WHEN ARE AGENTS IRREPLACEABLE? A TASK-LEVEL THEORY OF ACQUISITION AND COMPUTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

The success of an agent does not establish which capabilities the task requires. We introduce task-level resource diagnosis to determine when an agent is irreplaceable and when simpler configurations suffice. For each fixed-target task, we compare class-optimal risks with the interface and query budget held fixed across solver classes, allowing solvers in each class to redesign their solutions within the specified resource constraints. Exact evidence-coverage frontiers and processing risk bounds, together with successful agent constructions, establish conditions under which acquisition and enhanced computation are jointly necessary. Computation shapes both evidence acquisition and the answers derived from observations. Broader coverage need not reduce risk, and complete evidence can still require stronger computation. These mechanisms explain why tasks requiring the same agent capabilities can benefit from different resource investments. Controlled tasks test these predictions. Natural-language experiments show how interventions on acquisition and processing improve answer quality and affect resource use. The framework connects agent irreplaceability to task structure, identifying indispensable capabilities and guiding interventions on the limitations that prevent success.

1 INTRODUCTION

Large language model (LLM) agents select actions from an interaction history and use the resulting observations to guide subsequent decisions. Tool use expands the actions available to an agent, while feedback memory and search-based planning support decisions across multiple steps (Schick et al., 2023; Shinn et al., 2023; Zhou et al., 2024a). Applications include web navigation, software repository repair, and multistage scientific research (Zhou et al., 2024b; Yang et al., 2024; Lu et al., 2026). Reasoning and Acting (ReAct) organizes reasoning traces and environmental actions in an alternating sequence, so evidence obtained through one action can inform the next decision (Yao et al., 2023b). Throughout this process, a system must decide both how to use available information and when and where to acquire new information. Several capabilities contribute to a successful execution, making their individual roles difficult to infer from the outcome. The expanding range of agent applications makes understanding **when a given task requires an agent and why** increasingly important.

Capability, however, does not imply necessity, because many tasks already admit a direct procedure that meets the required performance target. **Figure 1** illustrates three simpler solution procedures. A given flight record supports direct lookup (A), whereas checking the status of a known instrument requires a fixed external query (B). Even a computationally demanding task may be solved entirely from the initial input. A fixed graph algorithm computes the shortest path in a fully specified graph with nonnegative edge weights (C) (Dijkstra, 1959), while specialized methods solve convex optimization problems and provide optimality certificates (Boyd & Vandenberghe, 2004). ConstraintBench also reports substantial errors in direct LLM solutions to optimization problems with solver-verified optima (Tso et al., 2026). When a simpler procedure already meets the target, the benefit of a full agent system must be weighed against the additional cost, latency, and potential failure modes of that system (Kapoor et al., 2025). Panel D combines acquisition and computation by retrieving external measurements and fitting temperature trends to compare heating rates. The full agent configuration is necessary when the full configuration meets the target but every allowed

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

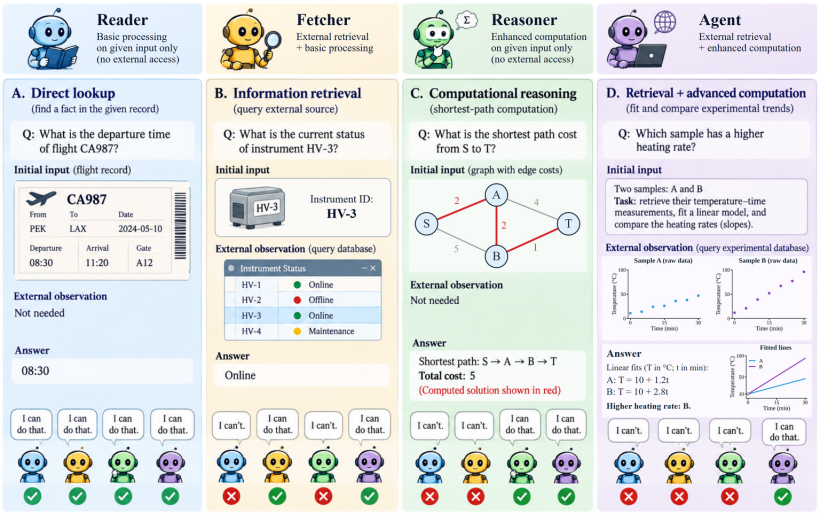


Figure 1: Illustrative tasks and sufficient capability configurations. Basic processing in these examples reads and copies fields. Measurements in panel D are synthetic.

alternative lacking acquisition or enhanced computation fails under the same task, interface, and query budget. Both fixed and adaptive workflows enter this comparison according to the capabilities they use.

Existing work approaches necessity through system efficiency and decisions under resource constraints. Simpler pipelines and optimized workflows show that changing the solver can improve accuracy and cost (Xia et al., 2025; Kapoor et al., 2025; Zhang et al., 2025). Training and inference-time control improve how models combine reasoning with external calls (Jin et al., 2025; Qian et al., 2025; Fang et al., 2026; Sun et al., 2026). Such gains make solver redesign central to necessity comparisons. Decision theory and usable information provide a basis for these comparisons by evaluating information and computation through the decision rules available to a solver (Blackwell, 1953; Russell & Wefald, 1991; Russell & Subramanian, 1995; Xu et al., 2020). Theory of Agent makes tool necessity explicit by identifying tasks beyond a model population's internal capabilities (Wang et al., 2026). A task-level certificate must extend this comparison to every solver permitted by the resource constraints, including alternative acquisition and processing strategies. The remaining challenge is therefore to identify task structures that make external acquisition and enhanced computation jointly necessary at a target risk when the interface and query budget are fixed across solver classes.

We introduce *task-level resource diagnosis* to determine when the full agent configuration is necessary and when fewer capabilities suffice. Class-optimal risks allow alternative solvers in each class to redesign their solutions within the specified resource limits. We connect these requirements to limitations in locating evidence and using acquired records. Computation governs both query selection and the conclusions drawn from observations. Tasks requiring the same agent capabilities can therefore call for different resource investments. The main contributions of this paper are as follows.

- **Task-level agent irreplaceability.** We compare class-optimal risks while fixing the interface and query budget across solver classes, allowing alternative solvers to redesign their entire solutions. The framework identifies minimal sufficient configurations and separates task requirements from the performance of a particular implementation.
- **Structural limits on agent simplification.** We establish how retained information and query budgets limit evidence coverage, and how computation limits the use of acquired evidence. Exact coverage frontiers and processing risk bounds, together with successful agent constructions, certify joint necessity through these distinct mechanisms.
- **Task-dependent resource investment.** We derive intervention predictions for retained context, query budgets, and processing resources. Controlled tasks and experiments on FinQA and HotpotQA connect these predictions to answer quality and resource use, show-

ing how processing can reduce evidence requirements and acquisition strategies can trade material use against model calls.

2 TASK-LEVEL RESOURCE DIAGNOSIS

We compare the full agent configuration with alternatives for a fixed task, interface, and query budget. Alternative solvers in each class may redesign their solutions within the specified resource limits. Class-optimal risk captures the best performance under each configuration. A capability is necessary when the remaining resources cannot meet the target risk.

2.1 FIXED TASKS AND MATCHED SOLVER CLASSES

Definition 1 (Fixed-target task)**.** A *fixed-target task* is specified by

$$\tau = (\Omega, P, O_0, \mathcal{I}, \mathcal{D}, \ell), \qquad \ell(d, \omega) = \ell_0(d, y^\star(\omega)). \tag{1}$$

Here $\omega \sim P$ is a latent environment and $O_0 : \Omega \to \mathcal{O}_0$ is the measurable initial-observation map. The interface $\mathcal{I}$ specifies allowed actions and the conditional distribution of each return given the environment, history, and action. A decision $d \in \mathcal{D}$ incurs loss $\ell(d, \omega) \in [0, 1]$. Acquisition changes visible evidence while the target $y^\star$ and evaluation remain fixed.

Within a task, a solver chooses what to query, when to stop, and how to answer. We represent these choices through three control rules.

