GAP-FREE ANYTIME-VALID COMPUTERIZED CLASSIFICATION TESTING UNDER THE 2PL MODEL

Anonymous authors
Paper under double-blind review

ABSTRACT

Computerized classification tests decide on which side of a cutoff an examinee lies. Classical sequential probability ratio tests calibrate errors at two separated abilities, leaving behavior inside a user-chosen indifference region unresolved. Under a known two-parameter logistic response model, we develop gap-free e-process procedures—requiring no user-chosen separated ability design points—whose finite-sample directional error control permits any predictable item policy. A shrinking-reference rule has local mean order $O(\delta^{-2}\log^2(1/\delta))$. Our main rule mixes over dyadic discrimination-standardized alternatives and attains the optimal worst-scale order $O(\delta^{-2}\log\log(1/\delta))$ in an ideal experiment with independent, identically parameterized cutoff-centered items. Known-truth simulations and logged-response comparisons show that the mixture substantially reduces exposure relative to shrinking. The logged study is observational and does not validate latent truth, item calibration, or counterfactual adaptive policies.

1 INTRODUCTION

Computerized adaptive testing (CAT) selects each item in response to an examinee's previous answers. By concentrating measurement where it is most informative, CAT can reduce burden, return decisions sooner, and make better use of a calibrated item bank than a common fixed form (Weiss & Kingsbury, 1984; Chang, 2015). A recent machine-learning survey organizes modern CAT around measurement, question selection, bank construction, and test control, and documents its growing use for evaluating both people and AI systems (Zhuang et al., 2026). This adaptivity is consequential in education, credentialing, placement, online learning, and model evaluation, where every additional item consumes time or computation.

Most CAT methodology targets a precise estimate of a latent ability. Yet many operational questions are discrete: has a learner mastered a skill, should an applicant pass, or which of two interventions is appropriate? When only the side of a policy cutoff matters, estimating ability more accurately than needed can waste items. Computerized classification testing (CCT) instead combines adaptive item selection with a stopping rule that terminates as soon as there is sufficient evidence for one category (Reckase, 1983; Thompson, 2009a).

The dominant frequentist construction places two design points on opposite sides of the cutoff and applies Wald's sequential probability ratio test (SPRT) (Wald, 1945; Reckase, 1983). For two simple hypotheses, the SPRT is optimal in expected sample size among tests with matched errors (Wald & Wolfowitz, 1948). Its likelihood-ratio process is also valid under optional stopping. The difficulty is therefore not that the SPRT itself fails under continuous monitoring. Rather, its CCT formulation requires an *indifference region*: two separated ability values at which errors are calibrated. Performance and even the interpretation of error control for abilities between those points depend on the chosen gap. This is undesirable when policy specifies one cutoff and no scientifically meaningful gray zone. The affected population is also unknown when that gap is chosen. In our logged benchmark, 7.5% of confidently reference-classified students still fall inside the SPRT design interval of half-width 0.5; the fraction is 32.4% before reference-confidence filtering. A seemingly moderate gap can therefore leave a non-negligible, prospectively unknown group without a pointwise error interpretation. Our known-truth experiment illustrates the same distinction: separated-design likelihood rules are fast but can cross in the wrong direction near the cutoff. They are not miscalibrated for their own targets; their targets simply differ from the touching-half-line guarantee we seek.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Modern anytime-valid inference gives a language for eliminating that gap. Nonnegative supermartingales—often called e-processes—can be monitored after every observation, under adaptive sampling, while retaining finite-sample error control (Vovk & Wang, 2021; Ramdas et al., 2023). These tools now support continuous A/B testing (Johari et al., 2022), bandit-style adaptive data collection (Karampatziakis et al., 2021), and time-uniform confidence sequences (Howard et al., 2021). The same principle is particularly natural for CCT because item choice can be based only on past responses.

We begin with a natural anytime-valid CCT for the two-parameter logistic (2PL) item-response model (Birnbaum, 1968): a discrimination-standardized reference on each side shrinks toward the cutoff. This construction exposes the cost of committing to one time-dependent scale. Our main procedure removes that commitment by mixing over a countable dyadic family of standardized alternatives. By *gap-free* we mean that neither procedure requires a user-chosen pair of ability design points; we do not mean that every examinee receives a finite-time decision. Our contributions are:

- two implementable gap-free CCTs—a shrinking-reference rule and a multiscale mixture rule—with finite-sample, time-uniform directional error control under any predictable item policy;
- an explicit high-probability stopping bound for the shrinking rule and an optimal worst-scale local mean order for the mixture in the ideal i.i.d. centered-item model, improving by one logarithmic factor over the shrinking baseline;
- a common-path, known-truth 2PL comparison with five representative CCT stopping rules, separating early crossings from terminal fallbacks and exposing the efficiency–calibration tradeoff near the cutoff; and
- a reproducible, student-disjoint logged-response evaluation on three educational datasets, including dataset-, direction-, and common-horizon analyses against shrinking, SPRT, and fixed-length rules.

Within each experiment, every method receives the same item sequence and response path, and evidence-threshold crossings are reported separately from terminal fallbacks. The known-truth benchmark generates responses from a self-contained heterogeneous 2PL bank. The logged experiments instead use recorded responses: a 75–25% student split separates item calibration from evaluation, and a 60–40% within-student split separates sequential evidence from the reference assessment. No logged evaluation response is simulated from a fitted model.

The closest precedents use sequential confidence intervals for adaptive mastery testing (Chang, 2005), generalized likelihood ratios under predictable CAT designs (Bartroff et al., 2008), or likelihoods integrated over both sides of the cutoff (Thompson & Ro, 2007). Unlike these methods, we establish finite-sample, time-uniform, pointwise directional error control over each ability half-line under any predictable item policy. Appendix A gives the detailed comparison and reviews the broader CCT and anytime-valid literature. The generic implication from a boundary likelihood ratio and monotone likelihood ratios to a one-sided e-process is not new. Our technical contribution is its predictable 2PL specialization together with the discrimination-standardized shrinking analysis, the computable multiscale construction and its local adaptation rate, and the explicit separation of certified and fallback behavior in CCT experiments.

These claims have a deliberate hierarchy. The error guarantees are finite-sample statements conditional on the specified 2PL response functions and known item parameters. The synthetic study checks finite-grid operating characteristics when that model is true. The logged study reports agreement with an estimated, response-disjoint reference label under observed exposure; it neither observes latent truth nor tests the assumptions behind the guarantees.

2 Problem Setup

We now formalize the adaptive 2PL setting and the one-sided guarantees that the two procedures must satisfy. This common setup will let us distinguish the choice of monitored evidence process from the choice of item-selection policy.

An item j has discrimination $a_j > 0$ and difficulty b_j. For latent ability $\theta \in \mathbb{R}$, the 2PL response model is $Y_t \mid \mathcal{F}_{t-1} \sim \mathrm{Bern}(p_{j_t}(\theta))$, where $p_j(\theta) = \sigma(a_j(\theta - b_j))$ and $\sigma(x) = (1 + e^{-x})^{-1}$. Here $\mathcal{F}_{t-1} = \sigma(j_1, Y_1, \ldots, j_{t-1}, Y_{t-1})$, and the next item j_t is $\mathcal{F}_{t-1}$-measurable. In plain words, j_t may use anything observed before time t—the earlier items and responses—but it cannot use the current

response Y_t or any future response. We call an item-choice rule with this property *predictable*. Responses are conditionally independent given θ and the selected items.

For a cutoff θ_c, let $H_0 : \theta < \theta_c$ ("fail") and $H_1 : \theta > \theta_c$ ("pass") denote the two *decision labels*. A test has a possibly infinite stopping time τ. On $\{\tau < \infty\}$ it returns $D \in \{H_0, H_1\}$; otherwise it abstains. We require

$$\sup_{\theta \geq \theta_c} \Pr_\theta(D = H_0) \leq \alpha_0, \qquad \sup_{\theta \leq \theta_c} \Pr_\theta(D = H_1) \leq \alpha_1. \tag{1}$$

Thus the closed null region for issuing H_0 is $[\theta_c, \infty)$, and the closed null region for issuing H_1 is $(-\infty, \theta_c]$. The symbols H_0, H_1 name outputs rather than a conventional disjoint null–alternative partition. Here D denotes only a decision certified by the test's evidence boundary: the monitored evidence for one side has reached its prescribed threshold. The constructions below specify the corresponding evidence processes—for example, $d_{\beta,t}$ for the shrinking rule and $M_{\beta,t}$ for the mixture—and their thresholds. If a finite bank is exhausted and an application requires a decision, we denote the separate terminal fallback by $\widetilde{D}$; its hybrid error is not covered by (1). The null regions overlap at the boundary. Consequently, a rule satisfying (1) can stop at θ_c with probability at most $\alpha_0 + \alpha_1$. A finite bank therefore needs an explicitly labeled terminal policy if every examinee must receive a decision; that fallback is not part of the anytime-valid guarantee. In particular, the experiments' choice $\alpha_0 = \alpha_1 = 0.05$ assigns 5% to each wrong direction and gives a 10% union bound on any certified decision at the cutoff; it is not a two-sided 5% bound. If both boundaries cross on the same response, either fixed measurable tie rule preserves these marginal bounds; our algorithms choose the larger normalized evidence and then break an exact tie toward H_1.

We use the following notation throughout. Index $\beta = 0$ denotes the lower ("fail") side and $\beta = 1$ the upper ("pass") side; set $s_\beta := (-1)^{\beta+1}$, $\alpha_\beta := \alpha_0$ for $\beta = 0$ and α_1 for $\beta = 1$, and $\ell_\beta := \log(1/\alpha_\beta)$. We assume throughout that $\alpha_0 + \alpha_1 < 1$. For local analyses, $\delta := |\theta - \theta_c|$ is always nonnegative and $\lambda := s_\beta a\delta = a(\theta - \theta_c)$ is the signed standardized gap for a repeated item of discrimination a. The symbol η is reserved for a tail probability; predictable reference displacements are denoted by ξ. For information comparisons, $I_j(\theta) := a_j^2 p_j(\theta)(1 - p_j(\theta))$ is the Fisher information of item j. For $p, q \in (0, 1)$, define $\mathrm{kl}(p, q) := p \log(p/q) + (1 - p) \log((1 - p)/(1 - q))$. The corresponding one-step information relative to the cutoff is $K_j(\theta, \theta_c) := \mathrm{kl}(p_j(\theta), p_j(\theta_c))$.

With the target guarantee now fixed, we turn to construction. We first study the simplest gap-free rule: a single alternative that moves toward the cutoff over time. Besides providing a valid baseline, this rule will expose the efficiency loss that motivates the multiscale method in Section 4.

3 A SHRINKING-ALTERNATIVE GAP-FREE CCT

The classical SPRT compares two fixed ability values (Wald, 1947; Reckase, 1983). That is efficient when those values describe the scientific question, but it creates the indifference region that we seek to avoid. Our first construction instead compares the cutoff with one reference ability on each side and moves both references toward the cutoff as evidence accumulates. We first prove validity for any predictable choice of references; we then select a shrinking reference sequence and analyze its stopping time.

The construction builds on adaptive likelihood-ratio tests, in which the numerator distribution at time t may be selected predictably from the information available before observing the next response (Robbins & Siegmund, 1972). In the present one-parameter setting, likelihood-ratio monotonicity extends validity at the cutoff to the corresponding composite half-line (Ramdas et al., 2023; Grünwald & Koolen, 2026). The reference sequence below adapts this principle to the 2PL model by standardizing the reference displacement by item discrimination and shrinking it at rate $t^{-1/2}$. Its validity holds for arbitrary predictable item selection, while its stopping behavior is quantified under the information conditions stated below. Further connections with the sequential-testing literature are given in Appendix A.3.

We now make this construction explicit under the 2PL response model. For the two decision directions, let $\theta_{0,t} < \theta_c < \theta_{1,t}$ be predictable reference abilities. For $\beta \in \{0, 1\}$ define $d_{\beta,0} = 1$ and

$$d_{\beta,t} = \prod_{s=1}^{t} \left(\frac{p_{j_s}(\theta_{\beta,s})}{p_{j_s}(\theta_c)} \right)^{Y_s} \left(\frac{1 - p_{j_s}(\theta_{\beta,s})}{1 - p_{j_s}(\theta_c)} \right)^{1 - Y_s}. \tag{2}$$

Proposition 3.1 (Anytime validity of predictable alternatives)**.** *Under any predictable item-selection rule and predictable reference-ability sequences, $(d_{0,t})_{t\geq 0}$ is a nonnegative supermartingale under every $\theta \geq \theta_c$, and $(d_{1,t})_{t\geq 0}$ is a nonnegative supermartingale under every $\theta \leq \theta_c$. Consequently,*

$$\Pr_\theta(\exists t : d_{0,t} \geq 1/\alpha_0) \leq \alpha_0 \quad (\theta \geq \theta_c), \qquad \Pr_\theta(\exists t : d_{1,t} \geq 1/\alpha_1) \leq \alpha_1 \quad (\theta \leq \theta_c).$$

If stopping occurs only when one of the two evidence processes $d_{\beta,t}$ reaches its threshold $1/\alpha_\beta$—what we call a certified boundary crossing—then $\Pr_{\theta_c}(\tau < \infty) \leq \alpha_0 + \alpha_1$.

This proposition treats every $p_j(\theta)$ as known and correctly specified. If fitted item parameters are substituted, sample splitting alone does not preserve the conditional mean inequality: calibration uncertainty, drift, local dependence, guessing, or multidimensionality can invalidate the exact e-process. "Any predictable item policy" concerns adaptivity of item choice inside the model, not robustness to response-model misspecification.

Appendix B.1 proves the result by reducing each composite half-line to the cutoff through 2PL monotonicity and applying Ville's inequality (Ville, 1939).

Proposition 3.1 deliberately leaves the reference sequence open. A fixed reference would restore a preferred gap, whereas a reference placed very near the cutoff from the outset would accumulate little evidence. We balance these effects by approaching the cutoff at the canonical $t^{-1/2}$ rate through the reference rule

$$\theta_{\beta,t} := \theta_c + (-1)^{\beta+1} \frac{\sqrt{\ell_\beta}}{a_{j_t}\sqrt{t}}, \qquad \beta \in \{0,1\}. \tag{3}$$

The factor $a_{j_t}^{-1}$ makes the log-odds displacement $\sqrt{\ell_\beta/t}$. Proposition 3.1 supplies the thresholds $1/\alpha_\beta$ in Algorithm 1.

Algorithm 1: Shrinking-alternative gap-free CCT

Input: Cutoff θ_c, levels α_0, α_1, predictable item policy
Output: A decision on boundary crossing; otherwise no decision

$d_{0,0} \leftarrow 1$, $d_{1,0} \leftarrow 1$;
for $t = 1, 2, \ldots$ **do**
  Select j_t using $\mathcal{F}_{t-1}$; then observe Y_t;
  for $\beta \in \{0,1\}$ **do**
    $\ell_\beta \leftarrow \log(1/\alpha_\beta)$ and Set $\theta_{\beta,t}$ according to (3);
    Update $d_{\beta,t}$ using the tth factor in (2);
  $\mathcal{C}_t \leftarrow \{\beta \in \{0,1\} : d_{\beta,t} \geq 1/\alpha_\beta\}$;
  if $\mathcal{C}_t \neq \varnothing$ **then**
    $\widehat{\beta}_t \leftarrow \arg\max_{\beta\in\mathcal{C}_t} \alpha_\beta d_{\beta,t}$, breaking ties toward 1;
    return $\tau = t$ and $H_{\widehat{\beta}_t}$;

The algorithm stops only at a certified crossing. To quantify stopping we analyze cutoff-centered items whose discriminations are bounded away from zero. Suppose $a_{j_t} \geq a_{\min} > 0$, put $\delta = |\theta - \theta_c|$ and $\gamma_\theta = a_{\min}\delta/(4 + 2a_{\min}\delta)$.

Proposition 3.2 (High-probability stopping time)**.** *Fix $\theta \neq \theta_c$ and let β be the correct index, so that $\operatorname{sign}(\theta - \theta_c) = (-1)^{\beta+1}$. If every administered item is centered at the cutoff, $b_{j_t} = \theta_c$, and $a_{j_t} \geq a_{\min} > 0$, let $\mathcal{A}_\beta < \infty$ and $D_\theta > 0$ be the explicit constants in Proposition B.4. Then, for every integer $t \geq \max(2, \mathcal{A}_\beta)$,*

$$\Pr_\theta(\tau \geq t + 1) \leq \exp(-t/D_\theta). \tag{4}$$

Thus the explicit horizon $\widehat{N}_\beta(\eta) = \lceil \max(2, \mathcal{A}_\beta, D_\theta \log(1/\eta)) \rceil$ satisfies $\Pr_\theta(\tau \geq \widehat{N}_\beta(\eta) + 1) \leq \eta$.

