$. Under $\theta = H$, the law of L_{t+1} has density $d\mathbb{P}_H = L_{t+1}\,d\mathbb{P}_L$ with respect to the law under L. Thus

$$\mathbb{E}_H[h(L_{t+1}) \mid \mathcal{F}_t] \;=\; \mathbb{E}_L\!\left[\frac{L_{t+1}(L_{t+1}-1)}{aL_{t+1}+1-a}\,\Big|\,\mathcal{F}_t\right].$$

The integrand is nonnegative because $x \mapsto \frac{x(x-1)}{ax+1-a}$ is nondecreasing and vanishes at $x = 1$. Hence the conditional expectation is ≥ 0, with equality only if $L_{t+1} = 1$ a.s. under L (and therefore under H). The statement under $\theta = L$ follows by symmetry (replace L_{t+1} with $1/L_{t+1}$). □

Proof of Theorem 12. The sub-/supermartingale property follows by Lemma 33 and boundedness of $\pi_t \in [0,1]$. For boundary hitting, note that if information arrives infinitely often with positive probability (i.e., $\mathbb{P}(L_{t+1} \neq 1 \text{ i.o.}) > 0$), then the sum of nonnegative conditional drifts $\sum_t \mathbb{E}_H[\pi_{t+1} - \pi_t \mid \mathcal{F}_t]$ diverges on that event unless π_t enters an arbitrarily small neighborhood of 1; Doob's submartingale convergence theorem then implies $\pi_t \to 1$ almost surely under H. A symmetric argument yields $\pi_t \to 0$ under L. If informative periods cease after some (random) time, the process is eventually constant and trivially hits a boundary region. In either case, the process reaches trust neighborhoods with probability one. □

C.2 LLR asymmetry and effort (for reference)

When the risky action is recommended, L_{t+1} decomposes as $L_{t+1} = L^{\text{rec}}_{t+1} \cdot L^{\text{out}}_{t+1}$, the product of the recommendation LLR and the outcome LLR. The latter has jumps J^+ and J^- as in (11)–(12); crucially, J^+ is independent of effort while J^- is strictly decreasing in the agent's effort $e^*(1, \pi_t; s_t)$. This is the microfoundation for the greater reputational downside of failures at high reputation used in Theorem 7.

D Comparative Statics Proofs

This appendix provides proofs for Theorem 17. We study how the equilibrium cutoff $s^*(\pi)$ moves with primitives.

D.1 Signal precision (Blackwell order)

Let the High type's signal distribution be indexed by a precision parameter κ, with $\kappa' \succcurlyeq \kappa$ in the Blackwell order (i.e., $f_H(\cdot \mid \omega; \kappa')$ is more informative about ω than $f_H(\cdot \mid \omega; \kappa)$ for each ω). The Low type's distribution is held fixed.

Lemma 34. *Fix (s, π). Under $\kappa' \succcurlyeq \kappa$, the posterior success probability $\mathbb{P}(\omega = 1 \mid s, a = 1, \pi; \kappa')$ is (weakly) higher than $\mathbb{P}(\omega = 1 \mid s, a = 1, \pi; \kappa)$ for all s in the MLR order. Consequently, $P_S(s; \pi)$ is (weakly) higher under κ' than under κ.*

Proof. Blackwell dominance implies that for any prior over ω and any likelihood ratio test based on s, the posterior places (weakly) more mass on $\omega = 1$ under the more informative experiment. As $e^*(1, \pi)$ is increasing in the success belief and the signal of effort is unchanged, the reduced-form success probability increases. □

Proposition 35. *If $\kappa' \succcurlyeq \kappa$ (High type more informative), then for all $\pi \in (0,1)$ the High type's cutoff satisfies $s^*(\pi; \kappa') \leq s^*(\pi; \kappa)$.*

Proof. From (10), $\Delta_H(s;\pi)$ is strictly increasing in $P_S(s;\pi)$. By Lemma 34, $P_S(s;\pi)$ shifts up under κ', so the zero set $\{s : \Delta_H(s;\pi) = 0\}$ weakly shifts left, and the monotone selection implies the cutoff falls. □

D.2 Good-state prior λ

Proposition 36. *For all $\pi \in (0,1)$, the cutoff $s^*(\pi)$ is (weakly) decreasing in the prior probability λ of a favorable state.*

Proof. A larger λ raises $\mathbb{P}(\omega = 1 \mid s, a = 1, \pi)$ for all s and hence raises $P_S(s;\pi)$ pointwise. The argument of Proposition 35 applies verbatim. □

D.3 Patience δ

We establish the comparative static under a sufficient curvature condition.

Proposition 37. *Suppose V is increasing and convex on $(0,1)$ (Assumption A3 plus convexity). Then for all $\pi \in (0,1)$, the cutoff $s^*(\pi)$ is (weakly) increasing in the discount factor δ.*

Proof. From (10), $\Delta_H(s;\pi) = \phi + \delta\big(P_S(s;\pi)\,V(\pi^+) + (1 - P_S(s;\pi))\,V(\pi^-) - V(\pi)\big)$. The derivative with respect to δ is the bracketed term, which is *negative* at high π under convexity because $V(\cdot)$ exhibits diminishing returns in π and, by Appendix B.3, failures become more damaging (reduce π more) as e^* rises with π. Globally, Appendix B shows that Δ_H has decreasing differences in $(1,\pi)$; multiplying a function with decreasing differences by a larger scalar δ preserves the monotone comparative static: the zero-crossing in s shifts (weakly) upward. A direct Topkis argument (decreasing differences in (a,δ)) delivers $s^*(\pi)$ nondecreasing in δ. □

D.4 A worked benchmark

This subsection fully develops the Gaussian–quadratic benchmark and derives the comparative statics in signal precision, the good-state prior, and patience. Throughout we fix $c(e) = \frac{1}{2}e^2$, so that

$$e^*(1,\pi;s) \;=\; \lambda(1,\pi;s), \tag{14}$$

where the dependence on the High type's cutoff $s^*(\pi)$ enters through the advice likelihoods below.

D.4.1 Signal distributions, advice likelihoods, and posteriors

Conditional on the state $\omega \in \{0,1\}$, the expert's private signal s is Gaussian:

$$s \mid (\theta = H, \omega) \sim \mathcal{N}(\mu_\omega, \sigma_H^2), \qquad s \mid (\theta = L, \omega) \sim \mathcal{N}(\mu_\omega, \sigma_L^2),$$

with $\mu_1 - \mu_0 > 0$ and $\sigma_L^2 > \sigma_H^2$ (so H is strictly more informative than L in the Blackwell sense). Fix a belief $\pi \in (0,1)$. In equilibrium, the High type uses a cutoff $s^*(\pi)$ (Theorem 5), recommending risk if and only if $s \geq s^*(\pi)$. Define for $\omega \in \{0,1\}$:

$$\begin{aligned} A(s) &\equiv r_H(1 \mid \omega = 1, \pi) = 1 - \Phi\Big(\frac{s - \mu_1}{\sigma_H}\Big), \\ B(s) &\equiv r_H(1 \mid \omega = 0, \pi) = 1 - \Phi\Big(\frac{s - \mu_0}{\sigma_H}\Big), \end{aligned} \tag{15}$$

where Φ is the standard normal cdf. Note that $A(s) > B(s)$ and both are strictly decreasing in s.

For the Low type, to keep on-path beliefs Bayes-consistent while keeping the algebra transparent, we adopt a simple frequency-matching *signal-independent* mixed strategy: upon any (s, π), the Low type recommends $a = 1$ with probability

$$p(\pi) \in (0,1), \qquad r_L(1 \mid \omega, \pi) = p(\pi) \quad \text{for} \;\; \omega \in \{0,1\},$$

so the event $\{a = 1\}$ carries no information about ω under L.[13]

Given a cutoff $s = s^*(\pi)$, the interim posterior that the *state* is good after observing $a = 1$ is

$$\lambda(1, \pi; s) \;=\; \mathbb{P}(\omega = 1 \mid a = 1, \pi) \;=\; \frac{\pi\,[\lambda A(s)] + (1-\pi)\,[\lambda p(\pi)]}{\pi\,[\lambda A(s) + (1-\lambda)B(s)] + (1-\pi)\,p(\pi)}. \tag{16}$$

Using (14), the success probability conditional on $a = 1$ is

$$P_S(1, \pi; s) \;=\; \lambda(1, \pi; s)^2. \tag{17}$$

For Bayesian updating about *type*, define, as in Appendix B,

$$\begin{aligned} p_H(\pi; s) &\equiv \mathbb{P}(\omega = 1 \mid \theta{=}H, a = 1, \pi) = \frac{\lambda A(s)}{\lambda A(s) + (1-\lambda)B(s)}, \\ p_L(\pi) &\equiv \mathbb{P}(\omega = 1 \mid \theta{=}L, a = 1, \pi) = \lambda. \end{aligned}$$

The log-likelihood jumps (Appendix C) are

$$\begin{aligned} J^+(\pi; s) &= \log \frac{p_H(\pi; s)}{p_L(\pi)} = \log \frac{A(s)}{\lambda A(s) + (1-\lambda)B(s)} \;\; > 0, \\ J^-(\pi; s) &= \log \frac{1 - e^* p_H(\pi; s)}{1 - e^* p_L(\pi)} \;\; < 0, \end{aligned} \tag{18}$$

with $e^* = \lambda(1, \pi; s)$.

D.4.2 Monotone effects of Gaussian precision

Let σ_H^2 be the High type's variance. Differentiating (15) at a fixed threshold s gives

$$\frac{\partial A}{\partial \sigma_H} = \varphi\Big(\frac{s-\mu_1}{\sigma_H}\Big)\frac{s-\mu_1}{\sigma_H^2} \;<\; 0, \qquad \frac{\partial B}{\partial \sigma_H} = \varphi\Big(\frac{s-\mu_0}{\sigma_H}\Big)\frac{s-\mu_0}{\sigma_H^2} \;>\; 0, \tag{19}$$

whenever $\mu_0 < s < \mu_1$, where φ is the standard normal pdf. Thus, a precision loss (higher σ_H) *reduces* the true-positive rate A and *increases* the false-positive rate B for the event $\{s \geq s^*(\pi)\}$.

