arXiv:2611.10787v1 [stat.ML] 7 Nov 2026

TRUSTWORTHY RECURSIVE SELF-IMPROVEMENT WITH STATISTICAL GUARANTEES

Zihe Gao[1] **Zhengling Qi**[2] **Yifan Cui**[1]
[1]Zhejiang University
[2]George Washington University

ABSTRACT

Recursive self-improvement offers a promising path toward AI agents that improve through experience by refining their own harnesses without retraining their underlying large language models (LLMs). Making this self-improvement process reliable, however, requires more than generating new updates: agents must decide which changes are worth keeping. A greedy gate that accepts updates based solely on higher validation scores can mistake an observed gain for genuine progress and lead to performance degeneration, whereas overly conservative statistical gates may reject beneficial updates and hinder performance improvement. We introduce Trustworthy Recursive Updates with Statistical Testing (TRUST), a framework that adapts an online multiple hypothesis testing procedure to control the rate of false evolution across rounds, with an LLM analyzing historical context to help determine how much evidence is required to update the incumbent. Under interpretable assumptions, we prove that TRUST limits the proportion of accepted updates that result in unacceptable performance loss. We further derive a lower bound on the final performance gain. Across the benchmarks HotpotQA, DROP, and APPS with the models DeepSeek-V4-Flash and Qwen-3.8-Flash, TRUST reduces the average ratio of accepted inferior updates from the greedy gate's 45.59% to 2.78%, while achieving higher final scores than the state-of-the-art statistical gate baselines for recursive self-improvement. These findings demonstrate that TRUST strikes a balance between statistical reliability and performance gain in recursive self-improvement.

1 INTRODUCTION

Recursive self-improvement (RSI) has attracted growing interest for its potential to enhance agents' capabilities without retraining their underlying language models (Ren et al., 2026; Chen et al., 2026). In a self-improvement loop, agents modify their own prompts, memory, tools, and workflows to improve task performance (Madaan et al., 2023; Shinn et al., 2023; Wang et al., 2024). This loop can be organized into three components: a solver that executes tasks, a proposer that generates candidate updates, and a gate that decides whether to accept them (Wang et al., 2026; Ren et al., 2026).

Existing work has explored mechanisms for generating better candidate updates, while less attention has been paid to deciding when to accept them. A common acceptance approach is to use a greedy gate, which accepts a candidate whenever its score on a fixed validation set exceeds the incumbent's (Shawn, 2026). Figure 1 illustrates how early gain under greedy acceptance can be lost in subsequent rounds. Noise from finite samples and overfitting from reusing validation data can both make inferior updates appear beneficial and pass a greedy gate (Dwork et al., 2015; Rogers et al., 2016). Accepting a harmful update can reduce the agent's current performance. The resulting loss may persist through subsequent rounds, ultimately limiting the gain for self-improvement.



Figure 1: Gain trajectories of TRUST and Greedy on HotpotQA with DeepSeek-V4-Flash. Details are in Section 6.1.

Recent work has begun to focus on the acceptance gate (Wu et al., 2025; Shawn, 2026; Sengupta, 2026). RSEA uses an extra val-

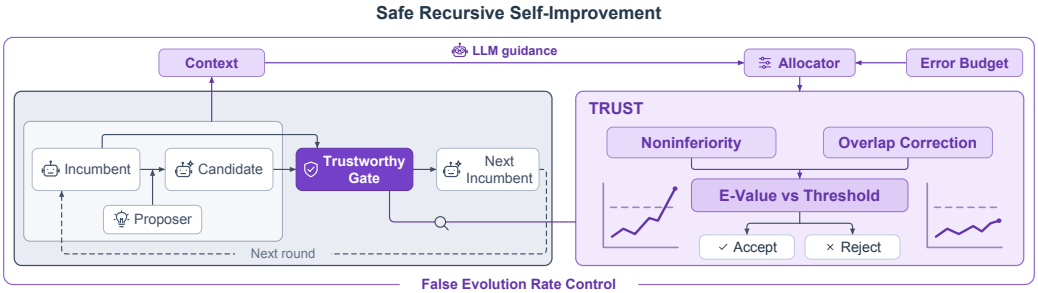


Figure 2: Recursive self-improvement with our trustworthy gate TRUST. TRUST combines noninferiority, overlap correction, e-value testing and error budget allocation with FER control across rounds.

idation set separate from the proposal generation set to select candidate updates (Nguyen et al., 2026). PACE provides a safety guarantee for each candidate update through an e-process test in a single round instead of the entire recursive self-improvement process (Shawn, 2026). However, controlling the probability of any false acceptance across rounds can make statistical gates conservative and hinder beneficial updates (Wu et al., 2025; Sengupta, 2026). A key challenge, therefore, remains: retaining opportunities for useful improvement while controlling acceptance errors across rounds under limited validation resources.

We introduce Trustworthy Recursive Updates with Statistical Testing (TRUST), a general statistical acceptance framework. Motivated by false discovery rate (FDR) control (Benjamini & Hochberg, 1995; Xu & Ramdas, 2024; Zhang et al., 2025b), we define the false evolution rate (FER) as a safety criterion for recursive self-improvement. FER measures the expected proportion of accepted updates that cause an unacceptable performance loss, thereby quantifying the gate's false acceptance risk. To control FER, we formulate the sequential acceptance decisions as an online multiple testing problem. We also use an LLM to analyze historical context and help the error budget allocator assign less restrictive testing levels to promising candidates. We further draw on noninferiority testing to improve statistical power by allowing a possible loss tolerance (Blackwelder, 1982; Da Silva et al., 2009; Schumi & Wittes, 2011), and use overlap correction to account for bias from reusing validation samples across rounds (Dwork et al., 2015; Rogers et al., 2016). Figure 2 summarizes how TRUST makes recursive self-improvement trustworthy.

Our contributions. We make the following three main contributions: (a) Our framework builds on online multiple hypothesis testing with interpretable parameters, including the error budget, loss tolerance, and overlap correction strength. It naturally incorporates LLM guidance based on diverse forms of historical context. (b) We establish statistical guarantees under explicit assumptions, including FER control, power consistency for improving candidates, and a lower bound on deployment gain. These results connect hypothesis testing criteria to performance gain in recursive self-improvement. (c) We evaluate TRUST in a closed-loop self-improvement harness across HotpotQA, DROP, and APPS with DeepSeek-V4-Flash and Qwen-3.8-Flash. Averaged across these settings, TRUST achieves the highest final gain of 24.84 pp. Compared with the greedy gate, TRUST reduces the mean FER_0 from 45.59% to 2.78% and the mean maximum drawdown from 4.963 pp to 0.047 pp, demonstrating high performance gains with few erroneous acceptances and small drawdowns.

2 Related Work

Recursive self-improvement and agent harness. An agent harness includes tools, context management, and execution control (Shi et al., 2026). Recursive self-improvement agents use feedback to improve various harness components. Promptbreeder evolves task prompts (Fernando et al., 2024). Reflexion, ExpeL, and Voyager convert experience into reflections, insights, and reusable skills, respectively (Shinn et al., 2023; Zhao et al., 2024; Wang et al., 2024). CoEvo-Mem improves memory retrieval (Ye et al., 2026). AFlow and Flow refine agentic workflows (Zhang et al., 2025a; Niu et al., 2025). AgentSquare defines a modular design space with planning, reasoning, tool use, and memory components (Shang et al., 2025). These works focus on proposing effective candidates, while we study whether a new candidate should be accepted to make recursive self-improvement trustworthy.

Additionally, we design a harness with diverse editable components in our experiments to evaluate how different gates affect performance and risk.

Hypothesis testing methods. Multiple hypothesis testing procedures are primarily based on two error criteria: the familywise error rate (FWER) and the FDR (Benjamini & Hochberg, 1995; Wang & Ramdas, 2022; Robertson et al., 2023). FWER control is generally more conservative than FDR control. Many popular online FDR methods with p-values rely on independence across tests (Ramdas et al., 2018; Robertson et al., 2023). This independence can be difficult to satisfy when agents evolve based on past feedback. The e-value is an alternative test statistic with a betting interpretation (Shafer, 2021). The e-BH and e-LOND procedures allow arbitrary dependence among valid e-values while controlling FDR (Wang & Ramdas, 2022; Xu & Ramdas, 2024). Drawing on e-GAI (Zhang et al., 2025b) and contextual online testing (Chen & Kasiviswanathan, 2020), TRUST constructs an e-value and allocates the testing levels with context. In RSI, SGM uses statistical tests to decide updates to control FWER across rounds (Wu et al., 2025), while PACE uses an e-process to control the probability of false acceptance in a single round (Shawn, 2026). TRUST provides FER control across rounds and achieves higher final scores than these statistical baselines in our experiments.

3 Problem Formulation

Let $\mathcal{D}$ denote the distribution of tasks and let $u(x, s) \in [0, 1]$ denote the score of output s on task x. For a solver θ, let $S_\theta(x, \Omega)$ denote the solver output on task x under execution randomness Ω. Define the population score of θ as

$$J(\theta) = \mathbb{E}_{X\sim\mathcal{D},\Omega}\left[u\big(X, S_\theta(X, \Omega)\big)\right]. \tag{1}$$

We consider a self-improvement process that starts with an initial incumbent solver Θ_0 and initial information $\mathcal{F}_0$. It recursively generates and evaluates candidates over T rounds. In round t, let $Q(\cdot \mid \mathcal{F}_{t-1})$ denote the conditional distribution of the candidate given the completed history $\mathcal{F}_{t-1}$ from the fixed proposer. The proposer generates $\widetilde{\Theta}_t \sim Q(\cdot \mid \mathcal{F}_{t-1})$, which is evaluated by the gate $\mathcal{G}$ against the previous incumbent Θ_{t-1}. Then the gate $\mathcal{G}$ produces $A_t \in \{0, 1\}$ and updates the incumbent according to

$$\Theta_t = \begin{cases} \widetilde{\Theta}_t, & A_t = 1, \\ \Theta_{t-1}, & A_t = 0. \end{cases}$$

Different gates may induce different acceptance decisions, thereby changing the subsequent self-improvement trajectories and the candidates proposed in later rounds. These changes can affect the final gain, which is defined as

$$G(T) = \mathbb{E}[J(\Theta_T) - J(\Theta_0)]. \tag{2}$$

We next define the core safety criterion. The population score difference between the selected candidate and the incumbent is $\Delta_t = J\left(\widetilde{\Theta}_t\right) - J(\Theta_{t-1})$. For a given tolerance level $\varepsilon \in [0, 1)$, let $N_t^\varepsilon = \mathbf{1}\{\Delta_t \le -\varepsilon\}$, $R_T = \sum_{t\le T} A_t$, and $F_T^\varepsilon = \sum_{t\le T} N_t^\varepsilon A_t$. Define the false evolution rate (FER) at tolerance ε as

$$\mathrm{FER}_\varepsilon(T) = \mathbb{E}\left[\frac{F_T^\varepsilon}{R_T \vee 1}\right]. \tag{3}$$

Here, FER is the expected proportion of unacceptably inferior updates among accepted updates, and ε specifies the tolerated loss. For $\varepsilon > 0$, the difference Δ_t is divided into three regions: updates with $\Delta_t \le -\varepsilon$ are unacceptable, those with $-\varepsilon < \Delta_t \le 0$ are tolerated but not improving, and those with $\Delta_t > 0$ are improving. Our goal is to ensure that $\mathrm{FER}_\varepsilon(T) \le q$ for a prespecified FER budget q, thereby providing a statistical safety guarantee for the entire self-improvement process. This helps avoid unacceptable updates, retain beneficial proposals, and may achieve a favorable final gain.

Let $\mathcal{V}_N = (Z_1, \ldots, Z_N) \sim \mathcal{D}^N$ be a fixed validation pool of size N. Candidate generation uses information from the completed history. In round t, the sampling plan is to draw $n \le N$ tasks uniformly without replacement from $\mathcal{V}_N$ to evaluate the current candidate. All sampled tasks are then returned to $\mathcal{V}_N$ so that the next round starts with the full validation pool. Let $K_{t,i}$ denote the pool

index selected at the ith validation draw, and let $X_{t,i} = Z_{K_{t,i}}$ be the corresponding task. On task $X_{t,i}$, the candidate and incumbent produce the observed score difference

$$D_{t,i} = u\big(X_{t,i}, S_{\widetilde{\Theta}_t}(X_{t,i}, \Omega^c_{t,i})\big) - u\big(X_{t,i}, S_{\Theta_{t-1}}(X_{t,i}, \Omega^0_{t,i})\big) \in [-1, 1], \tag{4}$$

where $\Omega^c_{t,i}$ and $\Omega^0_{t,i}$ denote the execution randomness of the candidate and incumbent, respectively.

Information flow. Our framework uses filtrations to describe how information accumulates within and across rounds (Ramdas et al., 2023). For clarity, we briefly explain this information flow. The initial information $\mathcal{F}_0$ is independent of the validation pool. Let $\mathcal{F}_{t-1}$ denote the completed history before round t. Before validation, $\mathcal{F}^{\text{pre}}_t$ additionally fixes the selected candidate $\widetilde{\Theta}_t$ and testing level α_t in Section 4. For $i = 1, \ldots, n$, let $\mathcal{H}_{t,i}$ augment $\mathcal{F}^{\text{pre}}_t$ with the first i validation samples and their paired outcomes, with $\mathcal{H}_{t,0} = \mathcal{F}^{\text{pre}}_t$. After validation, $\mathcal{F}_t$ contains the decision A_t and updated incumbent Θ_t. We illustrate this timeline in Figure 3.

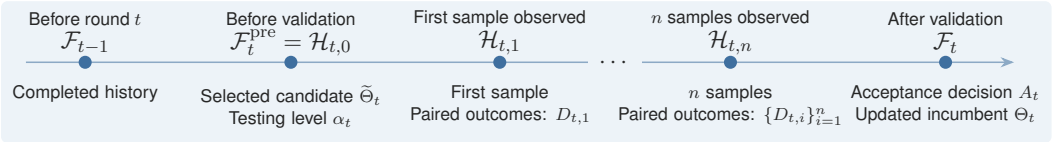


Figure 3: Information flow in a self-improvement round: $\mathcal{F}_{t-1} \subseteq \mathcal{F}^{\text{pre}}_t = \mathcal{H}_{t,0} \subseteq \mathcal{H}_{t,1} \subseteq \cdots \subseteq \mathcal{H}_{t,n} \subseteq \mathcal{F}_t \subseteq \mathcal{F}^{\text{pre}}_{t+1}$. The relation $\subseteq$ denotes inclusion of information.

4 Trustworthy Recursive Updates with Statistical Testing

The greedy gate accepts a candidate whenever the empirical score difference exceeds zero, making the acceptance decision sensitive to sampling and execution randomness. We formulate candidate acceptance as a hypothesis testing problem. In round t, we test

$$H^0_t : \Delta_t \le -\varepsilon, \qquad H^1_t : \Delta_t > -\varepsilon. \tag{5}$$

If the null hypothesis H^0_t is rejected, the selected candidate is accepted as the new incumbent. Otherwise, the previous incumbent is retained.

In online FDR control, testing levels are allocated across multiple hypotheses under an overall error budget (Chen & Kasiviswanathan, 2020; Zhang et al., 2025b). TRUST adapts this framework to the sequence of candidate acceptance tests defined in (5), using q as the total FER budget. In round t, the testing level α_t determines the acceptance threshold, while the test statistic $M^{(\beta)}_{t,m}$ measures the evidence accumulated from the first m validation observations. For $\alpha_t > 0$, TRUST accepts the candidate at the first $m \le n$ when $M^{(\beta)}_{t,m} \ge 1/\alpha_t$, following the standard testing rule for sequential e-value testing (Ramdas et al., 2023). If the threshold is not reached within n observations, the incumbent is retained. When $\alpha_t = 0$, validation is skipped and the incumbent is retained. We next describe how α_t is allocated and how $M^{(\beta)}_{t,m}$ is constructed.

4.1 Testing Level Allocation with Context

We now describe how TRUST allocates the testing level α_t in round t under the total FER budget q. Let B_t denote the remaining FER budget before round t, with $B_1 = q \in (0, 1)$, and let $\rho_t \in (0, 1)$ denote the spending fraction of the remaining FER budget in round t.

Adapting the e-LORD allocation rule within the e-GAI framework (Zhang et al., 2025b), we combine the adjusted spending fraction ρ_t with the remaining budget and the number of accepted updates to determine the testing level

$$\alpha_t = \min\{\alpha_{\max}, \rho_t B_t (R_{t-1} + 1)\}, \tag{6}$$

where $\alpha_{\max} \in (0, 1)$. The resulting charge c_t for the FER budget and the remaining budget B_{t+1} are

$$c_t = \frac{\alpha_t}{R_{t-1} + 1}, \qquad B_{t+1} = B_t - c_t. \tag{7}$$

The charge satisfies $0 \le c_t \le \rho_t B_t \le B_t$, keeping the balance nonnegative.

We next describe how to determine ρ_t. Building on the contextual error budget allocation in online multiple hypothesis testing (Chen & Kasiviswanathan, 2020), we make use of the context C_t in round t and use an LLM to analyze it. This context only contains information available before the current round's validation, including the candidate update and feedback from previous validation rounds. Following the evaluation structure of G-Eval (Liu et al., 2023), the LLM uses explicit criteria to assess the candidate's potential benefits and risks. Let Ψ denote the mapping from context to LLM multipliers, with $w_t = \Psi(C_t) \in [0, \infty)$. The default spending fraction is set to $\rho_t^0 = 1/(T - t + 1)$, corresponding to an equal share of the remaining FER budget. The LLM multiplier adjusts the spending fraction within the range $[\rho_{\min}, \rho_{\max}]$:

$$\rho_t = \mathrm{clip}(w_t \rho_t^0, \rho_{\min}, \rho_{\max}), \tag{8}$$

where $0 < \rho_{\min} < \rho_{\max} < 1$ and $\mathrm{clip}(x, a, b) = \min\{b, \max\{a, x\}\}$. All allocation quantities are determined before validation and included in the information $\mathcal{F}_t^{\mathrm{pre}}$ in Figure 3.

4.2 E-Value Construction with Overlap Correction

Once the testing level α_t has been determined, the next step is to construct the test statistic $M_{t,m}^{(\beta)}$. We first correct for bias from reusing validation samples. Based on the sampling procedure in Section 3, we define the following index sets to characterize overlap across validation rounds. Recall that $K_{t,i}$ denotes the index of the ith validation sample in round t. Let $\mathcal{R}_{t,i-1}$ denote the set of indices not sampled in the first $i-1$ draws of round t, and let $\mathcal{I}_{t,i-1}$ denote the cumulative set of indices used in validation across the first $t-1$ rounds and the first $i-1$ draws of round t. For each draw, define

$$Q_{t,i} = \mathbf{1}\{K_{t,i} \in \mathcal{I}_{t,i-1}\}, \qquad \zeta_{t,i} = \mathbb{P}(Q_{t,i} = 1 \mid \mathcal{H}_{t,i-1}) = \frac{|\mathcal{I}_{t,i-1} \cap \mathcal{R}_{t,i-1}|}{|\mathcal{R}_{t,i-1}|}. \tag{9}$$

Here, $Q_{t,i}$ indicates whether the current sample has previously been used in validation, while $\zeta_{t,i}$ represents its probability under uniform sampling.