Equivalently, with the error levels fixed and as $a_{\min}\delta \downarrow 0$, the explicit horizon has the simpler equivalent order

$$\widehat{N}_\beta(\eta) \asymp \frac{\log^2\big(1/(a_{\min}\delta)\big) + \log\big(1/(a_{\min}\delta)\big)\log(1/\eta)}{a_{\min}^2\delta^2}, \tag{5}$$

where the comparison constants may depend on α_β but not on δ or η. Thus, at fixed confidence the local cost is $\delta^{-2}\log^2(1/\delta)$, whereas at a fixed nonzero gap the confidence cost is only logarithmic in $1/\eta$.

Theorem 3.3 (Validity and mean stopping time)**.** *Algorithm 1 satisfies the one-sided error guarantees in (1). Under the assumptions of Proposition 3.2, the explicit constant T_β defined in Appendix B.2 satisfies $\mathbb{E}_\theta[\tau] \le T_\beta + \pi^2/6$ and, for fixed error levels,*

$$T_\beta + \frac{\pi^2}{6} \asymp \frac{\log^2(1/(a_{\min}\delta))}{a_{\min}^2\delta^2} \quad \text{as } a_{\min}\delta \downarrow 0. \tag{6}$$

Proof idea and interpretation. For the correct side, the cumulative log-likelihood is a predictable drift plus a mean-zero martingale. The centered-item assumptions make the drift grow on the $\sqrt{t}$ scale. Freedman's martingale inequality (Freedman, 1975), applied on a terminal block where the increments are of order $t^{-1/2}$, yields the exponential tail in (4) and hence the $D_\theta \log(1/\eta)$ horizon. Unlike Markov's inequality applied to a mean bound, this directly controls a chosen upper quantile of τ. The polynomial tail in Proposition B.4 is summable and gives Theorem 3.3, whose local order is $O(\delta^{-2}\log^2(1/\delta))$. Appendix B.2 contains the detailed bounds, their interpretation, and every drift and variance calculation; Appendix C proves the deterministic inversion needed for T_β.

Taken together, these results show that the shrinking rule is anytime valid and stops with exponentially decaying tails away from the cutoff, but its local mean contains an extra $\log^2(1/\delta)$ factor. That factor is the price of following a single prescribed scale. The next section removes this commitment by retaining evidence across many candidate scales simultaneously.

4 An improved multiscale mixture CCT

The shrinking rule invests in one prescribed scale although the useful scale depends on $|\theta - \theta_c|$. We instead mix fixed-scale likelihood-ratio processes $\{R_{\beta,k,t}\}_{k\ge 0}$ with summable weights. Geometric scales and polynomial weights connect this construction to discrete mixtures and stitching (Robbins, 1970; Fellouris & Tartakovsky, 2013; Howard et al., 2021); here they are standardized on the 2PL log-odds scale and implemented by certified finite partial sums. Appendix A.3 gives the detailed lineage. For the two decision directions, let

$$u_k = 2^{-k}, \qquad w_k = \frac{6}{\pi^2(k+1)^2}, \qquad k = 0, 1, \dots, \tag{7}$$

so that $\sum_{k\ge 0} w_k = 1$. Component k uses $\theta_{\beta,k,t} = \theta_c + (-1)^{\beta+1}u_k/a_{j_t}$, whose standardized log-odds displacement from the cutoff is always $\pm u_k$. Define

$$R_{\beta,k,t} = \prod_{s=1}^{t} \left(\frac{p_{j_s}(\theta_{\beta,k,s})}{p_{j_s}(\theta_c)}\right)^{Y_s} \left(\frac{1-p_{j_s}(\theta_{\beta,k,s})}{1-p_{j_s}(\theta_c)}\right)^{1-Y_s}, \tag{8}$$

$$M_{\beta,t} = \sum_{k=0}^{\infty} w_k R_{\beta,k,t}. \tag{9}$$

Because u_k is on the log-odds scale, discrimination is standardized across items. With $s_\beta = (-1)^{\beta+1}$, $z_t = a_{j_t}(\theta_c - b_{j_t})$, and $A(z) = \log(1+e^z)$, a component is updated by

$$\log R_{\beta,k,t} = \log R_{\beta,k,t-1} + s_\beta u_k Y_t - (A(z_t + s_\beta u_k) - A(z_t)). \tag{10}$$

For computation, use the partial sum $P_{\beta,t}^{(K)} = \sum_{k=0}^{K} w_k R_{\beta,k,t}$. Since the omitted terms are nonnegative,

$$P_{\beta,t}^{(K)} \le M_{\beta,t}. \tag{11}$$

A partial-sum crossing is therefore a genuine countable-mixture crossing.

Algorithm 2 uses a growing partial sum. The experiments use $K = 20$ and the following deterministic lower bound for the omitted tail:

$$\sum_{k=21}^{\infty} w_k R_{\beta,k,t} \ge \left(1 - \sum_{k=0}^{20} w_k\right) e^{-t2^{-21}}. \tag{12}$$

Adding this term to the 21-component sum remains a pathwise lower bound on $M_{\beta,t}$ and hence certifies only genuine crossings. It can stop at a different time from the dynamic rule analyzed below.

Proposition 4.1 (Anytime validity of the mixture)**.** *Under any predictable item-selection policy, $(M_{0,t})_{t\geq 0}$ is a nonnegative supermartingale under every $\theta \geq \theta_c$, and $(M_{1,t})_{t\geq 0}$ is a nonnegative supermartingale under every $\theta \leq \theta_c$. More generally, the same conclusion holds if the standardized references above are replaced, separately for every component, by any predictable processes satisfying $\theta_{0,k,t} \leq \theta_c \leq \theta_{1,k,t}$. Hence thresholding at $1/\alpha_0$ and $1/\alpha_1$ satisfies (1), and at the cutoff the probability of ever stopping is at most $\alpha_0 + \alpha_1$.*

As in Proposition 3.1, this is a conditional model guarantee for known response functions. The mixture does not by itself absorb item-calibration uncertainty or model misspecification.

Appendix B.3 gives the conditional mixture argument and Ville bound (Ville, 1939). The procedure below activates only logarithmically many components and reconstructs each new one from the complete history.

Algorithm 2: Multiscale-mixture gap-free CCT

Input: Cutoff θ_c, levels α_0, α_1, predictable item policy
Output: A decision on boundary crossing; otherwise no decision

Set u_k, w_k as in (7), $K_0 \leftarrow 4$, and initialize all active component evidence processes to one;
for $t = 1, 2, \ldots$ **do**
  Set $K_t \leftarrow 4 + \lceil \log_2(t+1) \rceil$ and reconstruct any newly active components by replaying stored observations through (10);
  Select j_t using $\mathcal{F}_{t-1}$; then observe Y_t;
  Update every active component using (10);
  Compute $P_{\beta,t}^{(K_t)}$ for $\beta \in \{0,1\}$;
  $\mathcal{C}_t \leftarrow \{\beta \in \{0,1\} : P_{\beta,t}^{(K_t)} \geq 1/\alpha_\beta\}$;
  if $\mathcal{C}_t \neq \varnothing$ **then**
    $\widehat{\beta}_t \leftarrow \arg\max_{\beta \in \mathcal{C}_t} \alpha_\beta P_{\beta,t}^{(K_t)}$, breaking ties toward 1;
    return $\tau = t$ and $H_{\widehat{\beta}_t}$;
  Continue otherwise;

Only $O(\log t)$ component pairs are active at time t. Because every fixed component is eventually activated and reconstructed from the full history, the implementation never permanently discards a fixed scale that may match the unknown gap.

Appendix B.4 gives the general heterogeneous-item stopping theorem; it requires cumulative conditional log-likelihood advantage because validity alone cannot prevent an uninformative item policy.

For comparison with shrinking, we use the ideal benchmark of independent cutoff-centered item clones with fixed discrimination. This isolates the cost of adapting to the unknown gap; it is not a recommendation for literal repeat exposure.

Theorem 4.2 (Mixture validity and mean stopping time)**.** *Let τ_{tr} denote the stopping time of the truncated rule in Algorithm 2, and let τ_∞ denote the ideal countable-mixture crossing time. Both rules satisfy (1). Suppose an unlimited supply of conditionally independent cutoff-centered item clones with discrimination $a > 0$ is available. Put $\lambda = a(\theta - \theta_c)$ and let β index the correct side. If $0 < |\lambda| \leq 1$, then the implemented stopping time satisfies*

$$\mathbb{E}_\theta[\tau_{\mathrm{tr}}] \leq \frac{16}{\lambda^2}\left[\log\frac{1}{\alpha_\beta} + \log\frac{\pi^2}{6} + 2\log\left(2 + \log_2\frac{1}{|\lambda|}\right) + \frac{1}{2}\right]. \tag{13}$$

The same bound holds for τ_∞, since $\tau_\infty \leq \tau_{\mathrm{tr}}$. For every $\lambda \neq 0$, both rules have finite mean stopping time, and the implemented rule has $O(\lambda^{-2}\log\log(1/|\lambda|))$ local order.

The grid has a component within a factor of two of each nonzero gap. It gains drift of order λ^2 and pays only $O(\log\log(1/|\lambda|))$ for its weight; the activation schedule makes it available before crossing. Appendix B.5 proves the bound.

For a finite-sample benchmark without matched items, put $K^\star(\theta, \theta_c) = \sup_j \mathrm{kl}(p_j(\theta), p_j(\theta_c))$. Every rule satisfying (1) and having finite mean at the displayed nonzero parameter must obey

$$\mathbb{E}_\theta \tau \geq \begin{cases} \mathrm{kl}(1-\alpha_0, \alpha_1)/K^\star(\theta, \theta_c), & \theta > \theta_c, \\ \mathrm{kl}(1-\alpha_1, \alpha_0)/K^\star(\theta, \theta_c), & \theta < \theta_c. \end{cases}$$

Appendix B.6 proves the stronger stopped-information and finite-bank versions. Since $\mathrm{kl}(p_j(\theta_c \pm \delta), p_j(\theta_c)) \leq a_j^2\delta^2/8$, bounded discrimination gives the usual δ^{-2} lower order. This result does not cover abstention, cap-charged exposure, or fallback labels.

For the i.i.d. centered-item experiment, Farrell's lower bound adds the cost of adapting one rule to every nonzero gap. Any directionally valid rule that stops almost surely at every $\theta \neq \theta_c$ must, with $q = \Pr_{\theta_c}(\tau = \infty) \geq 1 - \alpha_0 - \alpha_1 > 0$, satisfy on each side (Farrell, 1964),

$$\limsup_{\delta \downarrow 0} \frac{a^2\delta^2\, \mathbb{E}_{\theta_c \pm \delta}\, \tau}{\log\log(1/\delta)} \geq 8q.$$

Appendix B.7 verifies the assumptions and constant. The $\limsup$ forces this order only along a sequence approaching the cutoff. Thus the mixture improves shrinking by one logarithmic factor and attains the optimal worst-scale local order in the ideal experiment.

The theory has now separated the two main questions. Both rules are directionally valid under predictable item selection, while the multiscale rule pays a smaller local adaptation cost. We next examine whether this theoretical advantage is visible at finite sample sizes and contrast it with the speed of indifference-region and fixed-length benchmarks.

5 Experiments

We study three complementary questions: how the rules behave under known 2PL truth, how they stop on naturally occurring logged paths, and how sequential decisions compare with fixed rules at common horizons. Within each study all methods receive the same items and responses. This isolates stopping-rule behavior rather than the coupled effect of stopping and adaptive item choice. Appendix D gives the complete protocol, comparator definitions, calibration, splits, and supplementary results.

5.1 Known-truth comparison

We generate a heterogeneous 1,000-item 2PL bank and administer items without replacement in decreasing cutoff information. The cutoff is the bank-average ability corresponding to an expected score of 14 out of 20 ($\theta_c = 0.888$). At each registered distance $|\theta - \theta_c| \in \{0.4, 0.8, 1.2\}$, 1,000 paths are generated on either side and followed for at most 1,000 items. We compare shrinking and the fixed-21-component mixture at directional levels $\alpha_0 = \alpha_1 = 0.05$ with SPRT, a spending-based profile-LR interval, integrated CLR, practical GLR, and SCGLR. These comparators target different inferential criteria, so the study contrasts operating characteristics rather than equally leveled tests.

We distinguish adverse crossings (wrong-direction early decisions), certification yield (the fraction crossing by the cap), and cap-charged exposure (items used, charging 1,000 to a noncrossing path). A noncrossing is an abstention, not an error; fallback errors are reported separately in the appendix.

Table 1: Known-truth risk–coverage profile away from the cutoff. Certified is the percentage of paths crossing a valid e-process boundary by the 1,000-item cap. Exposure charges the cap to noncrossing paths; saving is the mixture's reduction in this cap-charged mean relative to shrinking. Adverse crossings are unconditional wrong-direction decisions. The complete count table, including all classical comparators, is in Table 3.

$\lvert\theta - \theta_c\rvert$	Certified (%)		Mixture exposure	Saving (%)	Adverse (%)
	Shrinking	Mixture	(items)	vs. shrinking	Shrinking / mixture
0.4	99.00	99.85	85.0	36.9	0.00 / 0.00
0.8	100.00	100.00	18.6	24.6	0.00 / 0.00
1.2	100.00	100.00	12.5	10.6	0.00 / 0.00

Neither gap-free rule makes an adverse crossing in the 2,000 pooled paths at any distance. The two-sided 95% exact-binomial upper endpoint for each zero count is 0.18%, so zero observed events do not establish zero risk. At gap 0.4, the mixture raises certification from 99.00% to 99.85% and reduces cap-charged exposure from 134.6 to 85.0 items, a 36.9% saving. At gaps 0.8 and 1.2, both certify every path, while the mixture reduces exposure from 24.7 to 18.6 items (24.6%) and from 13.9 to 12.5 items (10.6%).

The faster separated or integrated likelihood rules trade speed for adverse crossings. At gap 0.4, SPRT, integrated CLR, practical GLR, and SCGLR use 8.6, 11.8, 7.2, and 7.2 items, but have adverse rates 5.95%, 9.80%, 7.95%, and 8.00%. Those rates fall to 0.25–3.15% at gap 0.8 and 0.05–0.50% at gap 1.2. The profile-LR rule has no observed adverse crossing but uses 446.4 items at gap 0.4. Thus the mixture materially improves the conservative shrinking baseline without adopting the different error target that makes the aggressive comparators faster. These results are conditional on one generated bank and exclude item-model uncertainty.

5.2 Logged-response comparison

We use ASSISTments 2009–2010 (ASSISTments, 2009), EdNet KT1 (Choi et al., 2020), and NeurIPS/Eedi (Wang et al., 2021). Students are split 75–25% between item calibration and evaluation; within evaluation students, 60% of responses form the sequential path and 40% form a response-disjoint reference assessment. The primary analysis retains cases whose reference 95% Wald interval excludes the bank-derived 14/20 cutoff. Consequently, "wrong" means disagreement with a fitted reference label, not known latent error. The estimand is conditional reference disagreement under logged exposure, not counterfactual adaptive performance.

The indifference-region issue remains empirically material: 17,758 of 40,811 eligible students (43.5%) have reference estimates within $[\theta_c - 0.5, \theta_c + 0.5]$; even after confidence filtering, 2,701 of 19,436 (13.9%) remain there. The pooled sequential results are:

Table 2: Logged-path results on 19,436 confidently reference-classified students. Hybrid error applies the same cutoff-score fallback to unresolved paths.

Method	Mean items	Crossed (%)	Wrong crossing (%)	Hybrid error (%)
Shrinking CCT	160.2	61.98	0.60	2.57
Mixture CCT	119.3	74.75	0.89	2.78
SPRT ($\Delta = 0.5$)	32.6	95.21	4.98	5.90

The mixture uses 40.9 fewer responses than shrinking (95% paired-bootstrap interval 38.7–43.0), increases crossing by 12.78 percentage points, and increases hybrid disagreement by only 0.22 points (0.12–0.31). Relative to SPRT it uses 86.7 more responses (83.7–89.9) but reduces hybrid disagreement by 3.12 points (2.86–3.39). This is the same efficiency–calibration tradeoff seen under known truth: the mixture recovers much of shrinking's exposure cost, whereas SPRT is faster under its separated design. EdNet contributes 18,460 students (95.0%), so the pooled precision is primarily EdNet precision; dataset and directional results in Appendix D.3 are descriptive checks of heterogeneity.