Lemma 38. *Holding (s, π) fixed, $\partial_{\sigma_H} \lambda(1, \pi; s) < 0$ and $\partial_{\sigma_H} P_S(1, \pi; s) < 0$.*

Proof. Write (16) as $\lambda(1, \pi; s) = N/D$ with

$$N = \lambda\big[\pi A + (1-\pi)p(\pi)\big], \qquad D = \pi\big[\lambda A + (1-\lambda)B\big] + (1-\pi)p(\pi).$$

Using (19), $\partial_{\sigma_H} N = \lambda\pi\,\partial_{\sigma_H} A < 0$ while $\partial_{\sigma_H} D = \pi\big[\lambda\,\partial_{\sigma_H} A + (1-\lambda)\,\partial_{\sigma_H} B\big]$ has the sign of $(1-\lambda)\partial_{\sigma_H} B + \lambda \partial_{\sigma_H} A > 0$ because $|\partial_{\sigma_H} A|$ and $\partial_{\sigma_H} B$ are both positive and $(1-\lambda) \geq \lambda$ cannot hold for all λ, but, crucially, for any $\lambda \in (0,1)$ and $\mu_0 < s < \mu_1$, the Gaussian tails satisfy

$$(1-\lambda)\,\varphi\Big(\frac{s-\mu_0}{\sigma_H}\Big)\frac{s-\mu_0}{\sigma_H^2} \;+\; \lambda\,\varphi\Big(\frac{s-\mu_1}{\sigma_H}\Big)\frac{\mu_1-s}{\sigma_H^2} \;>\; 0,$$

[13] Any alternative measurable construction that maintains $r_L(1 \mid 1, \pi) \geq r_L(1 \mid 0, \pi)$ works as well; the comparative statics below are unchanged.

since each term is strictly positive. Hence $\partial_{\sigma_H} D > 0$. By the quotient rule,

$$\partial_{\sigma_H}\lambda(1,\pi;s) = \frac{D\,\partial_{\sigma_H}N - N\,\partial_{\sigma_H}D}{D^2} < 0.$$

Equation (17) then yields $\partial_{\sigma_H}P_S = 2\lambda(1,\pi;s)\,\partial_{\sigma_H}\lambda(1,\pi;s) < 0$. □

Lemma 39. *Holding (s,π) fixed, $\partial_{\sigma_H}J^+(\pi;s) < 0$ and $\partial_{\sigma_H}J^-(\pi;s) > 0$.*

Proof. Differentiate (18). For J^+,

$$\partial_{\sigma_H}J^+ = \frac{\partial_{\sigma_H}A}{A} - \frac{\lambda\,\partial_{\sigma_H}A + (1-\lambda)\,\partial_{\sigma_H}B}{\lambda A + (1-\lambda)B} < 0$$

by (19) and the same positivity argument as in Lemma 38. For J^-, holding p_H fixed for the moment,

$$\frac{\partial J^-}{\partial e^*} = -\Big(\frac{p_H}{1-e^*p_H} - \frac{p_L}{1-e^*p_L}\Big) < 0 \quad \text{since } p_H > p_L = \lambda,$$

so the fall in e^* from Lemma 38 raises J^-. In addition, $\partial_{\sigma_H}p_H < 0$ by the same calculus as for J^+, which further increases J^- (makes it less negative). Thus $\partial_{\sigma_H}J^- > 0$. □

Proposition 40. *For any $\pi \in (0,1)$ with interior experimentation, $\dfrac{\partial s^*(\pi)}{\partial\sigma_H^2} > 0$.*

Proof. Let $\Delta_H(s;\pi)$ denote the risky-minus-safe continuation difference (Appendix A). The cutoff satisfies $\Delta_H(s^*(\pi);\pi) = 0$ and $\partial_s\Delta_H > 0$ (single crossing). By the implicit function theorem,

$$\frac{\partial s^*(\pi)}{\partial\sigma_H^2} \;=\; -\frac{\partial_{\sigma_H^2}\Delta_H(s;\pi)}{\partial_s\Delta_H(s;\pi)}\bigg|_{s=s^*(\pi)}.$$

Using (17) and the update formulae for $\pi^\pm$ via $J^\pm$, one obtains

$$\partial_{\sigma_H^2}\Delta_H = \delta\Big[(\partial_{\sigma_H^2}P_S)\,(V(\pi^+) - V(\pi^-)) + P_S\,V'(\pi^+)\,\partial_{\sigma_H^2}\pi^+ + (1-P_S)\,V'(\pi^-)\,\partial_{\sigma_H^2}\pi^-\Big].$$

By Lemma 38, $\partial_{\sigma_H^2}P_S < 0$. By Lemma 39, $\partial_{\sigma_H^2}\pi^+ < 0$ and $\partial_{\sigma_H^2}\pi^- > 0$ (success becomes less, and failure more, like the prior), and with $V'(\cdot) \geq 0$ and $V(\pi^+) > V(\pi^-)$, the bracket is strictly negative. Since $\partial_s\Delta_H > 0$, the sign is $\partial s^*/\partial\sigma_H^2 > 0$. □

D.4.3 Effect of the good-state prior λ

Lemma 41. *Holding (s,π) fixed, $\partial_\lambda\lambda(1,\pi;s) > 0$ and $\partial_\lambda P_S(1,\pi;s) > 0$.*

Proof. From (16), $N = \lambda[\pi A + (1-\pi)p]$ and $D = \pi[\lambda A + (1-\lambda)B] + (1-\pi)p$. Then

$$\begin{aligned}\partial_\lambda\lambda(1,\pi;s) &= \frac{D[\pi A + (1-\pi)p] - N\cdot\pi(A-B)}{D^2}\\ &= \frac{(1-\pi)p\cdot\pi(A-B) + \pi A\cdot[(1-\pi)p + \pi(1-\lambda)B]}{D^2} > 0,\end{aligned}$$

since $A > B$ and all terms are nonnegative with at least one strictly positive. Hence $\partial_\lambda P_S = 2\lambda(1,\pi;s)\,\partial_\lambda\lambda(1,\pi;s) > 0$. □

Lemma 42. *Holding (s, π) fixed, $\partial_\lambda J^+(\pi; s) < 0$ and $\partial_\lambda J^-(\pi; s) > 0$.*

Proof. From (18), $J^+ = \log \frac{A}{\lambda A + (1-\lambda)B}$, so $\partial_\lambda J^+ = -\frac{A-B}{\lambda A+(1-\lambda)B} < 0$. For J^-, $p_L = \lambda$ increases with λ while p_H also (weakly) increases with λ but remains $> p_L$; holding e^* fixed,

$$\partial_\lambda J^- = \frac{\partial}{\partial \lambda} \log \frac{1 - e^* p_H}{1 - e^* \lambda} = -\frac{e^* \, \partial_\lambda p_H}{1 - e^* p_H} + \frac{e^*}{1 - e^* \lambda} \; > \; 0,$$

because $\partial_\lambda p_H \in (0, 1)$ and the second term dominates when $p_H > p_L = \lambda$. The rise in e^* from Lemma 41 further increases J^- (less negative). □

Proposition 43. *For any $\pi \in (0, 1)$ with interior experimentation, $\dfrac{\partial s^*(\pi)}{\partial \lambda} < 0$.*

Proof. As in the proof of Proposition 40, apply the implicit function theorem. Using Lemma 41, $\partial_\lambda P_S > 0$ raises the first term in $\partial_\lambda \Delta_H$; Lemma 42 moves $\pi^\pm$ *toward* π (success less informative, failure less damaging), which increases the expected continuation from taking risk because V is increasing; thus $\partial_\lambda \Delta_H > 0$ and $\partial s^* / \partial \lambda < 0$. □

D.4.4 Effect of patience δ

In the benchmark, δ scales the expected continuation difference from taking risk:

$$\frac{\partial \Delta_H(s; \pi)}{\partial \delta} \;=\; P_S(1, \pi; s) \, V(\pi^+(\pi; s)) + \big(1 - P_S(1, \pi; s)\big) \, V(\pi^-(\pi; s)) - V(\pi). \tag{20}$$

The sign of (20) is governed by how the lottery over $\{\pi^-, \pi^+\}$ compares to the sure thing π. With V convex and π the prior mean of the posterior (law of total expectation for posteriors), Jensen's inequality yields

$$P_S \, V(\pi^+) + (1 - P_S) \, V(\pi^-) \;\geq\; V\big(P_S \, \pi^+ + (1 - P_S) \, \pi^-\big) \;=\; V(\pi), \tag{21}$$

with strict inequality whenever $\pi^+ \neq \pi^-$. However, as π rises the *downside* ($\pi^- < \pi$) becomes more likely to be triggered by higher effort (Appendix B), and the informational content of success does not rise with effort—so the convexity comparison increasingly favors the status quo at high π once we internally account for how $\pi^\pm$ depend on π (Appendix B.2).

The benchmark therefore delivers the same qualitative comparative static as Theorem 17(iii), and we can make it explicit by bounding $\pi^\pm$.

Lemma 44. *Fix π. There exist constants $0 < \underline{\eta}(\pi) < \overline{\eta}(\pi) < \infty$ (computable from A, B) such that*

$$\begin{aligned} \pi^+(\pi; s) &\in \Big[\pi + \underline{\eta}(\pi)\,(1 - \pi)\,, \quad \pi + \overline{\eta}(\pi)\,(1 - \pi)\Big], \\ \pi^-(\pi; s) &\in \Big[\pi - \overline{\eta}(\pi)\,\pi\,, \quad \pi - \underline{\eta}(\pi)\,\pi\Big], \end{aligned}$$

with $\underline{\eta}(\pi)$ decreasing and $\overline{\eta}(\pi)$ increasing in σ_H, and $\underline{\eta}(\pi)$ decreasing in λ.