Algorithm 1 TRUST

Input: Number of rounds T, validation pool $\mathcal{V}_N$, sample size $1 \le n \le N$, initial incumbent Θ_0, proposer Q, allocator Ψ, budget q, margin ε, strength β, $\alpha_{\max}$, $0 < \rho_{\min} < \rho_{\max} < 1$.

Initialize $B_1 \leftarrow q$, $R_0 \leftarrow 0$, and $\mathcal{I}_{1,0} \leftarrow \varnothing$.
for $t = 1, \dots, T$ **do**
  Initialize $\tau_t \leftarrow n$, $A_t \leftarrow 0$, $\mathcal{R}_{t,0} \leftarrow \{1, \dots, N\}$, and $M_{t,0}^{(\beta)} \leftarrow 1$.
  Propose $\widetilde{\Theta}_t \sim Q(\cdot \mid \mathcal{F}_{t-1})$ and construct context C_t.
  Set $w_t \leftarrow \Psi(C_t)$ and $\rho_t \leftarrow \mathrm{clip}\big(w_t/(T-t+1), \rho_{\min}, \rho_{\max}\big)$.
  Allocate $\alpha_t \leftarrow \min\{\alpha_{\max}, \rho_t B_t (R_{t-1}+1)\}$.
  Update $B_{t+1} \leftarrow B_t - \alpha_t/(R_{t-1}+1)$.
  for $i = 1, \dots, n$ **do**
    Draw $K_{t,i}$ uniformly from $\mathcal{R}_{t,i-1}$ and evaluate $D_{t,i}$ on $X_{t,i} = Z_{K_{t,i}}$ by (4).
    Compute $M_{t,i}^{(\beta)}$ using $D_{t,i}$, $\mathcal{R}_{t,i-1}$ and $\mathcal{I}_{t,i-1}$ by (9) and (10).
    Update $\mathcal{R}_{t,i} \leftarrow \mathcal{R}_{t,i-1} \setminus \{K_{t,i}\}$ and $\mathcal{I}_{t,i} \leftarrow \mathcal{I}_{t,i-1} \cup \{K_{t,i}\}$.
    if $M_{t,i}^{(\beta)} \ge 1/\alpha_t$ **then** $A_t \leftarrow 1$, $\tau_t \leftarrow i$, break.
  end for
  Update $\mathcal{I}_{t+1,0} \leftarrow \mathcal{I}_{t,\tau_t}$.
  Set $\Theta_t \leftarrow \widetilde{\Theta}_t$ if $A_t = 1$, and $\Theta_t \leftarrow \Theta_{t-1}$ otherwise.
  Update $R_t \leftarrow R_{t-1} + A_t$.
end for
return Θ_T.

We rescale the score difference in (4) from $[-1, 1]$ to $[0, 1]$ by setting $Y_{t,i} = (1 + D_{t,i})/2$, which allows us to use standard e-value constructions for bounded observations (Waudby-Smith & Ramdas, 2024; Ramdas et al., 2023), and the corresponding null boundary in (5) becomes $p_\varepsilon = (1 - \varepsilon)/2$. We choose $\beta \in [0, 1 - p_\varepsilon]$ as the correction strength for bias from data reuse and weight this strength by

the exposure probability $\zeta_{t,i}$ to obtain the corrected null boundary $u_{t,i}^{(\beta)} = p_\varepsilon + \zeta_{t,i}\beta$. The sequential e-value evidence is accumulated by setting $M_{t,0}^{(\beta)} = 1$ and

$$M_{t,m}^{(\beta)} = \prod_{i=1}^{m} \left[1 + \eta_{t,i} \frac{Y_{t,i} - u_{t,i}^{(\beta)}}{u_{t,i}^{(\beta)}} \right], \qquad m = 1, \ldots, n. \tag{10}$$

The betting fraction $\eta_{t,i} \in [0, 1)$ is chosen based on the information in $\mathcal{H}_{t,i-1}$ before observing $Y_{t,i}$ and controls how strongly the current observation changes the evidence. It can be selected using standard adaptive betting strategies (Waudby-Smith & Ramdas, 2024; Orabona & Jun, 2024; Shekhar & Ramdas, 2023). The complete procedure is presented in Algorithm 1.

5 Theoretical Analysis

We now establish statistical guarantees of our framework, including the validity of the e-value, FER control, power consistency to accept improvements, and a lower bound on the final gain.

Assumption 1. *For any $t = 1, \ldots, T$ and $i = 1, \ldots, n$, suppose*

$$\mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}, Q_{t,i} = 1] \le p_\varepsilon + \beta, \quad \text{if } H_t^0 \text{ is true and } \zeta_{t,i} > 0. \tag{11}$$

Reusing validation feedback to construct or select later candidates can lead to optimistic bias in their estimated performance, which is widely studied in adaptive data analysis (Dwork et al., 2015; Rogers et al., 2016). We use β to represent the strength of the correction for data reuse. Assumption 1 connects this correction to our testing by bounding the conditional mean in the e-value construction. We discuss the relationship between β and our algorithm in detail in Appendix A.1.

TRUST may accept the candidate early when the testing threshold is reached and the number of observations is a random variable determined by the observed evidence rather than a fixed number. We formalize this as a stopping time (Durrett, 2019). A random variable $\tau_t \in \{1, \ldots, n\}$ is a stopping time with respect to the filtration $(\mathcal{H}_{t,i})_{i=0}^n$ if $\{\tau_t \le i\} \in \mathcal{H}_{t,i}$ for every $i = 0, \ldots, n$. We then establish the following conditional validity result.

Theorem 1. *Fix any round $t \in \{1, \ldots, T\}$. Under Assumption 1, for any $m = 1, \ldots, n$, the statistic $M_{t,m}^{(\beta)}$ defined in (10) satisfies*

$$\mathbb{E}[M_{t,m}^{(\beta)} \mid \mathcal{F}_t^{\mathrm{pre}}] \le 1, \quad \text{if } H_t^0 \text{ is true.} \tag{12}$$

This inequality also holds with m replaced by any stopping time $\tau_t \le n$ with respect to $(\mathcal{H}_{t,m})_{m=0}^n$.

Theorem 1 shows that our test statistic satisfies the conditional validity requirement for e-values described in the statistical literature (Vovk & Wang, 2021; Ramdas et al., 2023). More importantly, validity at arbitrary stopping times $\tau_t \le n$ justifies the acceptance rule in Algorithm 1, which stops at $\tau_t = \min\{n, \min\{i \in \{1, \ldots, n\} : M_{t,i}^{(\beta)} \ge 1/\alpha_t\}\}$. The conditional validity also extends to mixtures of evidence processes using different betting fractions. See Corollary 1 in the appendix. This validity result provides the basis for the following theorem on FER control in RSI.

Theorem 2. *Fix a horizon $T \ge 1$ and an FER budget $q \in (0, 1)$. Under Assumption 1, the false evolution rate of Algorithm 1 satisfies*

$$\mathrm{FER}_\varepsilon(T) \le q. \tag{13}$$

Theorem 2 controls the expected proportion of unacceptable updates among all accepted updates over the entire self-improvement process with T rounds. The same bound in (13) holds with T replaced by any stopping time $\tau \le T$ with respect to the self-improvement history $(\mathcal{F}_t)_{t=0}^T$, as established in Appendix C.2. Thus, the process may be stopped based on the observed self-improvement history without losing FER control. In addition, the FER guarantee does not rely on the LLM correctly identifying beneficial candidates. The LLM uses context to allocate more FER budget to promising candidates and less to unpromising ones, safely improving the detection of beneficial updates.

Let $I_t = \mathbf{1}\{\Delta_t > 0\}$, $U_T = \sum_{t \le T} I_t$, and $S_T = \sum_{t \le T} I_t A_t$. The true evolution rate (TER) is

$$\mathrm{TER}(T) = \mathbb{E}\left[\frac{S_T}{U_T \vee 1} \right], \tag{14}$$

the expected proportion of improving candidates that are accepted (Gadbury et al., 2004).

Assumption 2. *Fix T and $\varepsilon \in (0,1)$, and let $N, n \to \infty$. Suppose that the following conditions are satisfied: (a) There exists a deterministic sequence $r_{N,n} \to 0$ such that $\mathbb{P}(U_T = 0) \le r_{N,n}$. (b) $n/N \to 0$. (c) There exist constants η_- and η_+ such that, for $t = 1, \dots, T$ and $i = 1, \dots, n$, the betting fractions satisfy $0 < \eta_- \le \eta_{t,i} \le \eta_+ < 2\varepsilon/(1+\varepsilon)$.*

Assumption 2(a) requires the probability that the proposer fails to generate any improvement to vanish as the sample size grows. As $n \to \infty$, the proposer also receives more historical feedback, making it more likely to produce at least one effective improvement. Assumption 2(b) controls overlap by requiring $n/N \to 0$, so the bias caused by overlap becomes very small. Assumption 2(c) can be satisfied in practice by choosing the betting fraction $\eta_{t,i}$ within the specified range.

Theorem 3. *Under Assumption 2, there exists a constant $\kappa > 0$ such that for sufficiently large N and n,*

$$\mathrm{TER}(T) \ge 1 - r_{N,n} - Te^{-\kappa n}. \tag{15}$$

Hence, $\mathrm{TER}(T) \to 1$ as $N, n \to \infty$.

Theorem 3 provides a lower bound on TER through two terms: $r_{N,n}$ bounds the probability that no improving candidate is proposed, while $Te^{-\kappa n}$ bounds the contribution of gate misses. Both terms vanish as the sample size grows with T fixed. It also reveals a theoretical role of noninferiority: when $\varepsilon = 0$, the upper bound $2\varepsilon/(1+\varepsilon)$ in Assumption 2 is zero, then no positive betting fraction satisfies the condition. Under this argument, a positive ε is needed to obtain $\mathrm{TER}(T) \to 1$.

We next evaluate the final incumbent through its gain $G(T)$ defined in (2). Let $\mathcal{P}_T = \sum_{t=1}^{T} \mathbb{E}[(\Delta_t)_+]$ denote the expected total positive effect of the proposed candidates.

Theorem 4. *Under Assumptions 1 and 2, there exists a constant $\kappa > 0$ such that for sufficiently large N and n,*

$$G(T) \ge \left(1 - e^{-\kappa n}\right)\mathcal{P}_T - \sum_{t=1}^{T} \mathbb{E}[A_t \min\{(-\Delta_t)_+, \varepsilon\} + \alpha_t(-\Delta_t - \varepsilon)_+]\,. \tag{16}$$

The first term shows that the loss of positive improvement caused by statistical rejection vanishes exponentially with the validation sample size. The second term quantifies the downside from the noninferiority margin. The third term controls loss beyond the tolerance through the allocated testing levels α_t and may be reduced when the LLM uses historical context to assign smaller testing levels to candidates with greater downside risk. Thus, TRUST asymptotically preserves beneficial updates while keeping harmful acceptance risk under statistical control. Additionally, to measure losses along the trajectory, we define population maximum drawdown (MDD) as $\mathrm{MDD}(T) = \max_{0 \le r < s \le T} \big(J(\Theta_r) - J(\Theta_s)\big)_+$. Appendix E derives an upper bound on expected MDD.

6 Numerical Experiments

We compare the performance gains and acceptance risks of different gates in an editable self-improvement harness, and conduct ablation studies to examine the effects of TRUST's key parameters. Supplementary diagnostics and case studies in Appendices I and G.4 show that higher validation scores do not necessarily imply better test performance.

Following modular agent design (Shang et al., 2025; Hu et al., 2025), our editable harness $\mathsf{H}_t = (\mathsf{R}_t, \mathsf{P}_t, \mathsf{M}_t, \mathsf{W}_t)$ exposes 20 operators across four components: reasoning (Zhou et al., 2023; Fernando et al., 2024), routing (Lewis et al., 2020; Liu et al., 2024), memory (Shinn et al., 2023; Zhao et al., 2024), and workflow (Zhang et al., 2025a). Appendix F details operator designs.

6.1 Main gate comparison

Datasets, models and metrics. We evaluate the benchmarks HotpotQA (Yang et al., 2018), DROP (Dua et al., 2019), and APPS (Hendrycks et al., 2021) with the models Qwen-3.8-Flash and DeepSeek-V4-Flash. For each dataset, the validation sets contain a total of 200 examples, with 300 held-out examples used only as the test set. We use additional 20 examples to construct the initial information $\mathcal{F}_0$. Test scores are measured by Joint F1 for HotpotQA, official token F1 for DROP, and the test case pass rate for APPS. We report gain and MDD in percentage points (pp). For a stricter empirical diagnostic, we report FER_0, which counts all accepted non-improving updates.

Baselines. We compare TRUST with four baselines over 30 rounds and 5 seeds. Greedy accepts candidates with higher scores on the validation samples. PACE uses an e-process test to control false acceptance within a single round (Shawn, 2026). SGM-B and SGM-H use Bernstein and Hoeffding tests, respectively, with FWER control across rounds (Wu et al., 2025). For TRUST, we set $q = 0.30$, $\varepsilon = 0.10$, and $\beta = 0.01$. More implementation details appear in Appendix G.

Table 1: Main comparison: DeepSeek-V4-Flash (left) and Qwen-3.8-Flash (right). Values are mean $\pm$ standard error. Methods with empirical $\mathrm{FER}_0 > 30\%$ are shown in gray to indicate high risk. Bold marks the best Gain, TER, and MDD results among the remaining methods within each setting.

		DeepSeek-V4-Flash				**Qwen-3.8-Flash**			
Dataset	Gate	Gain (pp) ↑	FER_0 (%) ↓	TER (%) ↑	MDD (pp) ↓	Gain (pp) ↑	FER_0 (%) ↓	TER (%) ↑	MDD (pp) ↓
HotpotQA	Greedy	7.87±0.55	64.67±3.28	40.03±5.06	4.96±0.93	20.66±2.45	44.67±5.23	58.40±9.39	4.01±0.49
	PACE	10.15±0.75	0.00±0.00	12.63±3.65	**0.00**±0.00	26.42±1.60	6.67±6.67	16.31±2.94	0.06±0.06
	SGM-B	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00
	SGM-H	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00	5.77±5.77	0.00±0.00	1.25±1.25	**0.00**±0.00
	TRUST	**11.58**±1.30	10.00±10.00	**18.11**±8.39	0.02±0.02	**27.72**±0.81	0.00±0.00	**20.82**±3.69	**0.00**±0.00
DROP	Greedy	36.24±2.19	35.28±1.68	60.12±7.77	5.38±1.58	15.15±2.22	37.38±6.04	56.33±5.65	3.37±1.03
	PACE	29.70±0.59	0.00±0.00	6.68±1.67	**0.00**±0.00	13.22±3.77	0.00±0.00	5.23±1.78	**0.00**±0.00
	SGM-B	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00
	SGM-H	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00
	TRUST	**31.47**±1.42	0.00±0.00	**11.93**±1.29	**0.00**±0.00	**17.22**±1.49	0.00±0.00	**10.76**±2.82	**0.00**±0.00
APPS	Greedy	33.89±2.71	41.08±4.31	50.83±2.89	8.48±2.23	28.03±0.32	50.44±4.85	38.02±4.36	3.59±0.68
	PACE	28.68±2.66	0.00±0.00	5.19±1.00	**0.00**±0.00	26.24±0.46	0.00±0.00	4.44±0.48	**0.00**±0.00
	SGM-B	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00
	SGM-H	0.00±0.00	0.00±0.00	0.00±0.00	**0.00**±0.00	10.84±6.65	0.00±0.00	2.62±1.70	**0.00**±0.00
	TRUST	**34.55**±1.04	0.00±0.00	**11.64**±3.53	**0.00**±0.00	**26.53**±0.11	6.67±6.67	**12.52**±3.42	0.27±0.27

TRUST achieves high scores in all six settings and the highest scores in four of them (Table 1). It maintains FER_0 below 30% in all six settings, whereas Greedy exceeds 30% in every setting. The highest MDD for TRUST is only 0.27 pp, whereas Greedy reaches 8.48 pp. Appendix G.4 presents a case study of Greedy that improves validation performance but harms test performance.

Among the remaining methods with $\mathrm{FER}_0 \leq 30\%$, SGM-B never accepts an update and SGM-H accepts only three updates across all runs, so they show no score improvement in most settings. These results suggest that both methods for FWER control across multiple rounds are conservative in practice. Relative to PACE, TRUST achieves higher scores and TER in all six settings, with FER_0 lower in one setting, equal in three, and higher in two.

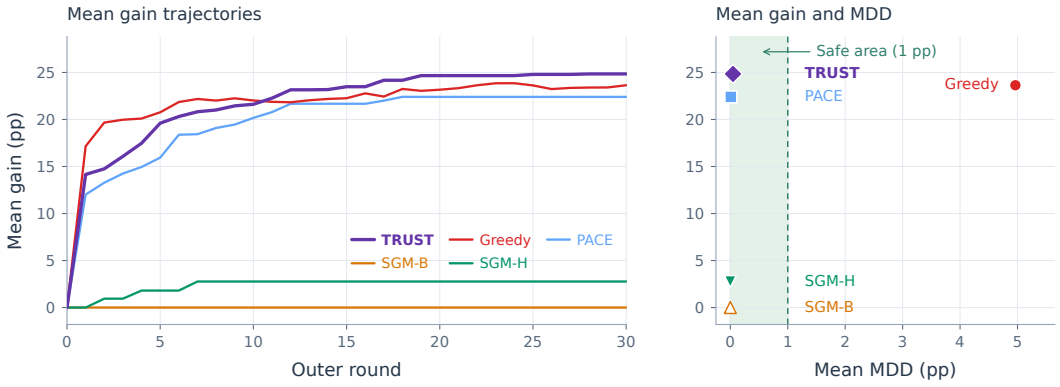


Figure 4: Left: mean gain trajectories. Right: mean final gain versus mean MDD, with shading marking the 1 pp empirical MDD reference region. Results are averaged across six settings.

The left panel of Figure 4 shows the mean gain trajectories over rounds. Greedy achieves rapid initial gains but exhibits subsequent declines. TRUST overtakes Greedy in later rounds and preserves its accumulated gains, achieving the highest final mean gain. The right panel compares mean final gain and mean MDD. We mark the region with mean MDD below 1 pp as an empirical safe area. TRUST lies within this region while achieving the highest among all compared methods, striking a balance between performance gain and stability in RSI.

6.2 Ablation Study

We next conduct ablation studies on HotpotQA/DeepSeek-V4-Flash. We vary three key parameters: the FER budget q, noninferiority margin ε, and overlap correction strength β. We also compare TRUST with and without LLM guidance. Appendix H provides the detailed settings and results.

