Relatively Robust Timing of Medical Tests under Ambiguous Prevalence Dynamics

Thomas A. Weber
Chair of Operations, Economics and Strategy
École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
thomas.weber@epfl.ch

***Abstract*—Medical tests are investments in information whose value depends on the prior probability of disease, on how that belief evolves over time, and on actions taken after evidence is observed. Deterministic prevalence dynamics are tractable, but policies calibrated under a single intensity can be unreliable when epidemiological speeds are misspecified. We model belief as evolving deterministically under an autonomous differential equation whose drift is scaled by an unknown multiplicative speed factor. The decision maker may wait, test at a cost, and treat (or not treat) after outcomes, with retesting possible after negative results. For stationary policies defined by two belief thresholds, we obtain a closed-form expression for the expected discounted value, yielding a fast evaluator for computing optimal thresholds. Relative robustness is measured by a performance ratio that compares the value of a candidate threshold pair with the ex-post optimal value under the realized speed factor; relatively robust thresholds maximize the worst-case ratio over an ambiguity interval. We provide a grid-search procedure and a numerical illustration motivated by post-exposure tuberculosis infection screening, where timing decisions unfold over weeks to years and intensity ambiguity is clinically consequential.**

***Keywords*—Medical Testing, Performance Index, Relative Robustness, Robust Optimization, Screening, Value of Information**

I. Introduction

Uncertainty is a central feature of the economics of healthcare [1], and screening for infectious diseases is a case in point. A diagnostic test is useful only insofar as it changes decisions, and the magnitude of that change depends on the prior probability of disease and on how that probability evolves over time with exposure, cohort conditions, and other drivers. In practice, the mechanisms governing prevalence accumulation are estimated with error, differ across cohorts, and may drift over time. A purely deterministic calibration can deliver a precise policy, but with misspecified intensity.

A standard response is to impose a fully stochastic model and optimize expected discounted payoffs. In many screening environments, however, the data needed to specify and validate a stochastic structure are limited, while recommended policies can be sensitive to distributional and horizon assumptions. Value-of-information methods in health decision science make this sensitivity explicit and have emphasized the option value of delay [2]–[7]. Moreover, information acquisition is itself a timing decision, and delaying evidence collection can be optimal when it is unlikely to alter the action choice. Cipriano and Weber [8] develop this logic in a dynamic health-policy setting, highlighting the value of timing prevalence measurement. A parallel operations-research literature formulates screening, monitoring, and treatment as sequential decision problems (often Markovian or partially observed), and derives structural policies such as thresholds and index rules [9]–[12]. These contributions are typically calibrated under a single epidemiological or clinical evolution law, which raises a distinct question: how fragile is an exact policy when the underlying evolution speed is misspecified?

We retain deterministic dynamics for the evolution of belief, but treat the *speed* of that evolution as ambiguous, modeled by an unknown multiplicative parameter in a compact interval. The decision maker chooses a testing policy *ex ante*, while the parameter is revealed only *ex post* through realized epidemiological conditions. The criterion is not an absolute worst-case payoff. Instead, we evaluate a chosen policy by a *performance ratio*, defined as the realized value relative to the value that could have been achieved under perfect information about the realized parameter. This relative comparison implements a robustness principle, familiar from competitive analysis in computer science [13], [14], which has been developed for economic decision problems under the term *relative robustness* [15]. The ratio criterion is appropriate for screening because it respects the fact that the attainable value varies with the speed of prevalence evolution; it compares the chosen policy with what was actually feasible in the realized environment.

II. Belief Dynamics

A. Law of Motion

Let $t \geq 0$ denote elapsed time since a baseline (index) assessment at which the initial belief $b(0) = b_0$ is formed.[1]Let $b(t)$ denote the decision maker's belief at time t that an individual is ill, conditional on the available information. The belief can be interpreted as the prevalence in a relevant cohort or risk stratum at time t. Under exchangeability, $b(t)$ is also the probability that a randomly selected individual from that stratum is ill. We assume that the belief state drifts

[1]In epidemiologic cohort designs, baseline (cohort entry) defines the start of follow-up; in event-history/survival settings, a common time scale is time since entry (time-on-study), though other scales such as chronological age may be appropriate in cohort applications.

deterministically and saturates at a ceiling $\bar{b} \in (0, 1]$, reflecting the fact that not all diseases progress to full incidence.[2]

Assumption 1 (**Monotone Drift with Ceiling**). *There exists $\bar{b} \in (0, 1]$ such that $f : [0, \bar{b}] \to \mathbb{R}$ is continuous on $[0, \bar{b}]$, continuously differentiable on $(0, \bar{b})$, and*

$$f(b) > 0 \ \text{ for } b \in (0, \bar{b}), \quad f(\bar{b}) = 0.$$

Between information arrivals, the decision maker's belief evolves according to

$$\dot{b} = (1 + \theta)f(b), \tag{1}$$

where θ is an unknown multiplicative parameter. In large populations, prevalence paths can be approximated by deterministic compartmental dynamics, so treating $b(t)$ as deterministic should be viewed here as a reduced-form approximation rather than as a claim that uncertainty is absent [16]–[20]. The nominal model corresponds to $\theta = 0$. The decision maker considers a compact ambiguity set, so

$$\theta \in \Theta = [\theta_1, \theta_2], \quad -1 < \theta_1 \leq 0 \leq \theta_2 < 1.$$

The initial belief $b(0) = b_0 \in (0, \bar{b})$ is taken as given.

Remark 1. If $\bar{b} < \bar{\alpha}$ (defined in Sec. III), then treatment based solely on the base-rate belief b is never optimal. Testing may still be valuable if a positive test can raise posterior belief above $\bar{\alpha}$. In the numerical illustration (Sec. VII) we focus on the empirically common case $\bar{b} > \bar{\alpha}$, in which both direct treatment (at high base rates) and test-triggered treatment may arise.

Example 1. A particularly tractable mechanism is logistic growth in the sense of Verhulst [21]. In epidemiology, the same functional form arises from a Susceptible–Infected (SI) mass-action term in the no-recovery case [16]–[18]. To accommodate the empirically common feature that not all diseases progress to full incidence, we allow for a carrying capacity $\bar{b} \in (0, 1]$ and set

$$f(b) = \eta \, b(\bar{b} - b), \quad \eta > 0. \tag{2}$$

Then $b(t)$ increases on $(0, \bar{b})$ and converges to $\bar{b}$. The multiplicative ambiguity in Eq. (1) corresponds to uncertainty in an effective incidence or contact intensity, a natural locus of misspecification across cohorts and environments [17], [19].

B. Belief Evolution

Under nominal dynamics ($\theta = 0$), the travel time from belief b to belief b' is given by

$$\tau_0(b, b') = \int_b^{b'} \frac{du}{f(u)}, \quad 0 < b < b' < \bar{b}. \tag{3}$$

The multiplicative ambiguity in Eq. (1) is naturally conserved in the general travel time.

[2]The time origin $t = 0$ is an index time at which the prior b_0 is assessed (e.g., at presentation, symptom onset, or a reported exposure). Thus, t measures the span since that index time; in cohort screening it may coincide with calendar age (cf. footnote 1), but in individual diagnostic timing problems it is more naturally "time since baseline" or "time since exposure."

Lemma 1 (**Time Scaling**). *Let $\tau_\theta(b, b')$ denote the travel time from b to b' under Eq. (1). Then*

$$\tau_\theta(b, b') = \frac{\tau_0(b, b')}{1 + \theta}, \quad 0 < b < b' < \bar{b}, \ \theta \in \Theta. \tag{4}$$

Proof. Along any trajectory under Eq. (1), $dt = db/((1 + \theta)f(b))$. Integrating from b to b' yields Eq. (4). □

Example 2. For the logistic drift in Eq. (2), one obtains the nominal travel time

$$\tau_0(b, b') = \frac{1}{\eta\bar{b}} \ln\left(\frac{b'/(\bar{b} - b')}{b/(\bar{b} - b)}\right), \quad 0 < b < b' < \bar{b}, \tag{5}$$

which then implies the ambiguous travel time $\tau_\theta(b, b')$ according to Lemma 1.