Proof. Immediate from the odds-ratio updates with $J^\pm$ in (18) and the monotonicity in Lemmas 39–42; the bounds follow by mapping log-odds intervals to probability space. □

Proposition 45. *Suppose V is increasing and convex. Then for any π at which experimentation is interior,*

$$\frac{\partial s^*(\pi)}{\partial \delta} \;\geq\; 0,$$

with strict inequality whenever π is sufficiently high.[14]

Proof. By (20), the sign of $\partial_\delta \Delta_H$ equals the sign of the Jensen gap in (21). Using Lemma 44, as π rises the spread $(\pi^+ - \pi^-)$ around π contracts asymmetrically toward the downside (failure becomes more anti-H with higher effort), so the convexity gain from taking risk shrinks and eventually becomes negative. Therefore $\partial_\delta \Delta_H \leq 0$ for all π (weakly) and < 0 for π high enough, implying $\partial s^*/\partial\delta \geq 0$ and > 0 at high π by the implicit function theorem. □

In the quadratic–Gaussian benchmark, for any interior π,

$$\frac{\partial s^*(\pi)}{\partial \sigma_H^2} > 0, \qquad \frac{\partial s^*(\pi)}{\partial \lambda} < 0, \qquad \frac{\partial s^*(\pi)}{\partial \delta} \geq 0 \;\; \text{(strict for large } \pi).$$

These conclusions align with the general results in Theorem 17 and make the channels explicit: precision shapes both the frequency and informativeness of successes/failures through (A, B); the prior λ raises success odds and softens failure; and patience scales a convex-reputation tradeoff that is tilted toward conservatism at high reputation.

E Exit Options and Equilibrium Selection

This appendix proves Proposition 19. The expert may irrevocably *exit* at the start of a period and secure a reservation value U_0 that is independent of reputation. Let $\widehat{V}(\pi)$ denote the expert's value function when exit is available.

E.1 Value recursion with an exit option

At belief π, after the High type observes s and chooses $a \in \{0, 1\}$, continuation value equals

$$\widehat{V}(\pi) = \max\Big\{U_0, \underbrace{\max_{a \in A_H(s,\pi)} \Big\{ \underbrace{\phi\, a}_{\text{stage payoff}} + \delta\, \mathbb{E}\big[\widehat{V}(\pi') \mid a, \pi\big]\Big\}}_{\equiv W(\pi)}\Big\}.$$

By the same arguments as in Section 3, $W(\pi)$ is well defined and bounded, and $\widehat{V}$ is the smallest bounded solution to the Bellman equation above.

Lemma 46. *$\widehat{V}(\pi) \geq V(\pi)$ for all π. If $U_0 \leq \min_\pi V(\pi)$, then $\widehat{V} \equiv V$.*

Proof. Trivial from $\max\{U_0, W(\pi)\} \geq W(\pi)$ and the definition of V as the fixed point without exit. □

[14] At very low π, the reputational upside can dominate under strong convexity of V, in which case $s^*(\pi)$ is locally flat in δ; the inequality is weak globally and strict for π above a model-dependent threshold.

E.2 Exit thresholds and reputation reinforcement

Define the (upper hemicontinuous) experimentation correspondence $\Gamma(\pi) \subseteq \{0,1\}$ for the High type as in Section 4. Let $\Pi_{\text{exp}} \equiv \{\pi : 1 \in \Gamma(\pi)\}$ and $\Pi_{\text{no}} \equiv \{\pi : 0 \in \Gamma(\pi)\}$.

Proposition 47. *There exists $\pi^{\text{exit}} \in [0,1)$ such that: (i) if $\pi \le \pi^{\text{exit}}$, then $\widehat{V}(\pi) = U_0$ and the High type exits with probability one; (ii) if $\pi > \pi^{\text{exit}}$ and $U_0 < \sup_\pi V(\pi)$, then exit does not occur on the equilibrium path.*

Proof. The map $\pi \mapsto W(\pi)$ is increasing and upper semicontinuous (Appendix B). The set $\{\pi : W(\pi) \le U_0\}$ is therefore a (possibly empty) closed interval $[0, \pi^{\text{exit}}]$. On this set, $\widehat{V}(\pi) = U_0$; off it, $\widehat{V}(\pi) = W(\pi) > U_0$ and exit is strictly dominated. Monotonicity of Γ implies that for $\pi > \pi^{\text{exit}}$ experimentation occurs with (weakly) higher probability, yielding strictly more convex reputation dynamics and reinforcing the no-exit region. □

Corollary 48. *If multiple MPBE exist at some π, the equilibrium with (weakly) higher experimentation probability yields (weakly) lower π^{exit} and strictly dominates in welfare for any $U_0 < \sup_\pi V(\pi)$.*

Proof. Follows from the sub/supermartingale arguments in Appendix C: higher experimentation raises expected log-likelihood drift under H and reduces the measure of belief states at which $W(\pi) \le U_0$. □

F Monitoring and Disclosure

We enrich the baseline with a public *monitoring signal* $m_t \in \{0,1\}$ realized after the recommendation $a_t \in \{0,1\}$ and observed by all players before the outcome. Conditional on the action and the expert's type,

$$\mathbb{P}(m_t = 1 \mid a_t = 1, \theta) = q_\theta, \qquad \mathbb{P}(m_t = 1 \mid a_t = 0, \theta) = \bar{q}_\theta, \tag{22}$$

with $q_H > q_L$ and $\bar{q}_H \ge \bar{q}_L$. The realized pair (y_t, m_t) updates the public reputation by Bayes' rule; let $\pi_{t+1} = \Phi(\pi_t; y_t, m_t)$ denote the posterior. We maintain Assumptions A1–A4.

F.1 Preliminaries and log-likelihood ratios

Fix a period with public reputation $\pi \in (0,1)$. When $a = 1$, the outcome $y \in \{0,1\}$ is generated as in the baseline: conditional on θ and on the induced effort $e^*(1,\pi)$, the success probability equals $e^*(1,\pi)\, p_\theta(\pi)$, where $p_H(\pi) > p_L(\pi)$ by Assumption A1. The monitoring signal m is conditionally independent of y given (a, θ) and has distribution (22).

It is convenient to work with the *log-likelihood ratio* (LLR) increments for the type posterior (Appendix C). Define

$$J^+(\pi) \equiv \log \frac{\mathbb{P}(y = 1 \mid \theta = H, a = 1, \pi)}{\mathbb{P}(y = 1 \mid \theta = L, a = 1, \pi)} = \log \frac{p_H(\pi)}{p_L(\pi)} \;>\; 0,$$
$$J^-(\pi) \equiv \log \frac{\mathbb{P}(y = 0 \mid H, a = 1, \pi)}{\mathbb{P}(y = 0 \mid L, a = 1, \pi)} = \log \frac{1 - e^* p_H(\pi)}{1 - e^* p_L(\pi)} \;<\; 0,$$

with $e^* = e^*(1,\pi)$, as in (11)–(12). The monitoring LLRs are

$$J_m^{(1)} \equiv \log \frac{q_H}{q_L} \;\text{ if } m = 1, \quad J_m^{(1)} = \log \frac{1 - q_H}{1 - q_L} \;\text{ if } m = 0 \qquad (a = 1), \tag{23}$$

$$J_m^{(0)} \equiv \log \frac{\bar{q}_H}{\bar{q}_L} \;\text{ if } m = 1, \quad J_m^{(0)} = \log \frac{1-\bar{q}_H}{1-\bar{q}_L} \;\text{ if } m = 0 \qquad (a = 0). \tag{24}$$

Since m and y are conditionally independent given (a, θ), the composite LLR for (y, m) is additive. If $a = 1$, the one-period log-odds innovation equals J^+ or J^- (from y) plus $J_m^{(1)}$ (from m); if $a = 0$, it equals $J_m^{(0)}$.

Lemma 49. *For each fixed action $a \in \{0, 1\}$, the composite observation (y, m) (with y vacuous when $a = 0$) is ordered by the likelihood ratio $\ell_a(y, m) \equiv \frac{\mathbb{P}(y,m|\theta=H,a,\pi)}{\mathbb{P}(y,m|\theta=L,a,\pi)}$. Then $\pi \mapsto \Phi(\pi; y, m)$ is strictly increasing in $\ell_a(y, m)$, and the High type's risky-minus-safe continuation difference $\Delta_H(s; \pi)$ preserves single crossing in the private signal s.*

Proof. By Bayes' rule, the posterior odds equal prior odds times the LLR, hence $\Phi(\pi; y, m)$ is strictly increasing in $\ell_a(y, m)$. Under Assumption A1, the High type's reduced-form success probability $P_S(s; \pi)$ is strictly increasing in s, while the incremental information from m does not depend on s. Therefore $\Delta_H(s; \pi)$ adds an s-independent term to the baseline expression (10), preserving strict single crossing in s and the cutoff representation of Theorem 5. □

F.2 Blackwell order and the option value of experimentation

Let $\mathcal{L}_a(\cdot; \pi)$ denote the law of the next-period posterior π' conditional on current reputation π and action a. (Formally, π' is a measurable function of (y, m) and current π.) The expert's continuation under a equals $\mathbb{V}_a(\pi) \equiv \mathbb{E}[V(\pi') \mid a, \pi]$, V is increasing and convex by Assumption A3 and Proposition 2. The *option value of experimentation* is $\mathbb{V}_1(\pi) - \mathbb{V}_0(\pi)$, the bracketed term in (10) up to the factor δ.

We parameterize the informativeness of the monitor by two scalars $(\kappa, \bar{\kappa}) \in [0, \infty)^2$, writing

$$q_H(\kappa) - q_L(\kappa) \;\text{ strictly increases in } \kappa, \qquad \bar{q}_H(\bar{\kappa}) - \bar{q}_L(\bar{\kappa}) \;\text{ strictly increases in } \bar{\kappa},$$

with $q_\theta(0) = \bar{q}_\theta(0) = \frac{1}{2}$. Larger $(\kappa, \bar{\kappa})$ thus represent Blackwell improvements of m under $a = 1$ and under $a = 0$, respectively.