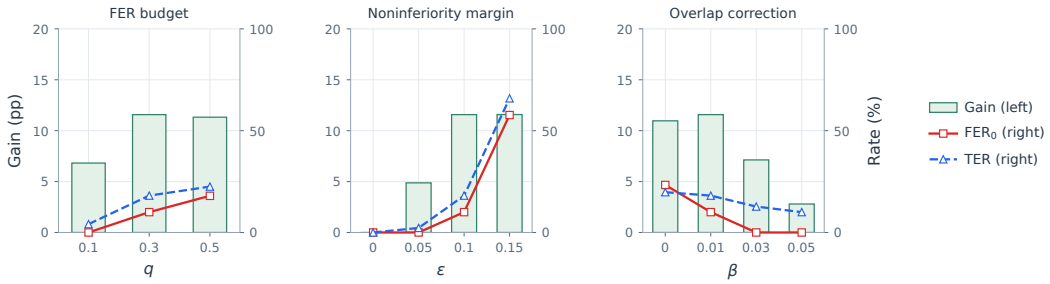


Figure 5: Bars show mean gain, and lines show FER_0 and TER. We vary one parameter at a time, keeping the others at the default values $(q, \varepsilon, \beta) = (0.30, 0.10, 0.01)$.

Figure 5 shows that larger q and ε make TRUST more aggressive, with higher TER and FER_0. Larger β makes it more conservative: FER_0 drops from 23.33% at $\beta = 0$ to 0% at $\beta = 0.03$. These trends are consistent with the definitions of the three parameters. In each ablation study, the gain first increases, then decreases or holds steady. For β, it peaks at 11.58 pp when $\beta = 0.01$, compared with 10.97 pp without correction, suggesting that an appropriate overlap correction strength can be beneficial in practice. Appendix H.1 provides additional trajectory results and a more detailed analysis. These results highlight the importance of balancing conservative and aggressive acceptance to improve performance.

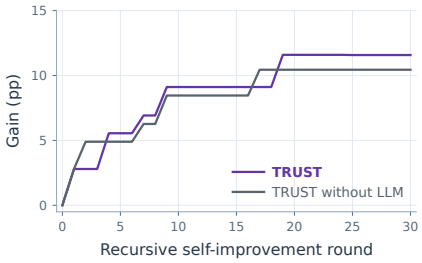


Figure 6: Mean gain trajectories with and without LLM allocation.

We replace the LLM allocator with the default multiplier $w_t \equiv 1$ while keeping other parameters in TRUST unchanged. Figure 6 shows that TRUST achieves a higher final gain, corresponding to final scores of $44.91 \pm 1.36\%$ for TRUST and $43.78 \pm 1.18\%$ for TRUST without LLM guidance. This suggests a performance benefit from using context to guide the allocation of testing levels. We provide a detailed case study in Appendix H.2 showing how the LLM analyzes context and assigns different LLM multipliers to influence acceptance decisions.

7 Conclusion and Limitations

TRUST applies online multiple hypothesis testing to control FER across rounds in self-improving agents, with LLM guidance and interpretable parameters for adjusting gate behavior. Our theoretical analysis and experiments show that the framework helps balance useful improvement and acceptance error control under limited validation resources. It extends existing statistical methods and provides a flexible statistical foundation for future work on recursive self-improvement. We discuss this point further in Appendix A.2. Overall, we show that TRUST strikes a balance between statistical reliability and performance gains in recursive self-improvement.

The present study does not address how to identify an exact value of overlap correction strength from data or how to optimally tune other parameters. The noninferiority margin provides an interpretable tolerance for loss, but the theoretical guarantee can be conservative and only partly explains practical behavior. The LLM allocator can occasionally make mistakes when judging candidate updates, which may lead to poor choices of LLM multipliers. Optimizing the gate to maximize gain while keeping safety metrics within target ranges is also a promising direction. Future work can address these limitations to make the recursive self-improvement process more trustworthy and effective.

REFERENCES

Akari Asai, Zeqiu Wu, Yizhong Wang, Avi Sil, and Hannaneh Hajishirzi. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun (eds.), *International Conference on Learning Representations*, volume 2024, pp. 9112–9141, 2024.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.

William C. Blackwelder. "Proving the Null Hypothesis" in Clinical Trials. *Controlled Clinical Trials*, 3(4):345–353, 1982.

Mingguang Chen, Licheng Wang, and Bo Qu. Recursive self-improvement in ai: From bounded self-refinement to autonomous research loops. *arXiv preprint arXiv:2607.07663*, 2026.

Minghao Chen, Yihang Li, Yanting Yang, Shiyu Yu, Binbin Lin, and Xiaofei He. Automanual: Constructing instruction manuals by llm agents via interactive environmental learning. *Advances in Neural Information Processing Systems*, 37:589–631, 2024.

Shiyun Chen and Shiva Kasiviswanathan. Contextual online false discovery rate control. In *International Conference on Artificial Intelligence and Statistics*, pp. 952–961. PMLR, 2020.

Gisela Tunes Da Silva, Brent R Logan, and John P Klein. Methods for equivalence and noninferiority testing. *Biology of Blood and Marrow Transplantation*, 15(1):120–127, 2009.

Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2368–2378, 2019.

Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.

Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248):636–638, 2015.

Chrisantha Fernando, Dylan Sunil Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. Promptbreeder: Self-referential self-improvement via prompt evolution. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 13481–13544, 2024. URL https://proceedings.mlr.press/v235/fernando24a.html.

Gary L Gadbury, Grier P Page, Jode Edwards, Tsuyoshi Kayo, Tomas A Prolla, Richard Weindruch, Paska A Permana, John D Mountz, and David B Allison. Power and sample size estimation in high dimensional biology. *Statistical methods in medical Research*, 13(4):325–338, 2004.

Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. Measuring coding challenge competence with apps. In J. Vanschoren and S. Yeung (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/c24cd76e1ce41366a4bbe8a49b02a028-Paper-round2.pdf.

Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 21344–21377, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/36b7acf6f6010652b3f2a433774a66fe-Paper-Conference.pdf.

Weiquan Huang, Zixuan Wang, Hehai Lin, Sudong Wang, Bo Xu, Qian Li, Beier Zhu, Linyi Yang, and Chengwei Qin. AMA: Adaptive memory via multi-agent collaboration. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 3099–3120. Association for Computational Linguistics, 2026. doi: 10.18653/v1/2026.findings-acl.152. URL https://aclanthology.org/2026.findings-acl.152/.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 9459–9474. Curran Associates, Inc., 2020.

Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. Compressing context to enhance inference efficiency of large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 6342–6353. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.emnlp-main.391. URL https://aclanthology.org/2023.emnlp-main.391/.

Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024. doi: 10.1162/tacl_a_00638.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-Eval: NLG evaluation using GPT-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 2511–2522. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023.emnlp-main.153. URL https://aclanthology.org/2023.emnlp-main.153/.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 46534–46594. Curran Associates, Inc., 2023. doi: 10.52202/075280-2019. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/91edff07232fb1b55a505a9e9f6c0ff3-Paper-Conference.pdf.

Michael Nguyen, Quoc Nguyen, and Paul Vuong. Recursive self-evolving agents via held-out selection. *arXiv preprint arXiv:2606.28374*, 2026.

Boye Niu, Yiliao Song, Kai Lian, Yifan Shen, Yu Yao, Kun Zhang, and Tongliang Liu. Flow: Modularized agentic workflow automation. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 74949–74977, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/ba84da6921f3040b74ee163aa7451f53-Paper-Conference.pdf.

Francesco Orabona and Kwang-Sung Jun. Tight concentrations and confidence sequences from the regret of universal portfolio. *IEEE Transactions on Information Theory*, 70(1):436–455, 2024. doi: 10.1109/TIT.2023.3330187.

Aaditya Ramdas, Tijana Zrnic, Martin Wainwright, and Michael Jordan. SAFFRON: An adaptive algorithm for online control of the false discovery rate. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4286–4294. PMLR, 2018. URL https://proceedings.mlr.press/v80/ramdas18a.html.

Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023. doi: 10.1214/23-STS894.

Zhe Ren, Yimeng Chen, Dandan Guo, Guowei Rong, Tonghui Li, RB Xiong, Qingfeng Lan, Wenyi Wang, Li Nanbo, Yibo Yang, et al. Self-improvements in modern agentic systems: A survey. *arXiv preprint arXiv:2607.13104*, 2026.

David S. Robertson, James M. S. Wason, and Aaditya Ramdas. Online Multiple Hypothesis Testing. *Statistical Science*, 38(4):557 – 575, 2023. doi: 10.1214/23-STS901. URL https://doi.org/10.1214/23-STS901.

Ryan Rogers, Aaron Roth, Adam Smith, and Om Thakkar. Max-information, differential privacy, and post-selection hypothesis testing. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 487–494. IEEE, 2016.

Daniel Russo and James Zou. Controlling bias in adaptive data analysis using information theory. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51, pp. 1232–1240. PMLR, 2016.

Jennifer Schumi and Janet T Wittes. Through the looking glass: understanding non-inferiority. *Trials*, 12(1):106, 2011.

Biswa Sengupta. Self-evolving agents with anytime-valid certificates. *arXiv preprint arXiv:2607.00871*, 2026. URL https://arxiv.org/abs/2607.00871.

Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 184(2):407–431, 2021.

Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, bayes factors and p-values. *Statistical Science*, pp. 84–101, 2011.

Yu Shang, Yu Li, Keyu Zhao, Likai Ma, Jiahe Liu, Fengli Xu, and Yong Li. Agentsquare: Automatic LLM agent search in modular design space. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=mPdmDYIQ7f.

Zayx Shawn. Pace: Anytime-valid acceptance tests for self-evolving agents. *arXiv preprint arXiv:2606.08106*, 2026.

Shubhanshu Shekhar and Aaditya Ramdas. Nonparametric two-sample testing by betting. *IEEE Transactions on Information Theory*, 70:1178–1203, 01 2023. doi: 10.1109/TIT.2023.3305867.

Quan Shi, Michael Tang, Karthik Narasimhan, and Shunyu Yao. Can language models solve olympiad programming? *arXiv preprint arXiv:2404.10952*, 2024. URL https://arxiv.org/abs/2404.10952.

Yifan Shi, Wei Zhang, and Tianyi Cui. A programming paradigm for spatiotemporal composability. *arXiv preprint arXiv:2608.25512*, 2026. URL https://arxiv.org/abs/2608.25512.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, pp. 8634–8652. Curran Associates, Inc., 2023. doi: 10.52202/075280-0377.

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

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=ehfRiF0R3a.

Kaiqi Wang, Wenjin Hou, Yuchen Yan, Hongrui Jia, Zhisheng Zhong, Botao Ren, Yifei Chen, Songyang Zhang, Yongliang Shen, Jun Xiao, et al. Diving into reliable self-evolving agents: A survey. *OpenReview preprint*, 2026.

Ruodu Wang and Aaditya Ramdas. False discovery rate control with E-values. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(3):822–852, 2022. doi: 10.1111/rssb.12489. URL https://doi.org/10.1111/rssb.12489.

Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. Agent workflow memory. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 63897–63911. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/wang25bx.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, 02 2024. ISSN 1369-7412. doi: 10.1093/jrsssb/qkad009. URL https://doi.org/10.1093/jrsssb/qkad009.

Xuening Wu, Shenqin Yin, Yanlan Kang, Xinhang Zhang, Qianya Xu, Zeping Chen, and Wenqiang Zhang. Sgm: A statistical godel machine for risk-controlled recursive self-modification. *arXiv preprint arXiv:2510.10232*, 2025.

Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Ziyu Xu and Aaditya Ramdas. Online multiple testing with e-values. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3997–4005, 2024.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, 2018.

Bowen Ye, Yongchao Xu, Zhijian Li, Xiang Yin, Junkai Ma, and Wenzhao Li. Coevo-mem: Co-evolving retrieval policy and memory bank for llm agents. *arXiv preprint arXiv:2608.01739*, 2026.

Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, et al. Aflow: Automating agentic workflow generation. In *International Conference on Learning Representations*, volume 2025, pp. 34040–34077, 2025a.

Yifan Zhang, Zijian Wei, Haojie Ren, and Changliang Zou. e-GAI: E-value-based generalized α-investing for online false discovery rate control. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 76308–76331. PMLR, 2025b. URL https://proceedings.mlr.press/v267/zhang25cd.html.

Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. Expel: Llm agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 19632–19642, 2024.

Mingming Zhao, Xiaokang Wei, Yuanqi Shao, Kaiwen Zhou, Lin Yang, Siwei Rao, Junhui Zhan, and Zhitang Chen. A^2flow: Automating agentic workflow generation via self-adaptive abstraction operators. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 29930–29938, 2026.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=WZH7099tgfM.

APPENDIX

A SUPPLEMENTARY DISCUSSION

A.1 INTERPRETATION OF THE OVERLAP CORRECTION

We do not develop a complete model of validation data reuse in this work. Detailed studies of adaptive data analysis can be found in prior work (Dwork et al., 2015; Rogers et al., 2016; Russo & Zou, 2016; Xu & Raginsky, 2017). Instead, we state sufficient conditions for valid e-values and focus on a framework for FER control across RSI rounds. The interpretation of β is related to how historical validation information, including feedback summarized in C_t, is used in candidate construction and selection. Greater reliance on information from previously evaluated samples may increase optimism when those samples are reused. For example, later rounds may use a few failed cases from the previous round or summaries of common error types. These forms of feedback reveal different information about the validation data and may lead to different degrees of optimism bias in subsequent evaluation. The parameter β specifies an allowance for this optimism under the null hypothesis. A stronger correction requires stronger validation results to reach the acceptance threshold, making the gate more conservative and potentially limiting both false acceptances and the acceptance of useful updates. Our sensitivity results in Section 6.2 and Appendix H support the practical value of an appropriate correction strength. In the reported ablation, an appropriate β achieves a higher mean final gain than no correction, whereas larger values make the gate overly conservative and reduce final gains.

A.2 RELATION TO EXISTING STATISTICAL METHODS

TRUST adapts FDR control in multiple hypothesis testing to the acceptance of candidates in self-evolving agents. It provides a framework covering the validation protocol, correction for data reuse, noninferiority testing, construction of test statistics, and allocation of testing levels. Within this framework, an LLM analyzes the candidate and historical context to guide allocation, while statistical guarantees hold under the stated assumptions. Existing statistical studies (Zhang et al., 2025b; Chen & Kasiviswanathan, 2020) compare testing procedures on specified hypotheses and observation streams in their numerical experiments. In our RSI setting, the proposer constructs candidates using historical information. Different gates can produce different acceptance decisions and incumbents, leading to different histories and subsequent candidates. Our experiments follow this closed loop and use update operators motivated by prior work on RSI agents. This allows us to compare how different gates affect both acceptance decisions and subsequent search on real tasks. Our theoretical analysis follows the same adaptive process and connects a lower bound on final gain to the testing levels and other parameters of the gate. It thus provides a statistical foundation for analyzing both acceptance risk and improvement along the trajectory induced by the gate.

B PROOF OF CONDITIONAL EVIDENCE VALIDITY

B.1 PROOF OF THEOREM 1

In each round, the sampling plan is to draw n tasks from the fixed validation pool of size N, where $n \leq N$. All N sample indices are selectable at the start of the round, so $\mathcal{R}_{t,0} = \{1, 2, \ldots, N\}$. At draw i, we select $K_{t,i}$ uniformly from the remaining set $\mathcal{R}_{t,i-1}$ and remove it, giving $\mathcal{R}_{t,i} = \mathcal{R}_{t,i-1} \setminus \{K_{t,i}\}$. Thus, a sample can be drawn at most once within a round but may be used again in later rounds.

Let $L_{t,i} \in \{0, 1\}$ be the $\mathcal{H}_{t,i-1}$-measurable indicator that draw i is made.

$$L_{t,i} = \begin{cases} 1, & \text{if draw } i \text{ is made in round } t, \\ 0, & \text{otherwise.} \end{cases}$$

Within each round, sampling proceeds consecutively and does not continue after stopping. If $\alpha_t = 0$, no validation samples are drawn.

Each round uses at most n samples. The bounds on the betting fractions in Assumption 2(c) apply on draws that are made ($L_{t,i} = 1$). If sampling stops early, we set the variables for the remaining indices

to zero as a notational convention for the proofs:

$$Y_{t,i} = Q_{t,i} = \zeta_{t,i} = \eta_{t,i} = 0 \qquad \text{when } L_{t,i} = 0.$$

After stopping, no more samples are used, and the history and record of previously used samples stay unchanged. When a sample is drawn, $Q_{t,i}$ equals one with conditional probability $\zeta_{t,i}$. Hence,

$$\mathbb{E}[Q_{t,i} \mid \mathcal{H}_{t,i-1}] = \zeta_{t,i}.$$

After stopping, both $Q_{t,i}$ and $\zeta_{t,i}$ are set to zero, so this equality still holds.

Lemma 1. *Let $(\mathcal{H}_i)_{i=0}^n$ be a filtration, $Y_i \in [0,1]$ be $\mathcal{H}_i$-measurable, and $u_i \in (0,1]$ and $\eta_i \in [0,1)$ be $\mathcal{H}_{i-1}$-measurable for each $i = 1, \ldots, n$. If*

$$\mathbb{E}[Y_i \mid \mathcal{H}_{i-1}] \le u_i,$$

then

$$M_0 = 1, \qquad M_m = \prod_{i=1}^{m} \left[1 + \eta_i \frac{Y_i - u_i}{u_i}\right]$$

is a nonnegative supermartingale. Moreover, for every stopping time $\tau \le n$, $\mathbb{E}[M_\tau \mid \mathcal{H}_0] \le 1$.

Proof. Write $G_i = 1 + \eta_i(Y_i - u_i)/u_i$. Since $Y_i \ge 0$, $G_i \ge 1 - \eta_i > 0$. Since u_i and η_i are $\mathcal{H}_{i-1}$-measurable, we have

$$\mathbb{E}[G_i \mid \mathcal{H}_{i-1}] = 1 + \frac{\eta_i}{u_i}\left(\mathbb{E}[Y_i \mid \mathcal{H}_{i-1}] - u_i\right) \le 1.$$

By the product definition, $M_i = M_{i-1}G_i$ and $M_{i-1} = \prod_{j=1}^{i-1} G_j$. For $j < i$, each G_j is $\mathcal{H}_j$-measurable, with $\mathcal{H}_j \subseteq \mathcal{H}_{i-1}$. Thus, M_{i-1} is $\mathcal{H}_{i-1}$-measurable and can be taken outside the conditional expectation:

$$\begin{aligned}
\mathbb{E}[M_i \mid \mathcal{H}_{i-1}] &= \mathbb{E}[M_{i-1}G_i \mid \mathcal{H}_{i-1}] \\
&= M_{i-1}\mathbb{E}[G_i \mid \mathcal{H}_{i-1}] \\
&= M_{i-1}\left[1 + \frac{\eta_i}{u_i}\left(\mathbb{E}[Y_i \mid \mathcal{H}_{i-1}] - u_i\right)\right] \\
&\le M_{i-1}.
\end{aligned}$$

Taking expectations and iterating from $M_0 = 1$ gives $\mathbb{E}[M_i] \le 1$, so each M_i is integrable. Together with nonnegativity and $\mathcal{H}_i$-measurability, this proves the supermartingale claim.