C. Discounting

All payoffs are discounted continuously at the constant rate $r > 0$. Thus, waiting for the belief $b > 0$ to evolve to $b' \in [b, \bar{b})$ yields an ambiguous discount factor,

$$D_\theta(b, b') = \exp\left(-\frac{r\,\tau_0(b, b')}{1 + \theta}\right), \quad \theta \in \Theta. \tag{6}$$

Example 3. For logistic drift, Eqs. (5) and (6) yield the ambiguous discount factor

$$D_\theta(b, b') = \left(\frac{b/(\bar{b} - b)}{b'/(\bar{b} - b')}\right)^{\hat{r}/(1+\theta)}, \quad \theta \in \Theta,$$

where we have set $\hat{r} = r/(\eta\bar{b})$.

Lemma 2 (**Closed-Form Flow under Logistic Drift**). *Suppose f is logistic as in Eq. (2). Define the odds transform*

$$\Omega(b) = \frac{b}{\bar{b} - b}, \quad 0 < b < \bar{b}.$$

Along any trajectory of Eq. (1), Ω evolves exponentially:

$$\Omega(b(t)) = \Omega(b_0) \exp\big((1 + \theta)\eta\bar{b}\,t\big), \quad t \geq 0. \tag{7}$$

Equivalently, the belief path admits the closed form

$$b(t) = \frac{\bar{b}\,\Omega(b_0) \exp\big((1 + \theta)\eta\bar{b}\,t\big)}{1 + \Omega(b_0) \exp\big((1 + \theta)\eta\bar{b}\,t\big)}, \quad t \geq 0. \tag{8}$$

Proof. For logistic f, one has $\dot{b} = (1 + \theta)\eta b(\bar{b} - b)$. A direct calculation gives $\dot{\Omega} = \dot{b}\,\bar{b}/(\bar{b} - b)^2 = (1 + \theta)\eta\bar{b}\,\Omega$, which implies Eq. (7) and therefore also Eq. (8). □

III. Test Technology and Treatment Payoff

A. Binary Test and Bayesian Updating

A binary diagnostic test yields outcome S (positive) or $\bar{S}$ (negative). Sensitivity and specificity satisfy $q_1, q_2 \in (1/2, 1)$:

$$P(S|X) = q_1, \quad P(\bar{S}|\bar{X}) = q_2,$$

where X denotes illness and $\bar{X}$ health. A test costs $C \geq 0$ each time it is administered. Moreover, given a current belief b, the probability of a positive outcome is

$$\ell(b) = q_1 b + (1 - q_2)(1 - b).$$

Bayes' rule yields the posterior after a positive test,

$$p(b) = \frac{\kappa b}{1 + (\kappa - 1)b}, \quad \kappa = \frac{q_1}{1 - q_2} > 1,$$

and after a negative test,

$$n(b) = \frac{(1/\nu)b}{1 + ((1/\nu) - 1)b}, \quad \nu = \frac{q_2}{1 - q_1} > 1.$$

For $b \in (0, 1)$, it is $n(b) < b < p(b)$.

B. Treatment Payoff and the Option Not to Treat

Let $B > 0$ denote the net benefit from treating an ill individual and let $H > 0$ denote the net harm from treating a healthy individual, where both include treatment cost. Treating at belief b yields the expected net payoff

$$U(b) = Bb - H(1 - b) = (B + H)b - H.$$

Not treating yields payoff 0. Hence the stopping payoff is

$$S(b) = \max\{0, U(b)\}.$$

This induces the treatment-indifference threshold

$$\bar{\alpha} = \frac{H}{B + H} \in (0, 1).$$

IV. Stationary Threshold Policies

A. Policy Class and Separation

Under the infinite-horizon specification, the value function depends only on the current belief and on θ. We focus on stationary two-threshold policies parameterized by (β, γ).

Definition 1 (Two-Threshold Policy). *Given (β, γ) with $0 \le \beta \le \gamma \le 1$, the policy is: (i)* wait *if $b < \beta$; (ii)* test *if $\beta \le b < \gamma$; (iii)* treat *if $b \ge \gamma$. After a negative outcome at belief b, the policy waits until the belief reaches β and then retests. After a positive outcome at belief b, the policy treats immediately (and, as always, may opt not to treat if the induced payoff is nonpositive).*

To maintain internal consistency, we restrict attention to threshold pairs for which positive outcomes jump into the treat region and negative outcomes jump into the wait region.

Assumption 2 (Separation). $n(\gamma) \le \beta \le \gamma \le p(\beta)$.

Once a solution has been determined, this assumption can be verified *ex post*.

B. Expected Value of Two-Threshold Policy

Let $V(b_0; \beta, \gamma|\theta)$ denote the expected discounted payoff generated by the policy (β, γ) starting from belief b_0. Define

$$D_\theta^-(b, \beta) = D_\theta(n(b), \beta).$$

Proposition 1 (Expected Value under (β, γ)). *Fix $\theta \in \Theta$ and (β, γ) satisfying Assumption 2 with $\gamma \ge \bar{\alpha}$. Let $V_\beta(\beta, \gamma|\theta)$ denote the continuation value when belief equals β under the policy. Then*

$$V_\beta(\beta, \gamma|\theta) = \frac{\ell(\beta)S(p(\beta)) - C}{1 - (1 - \ell(\beta))D_\theta^-(\beta, \beta)}. \tag{9}$$

Moreover, the value at the initial belief b_0 equals

$$V(b_0; \beta, \gamma|\theta) = D_\theta(b_0, \beta)V_\beta(\beta, \gamma|\theta), \quad 0 \le b_0 < \beta,$$

$$V(b_0; \beta, \gamma|\theta) = (1 - \ell(b_0))D_\theta^-(b_0, \beta)V_\beta(\beta, \gamma|\theta) + \ell(b_0)S(p(b_0)) - C, \quad \beta \le b_0 < \gamma,$$

and

$$V(b_0; \beta, \gamma|\theta) = S(b_0), \quad \gamma \le b_0 \le 1.$$

Proof. At $b = \beta$, the policy prescribes a test. If the outcome is positive, Assumption 2 implies $p(\beta) \ge \gamma \ge \bar{\alpha}$ and thus immediate treatment yields $S(p(\beta))$. If the outcome is negative, belief becomes $n(\beta) \le \beta$ and the policy waits deterministically until the belief returns back to β, yielding the continuation value $D_\theta(n(\beta), \beta)V_\beta$. Taking expectations and subtracting C yields a fixed-point equation in V_β, which rearranges to Eq. (9). The value at b_0 follows by conditioning on whether b_0 lies in the wait, test, or treat region (i.e., $b_0 \in (0, \beta) \cup [\beta, \gamma) \cup [\gamma, \bar{b})$). □

Remark 2. Under logistic drift, all discount factors in Proposition 1 are explicit rational powers of odds ratios by Eq. (5) and Eq. (6). In particular,

$$D_\theta(b, b') = \left(\frac{\Omega(b)}{\Omega(b')}\right)^{\hat{r}/(1+\theta)}, \quad \hat{r} = r/(\eta\bar{b}),$$

and therefore $D_\theta^-(b, \beta) = \big(\Omega(n(b))/\Omega(\beta)\big)^{\hat{r}/(1+\theta)}$. Hence the evaluator is closed form, and the precomputation of F in Sec. VI becomes unnecessary in the logistic case.

C. No-Test Benchmark and the Logic of Waiting

A policy with $\beta = \gamma$ never tests: it waits until the belief reaches γ and then treats. Since $S(b) = 0$ for $b \le \bar{\alpha}$, treating too early can only destroy value relative to waiting.

For a fixed θ, define the no-test value

$$V^{\mathrm{NT}}(b_0|\theta) = \sup_{b \in [b_0, \bar{b}]} D_\theta(b_0, b)\, S(b).$$

If the maximizer $b^{\mathrm{NT}}(\theta)$ exists in $(\bar{\alpha}, \bar{b})$, it solves the first-order condition

$$\frac{B + H}{(B + H)b - H} = \frac{r}{(1 + \theta)f(b)}. \tag{10}$$

Thus, absent testing, the treatment boundary depends on the discount rate r, the drift f, and the speed parameter θ.