Definition 2 (Solver)**.** A solver is a possibly randomized controller $G = (\pi, T, \psi)$ with acquisition policy π, stopping rule T, and terminal rule ψ. Starting from $H_0 = O_0$, the policy selects an allowed action A_s from H_{s-1} and receives O_s through $\mathcal{I}$. The resulting history is expressed as

$$H_t = (O_0, A_1, O_1, \ldots, A_t, O_t). \tag{2}$$

Actions, stopping, and the terminal decision D_G use only acquired information, accessible state, and internal randomness independent of ω. The rule ψ acts on the accessible stopped history and randomness. Let $\mathcal{G}_\tau$ contain all such solvers.

Access to the interface permits new observations. Computational capability governs how the controller uses them to query, stop, and answer. Let $a, c \in \{0, 1\}$ denote these two permissions. Under $a = 0$, the solver uses only O_0 and makes no acquisitions. Under $a = 1$, the solver can use $\mathcal{I}$.

For computation, fix the input representation, primitive operations, and state-access rules. Let $\mathcal{C}_0(\tau)$ contain the implementations allowed by specified resource bounds and $\mathcal{C}_1(\tau)$ those allowed after relaxing them. The bounds may limit operations, cumulative depth, or retained context and apply jointly to π, T, ψ. Thus computation can change acquisition choices as well as terminal processing. Both classes contain all zero-query constant solvers, and $\mathcal{C}_0 \subseteq \mathcal{C}_1 \subseteq \mathcal{G}_\tau$.

For both computational classes, fix a query limit B. The four *matched solver classes* are then defined by

$$\Gamma_{ac}(\tau, B) := \{G = (\pi, T, \psi) \in \mathcal{C}_c(\tau) : T \le B,\ a = 0 \Rightarrow T = 0\}, \tag{3}$$

where $T \le B$ holds almost surely. We suppress τ, B when fixed. The four classes $\Gamma_{00}, \Gamma_{10}, \Gamma_{01}, \Gamma_{11}$ are restricted passive, restricted acquisition, enhanced passive, and full, respectively.

Definition 3 (Agent configuration)**.** An *agent solver* $G_{\text{Agent}} = (\pi, T, \psi)$ has external acquisition access and enhanced computation for organizing acquisition and answering. These solvers form $\mathcal{G}_{\text{Agent}}(\tau, B) := \Gamma_{11}(\tau, B)$, the full configuration $(1, 1)$.

This configuration includes fixed query plans and plans that adapt to returned evidence, capturing the acquisition and computational capabilities used by agents (Russell & Norvig, 2010; Yao et al., 2023b). **Appendices A and B** give a running example and relate agent components to the two axes.

Additional permissions preserve every available solver. The coordinatewise order $(a, c) \preceq (a', c')$ means $a \le a'$ and $c \le c'$, and gives

$$\Gamma_{00} \subseteq \Gamma_{10} \subseteq \Gamma_{11}, \qquad \Gamma_{00} \subseteq \Gamma_{01} \subseteq \Gamma_{11}. \tag{4}$$

The two intermediate configurations offer different resources. Whether a configuration can replace an agent depends on the best performance across the corresponding solver class.

2.2 CLASS-OPTIMAL RISK AND AGENT IRREPLACEABILITY

The *class-optimal risk* is the infimum of expected task loss over a solver class, given by

$$R_{ac} := \inf_{G \in \Gamma_{ac}} \mathbb{E}[\ell(D_G, \omega)]. \tag{5}$$

Expectation includes the environment, interface returns, and internal randomness. Class inclusion gives $R_{11} \le R_{10} \le R_{00}$ and $R_{11} \le R_{01} \le R_{00}$. For target risk η, we retain the feasible configurations and identify those whose permissions cannot be reduced further, i.e.,

$$\mathcal{F}_\eta(\tau) := \{(a, c) \in \{0, 1\}^2 : R_{ac} \le \eta\}, \qquad \mathcal{M}_\eta(\tau) := \mathrm{Min}_{\preceq} \mathcal{F}_\eta(\tau). \tag{6}$$

Feasibility refers to the risk infimum. Concrete certificates use an explicit successful solver or a strict margin (**Appendix C.1**). For example, $\mathcal{M}_\eta = \{(1, 0), (0, 1)\}$ gives two sufficient alternatives, whereas $\mathcal{M}_\eta = \{(1, 1)\}$ makes both capabilities necessary.

Proposition 1 (Matched-risk criterion for agent irreplaceability)**.** *For a fixed task and the four matched classes, agent irreplaceability is characterized by*

$$\mathcal{M}_\eta(\tau) = \{(1, 1)\} \quad \iff \quad R_{11} \le \eta < \min\{R_{10}, R_{01}\}. \tag{7}$$

The criterion follows from class inclusion (**Appendix C.1**). The risk R_{01} tests whether computation on initial information suffices, while R_{10} tests acquisition with a restricted controller. Excluding both and constructing a successful full-class solver establishes agent irreplaceability. The next section obtains these exclusions from task structure.

3 INFORMATION, COMPUTATION, AND RESOURCE NECESSITY

Initial uncertainty can exclude enhanced computation without acquisition. Limits on locating or processing evidence can exclude restricted computation with acquisition. A successful agent under the same task and budget then completes the necessity certificate. We examine these limits and the resource changes they suggest.

3.1 INFORMATION REQUIREMENTS FOR EXTERNAL ACQUISITION

An initial request may identify an object without revealing the true state. We first ask whether the available information suffices to answer the request, even with unrestricted computation. The initial-information risk allows every measurable decision rule on O_0 and is given by

$$\rho_0(\tau) := \inf_{g : \mathcal{O}_0 \to \mathcal{D},\ g \text{ measurable}} \mathbb{E}[\ell(g(O_0), \omega)]. \tag{8}$$

Proposition 2 (Initial-information lower bound)**.** *Every enhanced passive class satisfies* $R_{01} \ge \rho_0(\tau)$.

Fixing the independent random seed gives an eligible rule in the infimum. Averaging over seeds preserves the bound (**Appendix C.1**). Thus $\rho_0 > \eta$ makes acquisition necessary. When $\rho_0 \le \eta$, passive feasibility still requires an implementable decision rule. Once acquisition is necessary, the next question is whether restricted computation can locate and use the evidence.

3.2 RESOURCE LIMITS ON EVIDENCE LOCALIZATION

To separate location clues from answer contents, consider M addressable records, each storing a binary state V_i. The initial observation identifies a target address $J = h(O_0) \in [M]$, and the answer is $Y = V_J$ under zero-one loss. Querying A_s returns the fixed record $O_s = V_{A_s}$. To isolate the need for external evidence, we use independent fair records distributed as

$$V_1, \ldots, V_M \overset{\text{iid}}{\sim} \text{Bernoulli}(1/2), \qquad V \perp O_0. \tag{9}$$

Independence makes an unread target unpredictable from the initial clue and other records. The fair prior fixes this uncertainty at risk $1/2$, while an exact target read supplies the answer. Thus $R_{00} = R_{01} = 1/2$ even when a passive solver computes J. To determine whether restricted

computation suffices with acquisition, we consider every query and stopping rule available in Γ_{10}. These pairs form the acquisition projection $\Pi_0 := \{(\pi, T) : \exists\psi,\ (\pi, T, \psi) \in \Gamma_{10}\}$. For each pair, let $S_{\pi,T} = \{A_s : 1 \le s \le T\}$ denote the visited addresses. The event $J \in S_{\pi,T}$ means that the target record is acquired. The supremal target-acquisition probability defines *addressability*, given by

$$\mathrm{Addr}_B(\Pi_0) := \sup_{(\pi,T)\in\Pi_0} \Pr\{J \in S_{\pi,T}\}. \tag{10}$$

To isolate acquisition limits, a *terminal relaxation* preserves $(\pi, T) \in \Pi_0$ but gives the terminal rule $\bar{\psi}$ unrestricted computation and full access to H_T, including O_0. This relaxation removes both retained-information and terminal-computation constraints. The relaxed risk is given by

$$\bar{R}_{10} := \inf_{(\pi,T)\in\Pi_0,\bar{\psi}} \mathbb{E}[\ell(\bar{\psi}(H_T), \omega)] \le R_{10}. \tag{11}$$

Theorem 1 (Localization risk for independent records)**.** *For independent fair records and any nonempty projection class Π_0, the risks satisfy*

$$R_{10} \ge \bar{R}_{10} = \tfrac{1}{2}(1 - \mathrm{Addr}_B(\Pi_0)). \tag{12}$$

The relaxed rule identifies an acquired target through O_0. If the target is omitted, complementing the target bit preserves the adaptive acquisition and stopping history, leaving error $1/2$ (**Appendix C.3**). Incomplete coverage therefore limits even unrestricted terminal computation on the full acquired history.