5.3 Common-horizon comparison

Among 8,523 confidently classified students with at least 200 sequential responses, we compare sequential rules with fixed endpoint decisions at seven common horizons. At horizon 200, the mixture uses 106.9 responses with 5.80% hybrid disagreement, versus 200 responses and 5.16% for the endpoint LR. At horizon 120, the corresponding values are 76.1 and 7.98% versus 120 and 7.59%. SPRT saturates near 40 responses at roughly 9% disagreement, while both gap-free rules use larger horizons to reduce terminal disagreement. Accordingly, the data show a resource–reference-agreement frontier, not an error-matched dominance result. The complete horizon grid and figure appear in Appendix D.4.

6 Conclusion

We developed two gap-free, anytime-valid CCTs for the 2PL model. The shrinking rule has local mean order $O(\delta^{-2}\log^2(1/\delta))$; the multiscale mixture reduces the adaptation cost to $\log\log(1/\delta)$ and attains the optimal worst-scale order in the ideal centered-item experiment. Both retain one-sided error control under predictable item selection and permit nontermination at the cutoff. In the known-truth study, the mixture substantially reduced shrinking's cap-charged exposure with no observed adverse crossing. On logged paths it used 25.5% fewer responses than shrinking for a 0.22-point increase in hybrid reference disagreement, whereas the faster separated-design SPRT disagreed more often.

The guarantees assume known, correctly specified item response functions; finite stopping additionally requires informative items, and the sharp rate comparison uses an ideal infinite bank. Logged responses neither reveal latent truth nor identify performance under counterfactual adaptive item selection, and the analysis does not propagate item-calibration uncertainty. Future work should develop calibration-robust evidence processes, sharpen heterogeneous finite-bank bounds, construct certified finite-horizon completion rules, and evaluate exposure-constrained adaptive policies prospectively.

Code availability

Code is available from the authors upon request.

Ethics statement

Educational response logs can reflect unequal opportunity, language, disability, prior instruction, and platform exposure. A CCT may reduce burden but can also automate consequential decisions; deployment requires population-specific validation, subgroup audits, accommodations, appeal mechanisms, and human oversight. Mathematical error budgets should not be claimed operationally unless model fit and item drift have been validated for the target population.

AI use statement

The theoretical developments and results concerning the first shrinking-alternative algorithm were obtained by the authors. The multiscale method and lower bounds were obtained using GPT-5.6 Sol. Generative AI tools subsequently assisted with implementation, proof checking, literature organization, figure and table code, and manuscript drafting. The authors reviewed the mathematics, code, citations, claims, and prose and take responsibility for the final work.

References

ASSISTments. ASSISTments 2009–2010 skill builder data. https://sites.google.com/site/assistmentsdata/home/2009-2010-assistment-data, 2009. Accessed 2026-09-04.

Juan Ramón Barrada, Julio Olea, Vicente Ponsoda, and Francisco José Abad. Incorporating randomness in the fisher information for improving item-exposure control in CATs. *British Journal of Mathematical and Statistical Psychology*, 61(2):493–513, 2008.

Juan Ramón Barrada, Julio Olea, Vicente Ponsoda, and Francisco José Abad. Item selection rules in computerized adaptive testing: Accuracy and security. *Methodology*, 5(1):7–17, 2009.

Jay Bartroff, Matthew Finkelman, and Tze Leung Lai. Modern sequential analysis and its applications to computerized adaptive testing. *Psychometrika*, 73(3):473–486, 2008.

Allan Birnbaum. Some latent trait models and their use in inferring an examinee's ability. In Frederic M. Lord and Melvin R. Novick (eds.), *Statistical Theories of Mental Test Scores*, pp. 397–479. Addison-Wesley, Reading, MA, 1968.

Hua-Hua Chang. Psychometrics behind computerized adaptive testing. *Psychometrika*, 80(1):1–20, 2015.

Hua-Hua Chang and Zhiliang Ying. A global information approach to computerized adaptive testing. *Applied Psychological Measurement*, 20(3):213–229, 1996.

Hua-Hua Chang and Zhiliang Ying. Nonlinear sequential designs for logistic item response theory models with applications to computerized adaptive tests. *The Annals of Statistics*, 37(3):1466–1488, 2009.

Yuan-Chin Ivan Chang. Application of sequential interval estimation to adaptive mastery testing. *Psychometrika*, 70(4):685–713, 2005.

Youngduck Choi, Youngnam Lee, Dongmin Shin, Junghyun Cho, Seoyon Park, Seewoo Lee, Jineon Baek, Chan Bae, Byungsoo Kim, and Jaewe Heo. EdNet: A large-scale hierarchical dataset in education. In *Artificial Intelligence in Education*, pp. 69–73, 2020. doi: 10.1007/978-3-030-52237-7_6.

Theo J. H. M. Eggen. Item selection in adaptive testing with the sequential probability ratio test. *Applied Psychological Measurement*, 23(3):249–261, 1999.

Theo J. H. M. Eggen and Guuske J. J. M. Straetmans. Computerized adaptive testing for classifying examinees into three categories. *Educational and Psychological Measurement*, 60(5):713–734, 2000.

Roger H. Farrell. Asymptotic behavior of expected sample size in certain one sided tests. *The Annals of Mathematical Statistics*, 35(1):36–72, 1964.

Georgios Fellouris and Alexander G. Tartakovsky. Almost optimal sequential tests of discrete composite hypotheses. *Statistica Sinica*, 23(4):1717–1741, 2013.

Matthew Finkelman. On using stochastic curtailment to shorten the SPRT in sequential mastery testing. *Journal of Educational and Behavioral Statistics*, 33(4):442–463, 2008.

David A. Freedman. On Tail Probabilities for Martingales. *The Annals of Probability*, 3(1):100–118, 1975.

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 998–1027. PMLR, 2016.

Aurélien Garivier and Emilie Kaufmann. Nonasymptotic sequential tests for overlapping hypotheses applied to near-optimal arm identification in bandit models. *Sequential Analysis*, 40(1):61–96, 2021.

Peter Grünwald, Rianne de Heide, and Wouter M. Koolen. Safe testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(5):1091–1128, 2024.

Peter D. Grünwald and Wouter M. Koolen. Supermartingales for one-sided tests: Sufficient monotone likelihood ratios are sufficient. *Statistics & Probability Letters*, 229:110574, 2026.

Kit-Tai Hau and Hua-Hua Chang. Item selection in computerized adaptive testing: Should more discriminating items be used first? *Journal of Educational Measurement*, 38(3):249–266, 2001.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.

Po-Hsien Hu, Ching-Lin Shih, and Cheng-Te Chen. A pre-rule for the sequential probability ratio test in a between-item grid multidimensional computerized classification test. *Behavior Research Methods*, 57:79, 2025.

Alan Huebner and Anthony D. Fina. The stochastically curtailed generalized likelihood ratio: A new termination criterion for variable-length computerized classification tests. *Behavior Research Methods*, 47(2):549–561, 2015.

Ramesh Johari, Pete Koomen, Leonid Pekelis, and David Walsh. Always valid inference: Continuous monitoring of A/B tests. *Operations Research*, 70(3):1806–1821, 2022.

Nikos Karampatziakis, Paul Mineiro, and Aaditya Ramdas. Off-policy confidence sequences. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5301–5310. PMLR, 2021.

Emilie Kaufmann and Wouter M. Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22(246):1–44, 2021.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.

Emilie Kaufmann, Wouter M. Koolen, and Aurélien Garivier. Sequential test for the lowest mean: From thompson to murphy sampling. In *Advances in Neural Information Processing Systems 31*, pp. 6335–6345. Curran Associates, Inc., 2018.

Sanghwa Kim, Dohyun Ahn, and Seungki Min. On the optimality of tracking fisher information in adaptive testing with stochastic binary responses, 2025.

Nick W. Koning and Sam van Meer. Anytime validity is free: Inducing sequential tests. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 88(4):1366–1384, 2026.

Tze Leung Lai. On optimal stopping problems in sequential hypothesis testing. *Statistica Sinica*, 7 (1):33–51, 1997.

Charles Lewis and Kathleen Sheehan. Using Bayesian decision theory to design a computerized mastery test. *Applied Psychological Measurement*, 14(4):367–386, 1990.

Tien-Hsiang Liu, Cheng-Te Chen, Chung-Ping Cheng, and Ching-Lin Shih. Improving test efficiency for a grid multidimensional computerized classification test by the application of a conditional latent-trait distribution to a sequential probability ratio test. *Behavior Research Methods*, 54(5): 2565–2578, 2022.

Steven W. Nydick. The sequential probability ratio test and binary item response models. *Journal of Educational and Behavioral Statistics*, 39(3):203–230, 2014.

Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter M. Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales, 2022.

Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.

Mark D. Reckase. A procedure for decision making using tailored testing. In David J. Weiss (ed.), *New Horizons in Testing: Latent Trait Test Theory and Computerized Adaptive Testing*, pp. 237–255. Academic Press, New York, 1983.

Herbert Robbins. Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, 41(5):1397–1409, 1970.

Herbert Robbins and David Siegmund. A class of stopping rules for testing parametric hypotheses. In Lucien M. Le Cam, Jerzy Neyman, and Elizabeth L. Scott (eds.), *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability*, volume 4, pp. 37–41, Berkeley, CA, 1972. University of California Press.

Nicki-Nils Seitz and Andreas Frey. The sequential probability ratio test for multidimensional adaptive testing with between-item multidimensionality. *Psychological Test and Assessment Modeling*, 55 (1):105–123, 2013.

James Sharpnack, Kevin Hao, Phoebe Mulcaire, Klinton Bicknell, Geoff LaFlair, Kevin Yancey, and Alina A. von Davier. BanditCAT and AutoIRT: Machine learning approaches to computerized adaptive testing and item calibration. In *Proceedings of Large Foundation Models for Educational Assessment*, volume 264 of *Proceedings of Machine Learning Research*, pp. 121–135. PMLR, 2025.

Judith A. Spray and Mark D. Reckase. Comparison of SPRT and sequential Bayes procedures for classifying examinees into two categories using a computerized test. *Journal of Educational and Behavioral Statistics*, 21(4):405–414, 1996.

Martha L. Stocking and Len Swanson. A method for severely constrained item selection in adaptive testing. *Applied Psychological Measurement*, 17(3):277–292, 1993.

Nathan A. Thompson. A practitioner's guide for variable-length computerized classification testing. *Practical Assessment, Research & Evaluation*, 12(1), 2007.

Nathan A. Thompson. Item selection in computerized classification testing. *Educational and Psychological Measurement*, 69(5):778–793, 2009a.

Nathan A. Thompson. Utilizing the generalized likelihood ratio as a termination criterion. In *Proceedings of the GMAC Conference on Computerized Adaptive Testing*, Minneapolis, MN, 2009b.

Nathan A. Thompson and Shungwon Ro. Computerized classification testing with composite hypotheses. In David J. Weiss (ed.), *Proceedings of the 2007 GMAC Conference on Computerized Adaptive Testing*, Minneapolis, MN, 2007.

Wim J. van der Linden and Bernard P. Veldkamp. Constraining item exposure in computerized adaptive testing with shadow tests. *Journal of Educational and Behavioral Statistics*, 29(3): 273–291, 2004.

Maaike M. van Groen, Theo J. H. M. Eggen, and Bernard P. Veldkamp. Multidimensional computerized adaptive testing for classifying examinees with within-dimensionality. *Applied Psychological Measurement*, 40(6):387–404, 2016.

Jean Ville. *Étude Critique de la Notion de Collectif*. Gauthier-Villars, 1939.

Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3):1736–1754, 2021.

Abraham Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16 (2):117–186, 1945.

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

Abraham Wald and Jacob Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948.

Zichao Wang, Angus Lamb, Evgeny Saveliev, Pashmina Cameron, Yordan Zaykov, José Miguel Hernández-Lobato, Richard E. Turner, Richard G. Baraniuk, Craig Barton, Simon Peyton Jones, Simon Woodhead, and Cheng Zhang. Results and insights from diagnostic questions: The NeurIPS 2020 education challenge. In *Proceedings of the NeurIPS 2020 Competition and Demonstration Track*, volume 133 of *Proceedings of Machine Learning Research*, pp. 191–205, 2021. URL https://proceedings.mlr.press/v133/wang21a.html.

Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(1):1–27, 2024.

David J. Weiss and G. Gage Kingsbury. Application of computerized adaptive testing to educational problems. *Journal of Educational Measurement*, 21(4):361–375, 1984.

R. A. Weitzman. Sequential testing for selection. *Applied Psychological Measurement*, 6(3):337–351, 1982.

Yan Zhuang, Qi Liu, Haoyang Bi, Zhenya Huang, Weizhe Huang, Jiatong Li, Junhao Yu, Zirui Liu, Zirui Hu, Yuting Hong, Zachary A. Pardos, Haiping Ma, Mengxiao Zhu, Shijin Wang, and Enhong Chen. Survey of computerized adaptive testing: A machine learning perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(8):8744–8763, 2026.

A Related work

This appendix expands the positioning in the introduction. It moves from adaptive item selection, through the closest CCT predecessors, to the sequential-inference principles underlying the two proposed procedures.

A.1 Adaptive testing and item selection

CAT alternates latent-trait estimation and item selection, typically using Fisher information at the current ability estimate (Weiss & Kingsbury, 1984; Chang, 2015; Barrada et al., 2009), while global and likelihood-weighted criteria offer alternatives to this local rule (Chang & Ying, 1996; Barrada et al., 2009). Maximum-information designs have consistency and asymptotic-normality guarantees in the Rasch model and, with modifications, in 2PL and 3PL models (Chang & Ying, 2009); recent work also establishes fixed-confidence and fixed-budget optimality for Fisher-information tracking in continuous-item models (Kim et al., 2025). These methods target accurate estimation, whereas CCT targets a certified decision relative to one or more cut scores.

Operational item selection cannot optimize information alone. Repeated use of high-discrimination items can create severe exposure imbalance; stratification, randomized information, constrained assembly, and shadow tests trade statistical information against pool use, content, and security (Hau & Chang, 2001; Barrada et al., 2008; Stocking & Swanson, 1993; van der Linden & Veldkamp, 2004). Recent machine-learning work casts CAT administration as a contextual-bandit problem, using Fisher information as reward while explicitly randomizing for exposure control (Sharpnack et al., 2025). Complementary bandit methods jointly optimize adaptive sampling, stopping, and decisions, with results on identification complexity, threshold testing, overlapping hypotheses, and time-uniform error control (Kaufmann et al., 2016; Garivier & Kaufmann, 2016; Kaufmann et al., 2018; Garivier & Kaufmann, 2021; Kaufmann & Koolen, 2021). Unlike CCT, these methods lack a shared latent ability and they do not provide directional validity across a cutoff.

A.2 Sequential and composite CCT

Reckase (1983) introduced an IRT-based sequential decision procedure, and subsequent work compared SPRT and Bayesian termination (Lewis & Sheehan, 1990; Spray & Reckase, 1996). Item selection can target Fisher information at the cutoff or Kullback–Leibler information between classification points (Eggen, 1999; Thompson, 2009a). Extensions cover more than two categories (Eggen & Straetmans, 2000), as well as multiple dimensions and correlation-aware termination (Seitz & Frey, 2013; van Groen et al., 2016; Liu et al., 2022; Hu et al., 2025). Our theory is unidimensional, but within the 2PL model validity permits any predictable item selection policy. Our experiments deliberately hold the item order and response path common across methods to isolate stopping-rule behavior; they do not compare adaptive item-selection policies.

Further CCT work improves efficiency through truncation, stochastic curtailment (Finkelman, 2008), generalized likelihood ratios (Thompson, 2009b; Huebner & Fina, 2015), and refinements for binary IRT models (Nydick, 2014); Thompson (2007) gives an operational overview. Mixture and generalized-likelihood-ratio procedures also arise from optimal stopping for composite sequential tests (Lai, 1997). Most of these constructions retain separated design points, a calibrated maximum length, or asymptotic boundaries.