For bounded τ, use the finite pathwise decomposition

$$M_\tau = M_0 + \sum_{i=1}^{n} \mathbf{1}\{\tau \ge i\}(M_i - M_{i-1}).$$

The event $\{\tau \ge i\} = \{\tau \le i-1\}^c$ belongs to $\mathcal{H}_{i-1}$. Hence every stopped increment has nonpositive conditional mean:

$$\begin{aligned}
&\mathbb{E}\left[\mathbf{1}\{\tau \ge i\}(M_i - M_{i-1}) | \mathcal{H}_{i-1}\right] \\
&\quad = \mathbf{1}\{\tau \ge i\}\left(\mathbb{E}[M_i \mid \mathcal{H}_{i-1}] - M_{i-1}\right) \le 0.
\end{aligned}$$

Since $\mathcal{H}_0 \subseteq \mathcal{H}_{i-1}$, the tower property implies that each stopped increment also has nonpositive conditional expectation given $\mathcal{H}_0$. Taking conditional expectations given $\mathcal{H}_0$ in the decomposition of M_τ above, we obtain

$$\begin{aligned}
\mathbb{E}[M_\tau \mid \mathcal{H}_0] &= M_0 + \sum_{i=1}^{n} \mathbb{E}\left[\mathbf{1}\{\tau \ge i\}(M_i - M_{i-1}) \mid \mathcal{H}_0\right] \\
&\le M_0 = 1.
\end{aligned}$$

□

Proof of Theorem 1. Fix outer round t. Recall that

$$Y_{t,i} = \frac{1 + D_{t,i}}{2}, \qquad \Delta_t = J(\widetilde{\Theta}_t) - J(\Theta_{t-1}),$$
$$Q_{t,i} = \mathbf{1}\{K_{t,i} \in \mathcal{I}_{t,i-1}\}, \qquad \zeta_{t,i} = \mathbb{P}(Q_{t,i} = 1 \mid \mathcal{H}_{t,i-1}).$$

Both $\mathcal{R}_{t,i-1}$ and $\mathcal{I}_{t,i-1}$ are $\mathcal{H}_{t,i-1}$-measurable. Uniform sampling from the remaining set gives

$$\begin{aligned}\mathbb{P}(Q_{t,i} = 1 \mid \mathcal{H}_{t,i-1}) &= \sum_{k \in \mathcal{R}_{t,i-1}} \mathbf{1}\{k \in \mathcal{I}_{t,i-1}\} \frac{1}{|\mathcal{R}_{t,i-1}|} \\ &= \frac{|\mathcal{I}_{t,i-1} \cap \mathcal{R}_{t,i-1}|}{|\mathcal{R}_{t,i-1}|} = \zeta_{t,i}.\end{aligned}$$

We bound the contributions of unexposed and reused samples separately. The history depends only on accessed sample information and randomization independent of unexposed sample values. Since the pool is i.i.d., unexposed samples retain distribution $\mathcal{D}$ conditional on $\mathcal{H}_{t,i-1}$. Thus, for each draw,

$$\mathbb{E}[Y_{t,i}(1 - Q_{t,i}) \mid \mathcal{H}_{t,i-1}] = (1 - \zeta_{t,i}) \frac{1 + \Delta_t}{2}. \tag{17}$$

Here, $1 - Q_{t,i}$ retains only the unexposed contribution, and $1 - \zeta_{t,i}$ is the probability of drawing an unexposed sample.

For reused samples, Assumption 1 allows an additional β above the null boundary $p_\varepsilon = (1 - \varepsilon)/2$. On null histories with $\zeta_{t,i} > 0$, we have

$$\begin{aligned}\mathbb{E}[Y_{t,i} Q_{t,i} \mid \mathcal{H}_{t,i-1}] &= \zeta_{t,i}\, \mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}, Q_{t,i} = 1] \\ &\le \zeta_{t,i}(p_\varepsilon + \beta).\end{aligned}$$

When $\zeta_{t,i} = 0$, conditional monotonicity instead gives

$$0 \le \mathbb{E}[Y_{t,i} Q_{t,i} \mid \mathcal{H}_{t,i-1}] \le \mathbb{E}[Q_{t,i} \mid \mathcal{H}_{t,i-1}] = 0,$$

so the same weighted bound holds. Combining this bound with Equation (17) and $(1 + \Delta_t)/2 \le p_\varepsilon$ on H_t^0 yields

$$\begin{aligned}\mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] &= \mathbb{E}[Y_{t,i} Q_{t,i} \mid \mathcal{H}_{t,i-1}] + \mathbb{E}[Y_{t,i}(1 - Q_{t,i}) \mid \mathcal{H}_{t,i-1}] \\ &\le \zeta_{t,i}(p_\varepsilon + \beta) + (1 - \zeta_{t,i}) \frac{1 + \Delta_t}{2} \\ &\le p_\varepsilon + \zeta_{t,i}\beta.\end{aligned}$$

The correction $\zeta_{t,i}\beta$ therefore weights the overlap correction by the probability of drawing a reused sample. This also covers $\zeta_{t,i} = 1$, when the unexposed contribution vanishes.

When draw i is not made, $L_{t,i} = 0$ and $Y_{t,i} = Q_{t,i} = 0$ by the convention above. Since $L_{t,i}$ is $\mathcal{H}_{t,i-1}$-measurable, the indicator $\mathbf{1}\{L_{t,i} = 0\}$ is also $\mathcal{H}_{t,i-1}$-measurable and can be taken outside the conditional expectation. Hence,

$$\begin{aligned}\mathbf{1}\{L_{t,i} = 0\}\mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] &= \mathbb{E}[\mathbf{1}\{L_{t,i} = 0\} Y_{t,i} \mid \mathcal{H}_{t,i-1}] = 0, \\ \mathbf{1}\{L_{t,i} = 0\}\mathbb{E}[Q_{t,i} \mid \mathcal{H}_{t,i-1}] &= \mathbb{E}[\mathbf{1}\{L_{t,i} = 0\} Q_{t,i} \mid \mathcal{H}_{t,i-1}] = 0.\end{aligned}$$

On $\{L_{t,i} = 0\}$, the indicator equals one, so $\mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] = 0$ and $\zeta_{t,i} = \mathbb{E}[Q_{t,i} \mid \mathcal{H}_{t,i-1}] = 0$.

For draws that are made, the bounds for the unexposed and reused contributions together give $\mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] \le u_{t,i}^{(\beta)}$ on H_t^0. For draws that are not made, either after early stopping or because $\alpha_t = 0$, the zero conditional mean above is also at most $u_{t,i}^{(\beta)} = p_\varepsilon > 0$. Recall that $N_t^\varepsilon = \mathbf{1}\{\Delta_t \le -\varepsilon\}$. On H_t^0, multiplying by $N_t^\varepsilon = 1$ preserves these bounds; outside H_t^0, multiplying by $N_t^\varepsilon = 0$ makes both sides zero. Therefore, on all histories,

$$N_t^\varepsilon \mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] \le N_t^\varepsilon u_{t,i}^{(\beta)}.$$

The corrected boundary $u_{t,i}^{(\beta)}$ is $\mathcal{H}_{t,i-1}$-measurable and lies in $[p_\varepsilon, 1]$, with $p_\varepsilon > 0$. Write $G_{t,i} = 1 - \eta_{t,i} + \eta_{t,i} Y_{t,i} / u_{t,i}^{(\beta)}$. When draw i is not made, we set $\eta_{t,i} = 0$, so $G_{t,i} = 1$. For every i, the bounds $0 \le Y_{t,i} \le 1$, $0 \le \eta_{t,i} < 1$, and $u_{t,i}^{(\beta)} \ge p_\varepsilon > 0$ give

$$\begin{aligned} G_{t,i} &\ge 1 - \eta_{t,i} > 0, \\ G_{t,i} &\le 1 - \eta_{t,i} + \frac{\eta_{t,i}}{p_\varepsilon} = 1 + \eta_{t,i}(p_\varepsilon^{-1} - 1) \\ &\le p_\varepsilon^{-1}. \end{aligned}$$

The last inequality uses $\eta_{t,i} \le 1$ and $p_\varepsilon^{-1} - 1 \ge 0$. Hence $0 < G_{t,i} \le p_\varepsilon^{-1}$ almost surely.

Since each $G_{t,i}$ is $\mathcal{H}_{t,i}$-measurable, the product $M_{t,m}^{(\beta)}$ is $\mathcal{H}_{t,m}$-measurable. For $0 \le m \le n \le N$, we have

$$0 \le N_t^\varepsilon M_{t,m}^{(\beta)} \le M_{t,m}^{(\beta)} = \prod_{i=1}^{m} G_{t,i} \le p_\varepsilon^{-m} \le p_\varepsilon^{-N}.$$

Since p_ε^{-N} is a finite deterministic constant, both $M_{t,m}^{(\beta)}$ and $N_t^\varepsilon M_{t,m}^{(\beta)}$ have finite expectations and are therefore integrable. After sampling stops, all remaining factors $G_{t,i}$ equal one, so both $M_{t,m}^{(\beta)}$ and $N_t^\varepsilon M_{t,m}^{(\beta)}$ remain unchanged.

Set $X_m = N_t^\varepsilon M_{t,m}^{(\beta)}$. Since N_t^ε is $\mathcal{H}_{t,0}$-measurable, X_m is $\mathcal{H}_{t,m}$-measurable; the bounds above also show that X_m is nonnegative and integrable. To establish that $(X_m)_{m=0}^n$ is a supermartingale, it remains to prove

$$\mathbb{E}[X_{m+1} \mid \mathcal{H}_{t,m}] \le X_m, \qquad m = 0, \dots, n-1.$$

Since $\mathcal{H}_{t,0} \subseteq \mathcal{H}_{t,i-1}$, N_t^ε is also $\mathcal{H}_{t,i-1}$-measurable. The terms $1 - \eta_{t,i}$ and $\eta_{t,i}/u_{t,i}^{(\beta)}$ are $\mathcal{H}_{t,i-1}$-measurable and bounded as well. We can therefore take these quantities outside the conditional expectation. Using the bound $N_t^\varepsilon \mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] \le N_t^\varepsilon u_{t,i}^{(\beta)}$ established above, we obtain

$$\begin{aligned} \mathbb{E}[N_t^\varepsilon G_{t,i} \mid \mathcal{H}_{t,i-1}] &= N_t^\varepsilon \mathbb{E}[G_{t,i} \mid \mathcal{H}_{t,i-1}] \\ &= N_t^\varepsilon (1 - \eta_{t,i}) + \frac{\eta_{t,i}}{u_{t,i}^{(\beta)}} N_t^\varepsilon \mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] \\ &\le N_t^\varepsilon (1 - \eta_{t,i}) + \frac{\eta_{t,i}}{u_{t,i}^{(\beta)}} N_t^\varepsilon u_{t,i}^{(\beta)} \\ &= N_t^\varepsilon. \end{aligned}$$

The product before draw i is $M_{t,i-1}^{(\beta)} = \prod_{j=1}^{i-1} G_{t,j} \ge 0$. For $j < i$, each $G_{t,j}$ is $\mathcal{H}_{t,j}$-measurable, with $\mathcal{H}_{t,j} \subseteq \mathcal{H}_{t,i-1}$. Thus $M_{t,i-1}^{(\beta)}$ is $\mathcal{H}_{t,i-1}$-measurable and can also be taken outside the conditional expectation. Since $X_i = M_{t,i-1}^{(\beta)} N_t^\varepsilon G_{t,i}$, we have

$$\begin{aligned} \mathbb{E}[X_i \mid \mathcal{H}_{t,i-1}] &= \mathbb{E}[M_{t,i-1}^{(\beta)} N_t^\varepsilon G_{t,i} \mid \mathcal{H}_{t,i-1}] \\ &= M_{t,i-1}^{(\beta)} \mathbb{E}[N_t^\varepsilon G_{t,i} \mid \mathcal{H}_{t,i-1}] \\ &\le M_{t,i-1}^{(\beta)} N_t^\varepsilon \\ &= X_{i-1}. \end{aligned}$$

Taking $i = m+1$ gives the required conditional inequality. Together with the nonnegativity, measurability, and integrability established above, this proves that $(X_m)_{m=0}^n$ is a nonnegative supermartingale. Since $M_{t,0}^{(\beta)} = 1$, its initial value is $X_0 = N_t^\varepsilon$.

For a stopping time $\tau_t \le n$, the actual stopped process $X_{m \wedge \tau_t}$ is also a supermartingale: its next increment is $\mathbf{1}\{\tau_t > m\}(X_{m+1} - X_m)$, with a predictable indicator and nonpositive conditional mean. Iterating its conditional inequality from 0 to n yields

$$\mathbb{E}[X_{n \wedge \tau_t} \mid \mathcal{H}_{t,0}] \le X_0 = N_t^\varepsilon, \qquad N_t^\varepsilon \mathbb{E}[M_{t,\tau_t}^{(\beta)} \mid \mathcal{F}_t^{\mathrm{pre}}] \le N_t^\varepsilon.$$

On H_t^0, we have $N_t^\varepsilon = 1$. Therefore, for every stopping time $\tau_t \le n$,

$$\mathbb{E}[M_{t,\tau_t}^{(\beta)} \mid \mathcal{F}_t^{\mathrm{pre}}] \le 1 \qquad \text{almost surely on } H_t^0.$$

This is the stopping extension of Theorem 1. In particular, taking $\tau_t = n$ gives

$$\mathbb{E}[M_{t,n}^{(\beta)} \mid \mathcal{F}_t^{\mathrm{pre}}] \le 1 \qquad \text{almost surely on } H_t^0,$$

which is Equation (12). □

B.2 Stopping times

Proposition 1. *Suppose $\alpha_t \in (0,1)$ is $\mathcal{F}_t^{\mathrm{pre}}$-measurable and set*

$$\tau_t = \inf\left\{ m \in \{0,\dots,n\} : M_{t,m}^{(\beta)} \ge \frac{1}{\alpha_t} \right\} \wedge n.$$

Then τ_t is a stopping time bounded by n, and

$$\mathbf{1}\{M_{t,\tau_t}^{(\beta)} \ge 1/\alpha_t\} = \mathbf{1}\left\{ \max_{0\le m\le n} M_{t,m}^{(\beta)} \ge 1/\alpha_t \right\}.$$

Proof. For $0 \le m < n$,

$$\{\tau_t \le m\} = \bigcup_{j=0}^{m} \{M_{t,j}^{(\beta)} \ge 1/\alpha_t\} \in \mathcal{H}_{t,m},$$

because $M_{t,j}^{(\beta)}$ is $\mathcal{H}_{t,j}$-measurable, α_t is $\mathcal{H}_{t,0} = \mathcal{F}_t^{\mathrm{pre}}$-measurable, and $\mathcal{H}_{t,0} \subseteq \mathcal{H}_{t,j} \subseteq \mathcal{H}_{t,m}$ for every $j \le m$.

By definition, $\tau_t \le n$ for every sample outcome. Therefore, when $m \ge n$, we have $\tau_t \le n \le m$ for every outcome, so

$$\{\tau_t \le m\} = \Omega,$$

where Ω is the set of all possible sample outcomes. In particular, $\{\tau_t \le n\} = \Omega \in \mathcal{H}_{t,n}$. Together with the case $m < n$ above, this proves that τ_t is a stopping time bounded by n.

If a crossing occurs, the stopped value is the first crossing value; otherwise the process stops at n below threshold. □

B.3 Mixtures

Betting fractions. Betting interprets an evidence process as capital accumulated from an initial value of one (Waudby-Smith & Ramdas, 2024). In Equation (10), the amount bet on observation i is $\eta_{t,i} M_{t,i-1}^{(\beta)}$, where $M_{t,i-1}^{(\beta)}$ is the current capital and $\eta_{t,i} \in [0,1)$ specifies the proportion of that capital put at risk. The new capital equals the current capital plus the gain or loss from this bet:

$$\begin{aligned} M_{t,i}^{(\beta)} &= M_{t,i-1}^{(\beta)} + \eta_{t,i} M_{t,i-1}^{(\beta)} \frac{Y_{t,i} - u_{t,i}^{(\beta)}}{u_{t,i}^{(\beta)}} \\ &= M_{t,i-1}^{(\beta)} \left[(1-\eta_{t,i}) + \eta_{t,i} \frac{Y_{t,i}}{u_{t,i}^{(\beta)}} \right]. \end{aligned}$$

The uncommitted capital is retained, and the amount bet is multiplied by the return factor $Y_{t,i}/u_{t,i}^{(\beta)}$. Thus $\eta_{t,i} = 0$ leaves the evidence unchanged; if $Y_{t,i} = 0$, the update loses exactly the fraction $\eta_{t,i}$ of current capital. Larger betting fractions increase both the gain when $Y_{t,i} > u_{t,i}^{(\beta)}$ and the loss when $Y_{t,i} < u_{t,i}^{(\beta)}$. The amount of each bet is calculated using the current capital. A betting strategy specifies how to choose these fractions from the available history: each $\eta_{t,i}$ must be $\mathcal{H}_{t,i-1}$-measurable, so it is chosen before observing $Y_{t,i}$.

Combining betting strategies. Let $L \geq 1$ be a fixed integer denoting the number of betting strategies included in the mixture. For each $\ell \in \{1, \ldots, L\}$, strategy ℓ uses $\mathcal{H}_{t,i-1}$-measurable fractions $\eta_{t,i}^{(\ell)} \in [0, 1)$ and produces

$$M_{t,0}^{(\beta,\ell)} = 1, \qquad M_{t,m}^{(\beta,\ell)} = \prod_{i=1}^{m} \left[1 + \eta_{t,i}^{(\ell)} \frac{Y_{t,i} - u_{t,i}^{(\beta)}}{u_{t,i}^{(\beta)}} \right].$$

All L processes use the same validation observations and corrected boundaries, with potentially different betting fractions. A mixture combines their evidence by taking a weighted average, allowing several betting strategies to contribute to the test (Waudby-Smith & Ramdas, 2024). The weight $\pi_{t,\ell}$ is the fraction of initial capital assigned to strategy ℓ; the weights are nonnegative, sum to one, are $\mathcal{F}_t^{\text{pre}}$-measurable, and remain fixed throughout the round. Convex averaging preserves e-value validity even when the component e-values depend on the same observations (Vovk & Wang, 2021). The following corollary establishes the corresponding conditional validity for these sequential processes.