Lemma 50. *Fix π. If $\bar{\kappa}' > \bar{\kappa}$, then $\mathcal{L}_0(\cdot; \pi, \bar{\kappa}')$ Blackwell-dominates $\mathcal{L}_0(\cdot; \pi, \bar{\kappa})$, hence $\mathbb{V}_0(\pi; \bar{\kappa}') > \mathbb{V}_0(\pi; \bar{\kappa})$ whenever V is strictly convex. Similarly, if $\kappa' > \kappa$, then $\mathcal{L}_1(\cdot; \pi, \kappa')$ Blackwell-dominates $\mathcal{L}_1(\cdot; \pi, \kappa)$ and $\mathbb{V}_1(\pi; \kappa') > \mathbb{V}_1(\pi; \kappa)$.*

Proof. For $a = 0$, the posterior is updated only with m, whose binary experiment becomes strictly more informative as $\bar{\kappa}$ increases (the LLR support expands by (24)). Blackwell's theorem implies a mean-preserving spread in the posterior, which strictly raises $\mathbb{E}[V(\pi')]$ under strict convexity. The case $a = 1$ is identical with (23). □

Proposition 51. *Fix π and let $s^*(\pi; \kappa, \bar{\kappa})$ be the High-type cutoff under monitor parameters $(\kappa, \bar{\kappa})$. Then:*

(i) *For all π at which experimentation is interior,*

$$\frac{\partial s^*(\pi; \kappa, \bar{\kappa})}{\partial \bar{\kappa}} \;\geq\; 0,$$

with strict inequality whenever V is strictly convex and the safe monitor is not degenerate.

(ii) *For all interior π,*

$$\frac{\partial s^*(\pi;\kappa,\bar{\kappa})}{\partial \kappa} \;\le\; 0,$$

with strict inequality under strict convexity and a nondegenerate risky monitor.

Proof. By Lemma 49, $s^*(\pi;\kappa,\bar{\kappa})$ solves the first-order condition $\Delta_H(s;\pi,\kappa,\bar{\kappa})=0$ with $\partial_s\Delta_H>0$. By the implicit function theorem,

$$\frac{\partial s^*(\pi;\kappa,\bar{\kappa})}{\partial \bar{\kappa}} = -\left.\frac{\partial_{\bar{\kappa}}\Delta_H(s;\pi,\kappa,\bar{\kappa})}{\partial_s\Delta_H(s;\pi,\kappa,\bar{\kappa})}\right|_{s=s^*}.$$

Using (10) and that m arrives under both actions,

$$\partial_{\bar{\kappa}}\Delta_H \;=\; \delta\left(\partial_{\bar{\kappa}}\mathbb{V}_1(\pi;\kappa,\bar{\kappa}) \;-\; \partial_{\bar{\kappa}}\mathbb{V}_0(\pi;\kappa,\bar{\kappa})\right).$$

The risky posterior law depends on m through $q_\theta(\kappa)$ but not through $\bar{q}_\theta(\bar{\kappa})$, so $\partial_{\bar{\kappa}}\mathbb{V}_1=0$, whereas $\partial_{\bar{\kappa}}\mathbb{V}_0>0$ by Lemma 50. Hence $\partial_{\bar{\kappa}}\Delta_H<0$, which yields $\partial_{\bar{\kappa}}s^*(\pi)\ge 0$, with strict inequality under strict convexity. The proof of (ii) is analogous: $\partial_\kappa\mathbb{V}_0=0$ and $\partial_\kappa\mathbb{V}_1>0$, hence $\partial_\kappa\Delta_H>0$ and $\partial_\kappa s^*(\pi)\le 0$, with strictness as stated. □

Corollary 52. *For any $\kappa\ge 0$, there exists $\bar{\kappa}_0>0$ such that for all $\bar{\kappa}\ge\bar{\kappa}_0$ the High-type cutoff satisfies $s^*(\pi;\kappa,\bar{\kappa})\ge s^*(\pi;\kappa,0)$ for all π, with strict inequality on a set of π of positive measure. In particular, along any equilibrium path that visits this set with positive probability, the ex-ante experimentation frequency strictly falls when the safe monitor is made sufficiently informative.*

Proof. Immediate from Proposition 51(i) and the strictness conclusion when V is strictly convex and the monitor is nondegenerate. □

F.3 Failure asymmetry with monitoring

Monitoring alters not only the option value but also the *composition* of diagnostic content. Under $a=1$, the composite LLR after success is $J^+(\pi)+J^{(1)}_m(m)$ while after failure it is $J^-(\pi)+J^{(1)}_m(m)$. Since J^+ is independent of effort but J^- becomes more negative as effort rises (Lemma 31), the *gap* in informativeness between success and failure narrows when a part of the signal (m) is available regardless of outcome or action; when the monitor under $a=0$ is precise (large $\bar{\kappa}$), the incremental informational advantage of $a=1$ comes primarily from y, which is exactly the channel penalized by higher effort at high reputation. The next proposition formalizes the resulting tilt toward conservatism.

Proposition 53. *Suppose V is increasing and convex and let $\bar{\kappa}>0$. Then the reputational-conservatism result (Theorem 7) continues to hold with monitoring, and the slope $\frac{d}{d\pi}s^*(\pi;\kappa,\bar{\kappa})$ is (weakly) larger for $\bar{\kappa}$ sufficiently large. In particular, for π in a neighborhood of 1, the High-type cutoff under monitoring exceeds the no-monitoring cutoff.*

Proof. By Lemma 49, single crossing and Topkis monotonicity in (a,π) are preserved, so $s^*(\pi)$ is (weakly) increasing in π as in Theorem 7. To compare slopes, differentiate $\Delta_H(s;\pi,\kappa,\bar{\kappa})$ with respect to π as in Appendix B.2. The effort channel (which makes failures more damaging at higher π) is unchanged, while the baseline continuation $V(\pi)$ is unaffected. The only new term is the safe-action information $\mathbb{V}_0(\pi;\bar{\kappa})$, which is independent of s but increases with $\bar{\kappa}$. Since this term reduces the convexity gain from switching to $a=1$, the decreasing-differences argument strengthens, yielding a (weakly) larger slope of s^* in π when $\bar{\kappa}$ is large. The local dominance near $\pi=1$ follows because at high reputation the downside (failure) dominates the upside (success), and subtracting a convexity gain from the safe action magnifies the precautionary motive. □

Proposition 51 shows a sharp *crowding-out*: information that arrives under the safe action raises the value of not experimenting and thus increases the risky-signal cutoff. Conversely, information that arrives only under the risky action encourages experimentation. Proposition 53 documents that, beyond levels, monitoring steepens the reputation–conservatism relation at the top, where the downside of a visible failure is largest due to higher implementation effort. These forces rationalize empirical situations in which transparency of process or effort reduces appetite for risky recommendations even when monitors are accurate.

G Committees and Advice Aggregation

This appendix strengthens the committee extension by deriving pivot probabilities explicitly, showing that the High type's problem preserves single crossing, and establishing how the cutoff depends on the aggregation threshold k.

G.1 Environment

There are $n \geq 2$ experts indexed by $i = 1, \ldots, n$. Each expert i has a fixed type $\theta^i \in \{H, L\}$, i.i.d. across i with prior $\pi \in (0,1)$ that a given expert is H. In period t, expert i privately observes a signal $s^i \in S$ about the period state $\omega_t \in \{0,1\}$ (i.i.d. over t with $\mathbb{P}(\omega_t = 1) = \lambda \in (0,1)$). Signal families satisfy Assumption A1 (MLRP; H Blackwell-dominates L). Experts simultaneously issue recommendations $a^i \in \{0,1\}$ (risky vs. safe). The implementer adopts the risky action iff at least k of the n recommendations are risky:

$$a \;=\; \mathbf{1}\left\{\sum_{i=1}^{n} a^i \geq k\right\}, \qquad k \in \{1, \ldots, n\}.$$

If $a = 1$, the implementer chooses effort $e \in [0,1]$ as in the baseline; success occurs with probability $e\,\mathbf{1}\{\omega = 1\}$; if $a = 0$, $y = 0$ with certainty. All recommendations and the outcome are publicly observed, and individual reputations about ability are updated by Bayes' rule. The period payoffs and the expert's continuation value are as in the baseline (Section 3); we focus on a symmetric MPBE in which all H-types use a cutoff $s^*(\pi)$ and all L-types use some (possibly different) measurable strategy.[15]

Let $r_\theta(1 \mid \omega, \pi)$ denote the probability that a type-θ expert issues $a^j = 1$ conditional on ω when the public reputation is π. From the viewpoint of a given expert i, the risky-recommendation probability of any other expert $j \neq i$ conditional on ω is

$$\rho_\omega(\pi) \;\equiv\; \pi\, r_H(1 \mid \omega, \pi) \;+\; (1-\pi)\, r_L(1 \mid \omega, \pi), \qquad \omega \in \{0,1\}. \tag{25}$$

Conditional on (ω, π), the number of other risky recommendations

$$S_{-i} \;\equiv\; \sum_{j \neq i} a^j$$

is binomial: $S_{-i} \sim \mathrm{Bin}(n-1, \rho_\omega(\pi))$ and is independent of i's private signal s^i.

[15] Nothing below requires the L-type to use a cutoff; we only use that for each ω and π the induced risky-recommendation probability of a generic other expert is well defined.