Chang (2005) applies one-sided sequential confidence intervals with β-protection to adaptive mastery testing, studying fixed-width and fixed-proportional-accuracy variants through classification accuracy, test length, and indifference-region width.

Bartroff et al. (2008) extends sequential generalized-likelihood-ratio theory to predictable CAT designs and applies modified Haybittle–Peto boundaries to finite-horizon 3PL mastery tests under limiting item-frequency and information conditions.

An earlier composite-hypothesis analogue is provided by Weitzman (1982), who adapted Wald's probability-ratio test by summing response-sequence probabilities over empirical criterion-score quantiles above and below a selection cutoff. Thompson & Ro (2007) extended this idea to a 3PL model by integrating likelihoods on both sides.

The present procedures address calibration differently. A cutoff-referenced e-process gives finite-sample, time-uniform wrong-direction control pointwise over each half-line under any predictable item policy. The countable multiscale construction does not interpret a population-averaged likelihood as a pointwise guarantee and requires neither an indifference width nor a tuned maximum length. It also yields an explicit stopping bound for the shrinking rule and the optimal worst-scale local mixture order in the ideal i.i.d. centered-item 2PL experiment.

A.3 Sequential likelihood ratios and anytime-valid inference

The statistical foundations of both procedures predate their CCT specialization. Likelihood-ratio martingales and Ville's maximal inequality provide time-uniform error control (Ville, 1939; Wald, 1947). Robbins and Siegmund extended this perspective to adaptive likelihood ratios whose numerator distribution may be selected predictably (Robbins & Siegmund, 1972), while mixture likelihood-ratio tests combine evidence across several alternatives (Fellouris & Tartakovsky, 2013).

Modern work gives a complementary, nonasymptotic toolkit for continuous monitoring. Exponential-supermartingale, mixture, and stitching arguments produce time-uniform concentration bounds and confidence sequences (Howard et al., 2020; 2021); betting constructions sharpen them for bounded observations (Waudby-Smith & Ramdas, 2024). The same framework organizes e-values, e-processes, and safe anytime-valid inference (Ramdas et al., 2023). Related e-value and safe-testing work develops calibration and testing principles for composite settings (Vovk & Wang, 2021; Grünwald et al., 2024). Nonnegative martingales are universal building blocks for admissible anytime-valid tests (Ramdas et al., 2022). At a designated maximum horizon, a valid terminal test can instead induce an anytime-valid procedure that matches the terminal test at that horizon (Koning & van Meer, 2026); the procedures here are constructed directly as open-ended evidence processes. For one-sided composite nulls, monotone likelihood ratios can extend supermartingale validity at the boundary to the complete null region (Grünwald & Koolen, 2026).

Accordingly, Proposition 3.1's one-step supermartingale logic is a transparent specialization, not a new general theorem about one-sided testing. The new work lies in choosing and analyzing standardized alternatives for predictable 2PL administration, implementing the countable mixture conservatively, and relating its stopping cost to both pointwise information and Farrell's worst-scale bound.

B Proofs

Roadmap. The proofs follow the order of the main text. Appendix B.1 proves anytime validity of the shrinking rule; Appendix B.2 proves its high-probability and mean stopping-time bounds; Appendix B.3 proves mixture validity and the conservative partial-sum implementation; Appendix B.4 gives the predictable-design concentration bound for heterogeneous, adaptively selected items; Appendix B.5 proves the repeated-item mean bound; and Appendix B.6 proves the information lower bound. Appendix C contains the deterministic inequalities used to invert the shrinking-rule tail.

B.1 Likelihood-ratio monotonicity and shrinking-rule validity

Lemma B.1 (Single-item likelihood-ratio monotonicity)**.** *Fix an item j and a reference $\theta_r \neq \theta_c$. Set $p = p_j(\theta)$, $p_c = p_j(\theta_c)$, $p_r = p_j(\theta_r)$, and*

$$L_j(y) = \left(\frac{p_r}{p_c}\right)^y \left(\frac{1-p_r}{1-p_c}\right)^{1-y}.$$

Then $f(p) = \mathbb{E}_{Y\sim\mathrm{Bern}(p)} L_j(Y)$ satisfies $f(p_c) = 1$ and is strictly monotone with $\operatorname{sign}(f') = \operatorname{sign}(\theta_r - \theta_c)$.

Proof. Because a 2PL probability lies strictly between zero and one, all ratios below are finite. Taking the expectation separately for $Y = 1$ and $Y = 0$ gives

$$f(p) = p\frac{p_r}{p_c} + (1-p)\frac{1-p_r}{1-p_c}, \qquad f'(p) = \frac{p_r - p_c}{p_c(1-p_c)}.$$

Substituting $p = p_c$ in the first display gives $f(p_c) = p_r + (1 - p_r) = 1$. Under the 2PL model, $p_j(x) = \sigma(a_j(x - b_j))$ with $a_j > 0$, and therefore

$$p_j'(x) = a_j p_j(x)(1 - p_j(x)) > 0.$$

Hence $p_r - p_c$ has the same sign as $\theta_r - \theta_c$. Since the denominator in $f'(p)$ is positive, this also is the sign of $f'(p)$, proving strict monotonicity and the lemma. □

Proof of Proposition 3.1. Write $L_{\beta,t} = d_{\beta,t}/d_{\beta,t-1}$ for the tth factor in (2). Predictability means that, conditional on $\mathcal{F}_{t-1}$, both the selected item j_t and the reference $\theta_{\beta,t}$ are fixed. The only new randomness is then

$$Y_t \mid \mathcal{F}_{t-1} \sim \mathrm{Bern}(p_{j_t}(\theta)).$$

Apply Lemma B.1 conditionally with $p = p_{j_t}(\theta)$ and $p_c = p_{j_t}(\theta_c)$. For the lower reference $\theta_{0,t} < \theta_c$, the function f in that lemma is decreasing. If $\theta \geq \theta_c$, monotonicity of the 2PL response curve gives $p \geq p_c$, and consequently

$$\mathbb{E}_\theta[L_{0,t} \mid \mathcal{F}_{t-1}] = f(p) \leq f(p_c) = 1.$$

For the upper reference $\theta_{1,t} > \theta_c$, f is increasing. If $\theta \leq \theta_c$, then $p \leq p_c$, and the same reasoning gives $\mathbb{E}_\theta[L_{1,t} \mid \mathcal{F}_{t-1}] \leq 1$.

The past evidence-process value $d_{\beta,t-1}$ is nonnegative and $\mathcal{F}_{t-1}$-measurable. Therefore

$$\mathbb{E}_\theta[d_{0,t} \mid \mathcal{F}_{t-1}] = d_{0,t-1}\,\mathbb{E}_\theta[L_{0,t} \mid \mathcal{F}_{t-1}] \leq d_{0,t-1} \qquad (\theta \geq \theta_c),$$

and analogously $\mathbb{E}_\theta[d_{1,t} \mid \mathcal{F}_{t-1}] \leq d_{1,t-1}$ for $\theta \leq \theta_c$. Thus the two processes are nonnegative supermartingales, each starting at one.

Ville's maximal inequality for a nonnegative supermartingale (E_t) with $E_0 = 1$ states that $\Pr(\sup_{t \geq 0} E_t \geq 1/\alpha) \leq \alpha$ (Ville, 1939). Applying it first to $d_{0,t}$ and then to $d_{1,t}$ yields exactly the two displayed crossing bounds in the proposition. Since Algorithm 1 can return H_0 only after $d_{0,t}$ crosses $1/\alpha_0$, and can return H_1 only after $d_{1,t}$ crosses $1/\alpha_1$, these crossing bounds imply (1). Finally, at $\theta = \theta_c$ both bounds apply. The event $\{\tau < \infty\}$ is contained in the union of the two crossing events, so the union bound gives $\Pr_{\theta_c}(\tau < \infty) \leq \alpha_0 + \alpha_1$. □

B.2 Stopping-time bounds for the shrinking rule

Lemma B.2 (Centered logistic separation)**.** *For $x \geq 0$,*

$$\sigma(x) - \frac{1}{2} = \frac{1}{2}\tanh(x/2) \geq \frac{x}{4 + 2x}.$$

Consequently, if $b_j = \theta_c$, then

$$|p_j(\theta) - p_j(\theta_c)| \geq \frac{a_j|\theta - \theta_c|}{4 + 2a_j|\theta - \theta_c|}.$$

Proof. The identity follows by writing

$$\sigma(x) = \frac{e^x}{1 + e^x} = \frac{1}{2}\left(1 + \frac{e^x - 1}{e^x + 1}\right) = \frac{1}{2}(1 + \tanh(x/2)).$$

For $u \geq 0$, all denominators in the following calculation are positive, and

$$\tanh u - \frac{u}{1 + u} = \frac{e^{2u} - 1 - 2u}{(e^{2u} + 1)(1 + u)} \geq 0,$$

where the last inequality is the elementary exponential bound $e^z \geq 1 + z$ with $z = 2u$. Taking $u = x/2$ proves

$$\sigma(x) - \frac{1}{2} = \frac{1}{2}\tanh(x/2) \geq \frac{x}{4 + 2x}.$$

If $b_j = \theta_c$, then $p_j(\theta_c) = \sigma(0) = 1/2$. With $x = a_j|\theta - \theta_c|$, the symmetry $\sigma(-x) = 1 - \sigma(x)$ shows that the absolute separation on either side of the cutoff equals $\sigma(x) - 1/2$. Substitution gives the second claim. □

We use the following scalar form of Freedman's inequality (Freedman, 1975, Theorem 1.6).

Fact B.3 (Freedman's inequality)**.** *Let (S_t) be a martingale with differences $X_t \le R$ almost surely and predictable quadratic variation $V_t = \sum_{s\le t} \mathbb{E}[X_s^2 \mid \mathcal{F}_{s-1}]$. For $q, v > 0$,*

$$\Pr(\exists t : S_t \ge q \text{ and } V_t \le v) \le \exp\left[-\frac{q^2}{2(v + Rq/3)}\right].$$

Proposition B.4 (Detailed stopping-time bounds for the shrinking rule)**.** *Fix $\theta \ne \theta_c$ and let β be the correct index, so that $\operatorname{sign}(\theta - \theta_c) = (-1)^{\beta+1}$. Suppose every administered item is centered at the cutoff, $b_{j_t} = \theta_c$, and $a_{j_t} \ge a_{\min} > 0$. Put $\delta = |\theta - \theta_c|$ and $\gamma_\theta = a_{\min}\delta/(4 + 2a_{\min}\delta)$. For integers $t \ge 1$, define*

$$q_\beta(t) = 2\gamma_\theta\sqrt{\ell_\beta}(\sqrt{t+1} - 1) - \frac{\ell_\beta}{8}(1 + \log t) - \ell_\beta, \tag{14}$$

$$v_\beta(t) = \frac{\ell_\beta}{4}(1 + \log t). \tag{15}$$

For every t such that $q_\beta(t) > 0$, Algorithm 1 satisfies the variance-sensitive bound

$$\Pr_\theta(\tau \ge t + 1) \le \exp\left[-\frac{q_\beta(t)^2}{2(v_\beta(t) + \sqrt{\ell_\beta}\, q_\beta(t)/3)}\right]. \tag{16}$$

Consequently, for $\eta \in (0, 1)$, let

$$N_\beta(\eta) = \min\left\{t \ge 1 : q_\beta(t) > 0, \quad \frac{q_\beta(t)^2}{2(v_\beta(t) + \sqrt{\ell_\beta}\, q_\beta(t)/3)} \ge \log\frac{1}{\eta}\right\}. \tag{17}$$

Then $N_\beta(\eta)$ is finite and

$$\Pr_\theta(\tau \ge N_\beta(\eta) + 1) \le \eta. \tag{18}$$

For the explicit inversion used in the main text, put

$$H_\theta = \log\frac{64}{\gamma_\theta^2}, \qquad C_\beta = 4\sqrt{\ell_\beta} + \ell_\beta\left(1 + \frac{H_\theta}{8}\right),$$
$$\mathcal{A}_\beta = \left\lceil\left(\frac{8C_\beta}{5\gamma_\theta\sqrt{\ell_\beta}}\right)^2\right\rceil, \qquad D_\theta = \frac{H_\theta/2 + 16/3}{\gamma_\theta^2}. \tag{19}$$

Then (4) holds for every integer $t \ge \max(2, \mathcal{A}_\beta)$. Define the explicit horizon

$$\widehat{N}_\beta(\eta) = \left\lceil \max\left(2, \mathcal{A}_\beta, D_\theta \log\frac{1}{\eta}\right)\right\rceil. \tag{20}$$

It satisfies

$$\Pr_\theta(\tau \ge \widehat{N}_\beta(\eta) + 1) \le \eta. \tag{21}$$

For the mean bound, set

$$\omega_\beta = \frac{13}{32}\sqrt{\ell_\beta} + \frac{4}{3} + \frac{\sqrt{5}}{2}, \qquad M_\beta = \frac{\omega_\beta}{2(1 - e^{-2})\gamma_\theta}, \tag{22}$$

$$\kappa(M) = \begin{cases} e^2, & M \le e/2, \\ \exp[-2W_{-1}(-1/(2M))], & M > e/2, \end{cases} \qquad T_\beta = \left\lceil \max(e^4, \kappa(M_\beta))\right\rceil. \tag{23}$$

Then

$$\Pr_\theta(\tau \ge k) \le \frac{1}{(k-1)^2} \qquad \text{for every integer } k \ge T_\beta + 1. \tag{24}$$

The constants above are deliberately conservative so that the confidence dependence is fully explicit.

We also verify the simpler order in (5). For fixed α_β, the quantity ℓ_β is a positive constant. As $a_{\min}\delta \downarrow 0$, we have

$$\gamma_\theta \asymp a_{\min}\delta, \qquad H_\theta \asymp L_\delta, \qquad C_\beta \asymp L_\delta.$$

Consequently, $\mathcal{A}_\beta \asymp (a_{\min}^2\delta^2)^{-1}L_\delta^2$ and $D_\theta \asymp (a_{\min}^2\delta^2)^{-1}L_\delta$. Substituting these two relations into (20), and using that the maximum and sum of two nonnegative quantities agree up to a factor of two, proves (5). The comparison constants are independent of η.

Interpretation. For the correct side of the cutoff, the cumulative log-likelihood is the sum of a predictable drift and a mean-zero martingale. Centering the administered items at the cutoff and imposing $a_{j_t} \geq a_{\min}$ make the cumulative drift at least $2\gamma_\theta \sqrt{\ell_\beta}(\sqrt{t+1}-1) - \ell_\beta(1+\log t)/8$. The martingale's predictable quadratic variation is at most $v_\beta(t)$. If the algorithm has not stopped by time t, that martingale must therefore make a downward deviation of at least $q_\beta(t)$. Freedman's martingale inequality (Freedman, 1975) gives (16); choosing the exponent to be at least $\log(1/\eta)$ gives (18). A blockwise application on the final fraction of the observations gives the explicit $D_\theta \log(1/\eta)$ horizon in (20). This is stronger than applying Markov's inequality to an expected stopping-time bound: it controls a chosen upper quantile of τ directly. Wald-type stopped-sum identities can yield mean stopping-time statements under suitable independence and integrability conditions (Wald, 1947), whereas the Freedman argument here handles the predictable, heterogeneous item sequence through conditional variances. Summing (24) gives Theorem 3.3; for fixed error levels its order is $O(\delta^{-2}\log^2(1/\delta))$ as $\delta \downarrow 0$.

The concentration step comes from Freedman's inequality. The additional work specific to the shrinking 2PL process is to identify its predictable drift and quadratic variation and then localize the argument to a terminal block. On this block the likelihood-ratio increments are uniformly of order $t^{-1/2}$, which produces the exponential tail in (4) and the logarithmic dependence on the requested tail probability in (20).