Proposition 2 (No-Test Boundary under Logistic Drift). *Suppose f is logistic as in Eq. (2). If the no-test maximizer $b^{\mathrm{NT}}(\theta)$ is interior in $(\bar{\alpha}, \bar{b})$, then it satisfies*

$$(B + H)\eta\,(1 + \theta)\,b(\bar{b} - b) = r\big((B + H)b - H\big), \tag{11}$$

and hence admits the closed form

$$b^{\mathrm{NT}}(\theta) = \frac{(B + H)\eta(1 + \theta)\bar{b} - r(B + H) + \sqrt{\Delta(\theta)}}{2(B + H)\eta(1 + \theta)}, \tag{12}$$

where

$$\Delta(\theta) = \big((B{+}H)\eta(1{+}\theta)\bar{b} - r(B{+}H)\big)^2 - 4(B{+}H)\eta(1{+}\theta)\,rH.$$

Proof. With logistic f, Eq. (10) is equivalent to Eq. (11). Solving the resulting quadratic equation and selecting the root in $(\bar{\alpha}, \bar{b})$ yields Eq. (12). □

V. Ex-Post Optimality and Relative Robustness

A. Ex-Post Optimal Value

For each realized $\theta \in \Theta$, define the ex-post optimal value within the policy class of Definition 1:

$$V^*(b_0|\theta) = \max_{(\beta,\gamma)\in\mathcal{D}} V(b_0;\beta,\gamma|\theta),$$

where $\mathcal{D}$ is the feasible set

$$\mathcal{D} = \Big\{(\beta,\gamma) \in [0,1]^2 : \ n(\gamma) \le \beta \le \gamma, \ \bar{\alpha} \le \gamma \le p(\beta)\Big\}.$$

If C is large or the test uninformative, the maximizer may lie on the degenerate boundary $\beta = \gamma$, corresponding to a no-test policy.

B. Performance Index and Relatively Robust Thresholds

For $(\beta,\gamma) \in \mathcal{D}$ and $\theta \in \Theta$ with $V^*(b_0|\theta) > 0$, the quality of the two-threshold policy relative to a given parameter realization is evaluated by the *performance ratio*,

$$\varphi(\beta,\gamma|\theta) = \frac{V(b_0;\beta,\gamma|\theta)}{V^*(b_0|\theta)} \in [0,1].$$

Definition 2 (Performance Index). *The performance index is defined as the worst-case performance ratio,*

$$\rho(\beta,\gamma) = \min_{\theta\in\Theta} \varphi(\beta,\gamma|\theta),$$

relative to the ambiguity set Θ.

A *relatively robust* threshold pair maximizes the performance index, so

$$(\hat{\beta},\hat{\gamma}) \in \arg\max_{(\beta,\gamma)\in\mathcal{D}} \rho(\beta,\gamma).$$

C. Boundary Reduction under Quasiconcavity

The practical simplification underlying relative robustness is that worst-case performance often occurs at $\partial\Theta$.

Proposition 3 (Boundary Representation). *If for each* $(\beta,\gamma) \in \mathcal{D}$*, the map* $\varphi(\beta,\gamma|\cdot)$ *is quasiconcave on* Θ*, then*

$$\rho(\beta,\gamma) = \min\Big\{\varphi(\beta,\gamma|\theta_1), \varphi(\beta,\gamma|\theta_2)\Big\}, \quad (\beta,\gamma) \in \mathcal{D}. \tag{13}$$

Proof. By quasiconcavity, $\varphi(\beta,\gamma|\cdot)$ attains its minimum over the compact interval Θ at an endpoint. □

When $(\hat{\beta},\hat{\gamma})$ lies in the interior of $\mathcal{D}$ and the minimum in Eq. (13) is attained through equality, one typically observes the equalization pattern

$$\varphi(\hat{\beta},\hat{\gamma}|\theta_1) = \varphi(\hat{\beta},\hat{\gamma}|\theta_2),$$

which can be viewed as a "balance at the extremes."

VI. Computation

Closed-form solutions for $(\hat{\beta},\hat{\gamma})$ are generally unavailable. Proposition 1, however, provides a fast evaluator for $V(b_0;\beta,\gamma|\theta)$ that enables brute-force computation on a grid.

Algorithm 1 Grid Search for Ex-Post Optimal Thresholds and Relatively Robust Thresholds

- **Input:** primitives $(b_0, \bar{b}, r, B, H, C, q_1, q_2)$, drift f, grids $\{b^i\}_{i=0}^N$, $\{\theta^j\}_{j=1}^M$
- Precompute $\bar{\alpha}$, $(p(b^i), n(b^i), \ell(b^i))$ on $\{b^i\}$, and the feasible grid set $\mathcal{D}_N = \mathcal{D} \cap (\{b^i\} \times \{b^i\})$
- Precompute $D_{\theta^j}(b^i, b^k)$ for all (i,k,j) as needed (closed form under logistic drift, Remark 2)
- **for** $j = 1$ to M **do**
  - Compute $V^*(b_0|\theta^j) = \max_{(\beta,\gamma)\in\mathcal{D}_N} V(b_0;\beta,\gamma|\theta^j)$ using Proposition 1
- **end for**
- Compute $\rho(\beta,\gamma) = \min_j V(b_0;\beta,\gamma|\theta^j)/V^*(b_0|\theta^j)$ for all $(\beta,\gamma) \in \mathcal{D}_N$
- Set $(\hat{\beta},\hat{\gamma}) \in \arg\max_{(\beta,\gamma)\in\mathcal{D}_N} \rho(\beta,\gamma)$
- **Output:** $(\hat{\beta},\hat{\gamma})$ and guarantee $\rho(\hat{\beta},\hat{\gamma})$

A. Grid Search on the Threshold Triangle

Let $\{b^i\}_{i=0}^N$ be a grid on $[0,1]$ with $b^0 = 0$ and $b^N = 1$, and let $\{\theta^j\}_{j=1}^M$ be a grid on Θ. Precompute $F(b^i) = \tau_0(b^0, b^i)$ by numerical quadrature in Eq. (3), then obtain $\tau_0(b^i, b^k) = F(b^k) - F(b^i)$ and $D_{\theta^j}(b^i, b^k)$ from Eq. (6). This makes the evaluation of Proposition 1 constant time per grid point.

B. Algorithm

The search complexity is $O(MN^2)$ with simple constant-time evaluators once F is tabulated. Coarse grids yield informative robust thresholds, and local refinement around the maximizer can tighten accuracy.

VII. Application

We illustrate the model on a screening problem motivated by post-exposure tuberculosis infection (TBI) testing among close contacts. A practical feature of contact investigation is that immunologic conversion is delayed, so retesting after an initially negative result is clinically meaningful, while effective exposure intensity is heterogeneous across outbreaks and settings. This combination maps naturally into (i) a monotone prevalence drift and (ii) multiplicative ambiguity in the speed of that drift.

A. Interpretation and Time Scale

We interpret $b(t)$ as the probability of TBI among identified close contacts at time t since an index exposure, conditional on the information available. A *test* is a TBI assay episode (screening plus confirmatory steps), and *treatment* is initiation of preventive therapy. The ceiling $\bar{b}$ captures the fact that not all exposed contacts become infected. Time is measured in months in the numerical illustration, consistent with the window-period logic in contact investigation and the operational cadence of follow-up.[3]

[3] If a different horizon is of interest (e.g., longer-term cohort screening), only the calibration changes; the structure of the policy and the relative robustness criterion are unchanged.

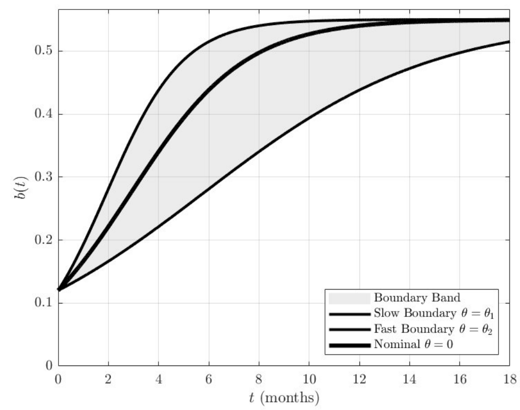

Fig. 1. Belief dynamics under speed ambiguity: $\dot{b} = (1 + \theta)\eta b(\bar{b} - b)$.