This identity turns a coverage bound into a risk bound. To isolate context retention, we allow unrestricted encoding and processing of returns, but pass initial context to acquisition through a bounded summary. Fewer distinguishable summaries leave more candidate addresses per summary.

Definition 4 (Context-retention bandwidth)**.** For nonnegative integers b, B, let $\Pi^{(b)}$ contain acquisition and stopping pairs that measurably encode O_0 as $C = f(O_0) \in [2^b]$ and make at most B queries. Subsequent rules use only C, address-labelled returns, the step count, and independent randomness. All later access to O_0 passes through C. Encoding, subsequent history, and workspace are otherwise unrestricted.

The bandwidth b limits retained localization information separately from the cost of computing h. To compare retention limits, let $\mathcal{C}_0$ retain b bits and $\mathcal{C}_1$ retain $b_\star = \lceil \log_2 M \rceil$ bits, with $0 \le b \le b_\star$. Both classes use this architecture, with terminal access only to the summary and query history. The enhanced class retains J and reads V_J, achieving $R_{11} = 0$ for $B \ge 1$.

Theorem 2 (Localization frontier for bandwidth and queries)**.** *For a target spread uniformly over the addresses and independent fair records, optimal coverage and terminally relaxed risk in the full summary class are given by*

$$\mathrm{Addr}_B(\Pi^{(b)}) = \min\{1, 2^b B/M\}, \quad \bar{R}_{10}(b, B) = \tfrac{1}{2}(1 - \min\{1, 2^b B/M\}). \tag{13}$$

Every subclass $\Pi_0 \subseteq \Pi^{(b)}$ satisfies $R_{10} \ge \bar{R}_{10} \ge \bar{R}_{10}(b, B)$.

For fixed records and randomness, each summary value determines an execution visiting at most B addresses. The 2^b summaries therefore cover at most $2^b B$ addresses in total. For $B \ge 1$, a matching construction partitions $\min\{M, 2^b B\}$ addresses into at most 2^b groups of size at most B. The encoder stores the index of the group containing J whenever possible, and the query stage reads that group (**Appendix C.4**). One additional summary bit and a doubled query budget therefore give the same optimal coverage.

To complete the necessity certificate, combine this frontier with passive risk $1/2$ and the enhanced solver with zero risk. For $B \ge 1$ and $0 \le \eta < 1/2$, the condition $2^b B/M < 1 - 2\eta$ places the restricted acquisition class above the target risk, giving $\mathcal{M}_\eta = \{(1, 1)\}$.

For other evidence structures, separate coverage from the uncertainty left when the target is omitted. With a coverage bound α and residual risk ρ, the risk bound is given by

$$R_{10} \ge \bar{R}_{10} \ge \rho(1 - \alpha), \tag{14}$$

provided both bounds hold across the acquisition class and all relaxed terminal rules. **Appendices C.2 and D** treat nonuniform targets, biased or correlated records, and multiple-record evidence. Answer copies and alternative sufficient evidence must be accounted for when bounding ρ.

3.3 COMPUTATIONAL LIMITS ON HISTORY PROCESSING

The same necessary configuration can arise when evidence is fully acquired but cannot be combined into an answer. Consider two known addresses containing L-bit blocks, with $N = 2L$ independent fair bits and query budget $B = 2$. The initial observation gives only sizes and addresses. Two fixed queries acquire all bits, separating the use of complete evidence from the difficulty of locating records. The target is the parity of these bits.

A prediction based on any fixed proper subset of these bits has error $1/2$. Improving on this risk therefore requires combining all bits. We use an online circuit with fan-in at most two, so depth limits how many bits can influence the answer. Ordinary depth accumulates across rounds and counts state updates, address selection, stopping, and output selection. History storage, fixed wires, and fan-out have zero depth cost. Block reads are interface primitives. The two computational classes differ only in nonnegative integer ordinary-depth bounds $d_0 \le d_1$. Circuit descriptions are fixed before the unknown input, and random seeds are independent of the input.

To bound every restricted interactive solver, we expand the interface reads into ordinary gates on the complete raw data. The resulting circuit preserves the solver output for each input and seed, with depth $d_{\mathrm{sim}} \le d_0 + 8$. Fan-in two then limits output dependence to at most 2^{d_0+8} bits. We impose the separation conditions

$$2^{d_0+8} < N, \qquad d_1 \ge \lceil \log_2 N \rceil, \tag{15}$$

under which flipping a bit outside this dependence set leaves the restricted output unchanged but reverses parity, forcing error $1/2$. An enhanced solver reads both blocks and uses a balanced exclusive-OR (XOR) tree. The class risks are therefore given by

$$R_{00} = R_{01} = R_{10} = \tfrac{1}{2}, \qquad R_{11} = 0, \qquad \bar{R}_{10} = 0. \tag{16}$$

Reading both blocks and outputting a constant is already admissible in the restricted class, so terminal relaxation achieves zero risk. Here sufficient raw evidence remains accessible, and the gap isolates a depth limitation. The bound includes adaptive queries and online computation. The constant 8 is a conservative interface-simulation overhead (**Appendices E.2–E.4**).

The complete history may contain all N bits, while depth limits how many influence the output. Other targets can admit useful predictions from fewer coordinates. For a binary target $F(X)$, any distribution on $X \in \{0,1\}^N$, and integer $0 \le k \le N$, the best risk using at most k fixed coordinates is given by

$$\varepsilon_F(k) := \inf_{S \subseteq [N],\ |S| \le k,\ g:\{0,1\}^{|S|} \to \{0,1\}} \Pr\{g(X_S) \ne F(X)\}. \tag{17}$$

The coordinates and rule are fixed before the current input. Processing on selected coordinates is unrestricted, with $S = \varnothing$ allowing constant rules.

Proposition 3 (Processing bound from output dependence)**.** *Let the internal seed be independent of X. If every restricted solver at every fixed seed has an output function depending on at most k raw coordinates, then $R_{10} \ge \varepsilon_F(k)$. For the two-block online model, the bound is expressed as*

$$R_{10} \ge \varepsilon_F(k_0), \qquad k_0 = \min\{N, 2^{d_0+8}\}. \tag{18}$$

Each fixed-seed output is an eligible predictor, and averaging preserves the bound. Fair parity has $\varepsilon_F(k) = 1/2$ for $k < N$. Boolean disjunction (OR) instead asks whether any bit flags an anomaly. Independent $\mathrm{Bernoulli}(1 - 2^{-1/N})$ flags make the two OR answers equally likely, giving $\varepsilon_F(k) = 2^{-k/N} - 1/2$. Thus partial dependence can progressively reduce risk for this target (**Appendix E.5**).

For a new implementation model, first establish a full-data simulation or output-dependence bound, then bound the corresponding prediction risk. Initial-information risk and a successful agent complete the certificate. **Appendix E.1** gives the general simulation transfer.

3.4 RESOURCE NECESSITY AND INTERVENTION PREDICTIONS

A structural bound $L_{10} \le R_{10}$ and an allowed agent with risk at most U_{11} certify joint necessity when $U_{11} \le \eta < \min\{\rho_0, L_{10}\}$, by **(7)**. A lower bound that does not exceed the target leaves the restricted class unresolved. The certificate identifies indispensable capabilities, while the limiting mechanism suggests resource interventions.

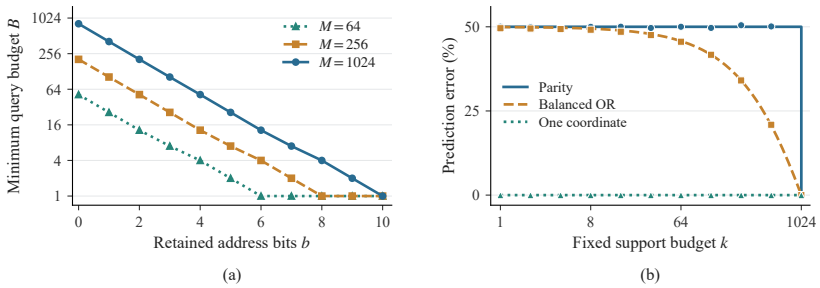


Figure 2. Task structure and resource limits. (a) Minimum query budget for 10% terminally relaxed risk. (b) Fixed-support prediction risk with complete raw evidence. Lines show analytical values, and points in (b) show measured errors of specified programs.