Proof of Proposition B.4. Let $s_\beta = (-1)^{\beta+1} = \operatorname{sign}(\theta - \theta_c)$ and $\ell = \ell_\beta$. At time s, abbreviate the true, cutoff, and reference response probabilities by

$$p_s = p_{j_s}(\theta), \qquad p_s^{(c)} = p_{j_s}(\theta_c), \qquad p_s^\star = p_{j_s}(\theta_c + \xi_s), \qquad \xi_s = s_\beta \frac{\sqrt{\ell}}{a_{j_s}\sqrt{s}}.$$

Define the log-likelihood increment and its conditional mean by

$$Z_s = Y_s \log \frac{p_s^\star}{p_s^{(c)}} + (1 - Y_s) \log \frac{1 - p_s^\star}{1 - p_s^{(c)}}, \tag{25}$$

$$m_s = \mathbb{E}_\theta[Z_s \mid \mathcal{F}_{s-1}]. \tag{26}$$

The cumulative sum of the increments Z_s is the cumulative log-likelihood evidence for the corresponding reference sequence. The log-odds difference between reference and cutoff is exactly

$$\log\left(\frac{p_s^\star}{1-p_s^\star}\frac{1-p_s^{(c)}}{p_s^{(c)}}\right) = a_{j_s}\xi_s = s_\beta\sqrt{\frac{\ell}{s}}. \tag{27}$$

For $g_s(x) = \log(1 - p_{j_s}(x))$,

$$g_s'(x) = -a_{j_s} p_{j_s}(x), \qquad g_s''(x) = -a_{j_s}^2 p_{j_s}(x)(1 - p_{j_s}(x)) \geq -\frac{a_{j_s}^2}{4}.$$

Taylor's theorem and (27) imply

$$\begin{aligned}
\log \frac{1-p_s^\star}{1-p_s^{(c)}} &= g_s(\theta_c + \xi_s) - g_s(\theta_c) \\
&\geq g_s'(\theta_c)\xi_s - \frac{a_{j_s}^2 \xi_s^2}{8} \\
&= -s_\beta p_s^{(c)} \sqrt{\frac{\ell}{s}} - \frac{\ell}{8s}.
\end{aligned}$$

The middle inequality is the second-order Taylor formula with a remainder at an intermediate point, followed by the uniform bound $g_s'' \geq -a_{j_s}^2/4$. Using $Z_s = s_\beta Y_s \sqrt{\ell/s} + \log((1-p_s^\star)/(1-p_s^{(c)}))$ and taking the conditional expectation in (25) therefore gives

$$\begin{aligned}
m_s &= s_\beta p_s \sqrt{\frac{\ell}{s}} + \log \frac{1-p_s^\star}{1-p_s^{(c)}} \\
&\geq s_\beta (p_s - p_s^{(c)}) \sqrt{\frac{\ell}{s}} - \frac{\ell}{8s} = |p_s - p_s^{(c)}| \sqrt{\frac{\ell}{s}} - \frac{\ell}{8s}.
\end{aligned} \tag{28}$$

Because $b_{j_s} = \theta_c$ and $a_{j_s} \geq a_{\min}$, Lemma B.2 gives $|p_s - p_s^{(c)}| \geq \gamma_\theta$: indeed, $x \mapsto x/(4+2x)$ is increasing on $[0,\infty)$, so its value at $a_{j_s}|\theta - \theta_c|$ is at least its value at $a_{\min}|\theta - \theta_c|$. Summing (28) and using the elementary integral comparisons

$$\sum_{s=1}^{t} \frac{1}{\sqrt{s}} \geq \int_0^t \frac{dx}{\sqrt{x+1}} = 2(\sqrt{t+1} - 1), \qquad \sum_{s=1}^{t} \frac{1}{s} \leq 1 + \int_1^t \frac{dx}{x} = 1 + \log t$$

yields

$$\sum_{s=1}^{t} m_s \geq 2\gamma_\theta \sqrt{\ell}(\sqrt{t+1} - 1) - \frac{\ell}{8}(1 + \log t) =: \psi(t). \tag{29}$$

Let $X_s = Z_s - m_s$ and $S_t = \sum_{s=1}^{t} X_s$. Since the random part of Z_s is its response Y_s times the log-odds difference in (27),

$$X_s = s_\beta \sqrt{\frac{\ell}{s}}(Y_s - p_s).$$

Thus (S_t) is a martingale. Moreover, $-X_s \leq \sqrt{\ell/s} \leq \sqrt{\ell}$ because $s_\beta(Y_s - p_s) \geq -1$. Its conditional variance is $\mathbb{E}[X_s^2 \mid \mathcal{F}_{s-1}] = (\ell/s)p_s(1-p_s)$; because $p_s(1-p_s) \leq 1/4$ and the harmonic-sum bound above applies,

$$V_t := \sum_{s=1}^{t} \mathbb{E}[X_s^2 \mid \mathcal{F}_{s-1}] = \sum_{s=1}^{t} \frac{\ell}{s} p_s(1-p_s) \leq \frac{\ell}{4}(1 + \log t). \tag{30}$$

Let $\tau_\beta = \inf\{t : d_{\beta,t} \geq 1/\alpha_\beta\}$. Since the algorithm stops when either side crosses, $\tau \leq \tau_\beta$. If $\tau > t$, neither side has crossed by time t, so in particular $\log d_{\beta,t} < \log(1/\alpha_\beta) = \ell$. Using $\log d_{\beta,t} = S_t + \sum_{s \leq t} m_s$ and (29),

$$-S_t > 2\gamma_\theta \sqrt{\ell}(\sqrt{t+1} - 1) - \frac{\ell}{8}(1 + \log t) - \ell = q_\beta(t). \tag{31}$$

Whenever $q_\beta(t) > 0$, Fact B.3, applied to the martingale $-S_t$ with increment bound $R = \sqrt{\ell}$ and the deterministic variance bound $v_\beta(t)$ from (15), gives

$$\Pr_\theta(\tau > t) \leq \Pr_\theta(-S_t > q_\beta(t)) \leq \exp\left[-\frac{q_\beta(t)^2}{2(v_\beta(t) + \sqrt{\ell}\, q_\beta(t)/3)}\right]. \tag{32}$$

The first event is contained in the second by (31); the last inequality is the fixed-time consequence of the maximal inequality in Fact B.3, because $V_t \leq v_\beta(t)$ deterministically. This proves (16). Substitution of the definition (17) proves (18) whenever the defining set is nonempty.

We next prove the explicit logarithmic dependence in (20). The preceding application used the largest increment over the entire history. That is convenient but wasteful: the increments at times of order t are only of order $t^{-1/2}$. We isolate such a terminal block. Since $0 < \gamma_\theta < 1/2$, put

$$\rho = \frac{\gamma_\theta^2}{64} \in (0,1), \qquad H_\theta = \log(1/\rho), \qquad m = \lceil \rho t \rceil.$$

For every $t \geq 2$, $1 \leq m < t$. Because $p_s^{(c)} = 1/2$, the increment in (25) has the exact representation

$$Z_s = \frac{\sqrt{\ell/s}}{2}\, s_\beta(2Y_s - 1) - \log\cosh\left(\frac{\sqrt{\ell/s}}{2}\right).$$

The elementary inequality $\log\cosh x \leq x$ for $x \geq 0$ therefore gives $Z_s \geq -\sqrt{\ell/s}$ and hence

$$\log d_{\beta,m} \geq -\sqrt{\ell} \sum_{s=1}^{m} s^{-1/2} \geq -2\sqrt{\ell m}. \tag{33}$$

On the terminal block, (28) and the integral bounds for the decreasing functions $x^{-1/2}$ and x^{-1} yield

$$\sum_{s=m+1}^{t} m_s \geq 2\gamma_\theta \sqrt{\ell}(\sqrt{t+1} - \sqrt{m+1}) - \frac{\ell}{8} \log \frac{t}{m}. \tag{34}$$

Specifically, $\sum_{s=m+1}^{t} s^{-1/2} \geq 2(\sqrt{t+1} - \sqrt{m+1})$ and $\sum_{s=m+1}^{t} s^{-1} \leq \log(t/m)$. If $\tau > t$, combine (33) and (34) with $\log d_{\beta,t} < \ell$. It follows that

$$-\sum_{s=m+1}^{t} X_s > B_t, \quad B_t := 2\gamma_\theta\sqrt{\ell}(\sqrt{t+1} - \sqrt{m+1}) - \frac{\ell}{8}\log\frac{t}{m} - \ell - 2\sqrt{\ell m}. \tag{35}$$

Now $m \geq \rho t$, $m \leq \rho t + 1$, and $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$. Consequently,

$$\begin{aligned} B_t &\geq \sqrt{\ell t}(2\gamma_\theta(1 - \sqrt{\rho}) - 2\sqrt{\rho}) - 2\sqrt{\ell}(1 + \sqrt{2}\,\gamma_\theta) - \ell\left(1 + \frac{H_\theta}{8}\right) \\ &\geq \frac{13}{8}\gamma_\theta\sqrt{\ell t} - C_\beta. \end{aligned}$$

Here we used $\sqrt{\rho} = \gamma_\theta/8$, $\gamma_\theta < 1/2$, and $2(1 + \sqrt{2}\,\gamma_\theta) < 4$; the resulting C_β is exactly the one in (19). Thus $t \geq \mathcal{A}_\beta$ implies

$$B_t \geq \gamma_\theta\sqrt{\ell t} =: q_t > 0. \tag{36}$$

Condition on $\mathcal{F}_m$ and apply Fact B.3 to the shifted martingale $-\sum_{s=m+1}^{m+r} X_s$. On this block its increments are bounded above by

$$R_t = \sqrt{\frac{\ell}{m+1}} \leq \frac{8\sqrt{\ell}}{\gamma_\theta\sqrt{t}},$$

and its predictable quadratic variation is at most

$$v_t \leq \frac{\ell}{4}\sum_{s=m+1}^{t}\frac{1}{s} \leq \frac{\ell}{4}\log\frac{t}{m} \leq \frac{\ell}{4}H_\theta.$$

Equations (35)– (36) and Freedman's inequality therefore give, for $t \geq \max(2, \mathcal{A}_\beta)$,

$$\begin{aligned} \Pr_\theta(\tau > t) &\leq \exp\left[-\frac{q_t^2}{2(v_t + R_t q_t/3)}\right] \\ &\leq \exp\left[-\frac{\gamma_\theta^2 t}{H_\theta/2 + 16/3}\right]. \end{aligned} \tag{37}$$

For the second line, substitute $q_t^2 = \gamma_\theta^2\ell t$, $v_t \leq \ell H_\theta/4$, and $R_t q_t/3 \leq 8\ell/3$; cancellation of ℓ gives the displayed exponent. This is (4), because $D_\theta = (H_\theta/2 + 16/3)/\gamma_\theta^2$. Taking $t = \widehat{N}_\beta(\eta)$ proves (21). Notice that the exponent in (37) is linear in t: restricting attention to a terminal block makes both the accumulated conditional variance and the bounded-increment correction independent of t after normalization.

Lemma C.1, with $a = \ell$ and $\gamma = \gamma_\theta$, states that for every $t \geq T_\beta$ the right-hand side is positive and Freedman's exponent is at least $2\log t$. Therefore, (32) gives

$$\Pr_\theta(\tau > t) \leq t^{-2}, \qquad t \geq T_\beta.$$

Replacing t by $k - 1$ proves (24). This also shows that the set in (17) is nonempty, and hence that $N_\beta(\eta)$ is finite. □

Proof of Proposition 3.2. Proposition B.4 gives (4). By the definition in (20), $\widehat{N}_\beta(\eta) \geq \max(2, \mathcal{A}_\beta)$ and $\widehat{N}_\beta(\eta)/D_\theta \geq \log(1/\eta)$. Substitution into (4) proves (21). □

Proof of Theorem 3.3. The two directional error guarantees are exactly those proved in Proposition 3.1; no stopping-time assumption is needed for that part. For the mean bound, because τ is positive and integer valued, $\tau = \sum_{k\geq 1}\mathbf{1}\{\tau \geq k\}$ pointwise; monotone convergence and Proposition B.4 therefore give

$$\mathbb{E}_\theta\,\tau = \sum_{k=1}^{\infty}\Pr_\theta(\tau \geq k) \leq T_\beta + \sum_{t=T_\beta}^{\infty} t^{-2} \leq T_\beta + \frac{\pi^2}{6}.$$

The first T_β probabilities were bounded by one, and for $k = T_\beta + 1, T_\beta + 2, \dots$ we used $\Pr(\tau \geq k) \leq (k-1)^{-2}$. Finally, $\sum_{t=T_\beta}^{\infty} t^{-2} \leq \sum_{t=1}^{\infty} t^{-2} = \pi^2/6$. The two inequalities in Lemma C.2 show that $\kappa(M)$ is of the same order as

$$M^2\big(\log(2M) + \log(2\log(2M))\big)^2.$$

The expression in parentheses is $\Theta(\log M)$ on $M \geq e/2$: its ratio to $\log M$ is continuous there, has a finite value at $e/2$, and converges to one as $M \to \infty$. Hence $\kappa(M) = \Theta(M^2 \log^2 M)$.

It remains to justify the equivalent order in (6). When the error levels are fixed, ω_β is a positive constant independent of δ. Moreover, as $a_{\min}\delta \downarrow 0$,

$$\gamma_\theta = \frac{a_{\min}\delta}{4 + 2a_{\min}\delta} \sim \frac{a_{\min}\delta}{4}, \qquad M_\beta = \Theta\left(\frac{1}{a_{\min}\delta}\right).$$

Therefore

$$\kappa(M_\beta) = \Theta\left(\frac{\log^2\big(1/(a_{\min}\delta)\big)}{a_{\min}^2\delta^2}\right).$$

As this quantity diverges, taking the maximum with e^4, applying the ceiling, and adding $\pi^2/6$ do not change its asymptotic order. This proves (6). □

B.3 Mixture validity and conservative finite computation

Proof of Proposition 4.1. For a fixed side β and component k, denote the new likelihood-ratio factor by

$$L_{\beta,k,t} := \frac{R_{\beta,k,t}}{R_{\beta,k,t-1}}.$$

Conditioning on $\mathcal{F}_{t-1}$ fixes j_t and the predictable reference $\theta_{\beta,k,t}$. Exactly as in the proof of Proposition 3.1, Lemma B.1 gives

$$\mathbb{E}_\theta[L_{0,k,t} \mid \mathcal{F}_{t-1}] \leq 1 \quad (\theta \geq \theta_c), \qquad \mathbb{E}_\theta[L_{1,k,t} \mid \mathcal{F}_{t-1}] \leq 1 \quad (\theta \leq \theta_c).$$

This remains true for the generalized references in the proposition because only their sidedness, not the value of u_k, is used. If a reference equals θ_c, its factor is exactly one. Since $R_{\beta,k,t-1}$ is nonnegative and $\mathcal{F}_{t-1}$-measurable,

$$\mathbb{E}_\theta[R_{\beta,k,t} \mid \mathcal{F}_{t-1}] = R_{\beta,k,t-1}\,\mathbb{E}_\theta[L_{\beta,k,t} \mid \mathcal{F}_{t-1}] \leq R_{\beta,k,t-1}$$

under the corresponding composite null. Thus every component is a nonnegative supermartingale starting at one.

For the finite partial mixture $P_{\beta,t}^{(K)} = \sum_{k=0}^{K} w_k R_{\beta,k,t}$, linearity of conditional expectation and $w_k \geq 0$ give

$$\mathbb{E}_\theta[P_{\beta,t}^{(K)} \mid \mathcal{F}_{t-1}] \leq P_{\beta,t-1}^{(K)}.$$

As $K \to \infty$, these partial sums increase pointwise to $M_{\beta,t}$. Conditional monotone convergence therefore yields

$$\mathbb{E}_\theta[M_{\beta,t} \mid \mathcal{F}_{t-1}] = \lim_{K\to\infty} \mathbb{E}_\theta[P_{\beta,t}^{(K)} \mid \mathcal{F}_{t-1}] \leq \lim_{K\to\infty} P_{\beta,t-1}^{(K)} = M_{\beta,t-1}.$$

Moreover, $M_{\beta,0} = \sum_{k\geq 0} w_k = 1$. Hence the countable mixtures are nonnegative supermartingales starting at one.