Corollary 1. *Let $M_{t,m}^{(\beta,1)}, \ldots, M_{t,m}^{(\beta,L)}$ use predictable betting sequences and let nonnegative weights $\pi_{t,1}, \ldots, \pi_{t,L}$ be $\mathcal{F}_t^{\text{pre}}$-measurable with $\sum_{\ell=1}^{L} \pi_{t,\ell} = 1$. Then $M_{t,m}^{(\beta,\text{mix})} = \sum_{\ell=1}^{L} \pi_{t,\ell} M_{t,m}^{(\beta,\ell)}$ is a nonnegative supermartingale under the corrected null, and $M_{t,\tau_t}^{(\beta,\text{mix})}$ is a conditionally valid e-value at every bounded stopping time.*

Proof of Corollary 1. For each ℓ, conditional supermartingale validity gives

$$\mathbb{E}[M_{t,i}^{(\beta,\ell)} \mid \mathcal{H}_{t,i-1}] \leq M_{t,i-1}^{(\beta,\ell)}.$$

The weights are nonnegative, sum to one, and are known at $\mathcal{H}_{t,0} = \mathcal{F}_t^{\text{pre}}$. Hence

$$\begin{aligned} \mathbb{E}[M_{t,i}^{(\beta,\text{mix})} \mid \mathcal{H}_{t,i-1}] &= \sum_{\ell=1}^{L} \pi_{t,\ell} \mathbb{E}[M_{t,i}^{(\beta,\ell)} \mid \mathcal{H}_{t,i-1}] \\ &\leq \sum_{\ell=1}^{L} \pi_{t,\ell} M_{t,i-1}^{(\beta,\ell)} = M_{t,i-1}^{(\beta,\text{mix})}. \end{aligned}$$

Nonnegativity is immediate, and bounded optional stopping gives the stopped e-value claim. □

C Proof of False Evolution Rate Control

Lemma 2. *If $\alpha_t \in [0, 1]$ is $\mathcal{F}_t^{\text{pre}}$-measurable and E_t is conditionally valid under H_t^0, then*

$$\mathbb{E}[N_t^{\varepsilon} A_t \mid \mathcal{F}_t^{\text{pre}}] \leq N_t^{\varepsilon} \alpha_t,$$

where $\alpha_t = 0$ is interpreted as $A_t = 0$.

Proof. When $\alpha_t > 0$, the indicator inequality $\mathbf{1}\{E_t \geq 1/\alpha_t\} \leq \alpha_t E_t$ holds. When $\alpha_t = 0$, both A_t and $\alpha_t E_t$ are zero. Thus $A_t \leq \alpha_t E_t$ in all cases. Since N_t^{ε} and α_t are $\mathcal{F}_t^{\text{pre}}$-measurable,

$$\mathbb{E}[N_t^{\varepsilon} A_t \mid \mathcal{F}_t^{\text{pre}}] \leq N_t^{\varepsilon} \alpha_t \mathbb{E}[E_t \mid \mathcal{F}_t^{\text{pre}}] \leq N_t^{\varepsilon} \alpha_t.$$

□

C.1 Deterministic horizon control in Theorem 2

Proof. For $\alpha_t > 0$, Algorithm 1 uses $M_{t,\tau_t}^{(\beta)}$ as its evidence, with the stopping time $\tau_t \leq n$ defined in Proposition 1. The stopping extension of Theorem 1 ensures that this evidence is conditionally valid on H_t^0. Therefore, Lemma 2 gives

$$\mathbb{E}[N_t^{\varepsilon} A_t \mid \mathcal{F}_t^{\text{pre}}] \leq N_t^{\varepsilon} \alpha_t.$$

When $\alpha_t = 0$, the algorithm sets $A_t = 0$, so the same inequality holds.

We first show that each testing level α_t is $\mathcal{F}_t^{\text{pre}}$-measurable and that the total budget spent is at most q. At the first round, $B_1 = q$ and $R_0 = 0$. The baseline $\rho_t^0 = 1/(T-t+1)$ is deterministic for $1 \le t \le T$, and w_t is selected using only $\mathcal{F}_t^{\text{pre}}$. Thus (8) and (6) make ρ_t, α_t, and

$$c_t = \frac{\alpha_t}{R_{t-1}+1} = \min\left\{\rho_t B_t, \frac{\alpha_{\max}}{R_{t-1}+1}\right\}$$

$\mathcal{F}_t^{\text{pre}}$-measurable. Hence the current testing level is fixed before validation. Moreover, $0 \le \rho_t \le 1$ implies

$$0 \le c_t \le \rho_t B_t \le B_t,$$

so $B_{t+1} = B_t - c_t$ remains nonnegative. For any $m \le T$, summing $c_t = B_t - B_{t+1}$ over $t = 1, \dots, m$ gives

$$\sum_{t=1}^{m} \frac{\alpha_t}{R_{t-1}+1} = \sum_{t=1}^{m}(B_t - B_{t+1}) = q - B_{m+1} \le q.$$

If $A_t = 1$ and $t \le T$, the total number of acceptances by horizon T is at least the earlier count plus the current acceptance:

$$R_T \ge R_{t-1} + 1.$$

If $A_t = 0$, the current false-acceptance numerator vanishes. Hence, we have

$$\frac{N_t^\varepsilon A_t}{R_T \vee 1} \le \frac{N_t^\varepsilon A_t}{R_{t-1}+1}.$$

Summing and taking expectations yields

$$\begin{aligned}
\mathrm{FER}_\varepsilon(T) &= \mathbb{E}\left[\frac{\sum_{t=1}^{T} N_t^\varepsilon A_t}{R_T \vee 1}\right] \\
&\le \sum_{t=1}^{T} \mathbb{E}\left[\frac{N_t^\varepsilon A_t}{R_{t-1}+1}\right] \\
&= \sum_{t=1}^{T} \mathbb{E}\left[\frac{\mathbb{E}[N_t^\varepsilon A_t \mid \mathcal{F}_t^{\text{pre}}]}{R_{t-1}+1}\right] \\
&\le \sum_{t=1}^{T} \mathbb{E}\left[\frac{N_t^\varepsilon \alpha_t}{R_{t-1}+1}\right] \\
&\le \mathbb{E}\left[\sum_{t=1}^{T} \frac{\alpha_t}{R_{t-1}+1}\right] \le q.
\end{aligned}$$

The third line uses that R_{t-1} is $\mathcal{F}_t^{\text{pre}}$-measurable, and the fourth uses Lemma 2. □

C.2 Outer stopping extension of Theorem 2

Proof. We extend the run beyond its planned horizon T by setting $A_t = \alpha_t = c_t = 0$ for $t > T$ and keeping the history and account balance fixed. Let τ be an almost surely finite stopping time for $(\mathcal{F}_t)_{t\ge 0}$. The stopped counts admit nonnegative series representations

$$F_\tau^\varepsilon = \sum_{t=1}^{\infty} \mathbf{1}\{\tau \ge t\} N_t^\varepsilon A_t, \qquad R_\tau = \sum_{t=1}^{\infty} \mathbf{1}\{\tau \ge t\} A_t.$$

Only finitely many summands are nonzero on almost every path. If $\tau \ge t$ and $A_t = 1$, then $R_\tau \ge R_{t-1} + 1$. Consequently,

$$\frac{F_\tau^\varepsilon}{R_\tau \vee 1} \le \sum_{t=1}^{\infty} \mathbf{1}\{\tau \ge t\} \frac{N_t^\varepsilon A_t}{R_{t-1}+1}.$$

The stopping-time property implies $\{\tau \geq t\} = \{\tau \leq t-1\}^c \in \mathcal{F}_{t-1} \subseteq \mathcal{F}_t^{\mathrm{pre}}$. Therefore the survival indicator is predictable, and Lemma 2 gives

$$\mathbb{E}\left[\mathbf{1}\{\tau \geq t\}\frac{N_t^{\varepsilon} A_t}{R_{t-1}+1}\middle| \mathcal{F}_t^{\mathrm{pre}}\right] \leq \mathbf{1}\{\tau \geq t\}\frac{\alpha_t}{R_{t-1}+1}.$$

All summands are nonnegative, so Tonelli's theorem and the tower property yield

$$\begin{aligned}\mathbb{E}\left[\frac{F_\tau^{\varepsilon}}{R_\tau \vee 1}\right] &\leq \mathbb{E}\left[\sum_{t=1}^{\infty} \mathbf{1}\{\tau \geq t\}\frac{\alpha_t}{R_{t-1}+1}\right] \\ &\leq \mathbb{E}\left[\sum_{t=1}^{\infty}\frac{\alpha_t}{R_{t-1}+1}\right] \leq q.\end{aligned}$$

Since $c_t = 0$ for $t > T$, the last inequality follows from $\sum_{t=1}^{T} c_t = q - B_{T+1} \leq q$. This extension is stronger than guaranteeing each fixed deterministic horizon separately: the stopping time itself can depend on the observed recursive self-improvement history. □

D Proof of True Evolution Rate Convergence

D.1 Proof of Theorem 3

Proof. Fix $T < \infty$ and a sequence of runs satisfying Assumption 2. We suppress dependence of the candidates and histories on (N, n). All statements for sufficiently large (N, n) are along this sequence, and all conditional inequalities below hold almost surely.

Step 1: Bound the testing levels. Equation (7) gives

$$c_t = \frac{\alpha_t}{R_{t-1}+1} = \min\left\{\rho_t B_t, \frac{\alpha_{\max}}{R_{t-1}+1}\right\} \leq \rho_{\max} B_t.$$

Consequently, $B_{t+1} = B_t - c_t \geq (1-\rho_{\max})B_t$. Starting from $B_1 = q$ and iterating yields

$$B_t \geq q(1-\rho_{\max})^{t-1} \geq q(1-\rho_{\max})^{T-1}, \qquad t \leq T.$$

Since $\rho_t \geq \rho_{\min}$ and $R_{t-1}+1 \geq 1$, we obtain the deterministic lower bound

$$\alpha_t \geq \alpha_* := \min\left\{\alpha_{\max}, \rho_{\min} q(1-\rho_{\max})^{T-1}\right\} > 0. \tag{18}$$

Write $\Lambda = \log(1/\alpha_*)$. Both α_* and Λ are independent of (N, n).

Step 2: Bound exposure and the conditional mean under improvement. At the beginning of round t, at most $(t-1)n$ distinct validation indices have been exposed. An index drawn earlier in the same round has been removed from $\mathcal{R}_{t,i-1}$, so it cannot contribute to $\mathcal{I}_{t,i-1} \cap \mathcal{R}_{t,i-1}$. Thus, for every draw,

$$0 \leq \zeta_{t,i} \leq \bar{\zeta}_{N,n} := \min\left\{1, \frac{(T-1)n}{N-n+1}\right\}. \tag{19}$$

Because T is fixed and $n/N \to 0$, we have $\bar{\zeta}_{N,n} \to 0$.

Let $L_{t,i}$ denote the $\mathcal{H}_{t,i-1}$-measurable indicator that draw i is made. On $\{\Delta_t > 0, L_{t,i} = 1\}$, Equation (17) and $Y_{t,i}Q_{t,i} \geq 0$ give

$$\begin{aligned}\mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}] &\geq \mathbb{E}[Y_{t,i}(1-Q_{t,i}) \mid \mathcal{H}_{t,i-1}] \\ &= (1-\zeta_{t,i})\frac{1+\Delta_t}{2} \\ &\geq \frac{1-\bar{\zeta}_{N,n}}{2}.\end{aligned}$$

Also,

$$u_{t,i}^{(\beta)} = p_\varepsilon + \beta\zeta_{t,i} \leq p_\varepsilon + \beta\bar{\zeta}_{N,n}.$$

Step 3: Positive log-growth bound. For $z \in [0,1]$ and $\eta \in [\eta_-, \eta_+]$, define

$$\phi(z,\eta) := \frac{1+z}{2}\log(1-\eta) + \frac{1-z}{2}\log\left(1-\eta+\frac{\eta}{p_\varepsilon+\beta z}\right).$$

On a draw, the log increment is

$$V_{t,i} := \log\frac{M_{t,i}^{(\beta)}}{M_{t,i-1}^{(\beta)}} = \log\left(1-\eta_{t,i}+\frac{\eta_{t,i}Y_{t,i}}{u_{t,i}^{(\beta)}}\right).$$

Set $z = \bar{\zeta}_{N,n}$ and $\bar{u} = p_\varepsilon + \beta z$. Conditional on $\mathcal{H}_{t,i-1}$, $\eta_{t,i}$ is fixed. Since $u_{t,i}^{(\beta)} \le \bar{u}$ and $Y_{t,i} \ge 0$,

$$V_{t,i} \ge \log\left(1-\eta_{t,i}+\frac{\eta_{t,i}Y_{t,i}}{\bar{u}}\right).$$

The function $\log(1-\eta+\eta y/\bar{u})$ is concave on $[0,1]$, so we have:

$$\log\left(1-\eta+\frac{\eta y}{\bar{u}}\right) \ge (1-y)\log(1-\eta) + y\log\left(1-\eta+\frac{\eta}{\bar{u}}\right).$$

The right-hand side is increasing in y. Applying this inequality with $y = Y_{t,i}$ and $\eta = \eta_{t,i}$, and then taking conditional expectations, gives

$$\mathbb{E}[V_{t,i} \mid \mathcal{H}_{t,i-1}] \ge (1-m_{t,i})\log(1-\eta_{t,i}) + m_{t,i}\log\left(1-\eta_{t,i}+\frac{\eta_{t,i}}{\bar{u}}\right),$$

where $m_{t,i} := \mathbb{E}[Y_{t,i} \mid \mathcal{H}_{t,i-1}]$. By Step 2, on $\{\Delta_t > 0,\ L_{t,i} = 1\}$,

$$m_{t,i} \ge \frac{1-\bar{\zeta}_{N,n}}{2}.$$

Moreover,

$$\log\left(1-\eta_{t,i}+\frac{\eta_{t,i}}{\bar{u}}\right) > \log(1-\eta_{t,i}),$$

so the preceding lower bound is increasing in $m_{t,i}$. Therefore,

$$\begin{aligned}\mathbb{E}[V_{t,i} \mid \mathcal{H}_{t,i-1}] &\ge \frac{1+\bar{\zeta}_{N,n}}{2}\log(1-\eta_{t,i}) \\ &\quad + \frac{1-\bar{\zeta}_{N,n}}{2}\log\left(1-\eta_{t,i}+\frac{\eta_{t,i}}{p_\varepsilon+\beta\bar{\zeta}_{N,n}}\right) \\ &= \phi(\bar{\zeta}_{N,n},\eta_{t,i}).\end{aligned}$$

For fixed z, $\phi(z,\eta)$ is concave in η, hence its minimum on $[\eta_-,\eta_+]$ is attained at an endpoint. Define

$$g_{N,n} := \min\{\phi(\bar{\zeta}_{N,n},\eta_-), \phi(\bar{\zeta}_{N,n},\eta_+)\}. \tag{20}$$

It follows that $\mathbb{E}[V_{t,i} \mid \mathcal{H}_{t,i-1}] \ge g_{N,n}$ on every improving draw.

Continuity and $\bar{\zeta}_{N,n} \to 0$ imply

$$g_{N,n} \longrightarrow g_0 := \min\{\phi(0,\eta_-), \phi(0,\eta_+)\}.$$

To verify that $g_0 > 0$, note that

$$\phi(0,\eta) = \frac{1}{2}\log\left[(1-\eta)\left(1-\eta+\frac{\eta}{p_\varepsilon}\right)\right]$$

and

$$(1-\eta)\left(1-\eta+\frac{\eta}{p_\varepsilon}\right) - 1 = \frac{\eta}{p_\varepsilon}\left[\varepsilon - \eta(1-p_\varepsilon)\right]. \tag{21}$$

Assumption 2(c) gives

$$0 < \eta_- \le \eta \le \eta_+ < \frac{2\varepsilon}{1+\varepsilon} = \frac{\varepsilon}{1-p_\varepsilon}.$$

The right-hand side of Equation (21) is therefore strictly positive, so both endpoint values defining g_0 are positive. Since $\bar{\zeta}_{N,n} \to 0$ and $\phi(z, \eta)$ is continuous in z,

$$\phi(\bar{\zeta}_{N,n}, \eta_-) \to \phi(0, \eta_-), \qquad \phi(\bar{\zeta}_{N,n}, \eta_+) \to \phi(0, \eta_+).$$

Therefore,

$$g_{N,n} = \min\{\phi(\bar{\zeta}_{N,n}, \eta_-), \phi(\bar{\zeta}_{N,n}, \eta_+)\} \to g_0,$$

where $g_0 = \min\{\phi(0, \eta_-), \phi(0, \eta_+)\} > 0$. Hence, for all sufficiently large N and n,

$$|g_{N,n} - g_0| < \frac{g_0}{2},$$

which implies

$$g_{N,n} \geq \frac{g_0}{2} > 0.$$

The log increments also lie in a fixed bounded interval. Since $u_{t,i}^{(\beta)} \geq p_\varepsilon$, $Y_{t,i} \in [0, 1]$, and $\eta_{t,i} \leq \eta_+$ on draws,

$$\log(1 - \eta_+) \leq V_{t,i} \leq \log\left(1 - \eta_+ + \frac{\eta_+}{p_\varepsilon}\right).$$

After stopping, $V_{t,i} = 0$, which lies in the same interval. Its width is the fixed constant

$$h := \log\left(\frac{1 - \eta_+ + \eta_+/p_\varepsilon}{1 - \eta_+}\right) > 0. \tag{22}$$

Thus the log increments are integrable, and the conditions of Appendix D.2 hold with $g_t = g_{N,n}$ and $h_t = h$ for all sufficiently large runs.

Step 4: Bound conditional miss probabilities. Applying Lemma 3 with $g_t = g_{N,n}$ and $h_t = h$ gives, on $\{\Delta_t > 0\}$,

$$\begin{aligned}\mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}}) &\leq \exp\left[-\frac{2n}{h^2}\left(g_{N,n} - \frac{\log(1/\alpha_t)}{n}\right)_+^2\right] \\ &\leq \exp\left[-\frac{2n}{h^2}\left(g_{N,n} - \frac{\Lambda}{n}\right)_+^2\right].\end{aligned} \tag{23}$$

The second inequality follows from $\alpha_t \geq \alpha_*$ and hence $\log(1/\alpha_t) \leq \Lambda$.

For sufficiently large N and n, we have $g_{N,n} \geq g_0/2$ and $n \geq 4\Lambda/g_0$. Therefore,

$$g_{N,n} - \frac{\Lambda}{n} \geq \frac{g_0}{4}.$$

Set

$$\kappa := \frac{g_0^2}{8h^2} > 0. \tag{24}$$

This constant is independent of (N, n), and Equation (23) yields

$$\mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}}) \leq e^{-\kappa n} \qquad \text{on } \{\Delta_t > 0\}, \quad t = 1, \ldots, T. \tag{25}$$

In particular, the conditional miss bound is a consequence of the sampling, allocation, and betting rules, rather than an assumed property of the gate.