More retained context or queries can improve coverage, but accurate answers also require target identity. In the uniform fair two-address task with $b = 0$, increasing B from 1 to 2 raises optimal coverage from $1/2$ to 1 while actual optimal risk stays $1/4$. Retaining one address bit instead gives zero risk with one query (**Appendix D.4**). Without the address bit, two different record values do not reveal which value the request designates. Extra reads cannot restore that identity.

When the localization coverage bound gives $\bar{R}_{10} > \eta$, even the full history leaves too much uncertainty, motivating better localization or more queries. The processing construction instead has $\bar{R}_{10} = 0 < R_{10}$ with all raw evidence accessible. More depth enables the required combination of that evidence. The same necessary agent configuration thus points to different resource investments. Both constructions admit successful fixed query plans. **Section 4** examines these interventions.

4 Experiments

We examine which agent capabilities different tasks require and how interventions on resource limits improve their solution procedures. Controlled tasks test the risk characterizations in **Section 3**. Natural-language tasks compare external acquisition and evidence processing in terms of answer quality and resource use. **Appendices F–K** provide the complete results and application examples of the four capability configurations.

4.1 Experimental setup

The controlled tasks vary the location of raw records and the target function, giving single-coordinate and parity tasks with either initial or external evidence. We use $N = 1024$ record bits and a query budget of $B = 2$. Restricted and enhanced computation have ordinary gate depths of 1 and 10, respectively, and each specified program is evaluated on 20,000 inputs. The localization experiments further vary retention bandwidth, query budget, the target-location distribution, and record correlation to compare analytical risks with measured program errors.

The natural-language experiments use 1,000 questions each from FinQA for financial question answering (Chen et al., 2021) and HotpotQA for multi-hop question answering (Yang et al., 2018), with Qwen2.5-14B-Instruct (Qwen Team, 2024). Acquisition procedures include fixed Best Matching 25 (BM25) retrieval (Robertson & Zaragoza, 2009), adaptive acquisition, and all-evidence input. Fixed retrieval selects 2, 4, or 8 blocks in one step, denoted B2, B4, and B8. Adaptive acquisition uses returned records to select subsequent queries and returns at most 8 blocks. All-evidence input contains every candidate record for the question. Different readouts reuse the same acquisition history. FinQA compares direct and program readouts, with and without an intermediate draft that organizes evidence and computational steps. The program readout generates and executes operations in a fixed domain-specific language (DSL), using only visible records. HotpotQA compares direct answering with evidence organization followed by answering. We report accuracy and exact match (EM), respectively. Resource accounting includes acquisition control, drafts, and final answers. Primary paired intervals resample report or article clusters, as detailed in **Appendix F**.

Table 1. Answer quality on natural-language tasks (%). FinQA reports accuracy and HotpotQA reports exact match (EM), with 1,000 questions per configuration. Bold marks the highest score in each column. Underlining marks the best $B \leq 8$ score when that score differs from the overall best. Dashes indicate configurations not evaluated.

Acquisition	FinQA accuracy				HotpotQA EM	
	Direct	Draft + direct	Program	Draft + program	Direct	Organized
No acquisition	0.2	0.2	0.8	0.5	14.6	24.0
Fixed B2	12.6	–	32.1	–	33.9	33.7
Fixed B4	15.3	–	39.7	–	42.5	42.4
Fixed B8	17.1	32.9	50.3	54.2	48.0	48.8
Adaptive $B \leq 8$	**18.8**	<u>34.0</u>	<u>51.6</u>	<u>55.2</u>	<u>52.2</u>	<u>52.8</u>
All evidence	17.5	**37.9**	**57.3**	**64.9**	**52.3**	**56.4**

4.2 Task structure and necessary capabilities

For each controlled task, we first compare the risks of the four solver classes. For $0 \leq \eta < 1/2$, the minimal sufficient configurations for initial single-coordinate, external single-coordinate, initial parity, and external parity tasks are $(0, 0)$, $(1, 0)$, $(0, 1)$, and $(1, 1)$, respectively. Measurements of the specified programs agree with the class-risk certificates in **Appendix G**. **Figure 2** further examines how task structure changes resource requirements.

In **Figure 2(a)**, increasing retention bandwidth from 4 to 5 bits at $M = 1024$ reduces the queries needed for a terminally relaxed risk of 10% from 52 to 26. Coverage alone does not ensure accurate answers. At $b = 4, B = 64$, a specified program covers every target but has 44.76% actual error. Retaining the complete address gives zero measured error at $b = 10, B = 1$ (**Appendix G.2**). This contrast reflects the target-identity limitation in **Section 3.4**.

With complete raw evidence, processing requirements depend on the target function. In **Figure 2(b)**, fair parity retains risk $1/2$ for every $k < N$, whereas balanced OR risk decreases with support capacity. These mechanisms distinguish investments in retaining localization information, acquiring records, and combining available evidence. Nonuniform targets and correlated records further vary the acquisition requirement, as examined in **Appendix G.3**.

4.3 Acquisition and processing on natural-language tasks

Table 1 compares combinations of acquisition procedures and readouts. Comparing readouts on a fixed history examines how the solver uses previously acquired records.

On FinQA with fixed B8 retrieval, the program readout improves over direct answering by 33.2 percentage points, with a 98.75% paired cluster interval of $[28.30, 37.93]$. Adding a draft to the all-evidence program readout yields a further 7.6 points, with interval $[3.93, 11.34]$. Complete input still leaves room to improve evidence organization and the execution of operations.

On HotpotQA, adaptive acquisition with direct answering improves over fixed B8 by 4.2 points, with a 97.5% interval of $[1.51, 6.93]$. Evidence organization improves all-evidence answering by 4.1 points, with interval $[1.19, 7.01]$, but improves B8 answering by only 0.8 points. The benefit of additional processing depends on the evidence obtained.

To examine where adaptive acquisition helps, we group HotpotQA questions by complete annotated-support coverage under the two procedures. The 231 questions with complete support only under adaptive acquisition contribute a net gain of 57 correct answers, while the 110 with complete support only under fixed retrieval contribute a net loss of 15. The groups with complete support under both or neither procedure contribute $+3$ and -3, respectively. The four groups sum to 42 additional correct answers, with the net gain concentrated in questions whose support coverage changes (**Appendix H.2**).

We next test whether additional processing can compensate for missing support. The FinQA comparison uses 961 questions meeting the fragment-matching criteria in **Appendix I**. We replace one

annotated supporting record with an unrelated fragment of similar length and add the same distractors to the paired histories. With no added distractors, program accuracy is 61.19% with complete support and 9.26% with missing support. Adding a draft gives 65.04% and 11.86%, respectively. The small gain with missing support points to the information gap as the first constraint to address. With complete support, processing gains are larger in histories with more distractors (**Appendix I**).

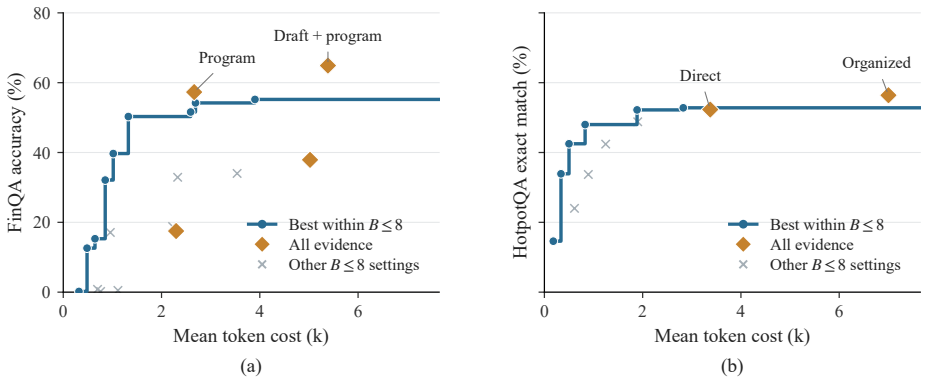


Figure 3. Answer quality and mean token use. (a) FinQA. (b) HotpotQA. The blue line gives the highest score within each mean token budget among evaluated $B \leq 8$ configurations. Gray crosses denote other configurations, and orange diamonds denote all-evidence configurations. Tokens include inputs, outputs, acquisition control, and drafts, with k denoting thousands.