Ville's maximal inequality now gives $\Pr_\theta(\sup_t M_{0,t} \geq 1/\alpha_0) \leq \alpha_0$ for $\theta \geq \theta_c$, and the analogous upper-side bound for $\theta \leq \theta_c$ (Ville, 1939). These are the directional guarantees in (1). Since all omitted terms are nonnegative, (11) holds. Therefore Algorithm 2 stops only when the corresponding full mixture has crossed, so the same guarantees hold for the implementation. At θ_c both supermartingale bounds apply, and a union bound gives the same bound $\alpha_0 + \alpha_1$ on the probability that either rule stops. □

B.4 Predictable-design concentration for the mixture

We now state a stopping result that permits heterogeneous items and adaptive item selection. Fix $\theta \neq \theta_c$, let β index its correct side, and fix a correctly sided component r of the generalized mixture in Proposition 4.1. At time t, write

$$p_t = p_{j_t}(\theta), \qquad p_t^{(c)} = p_{j_t}(\theta_c), \qquad p_{r,t}^\star = p_{j_t}(\theta_{\beta,r,t}),$$

and define

$$m_{r,t} = p_t \log \frac{p_{r,t}^\star}{p_t^{(c)}} + (1-p_t)\log \frac{1-p_{r,t}^\star}{1-p_t^{(c)}} = \mathrm{kl}(p_t, p_t^{(c)}) - \mathrm{kl}(p_t, p_{r,t}^\star), \tag{38}$$

$$\Delta_{r,t} = \mathrm{logit}(p_{r,t}^\star) - \mathrm{logit}(p_t^{(c)}) = a_{j_t}(\theta_{\beta,r,t} - \theta_c).$$

Set

$$G_{r,n} = \sum_{t=1}^n m_{r,t}, \qquad Q_{r,n} = \sum_{t=1}^n \Delta_{r,t}^2. \tag{39}$$

Here $G_{r,n}$ is the component's cumulative conditional log-likelihood advantage over the cutoff, while $Q_{r,n}$ controls the scale of its random fluctuations. Both may be random because item and reference selection may depend on earlier responses.

Theorem B.5 (Predictable-design mixture stopping bound)**.** *Let $h_{\beta,r} = \log(1/(\alpha_\beta w_r))$. Under any predictable item-selection rule and correctly sided predictable reference-ability sequences, the countable-mixture stopping time satisfies, for every $n \geq 1$ and $x, v > 0$,*

$$\Pr_\theta(\tau \geq n+1) \leq \Pr_\theta(G_{r,n} < h_{\beta,r} + x \text{ or } Q_{r,n} > v) + \exp\left(-\frac{2x^2}{v}\right). \tag{40}$$

Equivalently,

$$\Pr_\theta(\tau < \infty) \geq 1 - \Pr_\theta(G_{r,n} < h_{\beta,r} + x \text{ or } Q_{r,n} > v) - \exp\left(-\frac{2x^2}{v}\right). \tag{41}$$

In particular, if $G_{r,n} \geq g_n$ and $Q_{r,n} \leq v_n$ almost surely for deterministic $g_n > h_{\beta,r}$ and $v_n > 0$, then

$$\Pr_\theta(\tau \geq n+1) \leq \exp\left(-\frac{2(g_n - h_{\beta,r})^2}{v_n}\right). \tag{42}$$

Consequently, the mixture stops almost surely if this exponent diverges along some sequence of horizons. The same bounds hold for Algorithm 2 whenever component r has been activated by time n.

The theorem allows discriminations, difficulties, selected items, and reference abilities all to vary predictably. No design-free finite-stopping guarantee is possible when a policy may select uninformative items or references. The result instead separates insufficient accumulated information, represented by $G_{r,n}$, from an adverse response fluctuation controlled by $Q_{r,n}$.

Corollary B.6 (Uniform information accumulation)**.** *In the setting of Theorem B.5, suppose one correct-side component r satisfies, almost surely for every t, $m_{r,t} \geq \mu > 0$ and $\Delta_{r,t}^2 \leq d^2 < \infty$. Then, for every $n > h_{\beta,r}/\mu$,*

$$\Pr_\theta(\tau \geq n+1) \leq \exp\left(-\frac{2(n\mu - h_{\beta,r})^2}{nd^2}\right).$$

Thus the mixture stops almost surely with an exponentially decreasing tail. The condition permits heterogeneous, adaptively selected items: it requires only that one mixture scale retain a uniform per-item likelihood advantage and bounded standardized log-odds displacement.

Proof of Theorem B.5. Fix the correct side β and component r, and let $Z_{r,t}$ denote its tth log-likelihood-ratio increment. The Bernoulli exponential-family form is $p(y;z) = \exp(yz - A(z))$ for $y \in \{0,1\}$ and $A(z) = \log(1+e^z)$. The cutoff log odds are $z_t^{(c)} = \mathrm{logit}(p_t^{(c)})$, while the reference log odds are $z_t^{(c)} + \Delta_{r,t}$. Taking their log-likelihood difference gives

$$Z_{r,t} = \Delta_{r,t} Y_t - (A(\mathrm{logit}(p_t^{(c)}) + \Delta_{r,t}) - A(\mathrm{logit}(p_t^{(c)}))).$$

Because $Y_t \mid \mathcal{F}_{t-1} \sim \mathrm{Bern}(p_t)$ and $\Delta_{r,t}$ is predictable, taking the conditional expectation gives $m_{r,t}$ in (38). Equivalently, direct expansion of the two binary KL divergences shows

$$m_{r,t} = \mathrm{kl}(p_t, p_t^{(c)}) - \mathrm{kl}(p_t, p_{r,t}^{\star}).$$

Hence, with $X_{r,t} = Z_{r,t} - m_{r,t}$,

$$X_{r,t} = \Delta_{r,t}(Y_t - p_t), \qquad \log R_{\beta,r,n} = G_{r,n} + \sum_{t=1}^{n} X_{r,t}. \tag{43}$$

Conditionally on $\mathcal{F}_{t-1}$, $X_{r,t}$ has mean zero and its two possible values lie in an interval of length $|\Delta_{r,t}|$. The conditional form of Hoeffding's lemma (Hoeffding, 1963, Lemma 1) therefore shows, for every $\lambda > 0$, that

$$\mathbb{E}_\theta\left[\exp\left(-\lambda X_{r,t} - \frac{\lambda^2 \Delta_{r,t}^2}{8}\right)\middle| \mathcal{F}_{t-1}\right] \le 1.$$

Iterating this one-step conditional inequality proves that

$$\exp\left(-\lambda \sum_{t=1}^{n} X_{r,t} - \frac{\lambda^2}{8} Q_{r,n}\right), \qquad n \ge 0,$$

is a nonnegative supermartingale starting at one. Markov's inequality then implies, for $x, v > 0$,

$$\Pr_\theta\left(-\sum_{t=1}^{n} X_{r,t} > x,\ Q_{r,n} \le v\right) \le \exp\left(-\lambda x + \frac{\lambda^2 v}{8}\right).$$

Indeed, on the event on the left, the displayed exponential supermartingale is larger than $\exp(\lambda x - \lambda^2 v/8)$; its expectation is at most one. The quadratic exponent is minimized at $\lambda = 4x/v$, which gives $-\lambda x + \lambda^2 v/8 = -2x^2/v$.

If $\tau > n$, then the correct-side mixture has not crossed, so $M_{\beta,n} < 1/\alpha_\beta$. Because $M_{\beta,n} \ge w_r R_{\beta,r,n}$,

$$\log R_{\beta,r,n} < \log \frac{1}{\alpha_\beta w_r} = h_{\beta,r}.$$

On the additional event $\{G_{r,n} \ge h_{\beta,r} + x,\ Q_{r,n} \le v\}$, (43) thus forces $-\sum_{t \le n} X_{r,t} > x$. Splitting according to whether that additional event holds gives the set inclusion

$$\begin{aligned}\{\tau > n\} \subseteq \{&G_{r,n} < h_{\beta,r} + x \text{ or } Q_{r,n} > v\} \\ &\cup \left\{-\sum_{t=1}^{n} X_{r,t} > x,\ Q_{r,n} \le v\right\}.\end{aligned}$$

Taking probabilities, applying the union bound, and using the concentration inequality just proved yields (40). Also $\{\tau \le n\} \subseteq \{\tau < \infty\}$, so subtracting the same upper bound from one proves (41).

Under the deterministic assumptions, set $x = g_n - h_{\beta,r} > 0$ and $v = v_n$. The first event on the right-hand side of (40) is then empty, which proves (42). If its exponent diverges along an increasing sequence n_j, the right-hand side tends to zero. Since $\{\tau = \infty\} \subseteq \{\tau > n_j\}$ for every j,

$$\Pr_\theta(\tau = \infty) \le \inf_j \Pr_\theta(\tau > n_j) = 0.$$

Thus $\Pr_\theta(\tau < \infty) = 1$.

For Algorithm 2, once component r is active, its term $w_r R_{\beta,r,n}$ is reconstructed from all observations and included in the explicit partial sum. If that term reaches $1/\alpha_\beta$, then the partial sum is at least $1/\alpha_\beta$. The implementation must therefore stop by time n. Consequently, its event $\{\tau > n\}$ implies the same component inequality used above, and the proof applies without change. □

Proof of Corollary B.6. Summing the pointwise assumptions over $t = 1, \dots, n$ gives

$$G_{r,n} = \sum_{t=1}^{n} m_{r,t} \ge n\mu, \qquad Q_{r,n} = \sum_{t=1}^{n} \Delta_{r,t}^2 \le nd^2$$

almost surely. When $n > h_{\beta,r}/\mu$, the quantity $x = n\mu - h_{\beta,r}$ is positive. Substitute $g_n = n\mu$ and $v_n = nd^2$ into (42) to obtain the displayed inequality in the corollary. Finally,

$$\frac{2(n\mu - h_{\beta,r})^2}{nd^2} = \frac{2\mu^2}{d^2}n - \frac{4\mu h_{\beta,r}}{d^2} + \frac{2h_{\beta,r}^2}{nd^2} \longrightarrow \infty.$$

The almost-sure stopping conclusion now follows from the last part of Theorem B.5. □

B.5 Mean stopping time under a repeated centered item

For the mixture mean, write $\lambda = a(\theta - \theta_c)$ and $X_t = 2Y_t - 1$. Under repeated administration of a centered item,

$$\Pr_\lambda(X_t = x) = \frac{\exp(\lambda x/2)}{2\cosh(\lambda/2)}, \qquad \mathbb{E}_\lambda X_t = \tanh(\lambda/2). \tag{44}$$

To verify (44), note that $\Pr_\lambda(Y_t = 1) = \sigma(\lambda)$ and $\Pr_\lambda(Y_t = 0) = \sigma(-\lambda)$. For $x \in \{-1, 1\}$, both cases can be written as

$$\Pr_\lambda(X_t = x) = \frac{e^{\lambda x/2}}{e^{\lambda/2} + e^{-\lambda/2}} = \frac{e^{\lambda x/2}}{2\cosh(\lambda/2)}.$$

Subtracting the probabilities of $X_t = 1$ and $X_t = -1$ gives the stated expectation. Dividing the joint law with standardized parameter v by the joint law at the cutoff $v = 0$ gives the likelihood ratio

$$R_n(v) = \exp\left(\frac{v}{2}\sum_{t=1}^{n} X_t - n\log\cosh(v/2)\right). \tag{45}$$

Lemma B.7 (Positive drift at a nearby scale)**.** *If $0 < \lambda \le 1$ and $\lambda/2 \le v \le \lambda$, then*

$$\mathbb{E}_\lambda \log \frac{R_n(v)}{R_{n-1}(v)} = \frac{v}{2}\tanh(\lambda/2) - \log\cosh(v/2) \ge \frac{\lambda^2}{16}.$$

Proof. We first justify the two elementary inequalities used below. For $x \ge 0$, $\tanh x \le x$ because the derivative of $x - \tanh x$ is $\tanh^2 x \ge 0$ and the difference is zero at the origin. Consequently,

$$\frac{d}{dx}\left(\tanh x - x + \frac{x^3}{3}\right) = x^2 - \tanh^2 x \ge 0,$$

which proves $\tanh x \ge x - x^3/3$. Similarly, $d(x^2/2 - \log\cosh x)/dx = x - \tanh x \ge 0$, so $\log\cosh x \le x^2/2$.

The logarithm of the one-step ratio in (45) is $vX_n/2 - \log\cosh(v/2)$. By (44), its expectation is $v\tanh(\lambda/2)/2 - \log\cosh(v/2)$. Put $v = r\lambda$ with $r \in [1/2, 1]$. Applying the two inequalities just proved at $\lambda/2$ and $v/2$ gives

$$\frac{v}{2}\tanh(\lambda/2) - \log\cosh(v/2) \ge \frac{v\lambda}{4} - \frac{v\lambda^3}{48} - \frac{v^2}{8} = \lambda^2\frac{r(2-r)}{8} - \frac{r\lambda^4}{48}.$$

On $r \in [1/2, 1]$, $r(2-r) \ge 3/4$. Moreover, $r\lambda^4 \le \lambda^2$ because $r \le 1$ and $0 < \lambda \le 1$. Therefore the last display is at least $3\lambda^2/32 - \lambda^2/48 = 7\lambda^2/96 \ge \lambda^2/16$, as required. □

Lemma B.8 (Dyadic scale and weight cost)**.** *For $0 < \lambda \le 1$, let $k = \lceil \log_2(1/\lambda) \rceil$. Then $\lambda/2 < u_k \le \lambda$ and*

$$\log\frac{1}{w_k} < \log\frac{\pi^2}{6} + 2\log(2 + \log_2(1/\lambda)).$$

Proof. Write $x = \log_2(1/\lambda) \ge 0$. The definition $k = \lceil x \rceil$ gives $k - 1 < x \le k$. Exponentiating these inequalities with the decreasing map $y \mapsto 2^{-y}$ yields

$$\frac{\lambda}{2} = 2^{-(x+1)} < 2^{-k} = u_k \le 2^{-x} = \lambda.$$

From the definition of w_k,

$$\log\frac{1}{w_k} = \log\frac{\pi^2}{6} + 2\log(k+1).$$

Because $k < x + 1$, we have $k + 1 < x + 2$; substituting $x = \log_2(1/\lambda)$ proves the strict upper bound. □

Lemma B.9 (Expected crossing time). *Let $Z_1, Z_2, \ldots$ be integrable and i.i.d. with $m = \mathbb{E}\, Z_1 > 0$ and $Z_1 \le c$ almost surely. For $b > 0$ and $T_b = \inf\{n \ge 1 : \sum_{t=1}^{n} Z_t \ge b\}$, $\mathbb{E}\, T_b \le (b + c)/m$.*

Proof. Let $S_n = \sum_{t=1}^{n} Z_t$ and $T_b^{(N)} = T_b \wedge N$. This is a bounded stopping time, so Wald's identity for bounded stopping times (Wald, 1947) gives

$$\mathbb{E}\, S_{T_b^{(N)}} = m\, \mathbb{E}\, T_b^{(N)}.$$

The stopped sum is always smaller than $b + c$. Indeed, if $T_b > N$, then $S_N < b$ by the definition of the first crossing time. If $T_b \le N$, then $S_{T_b - 1} < b$ and $Z_{T_b} \le c$, so $S_{T_b} < b + c$. Consequently,

$$m\, \mathbb{E}\, T_b^{(N)} = \mathbb{E}\, S_{T_b^{(N)}} \le b + c.$$

As $N \to \infty$, $T_b^{(N)} \uparrow T_b$. Monotone convergence gives $m\, \mathbb{E}\, T_b \le b + c$, which proves the claim and simultaneously shows that $T_b < \infty$ almost surely. □

Proof of Theorem 4.2. **Validity.** Proposition 4.1 proves validity of the countable mixture. It also proves validity of Algorithm 2, because the algorithm stops only when a partial sum, and hence the full mixture, crosses.