TABLE I
Illustrative Calibration for Post-Exposure TBI Screening (Time in Months)

Parameter	Value	Interpretation
b_0	0.120	Initial belief at baseline
$\bar{b}$	0.550	Ceiling belief
$f(b)$	$\eta b(\bar{b} - b)$	Logistic drift (Example 1)
η	0.800	Nominal speed (per month)
r	0.080	Discount rate (per month)
Θ	$[-0.5, 0.5]$	Speed ambiguity
q_1	0.750	Sensitivity of TBI assay (representative)
q_2	0.800	Specificity of TBI assay (representative)
C	150	Test cost
B	3000	Benefit if infected and treated
H	1200	Harm if uninfected and treated

B. Calibration

Table I reports the primitives used in our numerical implementation. The belief drift is logistic (cf. Example 1) and the ambiguity set is $\Theta = [-0.5, 0.5]$, so the speed $(1 + \theta)$ ranges from one-half to one-and-a-half of the nominal intensity; see Figure 1.

Remark 3 (Calibration Rationale). The calibration is intended to match a short-horizon screening environment in which retesting is plausible and intensity misspecification is material. The delayed detectability (window period) and retesting logic in contact investigation support the timing interpretation [22]–[24]. Values for test performance and costs are chosen to be representative of TBI assay episodes used in economic evaluations [25], while (B, H) are scaled in net-monetary (or monetized net-health) units, noting that the policy thresholds are homogeneous in the scaling of (B, H, C) [26].

C. Heuristic Policies and Robust Policy

To separate the effect of ambiguity from the effect of optimization, we compare four stationary threshold rules.

Definition 3 (Three Ex-Post Heuristics). *For $\theta \in \{\theta_1, 0, \theta_2\}$, let $(\beta^*(\theta), \gamma^*(\theta))$ be a maximizer of $V(b_0; \beta, \gamma|\theta)$ over $\mathcal{D}$. The* slow-optimal*,* nominal-optimal*, and* fast-optimal *heuristics are the policies $(\beta^*(\theta_1), \gamma^*(\theta_1))$, $(\beta^*(0), \gamma^*(0))$, and $(\beta^*(\theta_2), \gamma^*(\theta_2))$, respectively.*

The relatively robust policy $(\hat{\beta}, \hat{\gamma})$ maximizes $\rho(\beta, \gamma)$ over $\mathcal{D}$, with ρ defined in Definition 2. For the calibration in Table I, the grid search yields

$$\beta^*(\theta_1) = 0.390, \quad \beta^*(0) = 0.442, \quad \beta^*(\theta_2) = 0.467,$$

$$\hat{\beta} = 0.427, \quad \hat{\gamma} = 0.495, \quad \rho(\hat{\beta}, \hat{\gamma}) = 0.977,$$

and the corresponding canonical reporting of γ respects feasibility and separation.[4] Figure 2 plots the performance ratio $\varphi(\beta, \gamma \mid \theta)$ for the relatively robust policy and the three ex-post heuristics, evaluated at the initial belief b_0 (the state at which the robust thresholds are computed).

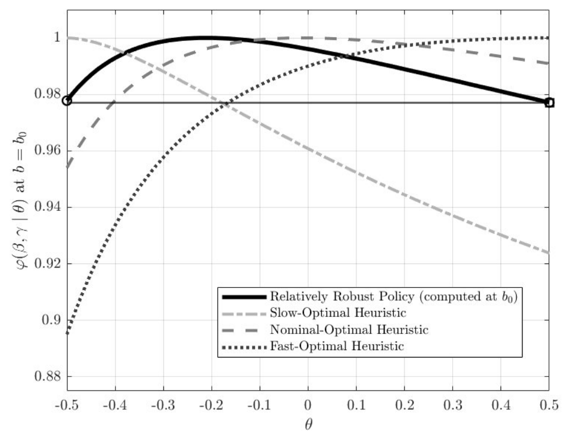

Fig. 2. Performance ratio as a function of $\theta \in \Theta$ for the relatively robust policy and the three ex-post heuristics, evaluated at the initial belief level $b_0 = 0.120$. The horizontal line indicates the robustness guarantee $\rho(\hat{\beta}, \hat{\gamma}) = \min_{\theta \in \Theta} \varphi(\hat{\beta}, \hat{\gamma} \mid \theta) = \varphi(\hat{\beta}, \hat{\gamma} \mid \theta_2) = 0.977$ for this calibration.

Moreover, Figure 3 shows typical belief trajectories under the robust policy with θ sampled in Θ. Each sample path evolves deterministically between test events and exhibits upward or downward belief jumps at test times. The shaded band highlights the test region $[\hat{\beta}, \hat{\gamma})$, and repeated negative outcomes generate visible retesting cycles, consistent with the window-period logic in short-horizon screening.

Remark 4 (Quasiconcavity). The boundary representation in Proposition 3 requires quasiconcavity of $\theta \mapsto \varphi(\beta, \gamma|\theta)$ on Θ. In the numerical illustration, quasiconcavity holds for the robust policy and for the three ex-post heuristics (cf. Figure 2). A proof can be obtained by establishing quasiconcavity of $\theta \mapsto V(b_0; \beta, \gamma|\theta)$ and $\theta \mapsto V^*(b_0|\theta)$ and combining these properties on Θ; operationally, quasiconcavity implies that the worst-case performance for a fixed policy occurs at θ_1 or θ_2, which is exactly what Figure 2 reveals.

[4]When $b_0 < \beta$, the value of a two-threshold policy depends on β and the continuation value at β; multiple γ values can be equivalent on a grid because $S(p(\beta))$ is constant in γ under separation. To avoid spurious "$\gamma = \bar{b}$" artifacts, we report a canonical feasible interior γ for each β, leaving the induced value unchanged at b_0.

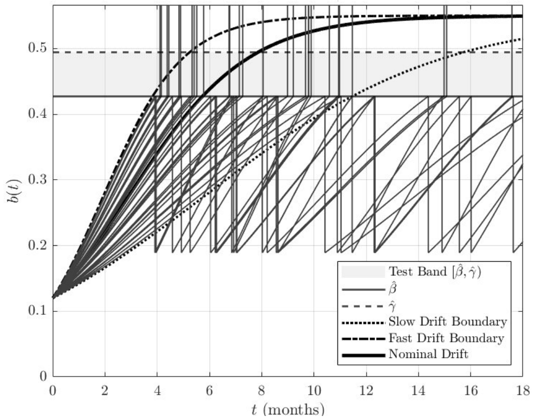

Fig. 3. Sample belief paths with update jumps (relatively robust policy; speed parameter $\theta \in \Theta$ sampled uniformly).

VIII. Discussion and Extensions

The framework treats belief drift as deterministic but speed-ambiguous, targeting intensity misspecification rather than shocks. Several extensions are natural. First, f may be derived from compartmental dynamics (with belief defined as an appropriate ratio), while ambiguity still enters through an intensity factor. Second, pooled or repeated evidence can replace the one-step Bayes updates (p, n). Third, primitives such as payoffs, test accuracy, drift (and even the ambiguity set Θ) may depend on patient features, yielding individualized thresholds and, under scarcity, triage. Capacity limits on tests or treatments can be modeled as constraints whose shadow prices produce higher "effective costs." Sampling or estimation bias can be handled by enlarging the ambiguity set (e.g., allowing the initial belief b_0 to vary in a plausible interval) and applying the same relative-robust benchmark.

A final motivation concerns outbreak settings (including fast-moving diseases after rescaling time units), where early data are thin and both baseline risk and effective intensity are poorly identified, naturally widening Θ and making robust tradeoffs operationally relevant.

IX. Conclusion

Screening decisions are about information acquisition and comparing the marginal costs and benefits of knowledge. Deterministic models clarify this tradeoff, yet they can embed parameter values whose stability is doubtful. We retain deterministic belief dynamics but allow ambiguity in their speed through a multiplicative factor. A two-threshold stationary policy then admits a closed-form value via a travel-time transform, keeping computation transparent. Relative robustness evaluates a policy by its value relative to the ex-post best policy under the realized speed and selects thresholds that maximize the worst-case performance ratio. The resulting rule is calibrated to perform well in both slow and fast regimes of prevalence accumulation, without presuming a probability model for which there is little defensible basis.

References

[1] K. J. Arrow, "Uncertainty and the welfare economics of medical care," *Amer. Econ. Rev.*, vol. 53, no. 5, pp. 941–973, Dec. 1963.

[2] H. Raiffa and R. Schlaifer, *Applied Statistical Decision Theory*. Cambridge, MA, USA: Harvard University Press, 1961.

[3] M. C. Weinstein, "Cost-effective priorities for cancer prevention," *Science*, vol. 221, no. 4605, pp. 17–23, Jul. 1983.