4.4 Resource allocation and answer quality

Figure 3 selects the highest score within each mean token budget among evaluated configurations with acquisition budget $B \leq 8$, and separately marks all-evidence configurations. This comparison examines how answer quality changes when resources are spent on acquiring more material or processing existing evidence.

The FinQA B4 program readout uses 4 evidence blocks and approximately 1.017 thousand tokens to achieve 39.7% accuracy, exceeding all-evidence direct answering at 17.5% with 31.60 blocks and 2.297 thousand tokens. Better processing can thus reduce the acquisition needed to obtain better answers. On HotpotQA, adaptive direct answering achieves 52.2% EM with 6.66 blocks and 1.886 thousand tokens. All-evidence direct answering uses 41.15 blocks and 3.375 thousand tokens for 52.3% EM. Adaptive acquisition reduces material and token use, while increasing mean calls from 1 to 3.60. Resource allocation in an agent workflow depends both on the task's limiting factor and on the relative costs of acquisition and sequential calls. **Appendices H and J** provide complete resource statistics, common-budget comparisons, and results across models.

5 Conclusion

We developed a task-level theory of agent irreplaceability by comparing class-optimal risks for a fixed task, interface, and query budget. Exact coverage frontiers and processing lower bounds establish conditions under which acquisition and enhanced computation are jointly necessary. The analysis distinguishes limits on obtaining evidence from limits on retaining and using acquired evidence. Broader coverage can leave risk unchanged, and complete evidence can still require stronger computation. Tasks requiring the same agent capabilities can therefore benefit from different resource investments.

The natural-language experiments show that improved processing can achieve higher accuracy with less evidence, while adaptive acquisition trades fewer records and tokens for more model calls. These findings connect the structural analysis to two decisions in agent design. Capability requirements identify which simpler configurations can replace an agent, while the mechanisms limiting performance guide where additional resources can help.

AI USE STATEMENT

We used generative AI tools to assist with language polishing, figure preparation, and checking mathematical proofs for potential errors or gaps. The authors reviewed all suggested edits and take full responsibility for the final manuscript.

REFERENCES

Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-RAG: Learning to retrieve, generate, and critique through self-reflection. In *International Conference on Learning Representations*, 2024.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions on Machine Learning Research*, 2023.

Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. FinQA: A dataset of numerical reasoning over financial data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3697–3711, 2021.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv 2110.14168*, 2021.

DeepSeek-AI. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv 2501.12948*, 2025.

E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1: 269–271, 1959.

Olive Jean Dunn. Multiple comparisons among means. *Journal of the American Statistical Association*, 56(293):52–64, 1961.

Zhengru Fang, Senkang Forest Hu, Zhonghao Chang, Yu Guo, Yihang Tao, Hongyao Liu, Mengzhe Ruan, Jun Huang, and Yuguang Fang. Inference-time budget control for LLM search agents. *arXiv preprint arXiv:2605.05701*, 2026.

C. A. Field and A. H. Welsh. Bootstrapping clustered data. *Journal of the Royal Statistical Society Series B*, 69(3):369–390, 2007.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. PAL: Program-aided language models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 10764–10799, 2023.

Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.

Gautier Izacard and Edouard Grave. Leveraging passage retrieval with generative models for open domain question answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 874–880. Association for Computational Linguistics, 2021.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Ö. Arık, Dong Wang, Hamed Zamani, and Jiawei Han. Search-R1: Training LLMs to reason and leverage search engines with reinforcement learning. In *Conference on Language Modeling*, 2025.

Florian Kaltenberger, Aloizio P. Silva, Abhimanyu Gosain, Luhan Wang, and Tien-Thinh Nguyen. OpenAirInterface: Democratizing innovation in the 5G era. *Computer Networks*, 176:107284, 2020.

Sayash Kapoor, Benedikt Stroebl, Zachary S. Siegel, Nitya Nadgir, and Arvind Narayanan. AI agents that matter. *Transactions on Machine Learning Research*, 2025.

Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

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 *Advances in Neural Information Processing Systems*, volume 33, 2020.

Yang Li and Junier Oliva. Active feature acquisition with generative surrogate models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 6450–6459, 2021.

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, 2024a.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. AgentBench: Evaluating LLMs as agents. In *International Conference on Learning Representations*, 2024b.

Chris Lu, Cong Lu, Robert Tjarko Lange, Yutaro Yamada, Shengran Hu, Jakob Foerster, David Ha, and Jeff Clune. Towards end-to-end automation of AI research. *Nature*, 651(8107):914–919, 2026.

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 *Advances in Neural Information Processing Systems*, volume 36, 2023.

Mathematical Association of America. MAA invitational competitions, n.d.

Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. GAIA: A benchmark for general AI assistants. In *International Conference on Learning Representations*, 2024.

Mistral AI. Mistral-7B-Instruct-v0.3. Model card, 2024.

Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

Cheng Qian, Emre Can Acikgoz, Qi He, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tur, Gokhan Tur, and Heng Ji. ToolRL: Reward is all tool learning needs. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv 2412.15115*, 2024.

Stephen Robertson and Hugo Zaragoza. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389, 2009.

Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3rd edition, 2010.

Stuart Russell and Eric Wefald. Principles of metareasoning. *Artificial Intelligence*, 49(1–3):361–395, 1991.

Stuart J. Russell and Devika Subramanian. Provably bounded-optimal agents. *Journal of Artificial Intelligence Research*, 2:575–609, 1995.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

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, 2023.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning. In *International Conference on Learning Representations*, 2025.

Chung-En Sun, Linbo Liu, Ge Yan, Zimo Wang, and Tsui-Wei Weng. LLM agents already know when to call tools—even without reasoning. *arXiv preprint arXiv:2605.09252*, 2026.

Joseph Tso, Preston Schmittou, Quan Huynh, and Jibran Hutchins. ConstraintBench: Benchmarking LLM constraint reasoning on direct optimization. *arXiv preprint arXiv:2602.22465*, 2026.

Hongru Wang, Cheng Qian, Manling Li, Jiahao Qiu, Boyang Xue, Mengdi Wang, Heng Ji, Amos Storkey, and Kam-Fai Wong. Position: Agent should invoke external tools ONLY when epistemically necessary. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Demystifying LLM-based software engineering agents. *Proceedings of the ACM on Software Engineering*, 2(FSE):801–824, 2025. doi: 10.1145/3715754.

Yilun Xu, Shengjia Zhao, Jiaming Song, Russell Stewart, and Stefano Ermon. A theory of usable information under computational constraints. In *International Conference on Learning Representations*, 2020.

John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *Advances in Neural Information Processing Systems*, volume 37, 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.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, volume 36, 2023a.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2023b.

Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. AFlow: Automating agentic workflow generation. In *International Conference on Learning Representations*, 2025.

Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language agent tree search unifies reasoning, acting, and planning in language models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 62138–62160, 2024a.

Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. WebArena: A realistic web environment for building autonomous agents. In *International Conference on Learning Representations*, 2024b.

A AN EXAMPLE OF BOOK AVAILABILITY

A.1 TARGET, EVIDENCE, AND EVALUATION

The task is to decide whether a specified book is available on the assigned shelf. The catalogue identifies the shelf whose record gives the book's availability. Consider library records fixed before interaction. The book is assigned to shelf 1 or 2 and is either checked out or available. We represent the environment space and corresponding elements by

$$\Omega = \{(1,0),(1,1),(2,0),(2,1)\}, \qquad \omega = (j,v), \tag{A.1}$$

where j is the shelf index and the states $v = 0$ and $v = 1$ denote checked out and available, respectively. The distribution P is uniform over the four environments. The initial observation contains the book title and the availability question. The interface permits a catalogue query, which returns j, and a query to shelf k, which returns v for a matching address $k = j$ and a missing-record symbol otherwise. Choose the decision and target spaces $\mathcal{D} = \mathcal{Y} = \{0,1\}$, the target map $y^\star(j,v) = v$, and the zero-one loss $\ell_0(d,y) = \mathbf{1}\{d \neq y\}$, giving the task loss $\ell(d,\omega) = \mathbf{1}\{d \neq v\}$. Reading these fixed records preserves the target throughout interaction.