Implemented rule for a small positive gap. Suppose $0 < \lambda \le 1$ and choose k as in Lemma B.8. Define

$$T_k = \inf\{n \ge 1 : \log R_n(u_k) \ge h_{1,k}\}, \qquad h_{1,k} = \log \frac{1}{\alpha_1 w_k}.$$

We first show that component k cannot reach this threshold while inactive. If $k > K_n$, then (45), $X_t \le 1$, and $K_n = 4 + \lceil \log_2(n+1) \rceil$ give

$$\log R_n(u_k) \le \frac{n u_k}{2} \le \frac{n}{2^{K_n + 2}} = \frac{n}{2^{6 + \lceil \log_2(n+1) \rceil}} \le \frac{n}{64(n+1)} < \frac{1}{64}.$$

In contrast, $\alpha_1 < 1$ and $w_k \le w_0 = 6/\pi^2$, so

$$h_{1,k} = \log \frac{1}{\alpha_1 w_k} > \log \frac{\pi^2}{6} > \frac{1}{64}.$$

Thus, on $\{T_k < \infty\}$, component k is active at time T_k and is included in the computed partial sum. Consequently,

$$P_{1,T_k}^{(K_{T_k})} \ge w_k R_{T_k}(u_k) \ge \frac{1}{\alpha_1},$$

so Algorithm 2 stops no later than T_k: $\tau_{\mathrm{tr}} \le T_k$. Moreover, $\tau_\infty \le \tau_{\mathrm{tr}}$ by (11). By (45), the log increments

$$Z_t = \log \frac{R_t(u_k)}{R_{t-1}(u_k)} = \frac{u_k}{2} X_t - \log \cosh(u_k/2)$$

are i.i.d. Lemmas B.8 and B.7 give $\mathbb{E}_\lambda\, Z_t \ge \lambda^2/16$. Furthermore, $Z_t \le u_k/2 \le 1/2$, because $X_t \le 1$ and $\log \cosh(u_k/2) \ge 0$. Lemma B.9, with $b = h_{1,k}$, $c = 1/2$, and $m \ge \lambda^2/16$, therefore yields

$$\mathbb{E}_\lambda\, \tau_{\mathrm{tr}} \le \mathbb{E}_\lambda\, T_k \le \frac{16}{\lambda^2} \left(\log \frac{1}{\alpha_1} + \log \frac{1}{w_k} + \frac{1}{2} \right).$$

Substitution of the weight bound in Lemma B.8 proves (13). If $-1 \le \lambda < 0$, then under parameter λ the variable $-X_t$ has the same distribution as X_t under parameter $-\lambda > 0$, while the correct lower-side component uses displacement $-u_k$. Repeating the argument with $|\lambda|$ and α_0 proves the negative-gap case.

Gaps larger than one. If $\lambda > 1$, component $u_0 = 1$ has mean log increment

$$\frac{1}{2} \tanh(\lambda/2) - \log \cosh(1/2) \ge \frac{1}{2} \tanh(1/2) - \log \cosh(1/2) \ge \frac{5}{48} > 0.$$

The first inequality uses monotonicity of $\tanh$; the second follows from $\tanh x \ge x - x^3/3$ and $\log \cosh x \le x^2/2$ at $x = 1/2$. The increments are bounded above by $1/2$, so Lemma B.9 gives

a finite expected time for component 0 to reach $h_{1,0}$. This component is active from the start of Algorithm 2; at its crossing time the computed partial sum is at least $w_0 R_n(1) \geq 1/\alpha_1$. Hence both τ_{tr} and τ_∞ have finite mean when $\lambda > 1$. The transformation $X_t \mapsto -X_t$ gives the same conclusion for $\lambda < -1$.

□

B.6 Information lower bounds for adaptive item selection

The next result applies to every stopping rule and every predictable item policy, not only to Algorithms 1 and 2. It measures sample size in the natural currency for adaptive experiments: the information supplied by the items that the algorithm actually selects.

Proposition B.10 (Adaptive stopped-information lower bound)**.** *Consider any possibly randomized predictable item-selection policy and any sequential decision rule satisfying (1), where $\alpha_0 + \alpha_1 < 1$. Define*

$$K_j(\theta, \theta_c) = \mathrm{kl}(p_j(\theta), p_j(\theta_c)).$$

If $\theta > \theta_c$ and $\mathbb{E}_\theta \tau < \infty$, then

$$\mathbb{E}_\theta \left[\sum_{t=1}^{\tau} K_{j_t}(\theta, \theta_c) \right] \geq \mathrm{kl}(1 - \alpha_0, \alpha_1). \tag{46}$$

If $\theta < \theta_c$ and $\mathbb{E}_\theta \tau < \infty$, then

$$\mathbb{E}_\theta \left[\sum_{t=1}^{\tau} K_{j_t}(\theta, \theta_c) \right] \geq \mathrm{kl}(1 - \alpha_1, \alpha_0). \tag{47}$$

Proof. We prove (46); the other side is identical after interchanging the decision labels. Let P_θ and P_c denote the laws under θ and θ_c. Finite mean implies $P_\theta(\tau < \infty) = 1$. Therefore (1) gives

$$P_\theta(D = H_1) \geq 1 - \alpha_0, \qquad P_c(D = H_1) \leq \alpha_1.$$

For $\tau_n = \tau \wedge n$, the chain rule for relative entropy under predictable experimental design gives

$$\mathrm{KL}(P_\theta^{\tau_n}, P_c^{\tau_n}) = \mathbb{E}_\theta \sum_{t=1}^{\tau_n} K_{j_t}(\theta, \theta_c). \tag{48}$$

Indeed, conditional on the past, the selected item is already known. For a randomized policy, it is drawn from the same parameter-free kernel under both laws. Item selection consequently contributes zero conditional KL, whereas the response contributes exactly $\mathrm{kl}(p_{j_t}(\theta), p_{j_t}(\theta_c))$.

Letting $n \to \infty$ in (48), monotone convergence on the right and the increasing-σ-field identity for relative entropy on the left yield

$$\mathrm{KL}(P_\theta^{\tau}, P_c^{\tau}) = \mathbb{E}_\theta \sum_{t=1}^{\tau} K_{j_t}(\theta, \theta_c).$$

Equivalently, this follows by applying the stopped-experiment change-of-measure lemma to τ_n and passing to the limit (Kaufmann et al., 2016, Lemma 1).

Data processing from the stopped transcript to the event $\{D = H_1\}$ now gives

$$\mathrm{KL}(P_\theta^{\tau}, P_c^{\tau}) \geq \mathrm{kl}(P_\theta(D = H_1), P_c(D = H_1)) \geq \mathrm{kl}(1 - \alpha_0, \alpha_1).$$

For the last inequality, binary relative entropy is increasing in its first argument and decreasing in its second whenever the first exceeds the second: the two derivatives are $\log(p(1-q)/(q(1-p))) > 0$ and $(q-p)/(q(1-q)) < 0$. This proves (46). Below the cutoff, use the event $\{D = H_0\}$ to obtain (47). □

The information inequality becomes a stopping-time inequality as soon as one upper-bounds how much information any feasible item sequence can accumulate.

Corollary B.11 (Information-envelope lower bound)**.** *Let*

$$h_\theta = \begin{cases} \mathrm{kl}(1-\alpha_0, \alpha_1), & \theta > \theta_c, \\ \mathrm{kl}(1-\alpha_1, \alpha_0), & \theta < \theta_c. \end{cases}$$

Suppose a nondecreasing concave function $\Gamma_\theta : [0,\infty) \to [0,\infty)$ satisfies, pathwise, for every feasible item sequence and every integer n,

$$\sum_{t=1}^{n} K_{j_t}(\theta, \theta_c) \le \Gamma_\theta(n).$$

Then every rule in Proposition B.10 satisfies

$$\mathbb{E}_\theta\, \tau \ge \Gamma_\theta^{-1}(h_\theta), \qquad \Gamma_\theta^{-1}(h) := \inf\{x \ge 0 : \Gamma_\theta(x) \ge h\}. \tag{49}$$

If the set defining the inverse is empty, no finite-mean rule with the stated error guarantees exists under that design constraint.

Proof. Proposition B.10, the pathwise envelope, and Jensen's inequality for the concave function Γ_θ give

$$h_\theta \le \mathbb{E}_\theta \sum_{t=1}^{\tau} K_{j_t}(\theta, \theta_c) \le \mathbb{E}_\theta\, \Gamma_\theta(\tau) \le \Gamma_\theta(\mathbb{E}_\theta\, \tau).$$

Taking the generalized inverse proves the claim. □

Two choices of the envelope make the result concrete. If items may be reused, put

$$K^\star(\theta, \theta_c) = \sup_j K_j(\theta, \theta_c).$$

Then $\Gamma_\theta(x) = xK^\star(\theta, \theta_c)$ and

$$\mathbb{E}_\theta\, \tau \ge \frac{h_\theta}{K^\star(\theta, \theta_c)}. \tag{50}$$

This already grants the algorithm an oracle that always chooses the most informative item for the true, unknown θ.

For a finite bank used without replacement, let $K_{(1)}(\theta, \theta_c) \ge \cdots \ge K_{(N)}(\theta, \theta_c)$ denote the nonincreasing rearrangement of $\{K_j(\theta, \theta_c) : 1 \le j \le N\}$. Define $C_\theta(0) = 0$, $C_\theta(m) = \sum_{r=1}^{m} K_{(r)}(\theta, \theta_c)$ for integer m, and interpolate linearly between integers; for $x \ge N$, set $C_\theta(x) = C_\theta(N)$. Its slopes are nonincreasing, so C_θ is concave. Any rule that uses each item at most once (and receives no further responses after exhausting the bank) satisfies

$$\sum_{t=1}^{n} K_{j_t}(\theta, \theta_c) \le C_\theta(n).$$

Consequently,

$$\mathbb{E}_\theta\, \tau \ge C_\theta^{-1}(h_\theta). \tag{51}$$

If $C_\theta(N) < h_\theta$, the bank does not contain enough information for any certified procedure that decides almost surely at θ with finite mean to meet the two directional guarantees. This does not preclude abstention or an uncertified terminal classifier after bank exhaustion.

For the 2PL model the item information is explicit. Put $z_j = a_j(\theta_c - b_j)$, let $s \in \{-1, 1\}$, and write $\theta = \theta_c + s\delta$. With $A(z) = \log(1 + e^z)$,

$$K_j(\theta_c + s\delta, \theta_c) = sa_j\delta A'(z_j + sa_j\delta) - A(z_j + sa_j\delta) + A(z_j), \tag{52}$$

$$\le \frac{a_j^2\delta^2}{8}. \tag{53}$$

To verify (53), differentiate the first line with respect to $x = sa_j\delta$ and integrate; the derivative is $xA''(z_j + x)$ and $A''(z) \le 1/4$. Hence, if every available item has $a_j \le a_{\max}$, every adaptive policy obeys the fully nonasymptotic bound

$$\mathbb{E}_{\theta_c + s\delta}\, \tau \ge \frac{8h_{\theta_c + s\delta}}{a_{\max}^2\delta^2}. \tag{54}$$

The sharper bank-specific bounds (50) and (51) should be preferred when the calibrated item parameters are known. Locally,

$$K_j(\theta_c + s\delta, \theta_c) = \frac{1}{2} I_j(\theta_c)\delta^2 + O(\delta^3), \qquad I_j(\theta_c) = a_j^2 p_j(\theta_c)(1 - p_j(\theta_c)),$$

which explains why high-Fisher-information items reduce the unavoidable sample size.

B.7 Farrell's worst-scale lower bound

The preceding results hold at each fixed gap, but they do not charge for using one rule that stops almost surely at every unknown nonzero gap. Farrell's theorem supplies that adaptation cost in the repeated-item model. We state the external theorem before specializing it so that its assumptions and its necessarily $\limsup$ conclusion are explicit.

Theorem B.12 (Farrell's lower bound for almost-sure finite stopping)**.** *Let $X_1, X_2, \ldots$ be i.i.d. from a regular canonical one-parameter exponential family*

$$dP_\lambda(x) = \exp(\lambda x - \psi(\lambda))\, dP_0(x),$$

where $\mathbb{E}_0 X = 0$, $\mathrm{Var}_0(X) = \sigma^2 \in (0, \infty)$, and the moment generating function is finite near zero. Let a two-sided sequential rule satisfy $\Pr_\lambda(\tau < \infty) = 1$ at every $\lambda \neq 0$ and obey fixed directional error bounds. With $\mu_\lambda = \mathbb{E}_\lambda X$ and $q = \Pr_0(\tau = \infty)$, on each side

$$\limsup_{\lambda \to 0^\pm} \frac{\mu_\lambda^2\, \mathbb{E}_\lambda \tau}{2\sigma^2 \log|\log|\mu_\lambda||} \geq q. \tag{55}$$

This is Theorem 1 of Farrell (1964). Its proof is the deep external input; the remaining calculation verifies every specialization step needed for the 2PL claim in the main text.

Corollary B.13 (Farrell bound for the repeated centered 2PL item)**.** *Suppose the same item with $b = \theta_c$ and discrimination $a > 0$ is repeatedly administered. Every directionally valid rule that stops almost surely for each $\theta \neq \theta_c$ satisfies, on each side,*

$$\limsup_{\delta \downarrow 0} \frac{a^2\delta^2\, \mathbb{E}_{\theta_c \pm \delta}\, \tau}{\log\log(1/\delta)} \geq 8q \geq 8(1 - \alpha_0 - \alpha_1), \qquad q = \Pr_{\theta_c}(\tau = \infty). \tag{56}$$

Proof. Put $\lambda = a(\theta - \theta_c)$ and $X = Y - 1/2$. Relative to the cutoff law $Y \sim \mathrm{Bern}(1/2)$,

$$\frac{dP_\lambda}{dP_0}(Y) = \exp(\lambda X - \psi(\lambda)), \qquad \psi(\lambda) = \log\cosh(\lambda/2).$$

Thus the repeated centered 2PL model is a regular canonical exponential family with a finite moment generating function, and

$$\sigma^2 = \mathrm{Var}_0(X) = \frac{1}{4}, \qquad \mu_\lambda = \mathbb{E}_\lambda X = \frac{1}{2}\tanh(\lambda/2) = \frac{\lambda}{4} + O(\lambda^3).$$

Both Algorithms 1 and 2 have finite mean and therefore stop almost surely at every nonzero λ in this model: $\Pr_\lambda(\tau < \infty) = 1$. The corollary also applies to any other rule with this almost-sure finite-stopping property. The assumptions of Theorem B.12 are therefore satisfied. Moreover,

$$\frac{2\sigma^2}{\mu_\lambda^2} = \frac{1/2}{\lambda^2/16 + O(\lambda^4)} = \frac{8}{\lambda^2}(1 + O(\lambda^2)), \qquad \log|\log|\mu_\lambda|| \sim \log\log(1/|\lambda|).$$

Substitution in (55) gives

$$\limsup_{\lambda \to 0^\pm} \frac{\lambda^2\, \mathbb{E}_\lambda \tau}{\log\log(1/|\lambda|)} \geq 8q.$$

Since $|\lambda| = a\delta$ and $\log\log(1/(a\delta)) \sim \log\log(1/\delta)$, this proves the first inequality in (56). Finally, at the cutoff either finite decision is a directional error, so

$$1 - q = \Pr_{\theta_c}(\tau < \infty) \leq \alpha_0 + \alpha_1,$$

which proves the second inequality. □

Unlike Proposition B.10, Farrell's result is not a pointwise nonasymptotic lower bound: a valid multiscale rule may favor selected gaps and be unusually fast there. The $\limsup$ says that it cannot be that fast at every sufficiently small scale. Extending the sharp constant in (56) from repeated i.i.d. items to general adaptive item selection would require a controlled-experiment analogue of Farrell's theorem; the stopped-information bounds above are the nonasymptotic statements that already hold without such an extension.

C TECHNICAL CLOSURE LEMMAS

This section collects the deterministic inequalities used to convert the shrinking rule's concentration bound into the explicit horizons stated in the main text. We begin with the polynomial-tail inversion and then establish the auxiliary bounds needed in its proof.

Lemma C.1 (Polynomial-tail closure)**.** *Let $\gamma, a > 0$, and set*

$$R = \sqrt{a}, \qquad v_k = \frac{a}{4}(1 + \log k), \qquad \underline{q}_k = 2\gamma\sqrt{a}(\sqrt{k+1} - 1) - \frac{a}{8}(1 + \log k) - a.$$

Define

$$\omega = \frac{13}{32}\sqrt{a} + \frac{4}{3} + \frac{\sqrt{5}}{2}, \qquad M = \frac{\omega}{2(1 - e^{-2})\gamma},$$

and let $T = \lceil \max(e^4, \kappa(M)) \rceil$, with κ as in (23). For every integer $k \geq T$ and every $q_k \geq \underline{q}_k$, one has $q_k > 0$ and

$$\exp\left[-\frac{q_k^2}{2(v_k + Rq_k/3)}\right] \leq k^{-2}.$$

Proof. Put $h(x) = \sqrt{x}/\log x$ for $x > 1$. Since

$$h'(x) = \frac{\log x - 2}{2\sqrt{x}(\log x)^2},$$

h is increasing on $[e^2, \infty)$. We first prove

$$\sqrt{k} \geq M \log k. \tag{57}$$

If $M \leq e/2$, then $k \geq e^4$ and monotonicity give $h(k) \geq h(e^4) = e/2 \geq M$, which is exactly (57). If $M > e/2$, put $z = -W_{-1}(-1/(2M))$. The lower real branch satisfies $z > 1$ and the defining identity $(-z)e^{-z} = -1/(2M)$ gives

$$ze^{-z} = \frac{1}{2M}, \qquad e^z = 2Mz.$$

Consequently,

$$h(e^{2z}) = \frac{e^z}{2z} = M.$$

Here $e^{2z} = \kappa(M) > e^2$. Since $k \geq T \geq \kappa(M)$ and h is increasing on this range, $h(k) \geq M$, again proving (57).