G.2 Pivot probability

Lemma 54. *Fix π and $k \in \{1, \dots, n\}$. The probability that expert i is pivotal—i.e., that the final decision switches from $a = 0$ to $a = 1$ when she switches her recommendation from 0 to 1—equals*

$$\begin{aligned} \zeta_k(\pi) &\equiv \sum_{\omega \in \{0,1\}} \mathbb{P}(\omega)\, \mathbb{P}(S_{-i} = k-1 \mid \omega, \pi) \\ &= \lambda\, b_{n-1,k-1}(\rho_1(\pi)) \;+\; (1-\lambda)\, b_{n-1,k-1}(\rho_0(\pi))\,, \end{aligned} \tag{26}$$

where $b_{N,j}(p) \equiv \binom{N}{j} p^j (1-p)^{N-j}$ is the binomial pmf.

Proof. For fixed ω, $a = 1$ under $a^i = 1$ iff $S_{-i} \geq k-1$, and under $a^i = 0$ iff $S_{-i} \geq k$. The only realizations that change the committee's decision are those with $S_{-i} = k-1$. The independence across experts conditional on ω and π yields $\mathbb{P}(S_{-i} = k-1 \mid \omega, \pi) = b_{n-1,k-1}(\rho_\omega(\pi))$. Averaging over ω gives (26). □

G.3 Decomposing the High type's marginal value

Let $\Delta_H^{\text{comm}}(s; \pi, k)$ be the High type's risky-minus-safe continuation gain at signal s, holding others to the symmetric profile. There are two channels:

(i) an *outcome channel* that operates only when the expert is pivotal (her recommendation changes the committee action and hence whether an outcome and reputational jump are realized);

(ii) a *signaling channel* coming from the informational content of her recommendation a^i about θ^i even when the committee's action is the same with $a^i = 0$ or 1.

Formally, write

$$\Delta_H^{\text{comm}}(s; \pi, k) \;=\; \underbrace{\zeta_k(\pi)\, \Delta_H(s; \pi)}_{\text{outcome channel}} + \underbrace{\Xi_H(s; \pi)}_{\text{signaling channel}} . \tag{27}$$

Here $\Delta_H(s; \pi)$ is the baseline (single-expert) risky-minus-safe continuation difference from (10)—it values the move from $a = 0$ to $a = 1$ *holding fixed* that the recommendation determines the action—and $\Xi_H(s; \pi)$ collects the pure signaling effect of a^i on the posterior about θ^i when the committee action is unchanged. Importantly, Ξ_H does not depend on k.

Lemma 55. *For any fixed (π, k), the map $s \mapsto \Delta_H^{\text{comm}}(s; \pi, k)$ has the strict single-crossing property. Consequently, the High type's best reply is a cutoff in s, and in a symmetric MPBE the High type uses a cutoff $s^*(\pi; k)$.*

Proof. By Lemma 3 in Appendix A, $s \mapsto \Delta_H(s; \pi)$ has strict single crossing. The pivot probability $\zeta_k(\pi)$ is constant in s. The signaling term $s \mapsto \Xi_H(s; \pi)$ inherits single crossing from Assumption A1 because, under MLRP, the LLR of the event $\{a^i = 1\}$ relative to $\{a^i = 0\}$ is increasing in s; hence the reputational benefit of choosing $a^i = 1$ rather than $a^i = 0$ is (weakly) increasing in s.[16] A sum of (strictly) single-crossing functions is (strictly) single crossing; thus (27) preserves the property and the cutoff characterization follows as in Theorem 5. □

[16] Formally, let $L_s \equiv \log \frac{f_H(s|\omega)}{f_L(s|\omega)}$ be the signal LLR (increasing in s by MLRP). The log posterior odds after observing the recommendation $a^i \in \{0,1\}$ but fixing the committee action is $L_0 + \log \frac{\mathbb{P}(a^i|H)}{\mathbb{P}(a^i|L)}$, and the difference between $a^i = 1$ and $a^i = 0$ is increasing in s.

G.4 Comparative statics in the threshold k

We now show how k shifts the cutoff. All k-dependence in (27) is through the pivot probability $\zeta_k(\pi)$.

Lemma 56. *Fix π and let $p_\omega \equiv \rho_\omega(\pi) \in (0,1)$ for $\omega \in \{0,1\}$. Then for any $k \in \{1,\dots,n-1\}$,*

$$\frac{\zeta_{k+1}(\pi)}{\zeta_k(\pi)} \;=\; \lambda\, R_{n-1,k}(p_1)\,\frac{b_{n-1,k-1}(p_1)}{\zeta_k(\pi)} \;+\; (1-\lambda)\, R_{n-1,k}(p_0)\,\frac{b_{n-1,k-1}(p_0)}{\zeta_k(\pi)},$$

where $R_{N,k}(p) \equiv \dfrac{b_{N,k}(p)}{b_{N,k-1}(p)} = \dfrac{N-k+1}{k}\cdot\dfrac{p}{1-p}$. In particular, if $p_0 \le \frac{k}{n}$ and $p_1 \le \frac{k}{n}$, then $\zeta_{k+1}(\pi) \le \zeta_k(\pi)$, with strict inequality unless $p_0 = p_1 = \frac{k}{n}$.

Proof. The identity is the binomial pmf ratio. If $p \le \frac{k}{n}$, then for $N = n-1$ one has $R_{N,k}(p) \le 1$ (because $\frac{N-k+1}{k} \le \frac{n-k}{k}$ and $\frac{p}{1-p} \le \frac{k/n}{1-k/n} = \frac{k}{n-k}$). The convex combination with nonnegative weights therefore does not exceed one; strictness is immediate unless both terms equal one. □

Proposition 57. *Fix π and suppose $p_\omega = \rho_\omega(\pi) \le \frac{k}{n}$ for $\omega \in \{0,1\}$ (in particular, for majority rules $k \ge \lceil n/2 \rceil$ this holds whenever $p_0, p_1 \le 1/2$). Then the High type's cutoff $s^*(\pi;k)$ is (weakly) increasing in k, with strict inequality whenever $\Delta_H(s^*(\pi;k);\pi) > 0$ (i.e., the outcome channel has strictly positive value at the threshold).*

Proof. By Lemma 55, $s^*(\pi;k)$ solves $\Delta_H^{\text{comm}}(s;\pi,k) = 0$ with $\partial_s \Delta_H^{\text{comm}} > 0$. Differencing in k,

$$\Delta_H^{\text{comm}}(s;\pi,k+1) - \Delta_H^{\text{comm}}(s;\pi,k) = \big(\zeta_{k+1}(\pi) - \zeta_k(\pi)\big)\,\Delta_H(s;\pi),$$

since the signaling term $\Xi_H(s;\pi)$ does not depend on k. By Lemma 56, $\zeta_{k+1}(\pi) \le \zeta_k(\pi)$ under the stated condition. At $s = s^*(\pi;k)$, we have $\Delta_H^{\text{comm}}(\cdot;\pi,k) = 0$, so

$$\Delta_H^{\text{comm}}(s^*(\pi;k);\pi,k+1) = (\zeta_{k+1} - \zeta_k)\,\Delta_H(s^*(\pi;k);\pi) \;\le\; 0,$$

with strict inequality if $\Delta_H(s^*(\pi;k);\pi) > 0$ and $\zeta_{k+1} < \zeta_k$. Because $\partial_s \Delta_H^{\text{comm}} > 0$, the implicit-function theorem implies $s^*(\pi;k+1) \ge s^*(\pi;k)$, strictly when the inequality above is strict. □

Proposition 57 shows that when other members are not "too risk-prone" (precisely, their risky probability p_ω does not exceed k/n), raising the threshold k shrinks the pivot probability ζ_k and thereby *raises* the cutoff. Intuitively, stricter aggregation makes a single member's recommendation less likely to change the committee's action, so the outcome-based option value of recommending risk is attenuated, while the pure signaling value of the recommendation is unaffected; to restore indifference at the margin, the High type requires a stronger private signal. The result applies in particular to majority rules when $p_0, p_1 \le 1/2$, a case that is empirically plausible at higher reputations where experimentation is less frequent (Theorem 7).

G.5 LLR calculus for own-recommendation signaling

For completeness, we record the likelihood-ratio contribution of the expert's own recommendation, which underlies the signaling term $\Xi_H(s;\pi)$ in (27). Let $R_\theta(\pi) \equiv \mathbb{P}(a^i = 1 \mid \theta, \pi)$ be the (unconditional-in-ω) risky frequency induced by the strategy at π. The LLR contribution of switching from $a^i = 0$ to $a^i = 1$ while holding the committee action fixed is

$$J^{\text{rec}}(\pi) \;\equiv\; \log\frac{R_H(\pi)}{R_L(\pi)} - \log\frac{1-R_H(\pi)}{1-R_L(\pi)} \;=\; \log\frac{R_H(\pi)\,[1-R_L(\pi)]}{R_L(\pi)\,[1-R_H(\pi)]}, \tag{28}$$

which is strictly positive whenever $R_H(\pi) > R_L(\pi)$. Under MLRP and cutoff strategies with (weakly) lower threshold for H than for L, $R_H(\pi) \geq R_L(\pi)$ holds, and the mapping $s \mapsto \Xi_H(s;\pi)$ is increasing because the event $\{a^i = 1\}$ is more likely for H at higher signals.[17]

Combining Lemma 54, Lemma 55, Proposition 57, and the recommendation-LLR (28) yields a complete characterization of the High type's policy in committees: it remains a reputation-dependent cutoff, and it is (weakly) more conservative in larger-k committees whenever others' risk propensity is below k/n, with strict conservatism whenever recommending risk has positive option value at the margin.

H Continuous-Time Approximation

Let periods have length $\Delta > 0$. Suppose $y_t \in \{0,1\}$ arrives with $\mathbb{P}(y_t = 1 \mid a_t = 1, \theta) = e^*(1,\pi_t)\, p_\theta\, \Delta + o(\Delta)$ and $\mathbb{P}(y_t = 1 \mid a_t = 0, \theta) = o(\Delta)$. Define the log-likelihood increment $\Delta L_{t+\Delta}$. Standard diffusion approximations yield:

Lemma 58. *As $\Delta \to 0$, under $a_t = 1$ the process $(L_t)_{t\geq 0}$ converges weakly to a Lévy process with drift $\mu_\theta(\pi_t) = e^*(1,\pi_t)\,[p_H(\pi_t) - p_L(\pi_t)]$ and bounded jumps at arrival times of $y = 1$. If additionally $e^*(1,\pi)$ and $p_\theta(\pi)$ are C^1, then L_t admits an Itô decomposition with locally Lipschitz coefficients on any compact set where $a_t \equiv 1$.*

We now give a functional central limit theorem (in the semimartingale sense) for the log-odds process.