Here (Ω, P) describes uncertainty about location and availability. The initial observation identifies the book whose location and status can be checked through the interface. The decision space and loss specify the answer and how the answer is evaluated. Because the same initial request occurs in all four environments, the request alone leaves uncertainty that queries can resolve.

A.2 SOLVERS AND THE FOUR MATCHED CLASSES

A zero-risk interactive solver first queries the catalogue and then queries the returned shelf. The acquisition policy queries the catalogue at H_0 and shelf O_1 at H_1, with $T = 2$ and $\psi(H_2) = O_2$. For $\omega = (2,0)$, the history records shelf 2 followed by the checked-out status, giving the final answer 0. This execution connects the task tuple, solver triple, and interface.

To instantiate computational capability, fix the query budget at $B = 2$. In $\mathcal{C}_0(\tau)$, the stopping step is fixed in advance. Each action argument and the terminal output is a predetermined constant or a copy of a predetermined history field. The class $\mathcal{C}_1(\tau)$ additionally permits operations combining fields and branches conditioned on returned values. Both contain constant solvers with zero acquisitions, while the extra operations available to the enhanced class give $\mathcal{C}_0(\tau) \subseteq \mathcal{C}_1(\tau)$. The two-query procedure above only copies the catalogue response into a query address and then copies the status into the answer. This procedure therefore belongs to $\mathcal{C}_0(\tau)$.

Table A.1: Matched risks for book availability.

Class	Allowed solution in this task	Optimal risk
Γ_{00}	Answer with a constant from the title and question.	$1/2$
Γ_{01}	Compute on the initial observation, which contains no availability information.	$1/2$
Γ_{10}	Query the catalogue, copy the shelf address, query that shelf, and output the status.	0
Γ_{11}	Use the same two-query procedure or more general computation and interaction.	0

Table A.1 summarizes the available procedures and their optimal risks. Thus the minimal sufficient set is given by $\mathcal{M}_\eta = \{(1,0)\}$ for risk targets $0 \leq \eta < 1/2$. This task requires acquisition but admits a simple controller. The need for enhanced computation depends on whether localization or processing exceeds the restricted class, as the task families in **Section 3** demonstrate.

B AGENT COMPONENTS AND THE TWO RESOURCE AXES

Agent systems combine reasoning, planning, memory, and tools. Their roles in our model depend on how they acquire external evidence and compute over the available history. **Table B.1** relates common components to the two resource axes.

Table B.1: Agent components and resource permissions.

Component or tool	Role in task τ and solver G
Document retrieval, database queries, and external search	Submit A_s through $\mathcal{I}$ and add the returned observation O_s to H_s. These operations use acquisition access $a = 1$.
Reasoning, planning, and reflection on an existing trajectory	Use H_t to choose acquisition actions, stopping, and terminal decisions through π, T, and ψ. The class $\mathcal{C}_c(\tau)$ specifies the allowed control rules.
Internal memory that stores or compresses observations	Retain and use information already present in H_t. State capacity and read/write operations may be restricted by $\mathcal{C}_c(\tau)$.
Retrieval from an external experience store	Introduce a previously unavailable record into H_s as O_s through action A_s. This uses acquisition access through $\mathcal{I}$.
Calculators or code operating on given data	Compute over O_0 or H_t. The class $\mathcal{C}_c(\tau)$ specifies the permitted operations and their resource bounds.

A component may contribute to both acquisition and computation by introducing external observations and performing control operations. The component's computational class depends on the implementation rules and resource bounds.

Retrieval-augmented generation combines a language model with an external passage store (Lewis et al., 2020); Fusion-in-Decoder aggregates information from multiple retrieved passages (Izacard & Grave, 2021). Self-reflective retrieval-augmented generation (Self-RAG) learns when to retrieve and how to assess retrieved content and generated responses (Asai et al., 2024). In the present framework, passage access changes the available observations, while the retrieval controller and reader determine how those observations are selected and used.

Chain-of-thought prompting produces intermediate reasoning steps (Wei et al., 2022), Tree of Thoughts searches over candidate reasoning paths (Yao et al., 2023a), and Self-Refine revises outputs using model-generated feedback (Madaan et al., 2023). These procedures can operate on a fixed input without acquiring new external facts. Their resource costs belong to the computation axis. Empirical test-time scaling also finds that the useful allocation of computation depends on the problem and base model (Snell et al., 2025).

Program-aided language models and Program of Thoughts generate intermediate programs and delegate execution to an interpreter (Gao et al., 2023; Chen et al., 2023). An interpreter operating only on supplied data changes processing permissions. If the generated program can also query external records, those calls must use the task's interface and count toward the task's acquisition budget.

AgentBench and the General AI Assistants (GAIA) benchmark evaluate performance across diverse interactive tasks (Liu et al., 2024b; Mialon et al., 2024). Applying resource diagnosis to such tasks additionally requires specifying the alternative solver classes. A benchmark score evaluates an implementation; excluding a capability requires a bound covering all permitted replacement solvers.

C PROOFS OF THE DIAGNOSTIC AND LOCALIZATION RESULTS

We first prove the matched-risk criterion, the initial-information bound, and the combination of structural certificates. We then establish the residual-risk reduction. The localization proofs calculate terminally relaxed risk for a fixed acquisition class and optimal coverage under the summary constraint. Identically distributed biased records and nonuniform targets are treated together. The

fair uniform case gives the two localization theorems in the main text. We then combine these results to establish the localization regimes in which both capabilities are necessary.

C.1 Matched-risk comparisons and certificates

Proof of Proposition 1. Class inclusion and risk monotonicity make $\mathcal{F}_\eta = \{(a, c) : R_{ac} \leq \eta\}$ upward closed. Under the condition $R_{11} \leq \eta < \min\{R_{10}, R_{01}\}$, configuration $(1, 1)$ is feasible and both intermediate configurations are infeasible. The bound $R_{00} \geq R_{10} > \eta$ also excludes configuration $(0, 0)$, so the feasible region reduces to $\mathcal{F}_\eta = \{(1, 1)\}$ with unique minimal element $(1, 1)$.

Conversely, suppose the minimal sufficient set is $\mathcal{M}_\eta = \{(1, 1)\}$. Feasibility gives $R_{11} \leq \eta$. Either intermediate configuration being feasible would contradict the minimality of $(1, 1)$. The intermediate risks must therefore satisfy $R_{10} > \eta$ and $R_{01} > \eta$. □

Possible sets of minimal configurations. The coordinatewise order on the four configurations is finite. The minimal elements of an upward-closed feasible region can form exactly the following six sets. Risk monotonicity gives the equivalent conditions. **Table C.1** lists the six possibilities and their equivalent risk conditions.

Table C.1: Minimal configurations and risk conditions.

$\mathcal{M}_\eta$	Equivalent risk condition
$\varnothing$	$R_{11} > \eta$
$\{(0, 0)\}$	$R_{00} \leq \eta$
$\{(1, 0)\}$	$R_{10} \leq \eta < R_{01}$
$\{(0, 1)\}$	$R_{01} \leq \eta < R_{10}$
$\{(1, 0), (0, 1)\}$	$\max\{R_{10}, R_{01}\} \leq \eta < R_{00}$
$\{(1, 1)\}$	$R_{11} \leq \eta < \min\{R_{10}, R_{01}\}$

For example, if $(1, 0)$ is feasible and $(0, 1)$ is infeasible, the bound $R_{00} \geq R_{01} > \eta$ also excludes the baseline, leaving $(1, 0)$ as the unique minimal configuration. If both intermediate configurations are feasible and $(0, 0)$ is infeasible, they form the minimal set together. The remaining cases follow from the same finite order.