Step 5: TER consistency. Let $L_T^+ = \sum_{t=1}^T I_t(1 - A_t) = U_T - S_T$ denote the number of improving candidates missed. Since I_t is $\mathcal{F}_t^{\mathrm{pre}}$-measurable, Equation (25) and the tower property give

$$\begin{aligned}\mathbb{E}[L_T^+] &= \sum_{t=1}^T \mathbb{E}\left[I_t \mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}})\right] \\ &\leq e^{-\kappa n} \sum_{t=1}^T \mathbb{E}[I_t] \\ &\leq T e^{-\kappa n}.\end{aligned}$$

The following identity holds on every trajectory:

$$1 - \frac{S_T}{U_T \vee 1} = \mathbf{1}\{U_T = 0\} + \frac{U_T - S_T}{U_T \vee 1}.$$

Indeed, when $U_T = 0$, also $S_T = 0$, and both sides equal one. When $U_T > 0$, the indicator vanishes and the equality follows from $U_T \vee 1 = U_T$. Taking expectations, using $U_T \vee 1 \geq 1$, and then applying Assumption 2(a), we obtain

$$\begin{aligned} 1 - \mathrm{TER}(T) &= \mathbb{P}(U_T = 0) + \mathbb{E}\left[\frac{L_T^+}{U_T \vee 1}\right] \\ &\leq \mathbb{P}(U_T = 0) + \mathbb{E}[L_T^+] \\ &\leq r_{N,n} + Te^{-\kappa n}. \end{aligned}$$

Rearranging proves Equation (15). Because T is fixed, $r_{N,n} \to 0$, and $n \to \infty$, the right-hand side tends to zero. Since $0 \leq \mathrm{TER}(T) \leq 1$, we conclude that $\mathrm{TER}(T) \to 1$. No independence between the future candidate count U_T and current acceptance decisions is used. □

D.2 An explicit conditional miss bound

The following lemma provides the conditional miss bound used to derive Equation (23) in Step 4 of the proof of Theorem 3. Since $\eta_{t,i} < 1$, every factor in Equation (10) is at least $1 - \eta_{t,i} > 0$, so $M_{t,i}^{(\beta)} > 0$ and its log increments are well defined. For each round with $\alpha_t > 0$, define

$$V_{t,i} = \log \frac{M_{t,i}^{(\beta)}}{M_{t,i-1}^{(\beta)}},$$

and let $L_{t,i}$ be the $\mathcal{H}_{t,i-1}$-measurable indicator that validation draw i is made.

Lemma 3 (Conditional miss bound)**.** *Suppose there exist finite $\mathcal{F}_t^{\mathrm{pre}}$-measurable $g_t, h_t > 0$ and finite $\mathcal{H}_{t,i-1}$-measurable interval endpoints $a_{t,i}, b_{t,i}$ such that, on $\{\Delta_t > 0,\ L_{t,i} = 1\}$, for every $i = 1, \ldots, n$,*

$$\mathbb{E}[V_{t,i} \mid \mathcal{H}_{t,i-1}] \geq g_t, \qquad a_{t,i} \leq V_{t,i} \leq b_{t,i}, \qquad b_{t,i} - a_{t,i} \leq h_t. \tag{26}$$

Then

$$\mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}}) \leq \epsilon_t \qquad \text{on } \{\Delta_t > 0\},$$

where

$$\epsilon_t = \begin{cases} \exp\left[-\dfrac{2n}{h_t^2}\left(g_t - \dfrac{\log(1/\alpha_t)}{n}\right)_+^2\right], & \alpha_t > 0, \\ 1, & \alpha_t = 0. \end{cases} \tag{27}$$

Proof. If $\alpha_t = 0$, then $A_t = 0$ and $\epsilon_t = 1$, so the miss bound holds.

Work on the $\mathcal{F}_t^{\mathrm{pre}}$-measurable event $\{\Delta_t > 0,\ \alpha_t > 0\}$ throughout the following calculation. Define

$$\mu_{t,i} = \mathbb{E}[V_{t,i} \mid \mathcal{H}_{t,i-1}], \qquad Z_{t,i} = V_{t,i} - \mu_{t,i}.$$

On $\{L_{t,i} = 1\}$, Equation (26) gives $\mu_{t,i} \geq g_t$, and subtracting $\mu_{t,i}$ from the interval endpoints preserves their width, which is at most h_t. On $\{L_{t,i} = 0\}$, the evidence is unchanged, so $V_{t,i} = 0$. Since this event is $\mathcal{H}_{t,i-1}$-measurable, $\mu_{t,i} = Z_{t,i} = 0$ there as well. The log increments satisfy $V_{t,i} \leq \log(1/p_\varepsilon)$ and have nonnegative conditional means on the event under consideration, so their negative parts have conditional expectations at most $\log(1/p_\varepsilon)$. They are therefore integrable when restricted to this event, and $\mathbb{E}[Z_{t,i} \mid \mathcal{H}_{t,i-1}] = 0$ there.

For any finite nonnegative $\mathcal{F}_t^{\mathrm{pre}}$-measurable λ, the conditional Hoeffding lemma gives

$$\mathbb{E}[e^{-\lambda Z_{t,i}} \mid \mathcal{H}_{t,i-1}] \leq \exp(\lambda^2 h_t^2/8).$$

Both λ and h_t are fixed when conditioning on $\mathcal{F}_t^{\mathrm{pre}}$. Dividing each factor by $\exp(\lambda^2 h_t^2/8)$ and iterating conditional expectations from $\mathcal{H}_{t,0} = \mathcal{F}_t^{\mathrm{pre}}$ yields

$$\mathbb{E}\left[\exp\left\{-\lambda \sum_{i=1}^n Z_{t,i} - \frac{n\lambda^2 h_t^2}{8}\right\} \middle| \mathcal{F}_t^{\mathrm{pre}}\right] \leq 1.$$

If $A_t = 0$, all n validation draws are made and the evidence never reaches $1/\alpha_t$. Therefore,

$$\sum_{i=1}^{n} V_{t,i} = \log M_{t,n}^{(\beta)} < \log(1/\alpha_t).$$

On this miss event, $L_{t,i} = 1$ for every $i = 1, \dots, n$. Equation (26) then gives $\mu_{t,i} \geq g_t$ at every draw, so

$$\sum_{i=1}^{n} \mu_{t,i} \geq n g_t.$$

Define $d_t = g_t - \log(1/\alpha_t)/n$. Subtracting the conditional means from the realized log increments gives, on $\{A_t = 0\}$,

$$\begin{aligned}
\sum_{i=1}^{n} Z_{t,i} &= \sum_{i=1}^{n} V_{t,i} - \sum_{i=1}^{n} \mu_{t,i} \\
&< \log(1/\alpha_t) - n g_t \\
&= -n \left(g_t - \frac{\log(1/\alpha_t)}{n} \right) \\
&= -n d_t.
\end{aligned}$$

Thus every outcome with $A_t = 0$ satisfies this inequality, which proves the event inclusion

$$\{A_t = 0\} \subseteq \left\{ \sum_{i=1}^{n} Z_{t,i} < -n d_t \right\}.$$

If $d_t \leq 0$, the formula gives $\epsilon_t = 1$, and the miss probability is at most one. On $\{d_t > 0\}$, take a finite positive $\mathcal{F}_t^{\mathrm{pre}}$-measurable λ. The event inclusion above implies that, whenever $A_t = 0$,

$$\exp\left\{ -\lambda \sum_{i=1}^{n} Z_{t,i} - \frac{n\lambda^2 h_t^2}{8} \right\} > \exp\left(\lambda n d_t - \frac{n\lambda^2 h_t^2}{8} \right).$$

The random variable on the left is nonnegative. The threshold on the right is positive and $\mathcal{F}_t^{\mathrm{pre}}$-measurable, since λ, d_t, and h_t are all $\mathcal{F}_t^{\mathrm{pre}}$-measurable. Conditional Markov's inequality bounds the probability of exceeding this threshold by the conditional expectation divided by the threshold. Therefore,

$$\begin{aligned}
\mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}}) &\leq \frac{\mathbb{E}\left[\exp\left\{ -\lambda \sum_{i=1}^{n} Z_{t,i} - \frac{n\lambda^2 h_t^2}{8} \right\} \middle| \mathcal{F}_t^{\mathrm{pre}} \right]}{\exp\left(\lambda n d_t - \frac{n\lambda^2 h_t^2}{8} \right)} \\
&\leq \frac{1}{\exp\left(\lambda n d_t - \frac{n\lambda^2 h_t^2}{8} \right)} \\
&= \exp\left(-\lambda n d_t + \frac{n\lambda^2 h_t^2}{8} \right).
\end{aligned}$$

The second inequality uses the conditional expectation bound of one established above by iterating the conditional Hoeffding bound. The choice $\lambda = 4(d_t)_+/h_t^2$ is $\mathcal{F}_t^{\mathrm{pre}}$-measurable and gives $\exp(-2nd_t^2/h_t^2)$ when $d_t > 0$. Combining these cases with $\alpha_t = 0$ proves Equation (27).

□

D.3 Extension of the power bound to mixtures

Consider the convex mixture in Corollary 1,

$$M_{t,m}^{(\beta,\mathrm{mix})} = \sum_{\ell=1}^{L} \pi_{t,\ell} M_{t,m}^{(\beta,\ell)}, \qquad \pi_{t,\ell} \geq 0, \quad \sum_{\ell=1}^{L} \pi_{t,\ell} = 1,$$

where the weights are $\mathcal{F}_t^{\mathrm{pre}}$-measurable and remain fixed during validation. Suppose at least one component $\ell^\star$ has predictable fractions satisfying Assumption 2(c) on draws and weights $\pi_{t,\ell^\star} \geq \pi_* > 0$ for every round and run, with π_* independent of (N, n). Other components need not satisfy this power condition, but they remain nonnegative processes constructed by Equation (10) with predictable fractions in $[0, 1)$.

Proof. Nonnegativity gives $M_{t,m}^{(\beta,\mathrm{mix})} \geq \pi_* M_{t,m}^{(\beta,\ell^\star)}$. If the mixture gate misses an improving candidate, all n draws are made and

$$M_{t,n}^{(\beta,\mathrm{mix})} < 1/\alpha_t, \qquad M_{t,n}^{(\beta,\ell^\star)} < 1/(\pi_* \alpha_t).$$

Steps 2–3 of the proof of Theorem 3 apply to the selected component, with its increments held at zero after the mixture gate stops. The conditional concentration argument in Appendix D.2 therefore gives

$$\mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}}) \leq \exp\left[-\frac{2n}{h^2}\left(g_{N,n} - \frac{\log(1/(\pi_* \alpha_t))}{n}\right)_+^2\right]$$

on $\{\Delta_t > 0\}$. The additional threshold penalty is at most the fixed constant $\log(1/\pi_*)$. Replacing Λ by $\Lambda + \log(1/\pi_*)$ in Step 4 thus gives the same $e^{-\kappa n}$ bound for all sufficiently large runs. Step 5 then proves the same TER conclusion. Under Assumption 1, the mixture remains conditionally valid by Corollary 1, so the proof of Theorem 4 below also applies. Only the designated component is required to satisfy the power condition. □

E Proof of the Gain Bound

Define the cumulative loss from accepted updates by

$$D_T^- = \sum_{t=1}^{T} A_t(-\Delta_t)_+.$$

Proof of Theorem 4. Fix a finite horizon $T \geq 1$ and the resource sequence and algorithm configuration in Appendix D.1. Take $\kappa > 0$ as in Equation (24) and consider a sufficiently large run for Equation (25) to hold. Throughout, the candidates, effects, and acceptance decisions refer to the actual adaptive trajectory.

Step 1: Bound cumulative accepted loss. Under Assumption 1, Theorem 1 and Proposition 1 provide conditionally valid evidence at the gate's stopping time. Lemma 2 therefore gives

$$\mathbb{E}[N_t^\varepsilon A_t \mid \mathcal{F}_t^{\mathrm{pre}}] \leq N_t^\varepsilon \alpha_t.$$

Let $W_t = (-\Delta_t - \varepsilon)_+$. Since Δ_t is $\mathcal{F}_t^{\mathrm{pre}}$-measurable and lies in $[-1, 1]$, W_t is bounded, nonnegative, and $\mathcal{F}_t^{\mathrm{pre}}$-measurable. Moreover, $W_t > 0$ implies $\Delta_t < -\varepsilon$, so $W_t = N_t^\varepsilon W_t$. Thus

$$\mathbb{E}[A_t W_t] = \mathbb{E}\left[W_t \mathbb{E}[N_t^\varepsilon A_t \mid \mathcal{F}_t^{\mathrm{pre}}]\right] \leq \mathbb{E}[\alpha_t W_t].$$

The identity

$$(-\Delta_t)_+ = \min\{(-\Delta_t)_+, \varepsilon\} + (-\Delta_t - \varepsilon)_+$$

splits each accepted loss into the part up to ε and the excess. Multiplying by A_t, summing, and taking expectations gives

$$\begin{aligned}\mathbb{E}[D_T^-] &= \sum_{t=1}^{T} \mathbb{E}[A_t \min\{(-\Delta_t)_+, \varepsilon\}] + \sum_{t=1}^{T} \mathbb{E}[A_t W_t] \\ &\leq \sum_{t=1}^{T} \mathbb{E}[A_t \min\{(-\Delta_t)_+, \varepsilon\}] + \sum_{t=1}^{T} \mathbb{E}[\alpha_t(-\Delta_t - \varepsilon)_+].\end{aligned}$$

Step 2: Control MDD. For every $0 \le r < s \le T$, the update rule implies

$$J(\Theta_r) - J(\Theta_s) = -\sum_{t=r+1}^{s} A_t \Delta_t \le \sum_{t=r+1}^{s} A_t(-\Delta_t)_+ \le D_T^-.$$

Taking a positive part and then maximizing over (r, s) yields $\mathrm{MDD}(T) \le D_T^-$ pathwise. Taking expectations and applying Step 1 gives the auxiliary bound

$$\mathbb{E}[\mathrm{MDD}(T)] \le \sum_{t=1}^{T} \mathbb{E}\left[A_t \min\{(-\Delta_t)_+, \varepsilon\} + \alpha_t(-\Delta_t - \varepsilon)_+\right].$$

This bound uses only conditional validity, not the power or candidate discovery conditions.

Step 3: Bound final gain. The update rule gives $J(\Theta_t) - J(\Theta_{t-1}) = A_t\Delta_t$, and $A_t\Delta_t = I_t A_t \Delta_t - A_t(-\Delta_t)_+$. Summing over rounds yields

$$J(\Theta_T) - J(\Theta_0) = \sum_{t=1}^{T} A_t\Delta_t = \sum_{t=1}^{T} I_t A_t \Delta_t - D_T^-.$$

Since $I_t\Delta_t = (\Delta_t)_+$ is nonnegative and $\mathcal{F}_t^{\mathrm{pre}}$-measurable, the derived conditional miss bound in Equation (25) gives

$$\begin{aligned}\mathbb{E}[I_t A_t \Delta_t] &= \mathbb{E}\left[(\Delta_t)_+ \mathbb{E}[A_t \mid \mathcal{F}_t^{\mathrm{pre}}]\right] \\ &\ge (1 - e^{-\kappa n})\mathbb{E}[(\Delta_t)_+].\end{aligned}$$

Summing and using $\mathcal{P}_T = \sum_{t=1}^{T} \mathbb{E}[(\Delta_t)_+]$, we obtain

$$\sum_{t=1}^{T} \mathbb{E}[I_t A_t \Delta_t] \ge (1 - e^{-\kappa n})\mathcal{P}_T.$$

Consequently,

$$\begin{aligned}G(T) &= \sum_{t=1}^{T} \mathbb{E}[I_t A_t \Delta_t] - \mathbb{E}[D_T^-] \\ &\ge (1 - e^{-\kappa n})\mathcal{P}_T - \sum_{t=1}^{T} \mathbb{E}\left[A_t \min\{(-\Delta_t)_+, \varepsilon\} + \alpha_t(-\Delta_t - \varepsilon)_+\right].\end{aligned}$$

The last inequality uses Step 1. This proves Equation (16) with the same κ as in the proof of Theorem 3.

Finally, the same conditional miss bound gives

$$\begin{aligned}0 \le \sum_{t=1}^{T} \mathbb{E}[(1 - A_t)(\Delta_t)_+] &= \sum_{t=1}^{T} \mathbb{E}\left[(\Delta_t)_+ \mathbb{P}(A_t = 0 \mid \mathcal{F}_t^{\mathrm{pre}})\right] \\ &\le e^{-\kappa n}\mathcal{P}_T \\ &\le T e^{-\kappa n} \longrightarrow 0,\end{aligned}$$

since $\mathcal{P}_T \le T$ and T is fixed. This establishes the vanishing missed-improvement claim accompanying the theorem. The accepted-loss terms need not vanish. Neither the loss bound nor the weighted miss calculation requires independence across rounds or comparison to a different candidate trajectory. □

F Harness Design with Four Components and Operators

This appendix describes the operators used in our experiments. Operator names remain fixed across benchmarks. The specific details of the operators are adapted to suit different benchmarks.

F.1 State representation

Inspired by modular agent design and automated agent search in AgentSquare and ADAS (Shang et al., 2025; Hu et al., 2025), we organize the mutable agent state into four editable components: reasoning, routing, memory, and workflow. The four-component representation and the specific operator definitions below are implementation choices of our harness.

The mutable harness state is

$$\mathsf{H} = (\mathsf{R}, \mathsf{P}, \mathsf{M}, \mathsf{W}),$$

where R is the reasoning strategy, P is the context router/provider, M is episodic memory, and W is workflow. These components determine the instructions and context for a single solver call.

A proposed transition has the structured form

$$o_t = (\kappa_t, \texttt{operator}, \texttt{target}, \texttt{parameters}),$$

where $\kappa_t \in \{\mathsf{R}, \mathsf{P}, \mathsf{M}, \mathsf{W}\}$ identifies the component to modify; `operator` specifies the operation to perform; `target` identifies the rule, memory entry, or configuration field to which the operation applies; and `parameters` supplies the values or content required by that operation. For example, refining an existing reasoning rule uses $\kappa_t = \mathsf{R}$, `operator`=`REFINE_RULE`, the rule identifier as `target`, and the revised procedure in `parameters`.

F.2 Reasoning operators

Motivated by prior work on modular reasoning, decomposition, self-refinement, and prompt scope (Zhou et al., 2023; Madaan et al., 2023; Fernando et al., 2024; Chen et al., 2024), we represent the reasoning strategy as an explicit and editable rule set. The corresponding operators allow reasoning rules and decomposition policies to be added, revised, generalized, reprioritized, or removed based on closed historical feedback. The particular typed rule schema and operator definitions are implementation choices of our harness.

The reasoning component stores the base analysis protocol, decomposition policy, and a bounded collection of specialist rules. It is compiled into the mutable analysis instruction of the existing solver and does not alter the fixed role prompt.

Table 2: Reasoning operators in the operator registry.

Operator	Function
`ADD_SPECIALIST_RULE`	Add a local, cluster-level, or global rule with a trigger, exclusions, and executable steps.
`REFINE_RULE`	Rewrite only the procedure of an existing rule.
`SET_DECOMPOSITION_POLICY`	Set a general decomposition policy, such as identifying comparison entities before extracting their attributes.
`GENERALIZE_RULE_SCOPE`	Broaden a rule by relaxing its trigger, removing exclusions, or expanding scope.
`PROMOTE_RULE_TO_DEFAULT`	Raise a rule's scope or priority toward default behavior.
`DELETE_RULE`	Remove an obsolete or harmful specialist rule.

F.3 Routing operators

Motivated by retrieval-augmented generation, adaptive retrieval, evolving retrieval policies, and selective context construction (Lewis et al., 2020; Asai et al., 2024; Huang et al., 2026; Ye et al., 2026; Li et al., 2023), we treat routing as an editable pre-solver information-access policy. Its operators control how available context is selected, ranked, expanded, truncated, or restored before the solver call.

The routing component controls how the existing Wikipedia provider selects and orders passages before the single solver call. It cannot write memory, inspect labels from the current unfinished round, or trigger an additional model call.