Theorem 59. *Fix a finite horizon $T > 0$ and a vanishing period length $\Delta \downarrow 0$. For each Δ, let $(\mathcal{F}^\Delta_{k\Delta})_{k\geq 0}$ be the public filtration. Define the (càdlàg, piecewise-constant) log-odds process*

$$L^\Delta_t \;=\; \sum_{k<t/\Delta} \Delta L^\Delta_{k+1}, \qquad \Delta L^\Delta_{k+1} \in \{0, J^+(\pi^\Delta_{k\Delta}), J^-(\pi^\Delta_{k\Delta})\},$$

where $J^\pm(\cdot)$ are the success/failure log-likelihood jumps (Appendix C), uniformly bounded and locally Lipschitz on $(0,1)$, and $\pi^\Delta_{k\Delta}$ is the public reputation just before time $k\Delta$. Write

$$m^\Delta_k \;\equiv\; \mathbb{E}\Big[\Delta L^\Delta_{k+1} \mid \mathcal{F}^\Delta_{k\Delta}\Big], \qquad \xi^\Delta_k \;\equiv\; \Delta L^\Delta_{k+1} - m^\Delta_k,$$

and define the martingale

$$M^\Delta(t) \;\equiv\; \sum_{k<t/\Delta} \xi^\Delta_k, \qquad t \in [0,T],$$

and the drift process $B^\Delta(t) \equiv \sum_{k<t/\Delta} m^\Delta_k$ so that $L^\Delta_t = L^\Delta_0 + B^\Delta(t) + M^\Delta(t)$.

Assume the following intensity scaling and regularity *hold:*

(H1) There exist bounded, locally Lipschitz maps $\rho, \Lambda : (0,1) \to [0,\infty)$ such that

$$\mathbb{P}(a_{k\Delta} = 1 \mid \mathcal{F}^\Delta_{k\Delta}) = \rho(\pi^\Delta_{k\Delta})\,\Delta + o_p(\Delta),$$
$$\mathbb{P}(y_{k\Delta} = 1 \mid a_{k\Delta} = 1, \mathcal{F}^\Delta_{k\Delta}) = \Lambda(\pi^\Delta_{k\Delta}) + o_p(1),$$

[17] In the special case in which the L-type mixes to exactly match the H-type's risky frequency ($R_H(\pi) = R_L(\pi)$), $J^{\text{rec}}(\pi) = 0$ and the signaling term vanishes; then $\Delta^{\text{comm}}_H(s;\pi,k) = \zeta_k(\pi)\,\Delta_H(s;\pi)$ and the cutoff is independent of k. Proposition 57 therefore isolates the comparative static in k to environments where recommendations carry some informational content about type.

uniformly on $[0,T]$ *in probability. In particular, the compensators*

$$A^{+,\Delta}(t) \equiv \sum_{k<t/\Delta} \rho(\pi^{\Delta}_{k\Delta})\Lambda(\pi^{\Delta}_{k\Delta})\,\Delta, \quad A^{-,\Delta}(t) \equiv \sum_{k<t/\Delta} \rho(\pi^{\Delta}_{k\Delta})\big(1-\Lambda(\pi^{\Delta}_{k\Delta})\big)\,\Delta$$

converge in probability, uniformly on $[0,T]$ *(u.c.p.), to*

$$A^{+}(t) = \int_0^t \rho(\pi_s)\Lambda(\pi_s)\,ds, \qquad A^{-}(t) = \int_0^t \rho(\pi_s)\big(1-\Lambda(\pi_s)\big)\,ds.$$

(H2) There is $C<\infty$ *such that* $|J^{\pm}(\pi)| \le C$ *for all* $\pi \in (0,1)$*; moreover* $J^{\pm}$ *are locally Lipschitz in* π*.*

(H3) The conditional means satisfy

$$B^{\Delta}(t) = \sum_{k<t/\Delta} m^{\Delta}_k = \int_0^t \mu(\pi^{\Delta}_s)\,ds + o_p(1) \quad u.c.p.\ on\ [0,T],$$

for a bounded, locally Lipschitz $\mu : (0,1) \to \mathbb{R}$*.*

Then, as $\Delta \to 0$*, the pair* $\big(L^{\Delta},\, \nu^{\Delta}\big)$ *converges weakly in* $\mathbb{D}([0,T]) \times \mathbb{M}$ *to a special semimartingale* (L,ν) *with* characteristics

$$\text{(drift)}\ \ B(t) = \int_0^t \mu(\pi_s)\,ds, \qquad \text{(continuous covariation)}\ \ C \equiv 0,$$

and jump compensator

$$\nu(dt,dx) = \rho(\pi_t)\,\Lambda(\pi_t)\,dt\,\delta_{J^+(\pi_t)}(dx) + \rho(\pi_t)\,\big(1-\Lambda(\pi_t)\big)\,dt\,\delta_{J^-(\pi_t)}(dx),$$

where δ_z *denotes the Dirac measure at* z*. In particular,*

$$L_t = L_0 + \int_0^t \mu(\pi_s)\,ds + J^+(\pi)\Big(N^+(t) - A^+(t)\Big) + J^-(\pi)\Big(N^-(t) - A^-(t)\Big),$$

with (N^+, N^-) *independent inhomogeneous Poisson processes conditionally on* $(\pi_s)_{s\le t}$ *and compensators* $A^{\pm}$ *as above. The predictable quadratic variation satisfies*

$$\langle M\rangle(t) = \int_0^t \sigma^2(\pi_s)\,ds, \qquad \sigma^2(\pi) = \rho(\pi)\Big[\Lambda(\pi)\big(J^+(\pi)\big)^2 + \big(1-\Lambda(\pi)\big)\big(J^-(\pi)\big)^2\Big].$$

Small-jump regime. *If, in addition,* $\sup_{\pi\in(0,1)}\big(|J^+(\pi)| \vee |J^-(\pi)|\big) \to 0$ *as* $\Delta \to 0$*, then* $M^{\Delta} \Rightarrow W$ *where* W *is a continuous Gaussian martingale with* $[W](t) = \int_0^t \sigma^2(\pi_s)\,ds$*; hence*

$$L^{\Delta} \Rightarrow L_0 + \int_0^t \mu(\pi_s)\,ds + W(t) \quad in\ \mathbb{D}([0,T]).$$

Proof. Step 1. For each Δ and k, condition on $\mathcal{F}^{\Delta}_{k\Delta}$. Given $a_{k\Delta} \in \{0,1\}$, the next-period outcome $y_{k\Delta} \in \{0,1\}$ yields $\Delta L^{\Delta}_{k+1} \in \{0, J^+, J^-\}$ with probabilities

$$\mathbb{P}(\Delta L^{\Delta}_{k+1} = J^+ | \mathcal{F}^{\Delta}_{k\Delta}) = \rho(\pi^{\Delta}_{k\Delta})\,\Lambda(\pi^{\Delta}_{k\Delta})\,\Delta + o_p(\Delta),$$

$$\mathbb{P}(\Delta L^{\Delta}_{k+1} = J^- | \mathcal{F}^{\Delta}_{k\Delta}) = \rho(\pi^{\Delta}_{k\Delta})\big(1-\Lambda(\pi^{\Delta}_{k\Delta})\big)\,\Delta + o_p(\Delta),$$

and $\mathbb{P}(\Delta L^{\Delta}_{k+1} = 0 | \mathcal{F}^{\Delta}_{k\Delta}) = 1 - O_p(\Delta)$, uniformly on $[0, T]$ in probability by (H1). By construction, M^{Δ} is a square-integrable martingale with respect to the càdlàg filtration $(\mathcal{F}^{\Delta}_t)$, with jumps $\Delta M^{\Delta}_{(k+1)\Delta} = \xi^{\Delta}_k$ and predictable quadratic variation

$$\langle M^{\Delta} \rangle(t) \; = \; \sum_{k<t/\Delta} \mathbb{E}\Big[(\xi^{\Delta}_k)^2 \mid \mathcal{F}^{\Delta}_{k\Delta}\Big].$$

Since $|m^{\Delta}_k| = O_p(\Delta)$ and $|\Delta L^{\Delta}_{k+1}| \le C$ by (H2), we have

$$\mathbb{E}\Big[(\xi^{\Delta}_k)^2 \mid \mathcal{F}^{\Delta}_{k\Delta}\Big] = \mathbb{E}\Big[(\Delta L^{\Delta}_{k+1})^2 \mid \mathcal{F}^{\Delta}_{k\Delta}\Big] + O_p(\Delta^2)$$

and thus, uniformly on $[0, T]$ in probability,

$$\langle M^{\Delta} \rangle(t) \; = \; \sum_{k<t/\Delta} \Big[(J^+)^2 \rho \Lambda \; + \; (J^-)^2 \rho(1 - \Lambda)\Big](\pi^{\Delta}_{k\Delta})\,\Delta + o_p(1) \; \xrightarrow[\Delta \to 0]{\mathbb{P}} \; \int_0^t \sigma^2(\pi_s)\, ds,$$

where σ^2 is as stated and we used Riemann-sum convergence plus boundedness/Lipschitz continuity (H1)–(H2). Similarly, by (H3) $B^{\Delta}(\cdot) \to \int_0^{\cdot} \mu(\pi_s)\, ds$ u.c.p. Hence the *predictable characteristics* of L^{Δ} (with truncation function $h(x) \equiv x$) are

$$B^{\Delta}(t) \to B(t), \qquad C^{\Delta}(t) \equiv 0 \to 0, \qquad \nu^{\Delta}(dt, dx) \Rightarrow \nu(dt, dx),$$

where

$$\nu^{\Delta}(dt, dx) \; = \; \sum_{k<t/\Delta} \Big\{ \rho(\pi^{\Delta}_{k\Delta})\Lambda(\pi^{\Delta}_{k\Delta})\,\Delta\,\delta_{J^+(\pi^{\Delta}_{k\Delta})}(dx) \\ + \rho(\pi^{\Delta}_{k\Delta})(1 - \Lambda(\pi^{\Delta}_{k\Delta}))\,\Delta\,\delta_{J^-(\pi^{\Delta}_{k\Delta})}(dx) \Big\}.$$