Proof of Proposition 2. Let $\mathcal{O}_0$ be the range of the initial observation. For the task loss $\ell(d, \omega) \in [0, 1]$, the initial-information risk is the infimum of expected loss over all measurable decision rules on this space, given by

$$\rho_0(\tau) := \inf_{g:\mathcal{O}_0 \to \mathcal{D} \text{ measurable}} \mathbb{E}[\ell(g(O_0), \omega)]. \tag{C.1}$$

This infimum allows unrestricted computation. A passive solver outputs $D = g_G(O_0, U)$ for a seed U independent of ω. At every fixed seed u, the section $o \mapsto g_G(o, u)$ is an eligible measurable rule. The fixed-seed risk is therefore bounded as

$$\mathbb{E}_\omega[\ell(g_G(O_0(\omega), u), \omega)] \geq \rho_0(\tau).$$

Independence expresses randomized risk as the average of fixed-seed risks. Nonnegative loss permits exchanging the order of integration, so the randomized risk is bounded as

$$\mathbb{E}[\ell(D, \omega)] = \int \mathbb{E}_\omega[\ell(g_G(O_0(\omega), u), \omega)]\, dP_U(u) \geq \rho_0(\tau).$$

Taking the infimum over either passive computational class proves $R_{00}, R_{01} \geq \rho_0(\tau)$, including cases where the infimum is unattained.

For binary decisions under zero-one loss, the initial-information risk has an explicit form. Consider the decision space $\mathcal{D} = \{0, 1\}$, binary target $Y = y^\star(\omega) \in \{0, 1\}$, and zero-one loss $\ell(d, \omega) =$

$\mathbf{1}\{d \neq Y\}$. Define the posterior probability by $p(O_0) = \Pr(Y = 1 \mid O_0)$. For every measurable initial decision g, the conditional error can be bounded as

$$\begin{aligned}\Pr(g(O_0) \neq Y \mid O_0) &= p(O_0)\mathbf{1}\{g(O_0) = 0\} + (1 - p(O_0))\mathbf{1}\{g(O_0) = 1\} \\ &\geq \min\{p(O_0), 1 - p(O_0)\}.\end{aligned} \tag{C.2}$$

The measurable rule $g^\star(O_0) = \mathbf{1}\{p(O_0) > 1/2\}$ attains this pointwise bound, so the initial-information risk is given by

$$\rho_0(\tau) = \mathbb{E}\big[\min\{p(O_0), 1 - p(O_0)\}\big]. \tag{C.3}$$

□

The condition $\rho_0 > \eta$ excludes every measurable initial-observation rule. The condition $R_{01} > \eta$ excludes the specified enhanced passive class and may also reflect the computational restrictions of that class.

A passive class that implements $g^\star$ has risk ρ_0. Equality can also hold through a sequence approaching the infimum. For general decision spaces and losses, the infimum over measurable rules defines ρ_0 directly and requires no measurable selection of pointwise optimal actions.

If a constant d satisfies $\mathbb{E}[\ell(d, \omega)] \leq \eta$, the constant-solver convention supplies a successful member of Γ_{00}. The resulting bound $R_{00} \leq \eta$ establishes the minimal sufficient set $\mathcal{M}_\eta = \{(0, 0)\}$. A bound $\rho_0 \leq \eta$ provides the unrestricted benchmark. Establishing feasibility of an actual computational class, however, additionally requires an admissible construction.

Valid intervals and partial diagnosis. For configuration $x = (a, c)$, suppose the risk lies in an interval satisfying $L_x \leq R_x \leq U_x$. An allowed solver $G \in \Gamma_x$ with expected loss at most U_x establishes the upper bound. A uniform bound on every member of Γ_x establishes the lower bound. Thus the upper-bound condition $U_x \leq \eta$ certifies feasibility, whereas the lower-bound condition $L_x > \eta$ certifies exclusion. A failed implementation alone leaves the remaining solvers in the class unresolved.

For ordered configurations $x \preceq y$, risk monotonicity gives $R_y \leq R_x$. Feasibility propagates to stronger configurations, whereas exclusion propagates to weaker ones. Certifying x and excluding all strictly smaller configurations establishes membership $x \in \mathcal{M}_\eta$. Identifying the entire minimal set also requires resolving configurations that could contribute other minimal elements. Unresolved comparisons leave the established partial conclusions intact.

These rules use feasibility at the class-risk infimum. At an unattained boundary $R_x = \eta$, an explicit solver with risk at most η still requires a construction. A concrete successful solver or a strict risk margin provides an implementation-level certificate.

Corollary 1 (Joint necessity from structural bounds)**.** *Fix the task, interface, query budget, computational classes, and target risk. Let L_{10} be a valid structural lower bound on R_{10}. Suppose an allowed $G^\star \in \Gamma_{11}$ satisfies $\mathbb{E}[\ell(D_{G^\star}, \omega)] \leq U_{11}$. Under the condition $U_{11} \leq \eta < \min\{\rho_0(\tau), L_{10}\}$ the minimal sufficient set is given by $\mathcal{M}_\eta(\tau) = \{(1, 1)\}$.*

Proof of Corollary 1. Fix the task, interface, query budget, nested computational classes, and risk target η. Let the bound $L_{10} \leq R_{10}$ hold over the entire restricted acquisition class. Suppose an allowed $G^\star \in \Gamma_{11}$ has expected loss at most U_{11}, i.e.,

$$\mathbb{E}[\ell(D_{G^\star}, \omega)] \leq U_{11}.$$

Under the condition $U_{11} \leq \eta < \min\{\rho_0(\tau), L_{10}\}$, **Proposition 2** gives $R_{01} \geq \rho_0(\tau) > \eta$, while the structural bound gives $R_{10} \geq L_{10} > \eta$. The construction gives the remaining upper bound

$$R_{11} \leq \mathbb{E}[\ell(D_{G^\star}, \omega)] \leq U_{11} \leq \eta.$$

The matched-risk criterion yields $\mathcal{M}_\eta(\tau) = \{(1, 1)\}$. The full-class upper bound comes from a concrete solver, so feasibility at the threshold is attained. □

Either localization or processing can supply L_{10}. A class-wide residual risk ρ under target omission and a bound α on the probability of acquiring the target record give $R_{10} \geq \bar{R}_{10} \geq \rho(1 - \alpha)$.

Alternatively, a loss-preserving full-data simulation for every restricted implementation transfers a processing-risk bound at the simulation cost.

Bounds on the same R_{10} can be combined through their maximum. Adding them requires an additional loss decomposition. Each bound and construction in a certificate must refer to the same task, interface, budget, and classes. These certificates give sufficient conditions for joint necessity. An unavailable certificate leaves necessity unresolved.

C.2 Residual risk under target omission

Let $E_{\pi,T} = \{J \notin S_{\pi,T}\}$ denote target omission.

Proposition 4 (Reduction from residual risk under target omission)**.** *Suppose every $(\pi, T) \in \Pi_0$ and relaxed terminal rule $\bar{\psi}$ have conditional risk under target omission bounded as*

$$\mathbb{E}[\ell(\bar{\psi}(H_T), \omega) \mid E_{\pi,T}] \geq \rho, \tag{C.4}$$

for every target omission event with positive probability $\Pr(E_{\pi,T}) > 0$. If the coverage bound $\mathrm{Addr}_B(\Pi_0) \leq \alpha$ holds for class-wide constants $\rho, \alpha \in [0, 1]$, the risks satisfy the bound in **(14)**.