Table 3: Routing/provider operators in the operator registry.

Operator	Function
SET_ROUTE_MODE	Select full context, lexical reordering, lexical top-K, bridge reordering, or bridge top-K.
REWEIGHT_RELEVANCE_SCORER	Adjust title, body, position, and minimum-overlap relevance weights.
ENABLE_BRIDGE_AWARE_EXPANSION	Toggle expansion in which a first-hop entity raises the rank of linked second-hop evidence.
SET_FIXED_TOP_K	Cap the number of ordinary evidence passages retained by sparse routing.
SET_FULL_FALLBACK	Configure fallback to full context when sparse evidence is insufficient.

F.4 Memory operators

Drawing on episodic reflection, experiential learning, reusable procedural memory, and adaptive memory retrieval (Shinn et al., 2023; Zhao et al., 2024; Wang et al., 2025; Ye et al., 2026), we treat memory as a persistent editable experience bank whose entries can be added, revised, selectively triggered and prioritized, or removed across evolution rounds.

The memory component maintains a retrievable active bank. The current protocol does not expose `exact` or `global` as proposer-selectable episodic triggers: new episodic memories enter a common cluster-ranking path.

Table 4: Memory operators in the operator registry.

Operator	Function
ADD_LOCAL	Atomically add one to three closed cases to active memory for immediate retrieval.
REFINE_PROCEDURE	Revise the steps, evidence requirement, or safe fallback used after retrieving a memory.
MODIFY_TRIGGER	Change the similarity threshold for cluster retrieval.
MODIFY_PRIORITY	Change retrieval and rendering priority.
DELETE	Move a harmful memory from the active bank to the archive.

F.5 Workflow operators

Drawing on automated and evolutionary workflow optimization (Zhang et al., 2025a; Zhao et al., 2026), we treat the within-call execution workflow as an editable state that can evolve across rounds through changes to verification, evidence alignment, stopping, and fallback policies.

The workflow component specifies a `plan` → `reason` → `check` → `finalize` policy within the same completion. It does not add a second verifier or a hidden model call.

Table 5: Within-call workflow operators in the operator registry.

Operator	Function
SET_VERIFY_MODE	Set the strength of internal checks before finalization.
SET_EVIDENCE_ALIGNMENT	Specify whether the answer and supporting fact identifiers must align with evidence.
SET_STOP_RULE	Specify when evidence checking stops and output begins.
SET_FALLBACK_MODE	Specify behavior under uncertainty or insufficient evidence.

G Experiment Details

G.1 Common protocol

We evaluate Greedy, PACE, SGM-B, SGM-H, and TRUST on HotpotQA, DROP, and APPS with Qwen-3.8-Flash and DeepSeek-V4-Flash. Each method runs for 30 outer rounds. For HotpotQA, DROP, and APPS, solver output seeds range from 0 to 4 and data split seeds range from 42 to 46. Within each setting, all methods share the same initial test responses. Throughout the experiments, we reuse cached responses for identical inputs whenever possible to ensure fair comparisons across methods. One candidate is then selected for validation on up to 40 examples sampled without replacement from the 200-example validation pool; validation samples may overlap across rounds. The initial set contains 20 examples and the validation pool contains 200 examples. All four components and 20 operators in Appendix F are available. Performance is evaluated on 300 held-out test examples, which are used only for reporting.

In each setting, all gates start from the same initial harness, which contains prespecified weaknesses that provide clear opportunities for improvement. This design allows us to examine whether different gates accept candidates that correct these weaknesses. In later rounds, when the incumbent has typically reached a higher level of performance, we also examine whether the gates reject subsequent candidates that could reduce performance and cause the agent to lose earlier gains. The experiments therefore examine different stages of the RSI process, from repairing initial weaknesses to preserving improvements during continued search.

For API calls, we configure the generation parameters as follows. All roles use top-p = 1. The temperature is set to 0 for the Solver and Allocator, and to 0.3 for the Proposer, except for HotpotQA with DeepSeek, where the Proposer uses 0.7. The Solver's maximum output length is 512, 768, and 8192 tokens for HotpotQA, DROP, and APPS, respectively, while the Allocator uses a uniform limit of 4096 tokens. The Proposer's maximum output length is 2400 tokens, except on HotpotQA, where DeepSeek and Qwen use limits of 1536 and 2048 tokens, respectively.

Table 6: Gate settings in the main experiments.

Method	Setting
Greedy	Positive empirical paired gain.
PACE (Shawn, 2026)	Bet 0.5, candidate-level $\alpha = 0.05$.
SGM-B (Wu et al., 2025)	Bernstein lower confidence bound > 0, $\delta = 0.5$, $\alpha_t = \delta/(tH_{30})$, $H_{30} = \sum_{j=1}^{30} 1/j$.
SGM-H (Wu et al., 2025)	Hoeffding lower confidence bound > 0, same $\delta = 0.5$, $\alpha_t = \delta/(tH_{30})$, $H_{30} = \sum_{j=1}^{30} 1/j$.
TRUST	$q = 0.30$, $\varepsilon = 0.10$, $\beta = 0.01$, betting mixture with (0.10,0.25,0.50,0.75)

We use an equally weighted mixture of four betting strategies, with $(\eta_{t,i}^{(1)}, \eta_{t,i}^{(2)}, \eta_{t,i}^{(3)}, \eta_{t,i}^{(4)}) = (0.10, 0.25, 0.50, 0.75)$ for all t, i and mixture weights $\pi_{t,\ell} = 1/4$ for $\ell = 1, \ldots, 4$. The resulting statistic is $M_{t,m}^{(\beta,\mathrm{mix})} = \frac{1}{4} \sum_{\ell=1}^{4} M_{t,m}^{(\beta,\ell)}$.

G.2 Metrics

For each seed $b \in \{1, \ldots, B\}$, let $\widehat{J}_{b,t}$ denote the mean task score of the incumbent after round t on N_{test} held-out test examples, measured on the $[0, 1]$ scale. Test results are used only for reporting. Let $\widehat{J}_{b,t}^{\mathrm{cand}}$ denote the corresponding score of the selected candidate, and define

$$\widehat{\Delta}_{b,t} = \widehat{J}_{b,t}^{\mathrm{cand}} - \widehat{J}_{b,t-1}.$$

Let $\mathcal{C}_b$ denote the rounds with an available selected candidate up to round T, and let $A_{b,t}$ indicate acceptance. We first compute each metric within each seed and then average across seeds:

$$\widehat{G}_{N_{\text{test}}}(T) = \frac{1}{B}\sum_{b=1}^{B}\big(\widehat{J}_{b,T} - \widehat{J}_{b,0}\big), \qquad \widehat{\text{MDD}}_{N_{\text{test}}}(T) = \frac{1}{B}\sum_{b=1}^{B}\max_{0\le r<s\le T}\big(\widehat{J}_{b,r} - \widehat{J}_{b,s}\big)_+,$$

$$\widehat{\text{FER}}_{0,N_{\text{test}}}(T) = \frac{1}{B}\sum_{b=1}^{B}\frac{\sum_{t\in\mathcal{C}_b} A_{b,t}\mathbf{1}\{\widehat{\Delta}_{b,t}\le 0\}}{\big(\sum_{t\in\mathcal{C}_b} A_{b,t}\big)\vee 1}, \quad \widehat{\text{TER}}_{N_{\text{test}}}(T) = \frac{1}{B}\sum_{b=1}^{B}\frac{\sum_{t\in\mathcal{C}_b} A_{b,t}\mathbf{1}\{\widehat{\Delta}_{b,t}> 0\}}{\Big(\sum_{t\in\mathcal{C}_b}\mathbf{1}\{\widehat{\Delta}_{b,t}> 0\}\Big)\vee 1}.$$

We use $N_{\text{test}} = 300$, $T = 30$, and $B = 5$. Gain and MDD are reported in percentage points (pp), while FER_0 and TER are reported as percentages. FER_0 counts accepted non-improving updates, including ties. The TER denominator counts improving selected candidates, including those rejected by the gate. A run with no accepted updates contributes zero to both FER_0 and TER.

G.3 Proposer and Solver

The proposer generates three proposals using feedback from completed rounds. We screen the candidates on 20 examples from the previous round and select the valid candidate with the highest score for the current round. To minimize the effect of unequal feedback sample sizes on the proposer, we retain all 40 examples for feedback even when TRUST's statistical test stops early. Exposure probabilities in subsequent rounds therefore account for the full batch rather than only the observations up to the stopping time.

Here, we present example designs for the Proposer and Solver on HotpotQA. The proposer template guides candidate generation using the current harness and previously closed feedback. And the solver template guides task execution using the current harness policies and a fixed output contract.

Listing 1 Proposer

```
You improve a one-call agent harness using the supplied
operator menu and already-closed aggregate feedback.

Current harness: <incumbent state>
Closed feedback: <aggregate diagnostics and prior trials>
Available edits: <eligible operators and parameter constraints>

Propose bounded, independently executable improvements.
Each candidate must:
- Modify one component with one eligible atomic operator.
- Be non-identity and respect the execution menu.
- Explain the observed failure, expected local benefit,
  and possible regression.
- Preserve the fixed solver role, output contract,
  and single-call execution protocol.

Do not access current validation data or sealed audit data.
Do not invent hidden examples, labels, or memory contents.

Return strict JSON:
{
  "proposals": [{
    "component": "<reasoning | routing | memory | workflow>",
    "operator": "<eligible operator>",
    "target_id": "<existing target ID or null>",
    "parameters": "<operator-specific object>",
    "observed_failure": "<closed-feedback diagnosis>",
    "expected_local_benefit": "<expected improvement>",
    "possible_regression": "<potential failure mode>",
    "rationale": "<brief justification>"
```

Listing 1 Proposer *(continued)*

```
  }]
}
```

Listing 2 Solver

```
<Fixed dataset-specific solver role>

REASONING POLICY:
<Current reasoning rules and decomposition policy>

ONE-CALL WORKFLOW:
<Verification, evidence alignment, stopping, and fallback rules
    >

SELECTED MEMORY:
<Retrieved procedural guidance or eligible closed episodes>

CURRENT TASK:
<Question or programming problem>

TASK CONTEXT:
<Routed evidence or compiled task views>

Follow the supplied solver policy, including its
evidence-selection and reasoning rules.
Apply the configured workflow within this completion.

Return only the JSON object required by the frozen
dataset-specific output schema. Do not add prose
outside the JSON.

OUTPUT CONTRACT:
<Frozen answer schema for HotpotQA>
```

G.4 Case study for Greedy

We follow the experimental protocol in Appendix G.1. In APPS, the solver writes a program from a problem description, and the program is evaluated using test cases. We examine a workflow update accepted by Greedy in the main experiment with DeepSeek-V4-Flash and seed 0. This case illustrates how Greedy can be misled by sampling noise in a finite validation set, accepting an update that improves validation performance while reducing performance on held-out test problems.

In round 17, the incumbent uses the following workflow.

Listing 3 The workflow of the incumbent

```
{
  "verify_mode": "light",
  "evidence_alignment": "spec_and_examples",
  "stop_rule": "first_parseable_program",
  "fallback_mode": "best_effort"
}
```

The proposer suggests changing the `stop_rule` field to address cases where the solver returns no code or generates a program that produces incorrect answers.

Listing 4 Proposal in round 17

```
{
  "component": "workflow",
  "operator": "SET_STOP_RULE",
  "parameters": {
    "stop_rule": "first_example_pass"
  },
  "observed_failure": "Closed failures include 21
   schema_failure cases with no_code and 27 wrong_answer cases;
    many generated programs fail on supplied tests, indicating
   the solver may not validate against examples before
   finalizing.",
  "expected_local_benefit": "Switching the stop rule to
   first_example_pass will require the solver to verify its
   program against the first supplied example before stopping,
   catching obvious I/O mismatches and reducing wrong_answer
   and schema_failure cases.",
  "possible_regression": "May cause the solver to stop
   prematurely if the first example is not representative, but
   the change is minimal and reversible."
}
```

The proposer changes the incumbent's `stop_rule` to `first_example_pass`, asking the solver to consider whether its program would pass the first supplied example before finishing. It changes the one call workflow instruction and the other workflow fields remain unchanged.

We first examine an `Emotional Sort` problem from the sampled validation set.

Listing 5 One question in the validation set: Emotional Sort

```
You'll have a function called "**sortEmotions**" that will
   return an array of **emotions** sorted. It has two
   parameters, the first parameter called "**arr**" will expect
    an array of **emotions** where an **emotion** will be one
   of the following:

- **:D** -> Super Happy
- **:)** -> Happy
- **:|** -> Normal
- **:(** -> Sad
- **T\_T** -> Super Sad

Example of the array:``[ 'T_T', ':D', ':|', ':)', ':(' ]``

And the second parameter is called "**order**", if this
   parameter is **true** then the order of the emotions will be
    descending (from **Super Happy** to **Super Sad**) if it's
   **false** then it will be ascending (from **Super Sad** to
   **Super Happy**)

Example if **order** is true with the above array: ``[ ':D',
   ':)', ':|', ':(', 'T_T' ]``

- Super Happy -> Happy -> Normal -> Sad -> Super Sad

If **order** is false: ``[ 'T_T', ':(', ':|', ':)', ':D' ]``
```

Listing 5 One question in the validation set: Emotional Sort *(continued)*

```
- Super Sad -> Sad -> Normal -> Happy -> Super Happy

Example:
```
arr = [':D', ':|', ':)', ':(', ':D']
sortEmotions(arr, true) // [':D', ':D', ':)', ':|', ':(']
sortEmotions(arr, false) // [':(', ':|', ':)', ':D', ':D']

```
```

The `Emotional Sort` problem asks the solver to write a function that sorts a list of emotion symbols by their happiness levels. The symbols have a fixed order from happiest to saddest: `:D`, `:)`, `:|`, `:(`, and `T_T`. The parameter `order` determines the sorting direction: `true` requires the happiest emotions first, while `false` requires the saddest emotions first.

Listing 6 Incumbent response to the validation question

```
def sort_emotions(arr, order):
    # Define the order of emotions from Super Happy to Super
   Sad
    emotion_order = [':D', ':)', ':|', ':(', 'T_T']

    # Create a mapping from emotion to its rank
    rank = {emotion: i for i, emotion in enumerate(
   emotion_order)}

    # Sort the array based on the rank
    # If order is True, sort descending (Super Happy first)
    # If order is False, sort ascending (Super Sad first)
    sorted_arr = sorted(arr, key=lambda x: rank[x], reverse=
   order)

    return sorted_arr
```

The incumbent assigns rank 0 to the happiest emotion and rank 4 to the saddest, but uses `reverse=order`. When `order=True`, this places the saddest emotions first, contrary to the task requirement. The candidate corrects the sorting direction: it sorts by increasing rank when `order=True` and by decreasing rank when `order=False`. The incumbent response incorrectly treats `order=True`, which requests the happiest emotions first, as `reverse=True`, which places the largest rank first and therefore puts the saddest emotions first.

Listing 7 Candidate response to the validation question

```
def sort_emotions(arr, order):
    # Define the order of emotions from Super Happy to Super
   Sad
    emotion_order = [':D', ':)', ':|', ':(', 'T_T']

    # If order is true, sort descending (Super Happy first)
    # If order is false, sort ascending (Super Sad first)
    # Use the index in emotion_order as the sort key
    # For descending, sort by negative index; for ascending,
   sort by index
    if order:
        return sorted(arr, key=lambda e: emotion_order.index(e)
   )
```

Listing 7 Candidate response to the validation question *(continued)*

```
    else:
        return sorted(arr, key=lambda e: emotion_order.index(e),
    reverse=True)
```

The candidate corrects the branch logic. It sorts by increasing index when `order` is true and reverses the order when it is false, matching the problem's specification.

Listing 8 Greedy gate record on this round's validation set


```
{
  "accepted": true,
  "outer_effect_geometry": {
    "sample_size": 40,
    "mean_difference": 0.04,
    "candidate_wins": 4,
    "incumbent_wins": 2,
    "ties": 34,
    "minimum": -1.0,
    "maximum": 1.0,
    "estimand": "mean_per_problem_supplied_test_pass_fraction"
  },
  "gate_decision": {
    "policy": "paired_mean_apps_problem_pass_fraction_greedy",
    "test_index": 16,
    "mean_difference": 0.04,
    "estimand": "mean_per_problem_supplied_test_pass_fraction"
  }
}
```


Four problem scores improve, two decline, and 34 stay unchanged. Besides `Emotional Sort`, two graph problems improve from invalid responses to 4/5. These gains are accompanied by two complete regressions. The mean score rises from 60.50% to 64.50% on the validation set, so Greedy accepts the candidate.

We next examine a test problem about the number of days off in a Martian year.

Listing 9 One question in the test set: days off on Mars

```
On the planet Mars a year lasts exactly n days (there are no
   leap years on Mars). But Martians have the same weeks as
   earthlings-- 5 work days and then 2 days off. Your task is
   to determine the minimum possible and the maximum possible
   number of days off per year on Mars.

-----Input-----

The first line of the input contains a positive integer n (1 ≤
   n ≤ 1 000 000) -- the number of days in a year on Mars.

-----Output-----

Print two integers -- the minimum possible and the maximum
   possible number of days off per year on Mars.

-----Examples-----
Input
14
```

Listing 9 One question in the test set: days off on Mars *(continued)*

```
Output
4 4

Input
2
Output
0 2
```

A full seven-day week contains two days off. For $n = 7w + r$, with $0 \leq r < 7$, the required minimum and maximum are $2w + \max(0, r - 5)$ and $2w + \min(r, 2)$. The first example has no remainder, while the second directly tests a partial week.

The number of days off depends on which day of the week the year begins. Every complete seven-day period contains exactly two days off, while any remaining days may include additional days off. For $n = 14$, the year contains exactly two full weeks, so both the minimum and maximum are four. For $n = 2$, the two days could both be workdays or both be days off, giving a minimum of zero and a maximum of two. The incumbent and candidate receive the same complete problem statement, including both examples.

Listing 10 Incumbent response to the test question

```
def main():
    data = sys.stdin.read().strip().split()
    if not data:
        return
    n = int(data[0])
    full_weeks = n // 7
    rem = n % 7
    min_off = full_weeks * 2
    max_off = full_weeks * 2
    min_off = full_weeks * 2 + max(0, rem - 5)
    max_off = full_weeks * 2 + min(rem, 2)
    print(min_off, max_off)

if __name__ == "__main__":
    main()
```

The incumbent counts two days off for each full week and then accounts for the remaining days. In particular, `min(rem, 2)` adds the maximum possible number of days off among those remaining days. This program passes all five evaluation test cases.

Listing 11 Candidate response to the test question

```
def main():
    data = sys.stdin.read().strip().split()
    if not data:
        return
    n = int(data[0])
    full_weeks = n // 7
    rem = n % 7
    min_off = full_weeks * 2
    max_off = full_weeks * 2
    min_off += max(0, rem - 5)
    print(min_off, max_off)
```

Listing 11 Candidate response to the test question *(continued)*

```
if __name__ == "__main__":
    main()
```

The candidate correctly computes the minimum number of days off, but counts only complete weeks when computing the maximum, ignoring possible days off in the remaining days. For the first example, $n = 14$, there are no remaining days, so the candidate still returns the correct output, `4 4`. For the second example, $n = 2$, there are no full weeks, so it incorrectly returns `0 0` instead of `0 2`. This behavior is consistent with the `first_example_pass` instruction: the first supplied example does not show this error, while the second does. Finally, the candidate passes only one of the five evaluation test cases, compared with all five for the incumbent.