[4] K. P. Claxton and M. J. Sculpher, "Using value of information analysis to prioritise health research: Some lessons from recent UK experience," *PharmacoEconomics*, vol. 24, no. 11, pp. 1055–1068, Nov. 2006.

[5] Z. Philips, K. Claxton, and S. Palmer, "The half-life of truth: What are appropriate time horizons for research decisions?" *Med. Decis. Making*, vol. 28, no. 3, pp. 287–299, May–Jun. 2008.

[6] S. Eckermann and A. R. Willan, "The option value of delay in health technology assessment," *Med. Decis. Making*, vol. 28, no. 3, pp. 300–305, May–Jun. 2008.

[7] S. Eckermann, J. Karnon, and A. R. Willan, "The value of value of information," *PharmacoEconomics*, vol. 28, no. 9, pp. 699–709, 2010.

[8] L. E. Cipriano and T. A. Weber, "Population-level intervention and information collection in dynamic healthcare policy," *Health Care Manag. Sci.*, vol. 21, no. 4, pp. 604–631, Dec. 2018.

[9] A. J. Schaefer, M. D. Bailey, S. M. Shechter, and M. S. Roberts, "Modeling medical treatment using Markov decision processes," in *Operations Research and Health Care: A Handbook of Methods and Applications*, M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, Eds. Boston, MA, USA: Springer, 2005, pp. 593–612.

[10] O. Alagoz, H. Hsu, A. J. Schaefer, and M. S. Roberts, "Markov decision processes: A tool for sequential decision making under uncertainty," *Med. Decis. Making*, vol. 30, no. 4, pp. 474–483, Jul.–Aug. 2010.

[11] M. Hauskrecht and H. Fraser, "Planning treatment of ischemic heart disease with partially observable Markov decision processes," *Artif. Intell. Med.*, vol. 18, no. 3, pp. 221–244, Mar. 2000.

[12] E. Kırkızlar, D. M. Faissol, P. M. Griffin, and J. L. Swann, "Timing of testing and treatment for asymptomatic diseases," *Math. Biosci.*, vol. 226, no. 1, pp. 28–37, Jul. 2010.

[13] D. D. Sleator and R. E. Tarjan, "Amortized efficiency of list update and paging rules," *Commun. ACM*, vol. 28, no. 2, pp. 202–208, Feb. 1985.

[14] S. Ben-David and A. Borodin, "A new measure for the study of on-line algorithms," *Algorithmica*, vol. 11, no. 1, pp. 73–91, Jan. 1994.

[15] T. A. Weber, "Relatively robust decisions," *Theory Decis.*, vol. 94, no. 1, pp. 35–62, Jan. 2023.

[16] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proc. Roy. Soc. Lond. A*, vol. 115, no. 772, pp. 700–721, Aug. 1927.

[17] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM Rev.*, vol. 42, no. 4, pp. 599–653, Dec. 2000.

[18] M. J. Keeling and P. Rohani, *Modeling Infectious Diseases in Humans and Animals*. Princeton, NJ, USA: Princeton University Press, 2008.

[19] R. M. Anderson and R. M. May, *Infectious Diseases of Humans: Dynamics and Control*. Oxford, UK: Oxford University Press, 1991.

[20] F. Brauer, C. Castillo-Chavez, and Z. Feng, *Mathematical Models in Epidemiology*. New York, NY, USA: Springer, 2019.

[21] P.-F. Verhulst, "Notice sur la loi que la population suit dans son accroissement," *Correspondance Mathématique et Physique*, vol. 10, pp. 113–121, 1838.

[22] Centers for Disease Control and Prevention, "Guidelines for the investigation of contacts of persons with infectious tuberculosis: Recommendations from the National Tuberculosis Controllers Association and CDC," *MMWR Recomm. Rep.*, vol. 54, no. RR-15, pp. 1–47, 2005.

[23] T. R. Sterling, G. Njie, D. Zenner, *et al.*, "Guidelines for the treatment of latent tuberculosis infection: Recommendations from the National Tuberculosis Controllers Association and CDC, 2020," *MMWR Recomm. Rep.*, vol. 69, no. RR-1, pp. 1–11, 2020.

[24] World Health Organization, *WHO Consolidated Guidelines on Tuberculosis: Module 1 – Prevention: Tuberculosis Preventive Treatment*. Geneva, Switzerland: World Health Organization, 2020.

[25] L. Goscé, K. Allel, Y. Hamada, *et al.*, "Economic evaluation of novel *Mycobacterium tuberculosis* specific antigen-based skin tests for detection of TB infection: A modelling study," *PLOS Glob. Public Health*, vol. 3, no. 12, p. e0002573, Dec. 2023.

[26] J. S. Pliskin, D. S. Shepard, and M. C. Weinstein, "Utility functions for life years and health status," *Oper. Res.*, vol. 28, no. 1, pp. 206–224, Jan.–Feb. 1980.

Appendix A: Direct Shape Results for Performance Ratios under Speed Ambiguity

1. Purpose and Standing Assumptions

This appendix provides a direct route to shape restrictions in the speed parameter θ for the performance ratio

$$\varphi(\beta,\gamma|\theta) = \frac{V(b_0;\beta,\gamma|\theta)}{V^*(b_0|\theta)},$$

tailored to the screening application in Sec. VII, where (i) belief drift is logistic, (ii) speed ambiguity is $\Theta = [-0.5, 0.5]$, and (iii) the relevant policies satisfy $b_0 < \beta$ (Table I has $b_0 = 0.12$ and the computed optimal/robust thresholds satisfy $\beta \approx 0.39$–0.47).

The goal is a simple sufficient condition that can be checked quickly and that holds in the numerical illustration. Under separation, θ enters the policy value through two primitives only: (i) discounting along deterministic belief trajectories and (ii) a single retesting factor that captures the discounted likelihood of returning to the test threshold after consecutive negative outcomes. Bounding that retesting factor yields curvature control and, in the present calibration, delivers the single-peaked performance-ratio shapes observed in Fig. 2 (evaluated at the initial belief b_0).

All standing assumptions introduced in the main part are maintained. In particular, belief evolves as $\dot{b} = (1+\theta)f(b)$ with $\theta \in \Theta \subset (-1,1)$ and discount rate $r > 0$. The test is binary with $(q_1,q_2) \in (1/2,1)^2$ and cost $C \geq 0$,

$$\ell(b) = q_1 b + (1-q_2)(1-b), \quad p(b) = \frac{\kappa b}{1+(\kappa-1)b},$$

$$n(b) = \frac{(1/\nu)b}{1+((1/\nu)-1)b},$$

where $\kappa = q_1/(1-q_2) > 1$ and $\nu = q_2/(1-q_1) > 1$. Treatment payoff is $S(b) = \max\{0,(B+H)b-H\}$ with $\bar{\alpha} = H/(B+H)$.

2. One-Threshold Reduction under Separation

The separation condition in the main part is

$$n(\gamma) \leq \beta \leq \gamma \leq p(\beta).$$

Under separation and $\gamma \geq \bar{\alpha}$, the continuation value at β depends on β but not on γ, because any positive test at β jumps to belief $p(\beta)$, which is already in the treat region.

Lemma A.1 (Value Depends on β Only under Separation). *Fix $\beta \in [0,1]$. Suppose γ satisfies $\beta \leq \gamma \leq p(\beta)$ and $\gamma \geq \bar{\alpha}$. Then for every $\theta \in \Theta$,*

$$V(b_0;\beta,\gamma|\theta) = V(b_0;\beta,p(\beta)|\theta).$$

In particular, within the feasible class $\mathcal{D}$, the map $(\beta,\gamma) \mapsto V(b_0;\beta,\gamma|\theta)$ is constant along the γ-interval $[\max\{\beta,\bar{\alpha}\},\, p(\beta)]$.

Proof. At belief $b = \beta$, a test is administered. Under separation, a positive result yields posterior $p(\beta) \geq \gamma$, so the policy treats immediately and the stopping payoff is $S(p(\beta))$, which is independent of γ. A negative result yields posterior $n(\beta) \leq \beta$, so the policy waits deterministically back to β and repeats. Thus, the fixed-point equation for the continuation value at β involves only β (through $\ell(\beta)$, $p(\beta)$, $n(\beta)$) and discounting along the path from $n(\beta)$ to β, but not γ. The remainder of the argument follows from Proposition 1. □

Remark A.1 (Canonical Choice of γ). Lemma A.1 formalizes the "without loss of generality" choice $\gamma = p(\beta)$ (or any other feasible interior γ) when attention is restricted to stationary two-threshold policies with separation. Accordingly, in the remainder of the appendix $V(b_0;\beta|\theta)$ denotes the value of the policy with threshold β and any feasible $\gamma \in [\max\{\beta,\bar{\alpha}\}, p(\beta)]$.