By (H1)–(H2), ν^{Δ} converges in probability (in the sense of measures on $[0, T] \times \mathbb{R}$ endowed with the vague topology) to

$$\nu(dt, dx) = \rho(\pi_t)\Lambda(\pi_t)\, dt\, \delta_{J^+(\pi_t)}(dx) + \rho(\pi_t)\big(1 - \Lambda(\pi_t)\big)\, dt\, \delta_{J^-(\pi_t)}(dx).$$

Step 2. By Aldous' criterion, it suffices to show that for any sequence of $(\mathcal{F}^{\Delta}_t)$-stopping times $\tau_{\Delta} \le T$ and any $\delta_{\Delta} \downarrow 0$,

$$L^{\Delta}_{\tau_{\Delta}+\delta_{\Delta}} - L^{\Delta}_{\tau_{\Delta}} \; \xrightarrow{\mathbb{P}} \; 0.$$

Since $L^{\Delta} = B^{\Delta} + M^{\Delta}$ and B^{Δ} converges u.c.p. to a continuous limit, it suffices to control M^{Δ}. For the martingale increments,

$$\mathbb{E}\Big[\big(M^{\Delta}_{\tau_{\Delta}+\delta_{\Delta}} - M^{\Delta}_{\tau_{\Delta}}\big)^2 \mid \mathcal{F}^{\Delta}_{\tau_{\Delta}}\Big] = \mathbb{E}\Big[\langle M^{\Delta} \rangle(\tau_{\Delta} + \delta_{\Delta}) - \langle M^{\Delta} \rangle(\tau_{\Delta}) \mid \mathcal{F}^{\Delta}_{\tau_{\Delta}}\Big].$$

By the u.c.p. convergence of $\langle M^{\Delta} \rangle$ established above and boundedness of σ^2, the right-hand side is bounded by $K\,\delta_{\Delta} + o_p(1)$ for some $K < \infty$. Hence the conditional second moment of the increment tends to 0 in probability, which implies Aldous tightness (see, e.g., Ethier and Kurtz (1986), Thm. 3.8.6). Therefore $\{L^{\Delta}\}$ is tight in $\mathbb{D}([0, T])$.

Step 3. Let L be any weak limit point along a subsequence $\Delta_n \downarrow 0$. By Step 1, the characteristics $(B^{\Delta_n}, C^{\Delta_n}, \nu^{\Delta_n})$ converge in probability to $(B, 0, \nu)$ uniformly on compacts. By the general convergence theorem for semimartingales (e.g., Jacod and Shiryaev (2003), Thm. IX.3.9), L is a special semimartingale with characteristics $(B, 0, \nu)$ and admits the decomposition

$$L_t \;=\; L_0 \;+\; \int_0^t \mu(\pi_s)\, ds \;+\; \int_0^t\!\!\int x\,(\mu^L(ds,dx) - \nu(ds,dx)),$$

where μ^L is the jump measure of L. Since ν is a.s. absolutely continuous w.r.t. dt with atoms at $J^{\pm}(\pi_t)$ and finite mass on bounded sets, L can be represented (conditionally on the path (π_s)) as a sum of two compensated inhomogeneous Poisson integrals with jump sizes $J^{\pm}(\pi_t)$ and intensities $\rho(\pi_t)\Lambda(\pi_t)$ and $\rho(\pi_t)(1-\Lambda(\pi_t))$, respectively. In particular, the martingale part has predictable quadratic variation $\langle M\rangle(t) = \int_0^t \sigma^2(\pi_s)\, ds$.

Step 4. Because the limiting characteristics are unique and continuous functionals of the (prelimit) characteristics, every subsequence has the same law for any limit point; hence the entire sequence converges to the unique limit characterized above.

If $\sup_\pi(|J^+(\pi)| \vee |J^-(\pi)|) \to 0$, then the jump compensator ν^Δ assigns vanishing mass to $\{|x| > \varepsilon\}$ for any fixed $\varepsilon > 0$. The martingale functional CLT (e.g., Rebolledo (1980); Hall and Heyde (1980), Thm. 3.2) applies: M^Δ converges to a continuous Gaussian martingale W with quadratic variation given by the limit of $\langle M^\Delta\rangle$, namely $\int_0^t \sigma^2(\pi_s)\, ds$. Adding the drift limit from (H3) yields the stated continuous-limit representation of L^Δ. □

Proposition 60. *Let $V(\pi)$ be C^1 and convex. In the diffusion regime with constant experimentation ($a \equiv 1$), the expert's value solves*

$$\rho V(\pi) = \phi + \mu(\pi)\, V'(\pi) + \sup_{s\in S}\{\delta\, \mathbb{E}[V(\pi')] - V(\pi)\},$$

where $\mu(\pi) = \pi(1-\pi)\, e^*(1,\pi)\, [p_H(\pi) - p_L(\pi)]$ is the expected log-odds drift mapped to belief space. The discrete-time cutoff characterization in Theorem 5 converges to the bang-bang policy of the HJB.

I Contingent Bonus on Success

We formalize the effect of a success–contingent transfer on the expert's policy and on learning. Suppose a bonus $\beta \geq 0$ is paid to the expert whenever $(a = 1, y = 1)$ occurs. Period utility is $\phi\, a + \beta\, \mathbf{1}\{a = 1, y = 1\}$; the implementer's problem and effort $e^*(1, \pi; s)$ are unchanged (the contract affects only the expert's utility). Standing Assumptions A1–A4 continue to hold.

I.1 Cutoff shift and comparative statics in β

Given public reputation $\pi \in (0, 1)$ and signal s, the High type's risky–minus–safe continuation difference with bonus β equals

$$\begin{aligned}\Delta_H^\beta(s;\pi) \;=\; \phi \;+\; \beta\, P_S(s;\pi) \;+\; \delta\Big\{ & P_S(s;\pi)\, V(\pi^+(\pi,s)) \\ & + (1 - P_S(s;\pi))\, V(\pi^-(\pi,s)) - V(\pi^{\text{rec}}(0;\pi,s))\Big\},\end{aligned} \tag{29}$$

where $P_S(s;\pi)$ is the success probability under $a = 1$ (Section 3), $\pi^{\text{rec}}(0;\pi,s)$ is the recommendation-only posterior after $a = 0$, and $\pi^{\pm}(\pi,s)$ are the outcome posteriors after $a = 1$ (Appendix A). Write the equilibrium cutoff as $s^*_\beta(\pi)$, defined by $\Delta_H^\beta(s^*_\beta(\pi);\pi) = 0$.

Proposition 61. *Fix $\pi \in (0, 1)$. Then:*

(i) For each $\beta \geq 0$, the map $s \mapsto \Delta_H^\beta(s;\pi)$ is continuous and strictly increasing (single crossing), hence there exists a unique cutoff $s^*_\beta(\pi)$ with $\Delta_H^\beta(s^*_\beta(\pi);\pi) = 0$.

(ii) $\beta \mapsto s^*_\beta(\pi)$ is continuous and (weakly) decreasing. If $P_S(s^*_\beta(\pi);\pi) \in (0,1)$ and $\partial_s \Delta_H^\beta(s^*_\beta(\pi);\pi) > 0$, then $s^*_\beta(\pi)$ is strictly *decreasing and differentiable in β, with*

$$\frac{\partial s^*_\beta(\pi)}{\partial \beta} \;=\; -\,\frac{P_S\Big(s^*_\beta(\pi);\pi\Big)}{\partial_s \Delta_H^\beta\Big(s^*_\beta(\pi);\pi\Big)} \;<\; 0.$$

(iii) If there exists $c(\pi) > 0$ with $\partial_s \Delta_H^\beta(s;\pi) \geq c(\pi)$ for all s in a neighborhood of $s^*_\beta(\pi)$ and all β in a compact set $B \subset \mathbb{R}_+$, then for any $\beta_1, \beta_2 \in B$,

$$\big|s^*_{\beta_2}(\pi) - s^*_{\beta_1}(\pi)\big| \;\leq\; \frac{|\beta_2 - \beta_1|}{c(\pi)}.$$

Proof. (i) The only change from Δ_H (Appendix A) is the additive term $\beta P_S(s;\pi)$; by Assumptions A1–A3 and Lemma 27, P_S is continuous and strictly increasing in s, while the remaining (baseline) term is continuous and strictly increasing by Lemma 28. Existence/uniqueness follows.

(ii) Monotone comparative statics: for fixed s, $\partial_\beta \Delta_H^\beta(s;\pi) = P_S(s;\pi) \geq 0$. Since $s \mapsto \Delta_H^\beta(s;\pi)$ has strict single crossing, Topkis implies that the smallest root $s^*_\beta(\pi)$ is (weakly) decreasing in β. When interior, the implicit function theorem applies and yields the derivative displayed.

(iii) Integrate the IFT formula along a path from β_1 to β_2 and bound the denominator below by $c(\pi)$; use that $0 \leq P_S \leq 1$. □

By the IFT formula, the *local* sensitivity satisfies $\big|\partial s^*_\beta(\pi)/\partial\beta\big| = P_S(s^*_\beta(\pi);\pi)/\partial_s \Delta_H^\beta(s^*_\beta(\pi);\pi)$. Since P_S rises with reputation via effort (Lemma 27), the absolute effect is largest at reputations where P_S is high (typically at higher π). Corners: if $P_S = 0$ at the threshold, the bonus has no first-order effect at that π; if $P_S = 1$ and $\partial_s \Delta_H^\beta > 0$, the formula still applies.