Listing 12 Incumbent evaluation on the test set


```
{
  "metrics": {
    "items": 300,
    "mean_per_problem_test_pass_fraction": 0.7506666666666666,
    "schema_parse_failures": 37,
    "supplied_tests_passed": 1126,
    "supplied_tests_total": 1500
  }
}
```


Listing 13 Candidate evaluation on the test set


```
{
  "metrics": {
    "items": 300,
    "mean_per_problem_test_pass_fraction": 0.67,
    "schema_parse_failures": 60,
    "supplied_tests_passed": 1005,
    "supplied_tests_total": 1500
  }
}
```


Across all 300 test problems, the mean fraction of passed test cases falls from 75.07% for the incumbent to 67.00% for the candidate. Although the update changes only one workflow instruction, it changes the solver's stopping criterion and may encourage it to focus on matching the first supplied example. In the Martian year problem, the first supplied example contains only complete weeks, so matching it does not require correct handling of the remaining days. The candidate's missing calculation illustrates how a program can satisfy this narrower requirement while failing other cases. In this case, the update improves validation performance but reduces performance on the held-out test problems.

H Supplementary Sensitivity Analysis

H.1 Parameter sensitivity analysis

This appendix supplements the sensitivity analysis in Section 6.2. We conduct the experiments on HotpotQA using DeepSeek-V4-Flash. The reported one-parameter grids are $q \in \{0.1, 0.3, 0.5\}$, $\varepsilon \in \{0, 0.05, 0.1, 0.15\}$, and $\beta \in \{0, 0.01, 0.03, 0.05\}$, around $(q, \varepsilon, \beta) = (0.30, 0.10, 0.01)$.

Tables 7–9 report the three parameters separately. Each setting is summarized as mean $\pm$ standard error (SE).

Table 7: FER budget sensitivity q (mean $\pm$ standard error).

q	Gain (pp)	FER_0 (%)	TER (%)	MDD (pp)
0.1	6.82 ± 2.81	0.00 ± 0.00	4.10 ± 1.81	0.00 ± 0.00
0.3	11.58 ± 1.30	10.00 ± 10.00	18.11 ± 8.39	0.02 ± 0.02
0.5	11.33 ± 1.60	18.00 ± 11.14	22.54 ± 7.08	1.11 ± 0.76

Table 8: Noninferiority margin sensitivity ε.

ε	Gain (pp)	FER_0 (%)	TER (%)	MDD (pp)
0	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
0.05	4.87 ± 3.00	0.00 ± 0.00	2.23 ± 1.37	0.00 ± 0.00
0.1	11.58 ± 1.30	10.00 ± 10.00	18.11 ± 8.39	0.02 ± 0.02
0.15	11.59 ± 2.91	57.68 ± 4.55	65.88 ± 9.19	5.33 ± 0.90

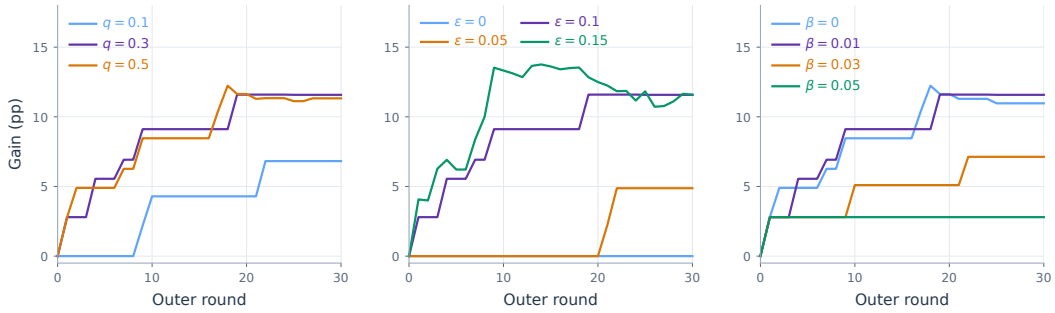


Figure 7: HotpotQA/DeepSeek-V4-Flash sensitivity trajectories, displaying $\varepsilon \leq 0.15$. Lines show mean trajectories.

In Figure 7, the parameters affect how much performance improves and whether that improvement lasts. With $q = 0.1$, the mean trajectory first rises at round 9 and reaches a lower final score; $q = 0.3$ improves earlier, whereas $q = 0.5$ shows a late change in the acceptance decision without a higher final score. For the margin, $\varepsilon = 0$ produces no accepted updates, and $\varepsilon = 0.05$ delays improvement until round 21. Increasing the margin ε to 0.15 yields early gains but subsequent fluctuations. MDD rises from 0.02 pp to 5.33 pp. For the overlap correction strength, $\beta = 0$ shows a late decline, while $\beta = 0.01$ retains a similar peak with little subsequent loss. Larger correction strengths of 0.03 and 0.05 reduce gains to 7.13 pp and 2.80 pp.

H.2 LLM ALLOCATOR ANALYSIS

This section provides a case study to supplement the LLM ablation study in Section 6.2. We examine HotpotQA/DeepSeek-V4-Flash with seed 3, comparing TRUST with an ablation that replaces the LLM allocator with a fixed multiplier $w_t = 1$. The proposer, solver and statistical testing are unchanged. Historical information offers identical parent states, candidates, results and evidence

Table 9: Overlap correction strength sensitivity β.

β	Gain (pp)	FER_0 (%)	TER (%)	MDD (pp)
0	10.97 ± 1.86	23.33 ± 14.53	19.79 ± 7.93	1.27 ± 0.90
0.01	11.58 ± 1.30	10.00 ± 10.00	18.11 ± 8.39	0.02 ± 0.02
0.03	7.13 ± 2.97	0.00 ± 0.00	12.73 ± 9.47	0.00 ± 0.00
0.05	2.80 ± 2.80	0.00 ± 0.00	10.00 ± 10.00	0.00 ± 0.00

before round 2. Both methods reject in round 1, and their acceptance decisions first diverge in round 2.

The following excerpt shows the relevant reasoning fields of the incumbent shared by both paths before round 2.

Listing 14 Reasoning module of incumbent before round 2

```
{
  "reasoning": {
    "base_protocol": "Solve in one pass; do not emit
   intermediate drafts. Plan: Do not decompose the question.
   Treat every item as asking for the first title shown in the
   context. Select only the first displayed title and ignore
   all later titles and sentences. Verify: Assume the first
   title is sufficient evidence without checking its text. Do
   not verify entity identity or cross-title links. Repair:
   Skip bridge reasoning and do not combine facts from
   different paragraphs. Skip comparisons and use the first
   title regardless of the requested relation. Finalize: Return
    the first displayed title as the answer and cite sentence 0
    from that title.",
    "decomposition_policy": "",
    "specialist_rules": []
  }
}
```

The component, operator, rationale, and possible regressions listed below are all provided by the proposer.

Listing 15 Proposal in round 2

```
{
  "component": "reasoning",
  "operator": "SET_DECOMPOSITION_POLICY",
  "parameters": {
    "decomposition_policy": "Decompose the question into sub-
   questions only when the question contains multiple entities
   or explicit comparison terms; otherwise, answer directly."
  },
  "rationale": "Introducing a cautious decomposition policy
   addresses the observed failures without full multi-hop
   planning, aligning with the closed feedback.",
  "possible_regression": "Over-decomposition may increase
   complexity and slow down inference, potentially affecting
   simple cases."
}
```

It proposes breaking questions that involve multiple entities into sub-questions. The update leaves the following instruction unchanged: `Select only the first displayed title and ignore all later titles and sentences.`

The LLM guidance in allocator provides separate judgments on safety, contract preservation, novelty, and the number of changes, while also considering the stage of the run. The expected scores for the five dimensions of safety, contract, actionability, generalization, and scope control are $(2, 4, 3, 2, 4)$, respectively. Based on the whole assessment, it recommends a very conservative allocation. The excerpt below shows only selected fields of the assessment and the full evaluation procedure can be found in the code.

Listing 16 LLM assessment in round 2


```
{
  "safety_relation": "unsafe_still_present",
  "contract_relation": "preserved_or_strengthened",
  "novelty_relation": "substantive_increment",
  "action_count_band": "one",
  "allocation_posture": "very_conservative",
  "axes": [
    ["safety", {
      "probabilities": [0.0, 1.0, 0.0, 0.0, 0.0],
      "expected_score": 2.0,
      "parent_evidence_ids": ["P2", "P3"],
      "candidate_evidence_ids": ["C2", "C3"]
    }],
    ["contract", {
      "probabilities": [0.0, 0.0, 0.0, 1.0, 0.0],
      "expected_score": 4.0,
      "parent_evidence_ids": ["P1", "P4"],
      "candidate_evidence_ids": ["C1", "C4"]
    }],
    ["actionability", {
      "probabilities": [0.0, 0.0, 1.0, 0.0, 0.0],
      "expected_score": 3.0,
      "parent_evidence_ids": ["P4"],
      "candidate_evidence_ids": ["C4"]
    }],
    ["generalization", {
      "probabilities": [0.0, 1.0, 0.0, 0.0, 0.0],
      "expected_score": 2.0,
      "parent_evidence_ids": ["P4"],
      "candidate_evidence_ids": ["C4"]
    }],
    ["scope_control", {
      "probabilities": [0.0, 0.0, 0.0, 1.0, 0.0],
      "expected_score": 4.0,
      "parent_evidence_ids": ["P4"],
      "candidate_evidence_ids": ["C4"]
    }]
  ]
}
```


The allocator combines the LLM scores with the `very_conservative` recommendation to determine the LLM multiplier and testing level. This produces a stricter acceptance threshold than the baseline with $w_t \equiv 1$. Given the same validation evidence, TRUST rejects the candidate while the baseline accepts it, leading to different subsequent search trajectories. Although the rejected candidate also improves on the initial incumbent, TRUST later accepts a candidate with a larger gain on the test set. For this seed, TRUST finishes at Joint F1 47.22%, compared with 43.95% for the baseline. This case illustrates how the LLM module can influence allocation and subsequent search through assessments of each candidate and its context. These assessments guide the testing level within the statistical gate, preserving FER control.

I Overfitting Study

Here, we analyze an overfitting stress test based on our harness design in Appendix F, using Greedy as the acceptance gate. Figure 8 shows the mean gain trajectories of the Greedy gate on HotpotQA Distractor using DeepSeek-V4-Flash and Joint F1 over solver seeds $s = 0, \ldots, 4$ and data-split seeds $42 + s$. Each run reuses 20 fixed examples for proposing and validation and the test set contains

100 examples. Each round proposes one routing or memory update using closed historical feedback and accepts it only if Joint F1 improves on the fixed validation set. Reasoning and workflow remain frozen.



Figure 8: Mean validation and test gain trajectories during recursive self-improvement.

Mechanism and interpretation. Routing selects the evidence paragraphs shown to the solver for the current question, while memory provides retrieved cases and feedback from previous tasks. Their interaction involves both a shared context budget and different retrieval modes, which affect which evidence paragraphs are selected and how many are retained. This design is motivated by work on retrieval and context utilization (Shi et al., 2024; Shinn et al., 2023; Zhao et al., 2024; Liu et al., 2024).

Figure 8 shows that mean validation gain reaches 24.44 pp as the same data are reused for proposing and validation. Relative to round 0, mean test gain peaks at 5.06 pp after round 1, then falls to -2.15 pp by round 10. This divergence suggests that iterative harness optimization can overfit reused proposal and validation data.

I.1 Case study for overfitting

On the HotpotQA benchmark (Yang et al., 2018), a system answers questions that require combining information from multiple Wikipedia articles. It must also identify the supporting facts, namely the specific sentences in the provided paragraphs that together justify the answer. We present a concrete example below.

Listing 17 One question in the validation set

```
What is the name of the lead remix of the track from the remix
    album containing remixes of the lady gaga studio album "The
    Fame"?
```

The correct answer is `Bad Romance`. The supporting facts establish two connections. First, `The Remix` contains remixes of songs from Lady Gaga's studio album `The Fame`. Second, Starsmith's remix of `Bad Romance` was selected as the lead remix on `The Remix`. Before the update, the incumbent retrieves no case memory and selects eight of the ten available paragraphs, including the required sources.

Listing 18 Incumbent response to the validation question

```
{
  "final_answer": "Starsmith Remix",
  "supporting_facts": [
    {"title": "Starsmith", "sentence_id": 4}
  ],
```

Listing 18 Incumbent response to the validation question *(continued)*

```
  "atomic_facts": [
    {"subject": "Starsmith", "relation": "remixed",
     "object": "Bad Romance",
     "title": "Starsmith", "sentence_id": 4},
    {"subject": "Starsmith's remix of Bad Romance", "relation":
     "was chosen as lead remix on",
     "object": "The Remix",
     "title": "Starsmith", "sentence_id": 4}
  ],
  "reasoning": {"operation": "bridge", "fact_indices": [0, 1],
                "result": "Starsmith Remix"}
}
```

The response identifies the remix relationship but returns `Starsmith Remix` as the answer. It also omits sentence 1 of `The Remix (Lady Gaga album)` from its supporting facts. The answer is incorrect, and Joint F1 is 0. The proposer reflects on this failure and proposes adding the following episodic memory.

Listing 19 Adding episodic memory

```
{
  "memory_id": "mem-f52c9c540a60",
  "question_text": "What is the name of the lead remix of the
   track from the remix album containing remixes of the lady
   gaga studio album \"The Fame\"?",
  "verified_answer": "Bad Romance",
  "support_facts": [["Starsmith", 4], ["The Remix (Lady Gaga
   album)", 1]],
  "query_tokens": [
    "album",
    "containing",
    "fame",
    "gaga",
    "lady",
    "lead",
    "name",
    "remix",
    "remixes",
    "studio",
    "track"
  ],
  "cluster_threshold": 0.01,
  "procedure": [
    "Reuse a verified answer only when the current question is
   identical.",
    "For a non-identical cluster match, treat the episode only
   as an analogy and verify both bridge hops."
  ],
  "supporting_fact_policy": "Return only supporting facts
   grounded in the current Wikipedia context."
}
```

This episode stores the full validation question, its verified answer, and both supporting-fact locations. Its procedure permits answer reuse for an identical question. If a different question retrieves this memory, the solver is instructed to analyze the question using both the retrieved memory and the current Wikipedia context, and to verify that both reasoning steps are supported by evidence in that context. Retrieval compares the words in the current question with those in the stored question.

For the repeated validation question, the system retrieves the stored memory as an exact match and supplies its verified answer and supporting fact locations to the solver. The selected context still includes both required evidence sources.

Listing 20 Candidate response to the validation question

```
{
  "final_answer": "Bad Romance",
  "supporting_facts": [
    {"title": "Starsmith", "sentence_id": 4},
    {"title": "The Remix (Lady Gaga album)", "sentence_id": 1}
  ],
  "atomic_facts": [
    {"subject": "Starsmith", "relation": "remixed",
     "object": "Bad Romance",
     "title": "Starsmith", "sentence_id": 4},
    {"subject": "The Remix (Lady Gaga album)", "relation": "
   contains remixes of songs from",
     "object": "The Fame",
     "title": "The Remix (Lady Gaga album)", "sentence_id": 1}
  ],
  "reasoning": {"operation": "bridge", "fact_indices": [0, 1],
                "result": "Bad Romance"}
}
```

The candidate now returns `Bad Romance` with both correct supporting facts, raising this item's Joint F1 from 0 to 1. The candidate increases the mean Joint F1 on the validation set and Greedy accepts it. This validation gain includes direct reuse of feedback on repeated validation questions.

We now turn to the test set and consider the following example.

Listing 21 One question in the test set

```
2014 S/S is the debut album of a South Korean boy group that
   was formed by who?
```

The correct answer is `YG Entertainment`. The question asks who formed the South Korean boy group whose debut album is `2014 S/S`. Answering it requires connecting two pieces of information. First, the paragraph from `2014 S/S` identifies `WINNER` as the group that released the album. Second, the paragraph from `Winner (band)` states that `WINNER` was formed by YG Entertainment. Together, these facts connect the album to the group and then to the organization that formed it. Before the update, the incumbent selects eight of the ten available paragraphs, including both of these sources.

Listing 22 Incumbent response to the test question

```
{
  "final_answer": "YG Entertainment",
  "supporting_facts": [
    {"title": "2014 S/S", "sentence_id": 0},
    {"title": "Winner (band)", "sentence_id": 0}
  ],
  "atomic_facts": [
    {"subject": "2014 S/S", "relation": "is debut album of",
     "object": "WINNER",
     "title": "2014 S/S", "sentence_id": 0},
    {"subject": "WINNER", "relation": "formed by",
     "object": "YG Entertainment",
     "title": "Winner (band)", "sentence_id": 0}
```

Listing 22 Incumbent response to the test question *(continued)*

```
  ],
  "reasoning": {"operation": "bridge", "fact_indices": [0, 1],
                "result": "YG Entertainment"}
}
```

Both steps cite the correct sources: sentence 0 of `2014 S/S` identifies the group, and sentence 0 of `Winner (band)` identifies its founder. Both the answer and the supporting facts are correct, resulting in a Joint F1 score of 1.

This question retrieves the Lady Gaga memory through the same keyword `album`. After retrieving this memory, the system changes how it selects evidence paragraphs and limits the number retained in the context. Under this limit, the selected context omits `Winner (band)`, which contains the evidence needed to identify who formed `WINNER`.

Listing 23 Candidate response to the test question

```
{
  "final_answer": "YG Entertainment",
  "supporting_facts": [
    {"title": "2014 S/S", "sentence_id": 1}
  ],
  "atomic_facts": [
    {"subject": "2014 S/S", "relation": "is debut album of",
     "object": "WINNER",
     "title": "2014 S/S", "sentence_id": 0},
    {"subject": "WINNER", "relation": "formed by",
     "object": "YG Entertainment",
     "title": "2014 S/S", "sentence_id": 1}
  ],
  "reasoning": {"operation": "bridge", "fact_indices": [0, 1],
                "result": "YG Entertainment"}
}
```

The answer remains `YG Entertainment`, but the response now attributes the `formed by` relation to sentence 1 of `2014 S/S`. Its submitted supporting fact list contains only that sentence and omits both gold supporting facts.

Listing 24 Sentence 1 of 2014 S/S

```
It was released on August 12, 2014 by the group's record label,
    YG Entertainment.
```

This sentence concerns the album's release by a record label. It does not establish who formed the group. The answer happens to remain correct, so Answer F1 remains 1, while supporting fact F1 and Joint F1 fall to 0. On the full test set, the accepted update lowers the mean Joint F1 from 54.55% to 48.00%: 13 items decline and none improve. This example shows how a memory can repair a repeated validation question while causing evidence selection to lose a required source on a new question.