3. Canonical Decomposition and the Retesting Factor z

For $b < b'$, define the nominal travel time and the ambiguous discount factor

$$\tau_0(b,b') = \int_b^{b'} \frac{du}{f(u)}, \quad D_\theta(b,b') = \exp\left(-\frac{r\,\tau_0(b,b')}{1+\theta}\right).$$

Fix β and define the (policy-dependent) primitives

$$a_0(\beta) = r\,\tau_0(b_0,\beta), \quad a_-(\beta) = r\,\tau_0(n(\beta),\beta),$$

$$k(\beta) = 1-\ell(\beta) \in (0,1), \quad N(\beta) = \ell(\beta)S(p(\beta)) - C.$$

For the policies of interest in Sec. VII, $b_0 < \beta$ and $N(\beta) > 0$, so Proposition 1 yields the compact representation

$$V(b_0;\beta|\theta) = \frac{N(\beta)\exp\left(-\frac{a_0(\beta)}{1+\theta}\right)}{1-k(\beta)\exp\left(-\frac{a_-(\beta)}{1+\theta}\right)}. \quad \text{(A.1)}$$

The key object is the retesting factor

$$\begin{aligned} z(\beta,\theta) &= k(\beta)\exp\left(-\frac{a_-(\beta)}{1+\theta}\right) \\ &= (1-\ell(\beta))\,D_\theta(n(\beta),\beta) \in (0,1), \end{aligned} \quad \text{(A.2)}$$

which is the product of (i) the probability of a negative outcome at β and (ii) the discount factor incurred while waiting from $n(\beta)$ back to β. The associated retesting multiplier is $M(\beta,\theta) = 1/(1-z(\beta,\theta))$.

It is convenient to use the monotone change of variables

$$x = \frac{1}{1+\theta}, \quad x \in \mathcal{X} = [x_2, x_1], \quad x_i = \frac{1}{1+\theta_i}. \quad \text{(A.3)}$$

Then (A.1) becomes

$$V(b_0;\beta|x) = \frac{N(\beta)e^{-a_0(\beta)x}}{1-k(\beta)e^{-a_-(\beta)x}} = \frac{N(\beta)e^{-a_0(\beta)x}}{1-z(\beta,x)}, \quad \text{(A.4)}$$

where $z(\beta,x) = k(\beta)e^{-a_-(\beta)x}$.

4. Log-Derivatives and a Uniform Bound under $z \le 1/2$

The following result provides the derivative that allows for much of the direct comparison.

Lemma A.2 (Log-Derivative Representation). *Assume that $b_0 < \beta$ and $N(\beta) > 0$. For $x \in \mathcal{X}$, let V be defined by* (A.4) *and set $z(\beta, x) = k(\beta)e^{-a_-(\beta)x}$. Then*

$$\frac{\partial}{\partial x} \log V(b_0; \beta|x) = -a_0(\beta) - \frac{a_-(\beta)\, z(\beta, x)}{1 - z(\beta, x)}. \tag{A.5}$$

Moreover, $\partial_x \log V < 0$ and

$$-a_0(\beta) - \frac{a_-(\beta)\, z(\beta, x)}{1 - z(\beta, x)} \le -a_0(\beta). \tag{A.6}$$

If, in addition, $z(\beta, x) \le 1/2$, then

$$0 \le \frac{z(\beta, x)}{1 - z(\beta, x)} \le 1,$$

and therefore

$$-a_0(\beta) - a_-(\beta) \le \partial_x \log V(b_0; \beta|x) \le -a_0(\beta). \tag{A.7}$$

Proof. Take logs in (A.4):

$$\log V = \log N(\beta) - a_0(\beta)x - \log\big(1 - z(\beta, x)\big).$$

Differentiate with respect to x and use $\partial_x z(\beta, x) = -a_-(\beta)z(\beta, x)$ to obtain (A.5). The bound (A.7) follows from $z/(1 - z) \le 1$ when $z \le 1/2$. □

Remark A.2 (Interpretation of $z \le 1/2$). The condition $z(\beta, \theta) \le 1/2$ is a "half-life" restriction on retesting: it implies $M(\beta, \theta) = 1/(1-z(\beta, \theta)) \le 2$, so the option value of repeated negative outcomes cannot inflate the continuation value at β by more than a factor of two. In turn, the sensitivity of $\log V$ to the speed parameter is dominated by the deterministic waiting penalty $a_0(\beta)$, which is the structure behind single-peaked performance ratios in the tuberculosis calibration.

5. Single-Crossing Based on Decreasing Differences

Fix two thresholds β and $\tilde\beta$ (both with $b_0 < \min\{\beta, \tilde\beta\}$ and $N(\beta), N(\tilde\beta) > 0$) and define, for $x \in \mathcal{X}$,

$$\Delta(x; \beta, \tilde\beta) = \log V(b_0; \beta|x) - \log V(b_0; \tilde\beta|x). \tag{A.8}$$

Single-crossing of the ratio $V(b_0; \beta|x)/V(b_0; \tilde\beta|x)$ is equivalent to single-crossing of Δ.

A convenient sufficient condition is decreasing differences of $\log V$ in (β, x), i.e., $\partial^2 \log V/(\partial\beta\, \partial x) \le 0$ on the relevant domain.

Proposition A.1 (Decreasing Differences ⇒ Single-Crossing). *Let $\mathcal{B} \subset (0, \bar b)$ be an interval on which $b_0 < \beta$ and $N(\beta) > 0$ for all $\beta \in \mathcal{B}$. If $\log V(b_0; \beta|x)$ has decreasing differences on $\mathcal{B} \times \mathcal{X}$, i.e.,*

$$\frac{\partial^2}{\partial\beta\, \partial x} \log V(b_0; \beta|x) \le 0 \quad \text{for all } (\beta, x) \in \mathcal{B} \times \mathcal{X},$$

then for any $\beta, \tilde\beta \in \mathcal{B}$, the map $x \mapsto \Delta(x; \beta, \tilde\beta)$ is monotone on $\mathcal{X}$. Consequently, $x \mapsto V(b_0; \beta|x)/V(b_0; \tilde\beta|x)$ is single-crossing on $\mathcal{X}$, and $\theta \mapsto V(b_0; \beta|\theta)/V(b_0; \tilde\beta|\theta)$ is single-crossing on Θ.

Proof. If $\log V$ has decreasing differences, then $\partial_x \log V(b_0; \beta|x)$ is nonincreasing in β for each x. Hence, for $\beta > \tilde\beta$,

$$\Delta'(x; \beta, \tilde\beta) = \partial_x \log V(b_0; \beta|x) - \partial_x \log V(b_0; \tilde\beta|x) \le 0,$$

for all $x \in \mathcal{X}$, so $\Delta(\cdot; \beta, \tilde\beta)$ is monotone and crosses any level at most once. The equivalence between x and θ follows from the monotone transform $x = 1/(1 + \theta)$. □

6. Sufficient Check for Decreasing Differences when $z \le 1/2$

Lemma A.3 (Half-Life Slope Dominance Bound). *Assume $b_0 < \beta$ and $N(\beta) > 0$. Suppose $a_0(\beta)$ and $a_-(\beta)$ are differentiable in β and define $z(\beta, x) = k(\beta)e^{-a_-(\beta)x}$. Then*

$$\frac{\partial^2 \log V(b_0;\beta|x)}{\partial\beta\, \partial x} = -a_0'(\beta) - \frac{a_-'(\beta)\, z(\beta,x)}{1-z(\beta,x)} - \frac{a_-(\beta)\, z_\beta(\beta,x)}{(1-z(\beta,x))^2}, \tag{A.9}$$

where

$$z_\beta(\beta, x) = e^{-a_-(\beta)x}\Big(k'(\beta) - k(\beta)a_-'(\beta)x\Big).$$

If, on an interval $\mathcal{B}$, the half-life bound $z(\beta, x) \le 1/2$ holds for all $(\beta, x) \in \mathcal{B} \times \mathcal{X}$ and

$$a_0'(\beta) \;\ge\; 4\, a_-(\beta) \sup_{x\in\mathcal{X}} \big|z_\beta(\beta, x)\big| \quad \text{for all } \beta \in \mathcal{B}, \tag{A.10}$$

then $\log V$ has decreasing differences on $\mathcal{B} \times \mathcal{X}$.