I.2 Targeting a desired threshold

A minimal bonus that implements any target cutoff $\tilde{s}(\pi)$ (with $\Delta_H(\tilde{s}(\pi);\pi) < 0$ at $\beta = 0$) is

$$\beta^{\min}(\pi, \tilde{s}) \;=\; -\frac{\Delta_H\big(\tilde{s}(\pi);\pi\big)}{P_S\big(\tilde{s}(\pi);\pi\big)}\,, \tag{30}$$

with the convention that if $P_S(\tilde{s}(\pi);\pi) = 0$ the local instrument is ineffective (no success can occur at the target). Any $\beta > \beta^{\min}$ strictly lowers the cutoff relative to $\tilde{s}(\pi)$.

I.3 Effect on the speed of learning (KL drift)

Let L_{t+1} be the one-step likelihood ratio between period-$(t{+}1)$ observations under H and L (Appendix C). Define the per-period expected Kullback–Leibler drift at reputation π under contract β by

$$\begin{aligned}\kappa(\pi;\beta) \;&\equiv\; \mathbb{E}[\log L_{t+1} \mid \pi_t = \pi; \beta]\\ &=\; \mathbb{P}(a_t = 1 \mid \pi;\beta)\cdot D_{\mathrm{KL}}\big(\mathcal{L}_H(y \mid a{=}1,\pi)\,\|\,\mathcal{L}_L(y \mid a{=}1,\pi)\big) \;+\; (\text{rec. term}),\end{aligned}$$

where the "rec. term" collects the (nonnegative) KL from observing the recommendation itself when it is informative about type. The outcome KL is strictly positive whenever $p_H(\pi,s) \neq p_L(\pi,s)$ on a set of positive probability (Assumption A1).

Proposition 62. *At $\beta = 0$, a marginal increase in the bonus strictly raises the expected KL drift:*

$$\frac{d}{d\beta}\,\mathbb{E}\big[\kappa(\pi_t;\beta)\big]\Big|_{\beta=0} \;=\; \mathbb{E}\Big[\frac{\partial\,\mathbb{P}(a_t=1\mid\pi_t;\beta)}{\partial\beta}\Big|_{\beta=0}\Big] \cdot \underbrace{D_{\mathrm{KL}}\big(\mathcal{L}_H(y\mid a{=}1,\pi_t)\,\|\,\mathcal{L}_L(y\mid a{=}1,\pi_t)\big)}_{>0} \;+\; (\text{rec. term}) \;>\; 0,$$

provided recommendations are not exactly *uninformative at $\beta = 0$ (in which case the "rec. term" vanishes).*

Proof. By Proposition 61, $\partial_\beta s^*_\beta(\pi)\big|_{\beta=0} < 0$ wherever $P_S(s^*(\pi);\pi) > 0$, so the risky region in s (hence $\mathbb{P}(a = 1 \mid \pi;\beta)$) expands at $\beta = 0$. Each experiment contributes strictly positive outcome KL by Assumption A1 (the Bernoulli success parameters under H and L differ on a set of positive probability). Linearity of expectation gives the display. The recommendation KL term is weakly positive and independent of β to first order unless the bonus also shifts recommendation informativeness directly; here it does not, since the agent's problem is unchanged. □

I.4 Affine success/failure contracts

More generally, consider an affine contract paying β_1 upon $(a{=}1, y{=}1)$ and $-\beta_0$ upon $(a{=}1, y{=}0)$, with $\beta_1, \beta_0 \geq 0$. Then

$$\Delta_H^{\beta_1,\beta_0}(s;\pi) \;=\; \phi \;+\; \big[\beta_1 P_S(s;\pi) - \beta_0(1 - P_S(s;\pi))\big] \;+\; \delta\Big\{\cdots\Big\}, \tag{31}$$

so all arguments above go through with the additive rotation $[\beta_1 P_S - \beta_0(1 - P_S)]$. In particular, increasing β_1 (holding β_0 fixed) lowers the cutoff, and increasing β_0 raises it; the IFT derivative becomes

$$\frac{\partial s^*_{\beta_1,\beta_0}(\pi)}{\partial\beta_1} = -\frac{P_S\Big(s^*_{\beta_1,\beta_0}(\pi);\pi\Big)}{\partial_s\Delta_H^{\beta_1,\beta_0}\Big(s^*_{\beta_1,\beta_0}(\pi);\pi\Big)}, \qquad \frac{\partial s^*_{\beta_1,\beta_0}(\pi)}{\partial\beta_0} = \frac{1 - P_S\Big(s^*_{\beta_1,\beta_0}(\pi);\pi\Big)}{\partial_s\Delta_H^{\beta_1,\beta_0}\Big(s^*_{\beta_1,\beta_0}(\pi);\pi\Big)}.$$

The bonus rotates the risky–safe margin by an amount proportional to P_S, so it is most potent where effort (hence P_S) is high. Because experiments carry positive KL, any policy that increases their frequency speeds up learning; the appendix's formulas make these effects transparent and easy to calibrate.

References

Backus, M. and A. T. Little (2020): "I Don't Know," *American Political Science Review*, 114, 724–743.

Birkhäuer, J., J. Gaab, J. Kossowsky, S. Hasler, P. Krummenacher, C. Werner, and H. Gerger (2017): "Trust in the Health Care Professional and Health Outcome: A Meta-Analysis," *PLOS ONE*, 12, e0170988.

Birkmeyer, J. D., J. F. Finks, A. O'Reilly, M. Oerline, A. M. Carlin, A. R. Nunn, J. Dimick, M. Banerjee, N. J. O. Birkmeyer, and M. B. S. Collaborative (2013): "Surgical Skill and Complication Rates after Bariatric Surgery," *The New England Journal of Medicine*, 369, 1434–1442.

Birkmeyer, J. D., A. E. Siewers, E. V. A. Finlayson, T. A. Stukel, F. L. Lucas, I. Batista, H. G. Welch, and J. E. Wennberg (2003): "Surgeon Volume and Operative Mortality in the United States," *The New England Journal of Medicine*, 349, 2117–2127.

Bolton, P., X. Freixas, and J. Shapiro (2012): "The Credit Ratings Game," *Journal of Finance*, 67, 85–111.

Bolton, P. and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67, 349–374.

Cripps, M., G. Mailath, and L. Samuelson (2004): "Imperfect Monitoring and Impermanent Reputation," *Econometrica*, 72, 407–432.

Deb, R., M. Mitchell, and M. M. Pai (2022): "(Bad) Reputation in Relational Contracting," *Theoretical Economics*, 17, 763–800.

Ethier, S. N. and T. G. Kurtz (1986): *Markov Processes: Characterization and Convergence*, Wiley Series in Probability and Mathematical Statistics, New York: John Wiley & Sons.

Halac, M., N. Kartik, and Q. Liu (2016): "Optimal Contracts for Experimentation," *The Review of Economic Studies*, 83, 1040–1091.

Hall, P. and C. C. Heyde (1980): *Martingale Limit Theory and Its Application*, New York: Academic Press.

Haskard Zolnierek, K. B. and M. R. DiMatteo (2009): "Physician Communication and Patient Adherence to Treatment: A Meta-Analysis," *Medical Care*, 47, 826–834.

Holmström, B. (1982): "Managerial Incentive Problems: A Dynamic Perspective," *Essays in Economics and Management in Honor of Lars Wahlbeck*, reprinted in *Review of Economic Studies* (1999), 66(1): 169–182.

Inderst, R. and M. Ottaviani (2012): "Financial Advice," *Journal of Economic Literature*, 50, 494–512.

Jacod, J. and A. N. Shiryaev (2003): *Limit Theorems for Stochastic Processes*, Grundlehren der Mathematischen Wissenschaften, Berlin: Springer, 2 ed.

Kamenica, E. and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68.

Loh, R. K. and R. M. Stulz (2011): "When Are Analyst Recommendation Changes Influential?" *Review of Financial Studies*, 24, 593–627.

Lukyanov, G., A. Vlasova, and M. Ziskelevich (2025): "Risky Advice and Reputational Bias," *arXiv preprint arXiv:2508.19707*, submitted.

Manso, G. (2011): "Motivating Innovation," *Journal of Finance*, 66, 1823–1860.

Morris, S. (2001): "Political Correctness," *Journal of Political Economy*, 109, 231–265.

Mylovanov, T. and N. Klein (2017): "Will Truth Out?—An Advisor's Quest to Appear Competent," *Journal of Mathematical Economics*, 72, 112–121.

Ottaviani, M. and P. N. Sørensen (2006): "Reputational Cheap Talk," *RAND Journal of Economics*, 37, 155–175.

Ozyurt, S. (2016): "A Fraudulent Expert and Short-Lived Customers," Tech. rep., Sabanci University Economics Working Paper, available at SSRN: 10.2139/ssrn.2671388.

Pavesi, F. and M. Scotti (2022): "Good lies," *European Economic Review*, 141, 103965.

Prat, A. (2005): "The Wrong Kind of Transparency," *American Economic Review*, 95, 862–877.

Prendergast, C. (1993): "A Theory of "Yes Men": Honesty and Opinion Herding in Organizations," *American Economic Review*, 83, 757–770.

Prendergast, C. and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104, 1105–1134.

Rebolledo, R. (1980): "Central Limit Theorems for Local Martingales," *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 51, 269–286.

Scharfstein, D. and J. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80, 465–479.

Schottmüller, C. (2019): "Too Good to Be Truthful: Why Competent Advisers are Fired," *Journal of Economic Theory*, 181, 333–360.

Tadelis, S. (2002): "The Market for Reputations as an Incentive Mechanism," *Journal of Political Economy*, 110, 854–882.

Thomas, C. (2019): "Experimentation with Reputation Concerns: Dynamic Signalling with Changing Types," *Journal of Economic Theory*, 179, 366–415.