Write $u = \log k \geq 4$. Since $k \geq e^4$,

$$\sqrt{k+1} - 1 \geq \sqrt{k} - 1 = \sqrt{k}(1 - k^{-1/2}) \geq (1 - e^{-2})\sqrt{k}.$$

Also

$$\frac{a}{8}(1 + u) + a = \frac{a}{8}(u + 9) \leq \frac{13a}{32}u,$$

where the last inequality is equivalent to $u \geq 4$. Using (57) and the definition of M now gives

$$\begin{aligned} \underline{q}_k &\geq 2\gamma\sqrt{a}(1 - e^{-2})Mu - \frac{13a}{32}u \\ &= \sqrt{a}\,u\left(\omega - \frac{13}{32}\sqrt{a}\right) = \sqrt{a}\left(\frac{4}{3} + \frac{\sqrt{5}}{2}\right)u > 0. \end{aligned}$$

Moreover, $1 + u \le 5u/4$ for $u \ge 4$, so

$$v_k \le \frac{5a}{16}u. \tag{58}$$

It remains to lower-bound the Freedman exponent. The positive root of

$$q^2 - \frac{4Ru}{3}q - 4uv_k = 0$$

is

$$q_\star = \frac{4Ru/3 + \sqrt{(4Ru/3)^2 + 16uv_k}}{2}.$$

Using $\sqrt{x+y} \le \sqrt{x} + \sqrt{y}$ for nonnegative x, y, followed by $R = \sqrt{a}$ and (58), gives

$$q_\star \le \frac{4Ru}{3} + 2\sqrt{uv_k} \le \sqrt{a}\,u\left(\frac{4}{3} + \frac{\sqrt{5}}{2}\right) \le \underline{q}_k \le q_k.$$

The quadratic is nonnegative to the right of its positive root; hence

$$q_k^2 \ge 4u\left(v_k + \frac{Rq_k}{3}\right).$$

Dividing by $2(v_k + Rq_k/3)$ shows that the exponent is at least $2u = 2\log k$. Exponentiating its negative gives the claimed k^{-2} bound. □

Lemma C.2 (Lambert-W inversion bounds)**.** *For $M \ge e/2$, let $\kappa(M) = \exp[-2W_{-1}(-1/(2M))]$. Then*

$$\frac{4}{(1+\log 2)^2}M^2(\log(2M) + \log(2\log(2M)))^2 \le \kappa(M) < 4M^2(\log(2M) + \log(2\log(2M)))^2.$$

Equality in the lower bound holds only at $M = e/2$, and the ratio of $\kappa(M)$ to the expression without its leading constant tends to 4.

Proof. Let $L = \log(2M) \ge 1$, $A = L + \log(2L)$, and $z = -W_{-1}(-1/(2M)) \ge 1$. From $(-z)e^{-z} = -1/(2M)$ we obtain $e^z = 2Mz$, and taking logarithms gives

$$z = L + \log z, \qquad \kappa(M) = e^{2z} = 4M^2z^2. \tag{59}$$

Because $\log z \ge 0$, the first identity implies $z \ge L$. Substituting this back into $\log z$ gives

$$z = L + \log z \ge L + \log L.$$

Furthermore,

$$(1+\log 2)(L + \log L) \ge L + \log L + \log 2 = A,$$

because $L + \log L \ge 1$ for $L \ge 1$. Thus $z \ge A/(1+\log 2)$. Squaring this inequality and substituting it into $\kappa(M) = 4M^2z^2$ proves the lower bound. Equality throughout requires $L + \log L = 1$, hence $L = 1$ and $M = e/2$.

For the upper bound, let $f(x) = x - \log x$. Its derivative $f'(x) = 1 - 1/x$ is nonnegative on $[1, \infty)$ and positive for $x > 1$. The inequality $e^L > 2L$ for $L \ge 1$ follows, for example, because $e^L/(2L)$ has its minimum over this range at $L = 1$, where it equals $e/2 > 1$. Hence $\log(2L) < L$ and $A < 2L$. Direct calculation yields

$$f(A) - L = A - \log A - L = \log(2L/A) > 0.$$

On the other hand, (59) says $f(z) = L$. Strict monotonicity therefore gives $z < A$. Substituting this into $\kappa(M) = 4M^2z^2$ proves the upper bound.

Finally,

$$\frac{A}{L} = 1 + \frac{\log(2L)}{L} \longrightarrow 1.$$

Since $L \le z < A$, division by A and the squeeze theorem give $z/A \to 1$. Therefore $\kappa(M)/(M^2A^2) = 4(z/A)^2 \to 4$, proving the final assertion. □

D EXPERIMENTAL DETAILS

Synthetic benchmark. A fixed seed generates the 1,000 item parameters and every Bernoulli response. For the fixed 21-component mixture, the calculation adds $(1 - \sum_{k=0}^{20} w_k)e^{-t2^{-21}}$ to the explicit partial sum, exactly as in (12); omitted component paths are not estimated. Because the administration order is determined before responses are drawn, it is a predictable heterogeneous design covered by the anytime-validity results.

The discriminations are drawn from $\mathrm{LogNormal}(0, 0.35)$ and clipped to $[0.5, 2.5]$; difficulties are drawn from $N(0, 0.8^2)$ and clipped to $[-2.5, 2.5]$. Items are ordered by decreasing Fisher information at the cutoff. The cutoff solves the equal-weight mean-probability equation for an expected score of 14 out of 20. We generate 1,000 paths on each side at every registered absolute gap and impose a 1,000-item cap.

Unified retained-method comparison. Every comparator is evaluated on the same item bank and response paths. Early, terminal, and hybrid decisions are kept separate, and equal absolute gaps are pooled using replication-count weights. The response-independent 2PL partition term is evaluated over the registered ability grid for the prefix MLE and GLR calculations. None of those implementation parameters is relabeled as a pointwise half-line error guarantee; only the two e-process procedures inherit (1) in this experiment.

The retained comparators are: SPRT between $\theta_c - 0.5$ and $\theta_c + 0.5$ with 0.05 Wald errors; a one-sided profile-likelihood interval with spending $\alpha/[t(t+1)]$; Thompson–Ro integrated CLR with 64-point Gauss–Legendre quadrature over four ability units on either side; a separated-composite GLR with constrained maxima beyond the same half-width; and SCGLR with 250 predictive simulations at each scheduled check. These settings define distinct inferential targets and are not treated as error-matched procedures.

D.1 COMPLETE SYNTHETIC COMPARATOR TABLE

The main text separates validity, certification yield, and cap-charged exposure because a fast forced decision and a slower certified decision are not interchangeable outcomes. For completeness, the table below gives the absolute cap-charged item count and adverse-crossing rate for every retained method at each registered distance $|\theta - \theta_c| \in \{0.4, 0.8, 1.2\}$.

Table 3: Unified comparison on the common synthetic 2PL paths at the three registered distances from the cutoff. Each cell is a Monte Carlo estimate of mean items / opposite-direction crossing (%). Results pool the two directions. Unresolved paths are charged the 1,000-item cap. Terminal fallbacks are excluded, so the table reports certified crossings only.

Method	$\lvert\theta - \theta_c\rvert = 0.4$	0.8	1.2
Shrinking CCT	134.6 / 0.00	24.7 / 0.00	13.9 / 0.00
Mixture CCT	85.0 / 0.00	18.6 / 0.00	12.5 / 0.00
SPRT/TSPRT	8.6 / 5.95	5.2 / 0.25	4.1 / 0.05
Profile-LR interval (spending)	446.4 / 0.00	37.4 / 0.00	14.4 / 0.00
Integrated CLR	11.8 / 9.80	4.0 / 3.15	2.7 / 0.50
Practical GLR	7.2 / 7.95	4.2 / 1.20	3.1 / 0.25
SCGLR	7.2 / 8.00	4.2 / 1.20	3.1 / 0.25

D.2 LOGGED-DATA CONSTRUCTION

Source releases. The analysis uses the untouched ASSISTments 2009–2010, EdNet KT1, and NeurIPS/Eedi public releases. Student response records are not redistributed.

Deterministic splits. For dataset d and student u, a SHA-256 digest of the master seed, d, and u determines the 75–25% calibration/evaluation allocation. A second deterministic hash assigns each

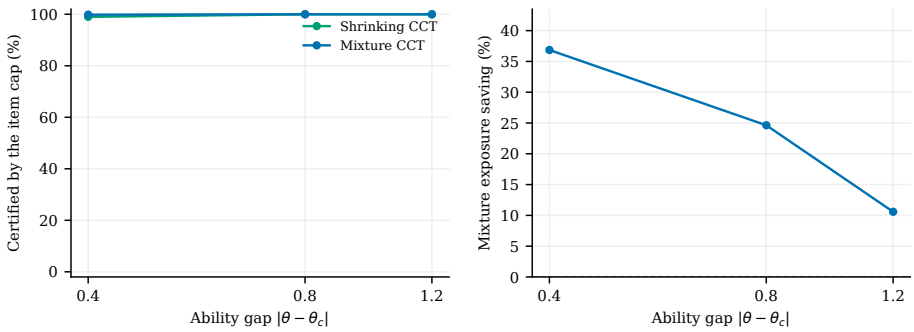

Figure 1: Finite-cap efficiency of the two gap-free CCTs. Left: certification by 1,000 items. Right: reduction in cap-charged exposure from using the mixture rather than shrinking. Points pool the two sides.

Table 4: Dataset-specific logged-path results. Cells report mean items/hybrid-disagreement percentage; n is the number of confidently reference-classified students.

Dataset (n)	Shrinking	Mixture	SPRT	Mixture crossed (%)
ASSISTments (62)	61.1/0.00	40.8/0.00	10.6/4.84	83.87
NeurIPS/Eedi (914)	52.2/0.22	37.6/0.22	15.3/1.09	93.33
EdNet (18,460)	165.9/2.69	123.6/2.92	33.6/6.14	73.80

usable evaluation response to the 60% sequential or 40% reference subset. Splits are stable across methods and runs.

Cross-fitted item calibration. Within a calibration student, responses are deterministically divided into two folds. The ability covariate for a response in one fold is computed from the opposite fold. Grouped, weakly regularized logistic regression then estimates the 2PL item intercept and discrimination. Items must satisfy the registered minimum observation and outcome-variation filters. Evaluation students never contribute to these fits.

Reference labels. The reference estimate solves the 2PL score equation on the 40% reference subset. Its label is the side of the cutoff, and its standard error is the inverse square root of observed Fisher information. The primary subset uses $|\widehat{\theta} - \theta_c| \geq 1.96\widehat{\mathrm{se}}(\widehat{\theta})$. Among all eligible students, 17,758 of 40,811 have $|\widehat{\theta} - \theta_c| < 0.5$; after applying the primary confidence filter, the corresponding count is 2,701 of 19,436. Within each dataset, the calibration-student ability proxy is centered at its median and scaled by its standard deviation before item fitting. The reported cutoff is then the solution of the equal-question bank-average 2PL equation for an expected score of 14 out of 20. Thus the reference label measures which side of a reproducible, bank-derived benchmark a student falls on; it is not a validated mastery label.

Recorded outcomes. For every student and sequential rule, the analysis records the available path length, first crossing time and direction, and fallback-augmented decision. The fixed-length comparison uses the same response paths at every registered horizon and is restricted to confidently reference-classified students with at least 200 sequential responses.

D.3 STRATIFIED AND PAIRED SEQUENTIAL RESULTS

Table 5: Pooled directional results. Brackets are 95% Wilson intervals across students, conditional on the fitted item banks and reference-label construction.

Reference side	Method	n	Mean items	Wrong crossing [95% CI] (%)	Hybrid error [95% CI] (%)
Below cutoff	Shrinking CCT	12,395	98.2	0.07 [0.04, 0.14]	2.06 [1.82, 2.32]
	Mixture CCT	12,395	75.7	0.22 [0.15, 0.32]	2.15 [1.91, 2.42]
	SPRT ($\Delta = 0.5$)	12,395	27.0	2.12 [1.88, 2.39]	3.18 [2.88, 3.50]
Above cutoff	Shrinking CCT	7,041	269.4	1.52 [1.26, 1.83]	3.47 [3.06, 3.92]
	Mixture CCT	7,041	196.1	2.07 [1.77, 2.43]	3.91 [3.48, 4.38]
	SPRT ($\Delta = 0.5$)	7,041	42.5	10.00 [9.32, 10.72]	10.69 [9.99, 11.44]

Table 6: Paired contrasts for the mixture CCT. Negative item differences favor the mixture; negative hybrid-error differences favor the mixture. Brackets are percentile intervals from 2,000 student-level paired bootstrap resamples.

Dataset	Comparator	n	Item difference [95% CI]	Hybrid-error difference [95% CI] (pp)
Pooled	Shrinking CCT	19,436	-40.9 [-43.0, -38.7]	0.22 [0.12, 0.31]
	SPRT ($\Delta = 0.5$)	19,436	86.7 [83.7, 89.9]	-3.12 [-3.39, -2.86]
ASSISTments	Shrinking CCT	62	-20.3 [-29.6, -11.0]	0.00 [0.00, 0.00]
	SPRT ($\Delta = 0.5$)	62	30.2 [20.8, 41.8]	-4.84 [-11.29, 0.00]
NeurIPS/Eedi	Shrinking CCT	914	-14.6 [-16.7, -12.6]	0.00 [0.00, 0.00]
	SPRT ($\Delta = 0.5$)	914	22.3 [19.3, 25.2]	-0.88 [-1.53, -0.33]
EdNet	Shrinking CCT	18,460	-42.2 [-44.6, -40.0]	0.23 [0.12, 0.33]
	SPRT ($\Delta = 0.5$)	18,460	90.1 [86.7, 93.5]	-3.22 [-3.51, -2.95]

D.4 COMMON-HORIZON NUMERICAL RESULTS

Table 7: Common-horizon comparison on confidently reference-classified students with at least 200 sequential responses. Sequential cells give mean items/hybrid-error percentage when unresolved paths receive the cutoff-score fallback at the displayed horizon; fixed rules use exactly the displayed number of items.

Horizon	Shrinking mean/error	Mixture mean/error	SPRT mean/error	Cutoff score error	Endpoint LR error	Gaussian-prior Bayes error
20	19.2/20.73	19.0/20.71	18.4/20.72	20.71	20.70	20.53
40	35.9/14.35	33.9/14.27	28.6/14.28	14.29	14.21	14.14
80	64.8/9.82	57.2/9.82	36.2/10.49	9.74	9.56	9.67
120	89.5/7.92	76.1/7.98	38.6/9.48	7.72	7.59	7.64
160	111.1/6.50	92.5/6.59	39.3/9.12	6.25	6.17	6.25
180	120.9/5.93	99.9/6.04	39.5/9.05	5.63	5.50	5.60
200	130.2/5.67	106.9/5.80	39.6/9.00	5.35	5.16	5.27

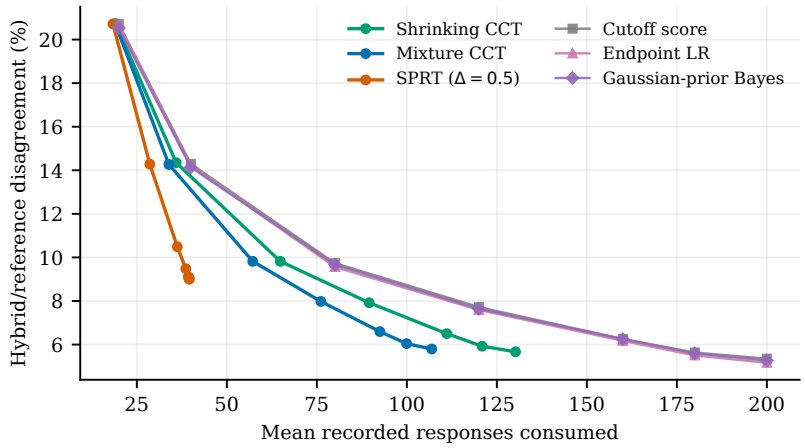

Figure 2: Resource–reference-agreement frontier for 8,523 confidently reference-classified students with at least 200 sequential responses. Sequential rules use the common fallback when unresolved; fixed rules consume the full horizon.