Proof. Equation (A.9) follows by differentiating (A.5) with respect to β. If $z \le 1/2$, then $(1 - z)^{-2} \le 4$, and therefore

$$-\frac{a_-(\beta)\, z_\beta(\beta, x)}{(1 - z(\beta, x))^2} \le 4\, a_-(\beta)\, \big|z_\beta(\beta, x)\big|.$$

Moreover, the middle term in (A.9) is ≤ 0 when $a_-'(\beta) \ge 0$ (as in the logistic verification below). Under (A.10), the right-hand side of (A.9) is ≤ 0 on $\mathcal{B} \times \mathcal{X}$. □

Remark A.3 (Simplicity of Half-Life Bound). Condition (A.10) compares the deterministic waiting slope $a_0'(\beta) = r/f(\beta)$ with the maximal sensitivity of $z(\beta, x)$. The half-life bound $z \le 1/2$ turns $(1 - z)^{-2}$ into the constant 4.

7. Implication for Quasiconcavity in the TBI Calibration

In the tuberculosis calibration of Sec. VII, the computed optimal thresholds satisfy

$$\beta^*(\theta_1) = 0.390,\; \beta^*(0) = 0.442,\; \beta^*(\theta_2) = 0.467,\; \hat\beta = 0.427,$$

so the testing thresholds β of the policies in Fig. 2 lie in

$$\mathcal{B}_{\mathrm{TBI}} = [0.390,\, 0.467].$$

On this interval, verification is minimal: (i) the half-life bound can be checked at $(\beta, \theta) = (0.390, 0.5)$ because $\beta \mapsto z(\beta, \theta)$ decreases while $\theta \mapsto z(\beta, \theta)$ increases; and (ii) the slope dominance check (A.10) is comfortably satisfied for Table I.

Combining Lemma A.3 with Proposition A.1 yields single-crossing of ratios $V(b_0; \beta|\theta)/V(b_0; \tilde\beta|\theta)$ for $\beta, \tilde\beta \in \mathcal{B}_{\mathrm{TBI}}$, supporting the boundary logic behind Proposition 3 for the policies used in the numerical illustration.

8. Logistic Drift Specialization for Table I

Under logistic drift $f(b) = \eta b(\bar{b} - b)$, define odds $\Omega(b) = b/(\bar{b} - b)$ and $\hat{r} = r/(\eta\bar{b})$. Then

$$\tau_0(b, b') = \frac{1}{\eta\bar{b}} \ln\left(\frac{\Omega(b')}{\Omega(b)}\right), \quad D_\theta(b, b') = \left(\frac{\Omega(b)}{\Omega(b')}\right)^{\hat{r}/(1+\theta)}.$$

Hence, for $b_0 < \beta$,

$$a_0(\beta) = \hat{r} \ln\left(\frac{\Omega(\beta)}{\Omega(b_0)}\right), \quad a_-(\beta) = \hat{r} \ln\left(\frac{\Omega(\beta)}{\Omega(n(\beta))}\right),$$

and

$$z(\beta, \theta) = k(\beta) \left(\frac{\Omega(n(\beta))}{\Omega(\beta)}\right)^{\hat{r}/(1+\theta)}.$$

Moreover, $a_0'(\beta) = r/f(\beta) = r/(\eta\beta(\bar{b} - \beta))$ is explicit and strictly positive on $(0, \bar{b})$, and

$$a_-'(\beta) = r\left(\frac{1}{f(\beta)} - \frac{n'(\beta)}{f(n(\beta))}\right),$$

$$n'(\beta) = \frac{1/\nu}{\left(1 + ((1/\nu) - 1)\beta\right)^2} \in (0, 1),$$

so $a_-'(\beta) > 0$ holds on the empirically relevant range. Together with $k'(\beta) = -\ell'(\beta) < 0$, this yields the monotonicity claims used in Subsec. 7 and reduces the half-life verification to an endpoint check.

Appendix B: Indirect Methods

We sketch an alternative route to quasiconcavity of $\varphi(\beta, \gamma|\cdot)$ based on total positivity (TP), as developed by Karlin [B1].

1. Mixture Representation for V and a Change of Variables

We begin with a series representation for the value of a fixed two-threshold policy $u = (\beta, \gamma)$ under the time-scaling structure.

Lemma B.1 (Discrete-Mixture Representation). *Fix a feasible two-threshold policy $u = (\beta, \gamma) \in \mathcal{D}$ with $\gamma \geq \bar{\alpha}$. Assume $b_0 < \beta$ and define*

$$a_0(u) = r\,\tau_0(b_0, \beta), \quad a_-(u) = r\,\tau_0(n(\beta), \beta),$$

$$k(u) = 1 - \ell(\beta) \in (0, 1), \quad N(u) = \ell(\beta)S(p(\beta)) - C.$$

Then for every $\theta \in \Theta$,

$$V(b_0; u|\theta) = N(u) \sum_{m=0}^{\infty} k(u)^m \exp\left(-\frac{a_0(u) + m a_-(u)}{1 + \theta}\right). \tag{B.1}$$

For $\beta \leq b_0 < \gamma$, $V(b_0; u|\theta)$ equals a θ-independent term $\ell(b_0)S(p(b_0)) - C$ plus the same type of series with $a_0(u)$ replaced by $r\tau_0(n(b_0), \beta)$. For $b_0 \geq \gamma$, $V(b_0; u|\theta) = S(b_0)$ is θ-independent.

Proof. For $b_0 < \beta$, Proposition 1 gives

$$V(b_0; u|\theta) = \frac{N(u)\, D_\theta(b_0, \beta)}{1 - k(u)\, D_\theta(n(\beta), \beta)}.$$

Since $0 \leq k(u) D_\theta(n(\beta), \beta) < 1$ on Θ,

$$\frac{1}{1 - k(u) D_\theta(n(\beta), \beta)} = \sum_{m=0}^{\infty} k(u)^m D_\theta(n(\beta), \beta)^m.$$

Using $D_\theta(b, b') = \exp(-(r\tau_0(b, b'))/(1 + \theta))$ and collecting exponents yields (B.1). The other cases follow from Proposition 1. □

Lemma B.1 expresses V as a positive combination of kernels $\exp(-a/(1 + \theta))$. Introduce the change of variables

$$x = \frac{1}{1 + \theta}, \quad \theta \in \Theta, \tag{B.2}$$

so $x \in \mathcal{X} = [x_2, x_1]$ with $x_i = 1/(1 + \theta_i)$ and $\exp(-a/(1 + \theta)) = \exp(-ax)$. For $b_0 < \beta$,

$$V(b_0; u|x) = \sum_{m=0}^{\infty} \alpha_m(u) \exp(-s_m(u)x), \tag{B.3}$$

with

$$\alpha_m(u) = N(u)\, k(u)^m, \quad s_m(u) = a_0(u) + m a_-(u). \tag{B.4}$$

Remark B.1 (Interpretation). The index m counts the number of negative retests at β before the first positive outcome. The weights $\alpha_m(u)$ have a geometric tail, while $s_m(u)$ is the associated effective discounted time.

Example B.1 (Logistic Drift). Under logistic drift $f(b) = \eta b(\bar{b} - b)$ (cf. Example 1), define odds $\Omega(b) = b/(\bar{b} - b)$ and $\hat{r} = r/(\eta\bar{b})$. Then

$$\tau_0(b, b') = \frac{1}{\eta\bar{b}} \ln\left(\frac{\Omega(b')}{\Omega(b)}\right),$$

and

$$s_m(u) = \hat{r}\left(\ln\frac{\Omega(\beta)}{\Omega(b_0)} + m \ln\frac{\Omega(\beta)}{\Omega(n(\beta))}\right).$$

Moreover, (B.3) sums to

$$V(b_0; u|x) = \frac{N(u)\left(\frac{\Omega(b_0)}{\Omega(\beta)}\right)^{\hat{r}x}}{1 - k(u)\left(\frac{\Omega(n(\beta))}{\Omega(\beta)}\right)^{\hat{r}x}},$$

consistent with Remark 2.

2. Pairwise Interval Property and Quasiconcavity of $\varphi(\beta, \gamma|\cdot)$

A function $g : \Theta \to \mathbb{R}$ is quasiconcave if all upper contour sets are intervals. Since $\Theta \subset \mathbb{R}$, this means: for every $\lambda \in \mathbb{R}$, the set $\{\theta \in \Theta :\ g(\theta) \geq \lambda\}$ is an interval (possibly empty).

Definition B.1 (Pairwise Interval Property). *Let $\{G_v : \Theta \to \mathbb{R}\}_{v\in\mathcal{V}}$ be a family of functions and fix $u \in \mathcal{V}$. We say that G_u has the pairwise interval property on Θ relative to $\{G_v\}$ if for every $v \in \mathcal{V}$ and every $\lambda \in \mathbb{R}$, the set*

$$\left\{\theta \in \Theta :\ G_u(\theta) \geq \lambda G_v(\theta)\right\}$$

is an interval (possibly empty).

Lemma B.2 (Pairwise Interval Property ⇒ Quasiconcavity). *Let $\{G_v\}_{v\in\mathcal{V}}$ be nonnegative functions on Θ, and define*

$$G^*(\theta) = \sup_{v\in\mathcal{V}} G_v(\theta), \quad \theta \in \Theta.$$

Fix $u \in \mathcal{V}$ and assume $G^*(\theta) > 0$ on Θ. If G_u has the pairwise interval property on Θ relative to $\{G_v\}$, then $\theta \mapsto G_u(\theta)/G^*(\theta)$ is quasiconcave on Θ.

Proof. Fix $\lambda \in \mathbb{R}$. Since $G^*(\theta) = \sup_v G_v(\theta)$,

$$\Big\{\theta \in \Theta : \tfrac{G_u(\theta)}{G^*(\theta)} \geq \lambda\Big\} = \bigcap_{v\in\mathcal{V}} \Big\{\theta \in \Theta : G_u(\theta) \geq \lambda G_v(\theta)\Big\}.$$

Each set in the intersection is an interval by Definition B.1, hence the intersection is an interval. □

Remark B.2 (On "Only Comparing to the Optimum"). If there exists $v^* \in \mathcal{D}$ such that $V^*(b_0|\theta) = V(b_0; v^*|\theta)$ for all $\theta \in \Theta$, then quasiconcavity of $\varphi(u|\cdot)$ reduces to the pairwise interval property of $\theta \mapsto V(b_0; u|\theta)/V(b_0; v^*|\theta)$. In general, the maximizer may depend on θ, and it suffices to check the property against policies that are optimal for some $\theta \in \Theta$.

3. Total Positivity, Single-Crossing, and TP Criterion

Under the transformation (B.2), the basic kernel is

$$K(s, x) = \exp(-sx), \quad s, x > 0,$$

which is TP_2 on $\mathbb{R}_+ \times \mathbb{R}_+$ (Karlin [B1]). This implies that monotone likelihood-ratio orderings of mixing weights yield single-crossing for ratios of the corresponding mixtures.

Proposition B.1 (TP_2 and MLR ⇒ Pairwise Interval Property). *Fix two policies u and $\tilde{u}$ and write their values in the x-variable as*

$$G_u(x) = \sum_{m\geq 0} \alpha_m \exp(-s_m x),\ G_{\tilde{u}}(x) = \sum_{m\geq 0} \tilde{\alpha}_m \exp(-s_m x),$$

for all $x \in \mathcal{X}$, on a common nondecreasing support $\{s_m\}_{m\geq 0}$ with nonnegative weights $\{\alpha_m\}_{m\geq 0}$ and $\{\tilde{\alpha}_m\}_{m\geq 0}$. If the likelihood-ratio sequence $\alpha_m/\tilde{\alpha}_m$ is monotone in m (either nondecreasing or nonincreasing), then $x \mapsto G_u(x)/G_{\tilde{u}}(x)$ is single-crossing on $\mathcal{X}$. Hence, for every $\lambda \in \mathbb{R}$, the set

$$\Big\{x \in \mathcal{X} : G_u(x) \geq \lambda G_{\tilde{u}}(x)\Big\}$$

is an interval, and equivalently $\{\theta \in \Theta : G_u(\theta) \geq \lambda G_{\tilde{u}}(\theta)\}$ is an interval.

Remark B.3 (Role of the Criterion). Proposition B.1 links the mixture representation to the pairwise interval property, and with Lemma B.2 yields quasiconcavity of $\theta \mapsto \varphi(\beta, \gamma|\theta)$ on Θ.

4. How the Sequences Arise in the Two-Threshold Model

We now turn to the question of how the sequences $\{s_m\}$ and $\{\alpha_m\}$ arise in the present model.

4.a. Support Points as Discounted Retest-Cycle Times: For $b_0 < \beta$, Lemma B.1 yields $s_m(u) = a_0(u) + ma_-(u)$, $a_0(u) = r\tau_0(b_0, \beta)$, and $a_-(u) = r\tau_0(n(\beta), \beta)$. Thus, $s_m(u)$ is discounted time to the first test at β plus m discounted retest-cycle times.

4.b. Weights as a Geometric Retesting Law: The weights are $\alpha_m(u) = N(u)k(u)^m$, $N(u) = \ell(\beta)S(p(\beta)) - C$, and $k(u) = 1 - \ell(\beta)$, so $k(u)^m$ is the probability of m consecutive negatives at β before the first positive outcome.

4.c. Canonical Common Support via Truncation and Union: Different policies generally generate different supports $\{a_0 + ma_-\}_{m\geq 0}$. For a truncation level M, define

$$\mathcal{S}_M(u) = \Big\{a_0(u) + ma_-(u) : m \in \{0, 1, \ldots, M\}\Big\}.$$

For two policies u and $\tilde{u}$, let $\{s_m\}_{m=0}^L$ be the sorted list of distinct points in $\mathcal{S}_M(u) \cup \mathcal{S}_M(\tilde{u})$. Define weights on this common support by aggregation,

$$\alpha(s) = \sum_{j:\ a_0(u)+ja_-(u)=s} N(u)k(u)^j,$$

$$\tilde{\alpha}(s) = \sum_{j:\ a_0(\tilde{u})+ja_-(\tilde{u})=s} N(\tilde{u})k(\tilde{u})^j,$$

and relabel $(s, \alpha(s), \tilde{\alpha}(s))$ as $(s_m, \alpha_m, \tilde{\alpha}_m)_{m=0}^L$.

5. MLR Ordering: Plausibility

The MLR condition in Proposition B.1 is sufficient and is natural when one policy induces stochastically more retesting than another.

5.a. Geometric Tails and MLR Characterization of Support Alignment: If two policies share the same (a_0, a_-), then

$$\frac{\alpha_m}{\tilde{\alpha}_m} = \frac{N}{\tilde{N}}\Big(\frac{k}{\tilde{k}}\Big)^m,$$

so $\alpha_m/\tilde{\alpha}_m$ is monotone in m whenever $k \geq \tilde{k}$ or $k \leq \tilde{k}$.

5.b. Ordering of k across Policies: Since $\ell(b)$ is affine and increasing, $k(\beta) = 1 - \ell(\beta)$ is decreasing in β. Thus, smaller β yields larger k and a heavier geometric tail in retesting.

5.c. Support Differences and Nearly Aligned Mixtures: When (a_0, a_-) differ, the supports do not coincide, but longer effective times arise from larger k, larger a_-, or both. These shifts are the channel through which MLR-type comparisons tend to hold on a common sorted support.

5.d. Operational Check: For fixed $u, \tilde{u}$, the truncation-and-union construction yields finite sequences $(s_m, \alpha_m, \tilde{\alpha}_m)_{m=0}^L$. One checks MLR by verifying

$$\frac{\alpha_{m+1}}{\tilde{\alpha}_{m+1}} \geq \frac{\alpha_m}{\tilde{\alpha}_m} \quad \text{for all } m \in \{0, 1, \ldots, L-1\},$$

or the reverse inequality, after removing indices with $\tilde{\alpha}_m = 0$. With M large enough that k^M is negligible, this provides a practical sufficient condition.

References

[B1] S. Karlin, *Total Positivity*, Stanford, CA, USA: Stanford University Press, 1968.