Proof of Proposition 4. Fix $(\pi, T) \in \Pi_0$ and a relaxed terminal rule accessing the complete stopped history. Express the terminal decision as $D = \bar{\psi}(H_T)$, allowing internal randomness independent of the environment. Define the target omission event as $E_{\pi,T} = \{J \notin S_{\pi,T}\}$. For a positive omission probability $P(E_{\pi,T}) > 0$, the hypothesis gives the same conditional-risk bound for all such rules. Nonnegative loss gives the following lower bound on total risk:

$$\mathbb{E}[\ell(D, \omega)] \geq \mathbb{E}[\ell(D, \omega)\mathbf{1}_{E_{\pi,T}}] = P(E_{\pi,T})\mathbb{E}[\ell(D, \omega) \mid E_{\pi,T}] \geq \rho P(E_{\pi,T}).$$

For the zero-probability case $P(E_{\pi,T}) = 0$, nonnegativity gives the same inequality. Membership $(\pi, T) \in \Pi_0$ and nonnegativity $\rho \geq 0$ bound the risk through class addressability as

$$\mathbb{E}[\ell(D, \omega)] \geq \rho\big(1 - P\{J \in S_{\pi,T}\}\big) \geq \rho\big(1 - \mathrm{Addr}_B(\Pi_0)\big).$$

Taking the infimum over relaxed terminal rules and then over $(\pi, T) \in \Pi_0$ yields the relaxed-risk bound

$$\bar{R}_{10} \geq \rho\big(1 - \mathrm{Addr}_B(\Pi_0)\big) \geq \rho(1 - \alpha), \tag{C.5}$$

where the last step uses the coverage bound $\mathrm{Addr}_B(\Pi_0) \leq \alpha$. Every actual restricted solver projects into Π_0 and has an allowed relaxed terminal rule. The relaxation therefore gives the bound $R_{10} \geq \bar{R}_{10}$. The argument holds at the infimum and supremum levels, including cases where optimal rules are unattained. □

A conditional bound established only for actual terminal rules yields a bound on R_{10}. Extending the bound to $\bar{R}_{10}$ requires the condition for every relaxed terminal rule. The independent-record and private-noise arguments below satisfy this stronger condition.

C.3 Localization risk for independent records

Fix any nonempty acquisition and stopping class Π_0. Use the localization rule $J = h(O_0)$ and record distribution $V_i \overset{\text{iid}}{\sim} \mathrm{Bernoulli}(\theta)$ with parameter $0 < \theta < 1$, subject to the independence condition $V \perp O_0$. Each query returns the exact value of a coordinate that remains fixed during interaction. Denote the optimal constant-decision risk by $\rho_\theta = \min\{\theta, 1 - \theta\}$.

Proof of Theorem 1 and the risk identity in Corollary 2. Combine all acquisition, stopping, and terminal randomness into U, independent of (O_0, V). Fix $j \in [M]$ and compare executions sharing (O_0, U, V_{-j}) with V_j set to 0 and 1. Before the first query to j, they receive the same addresses and returns, so their next-query and stopping decisions agree. Induction shows that both first query j at the same step or both stop before querying that address. The finite budget and finite content space let this comparison hold on a common event of probability one.

Define the event of omitting this address by $E_j = \{j \notin S_{\pi,T}\}$. Both E_j and the stopped history on this event are determined by (O_0, U, V_{-j}). The output D of any full-history terminal rule is also

determined by these variables on E_j. Because the localization rule $J = h(O_0)$ uses only the initial observation, the event $\{J = j\} \cap E_j$ is independent of V_j. Conditional on the other variables, V_j retains the $\mathrm{Bernoulli}(\theta)$ prior. The error contribution when target j is omitted is bounded as

$$\Pr\{D \neq V_j, J = j, E_j\} \geq \rho_\theta \Pr\{J = j, E_j\}.$$

Summing over j gives a bound valid for all adaptive queries, repeated queries, and stopping rules, expressed as

$$\Pr\{D \neq V_J, E_{\pi,T}\} \geq \rho_\theta \Pr(E_{\pi,T}). \tag{C.6}$$

For a fixed (π, T), the relaxed terminal rule recovers J from the O_0 component of the history. The rule outputs the labelled return V_J when the target record is acquired. If the target record is omitted, the output is 0 under the prior condition $\theta \leq 1/2$ and 1 otherwise. Exact reads give zero error when the target record is acquired, whereas independence gives conditional error ρ_θ when the target record is omitted. The best terminal risk for this acquisition procedure is therefore given by

$$\inf_{\bar{\psi}} \Pr\{\bar{\psi}(H_T) \neq V_J\} = \rho_\theta\big(1 - \Pr\{J \in S_{\pi,T}\}\big). \tag{C.7}$$

Taking the infimum over $(\pi, T) \in \Pi_0$ yields the class-risk relation

$$\bar{R}_{10} = \rho_\theta\big(1 - \mathrm{Addr}_B(\Pi_0)\big), \qquad R_{10} \geq \bar{R}_{10}. \tag{C.8}$$

This identity holds at the infimum even when an optimal acquisition policy is unavailable. The relaxed terminal rule identifies the target from O_0 and selects the corresponding return, so the acquisition stage need not identify whether the target record has been acquired. □

Given O_0 and independent internal randomness, V_J still follows the $\mathrm{Bernoulli}(\theta)$ prior of each record. Every passive rule therefore has risk at least ρ_θ, attained by the best constant output. This yields the identity $\rho_0(\tau) = R_{00} = R_{01} = \rho_\theta$. With a budget $B \geq 1$, an enhanced class that can compute and retain J, query that address, and output the return achieves zero risk, giving $R_{11} = 0$.

This proves the risk identity in **Corollary 2**. Under the fair-prior choice $\theta = 1/2$, every binary decision has conditional error $1/2$ when the target record is omitted, proving **Theorem 1**. For biased Bernoulli records, the optimal constant decision attains the risk bound under target omission.

C.4 The target-coverage frontier

Let $w_j = \Pr(J = j)$ denote the address probabilities, ordered as $w_{(1)} \geq \cdots \geq w_{(M)}$. The cumulative mass of the most probable addresses is given by $W_n = \sum_{j=1}^{n} w_{(j)}$, with $W_0 = 0$. For $M \geq 1$ and nonnegative integers b, B, the summary count and corresponding coverage limit are given by $K = 2^b$ and $n_\star = \min\{M, KB\}$, respectively. The coverage argument requires only the independence condition $V \perp O_0$ and allows dependent content coordinates.

Proof of Theorem 2 and the frontier in Corollary 2. Under the zero-query budget $B = 0$, both the probability of target acquisition and W_0 are zero. For a positive budget $B \geq 1$, fix an allowed encoding $C = f(O_0) \in [K]$ with the associated acquisition and stopping rules. Let U contain all internal randomness, independently of the environment. For each summary value with positive probability $\Pr(C = c) > 0$, write the visited-address set as $L_c(V, U)$. Denote these positive-probability summary values by $\mathcal{C}_+$.

Independence of (V, U) and C implies $|L_c(V, U)| \leq B$ almost surely under their common distribution for each $c \in \mathcal{C}_+$. Since $\mathcal{C}_+$ is finite, the bounds hold simultaneously on an event of probability one. On that event, the visited-address union and size bound are given by

$$L_\cup(V, U) = \bigcup_{c \in \mathcal{C}_+} L_c(V, U), \qquad |L_\cup(V, U)| \leq KB. \tag{C.9}$$

The actual visited set is $L_C(V, U)$. Acquiring the target record therefore implies membership $J \in L_\cup(V, U)$. Independence of (V, U) from J then bounds the acquisition probability as

$$\Pr\{J \in L_C(V, U)\} \leq \mathbb{E}\left[\sum_{j \in L_\cup(V,U)} w_j\right] \leq W_{n_\star}. \tag{C.10}$$

Writing executions as functions of the full V only couples the possible summary values for analysis. Each actual execution accesses queried returns. The bound consequently includes randomized, adaptive, and early-stopping policies.

For attainability, choose the $n_\star$ most probable addresses and partition them into at most K nonempty groups of size at most B. The encoder computes J and stores the corresponding group index when the target lies among the selected addresses, or the first group index otherwise. The query stage reads the entire indicated group. The size bound $n_\star \le KB$ ensures that this partition exists and respects the budget. The target record is acquired exactly when the target address belongs to the selected set. The optimal coverage is therefore given by

$$\mathrm{Addr}_B^{(b)} = W_{n_\star}. \tag{C.11}$$

Unequal group sizes and zero-probability addresses leave the construction valid. Attainability uses the full summary class's freedom to choose the encoder and subsequent rules. Subclasses inherit the upper bound and may impose further restrictions on this construction. □

The coverage formula and the i.i.d. Bernoulli record model in **Appendix C.3** give the optimal relaxed risk as

$$\bar{R}_{10}(b, B) = \rho_\theta(1 - W_{n_\star}). \tag{C.12}$$

For an actual projection class $\Pi_0 \subseteq \Pi^{(b)}$, class inclusion gives the risk bound

$$R_{10} \ge \bar{R}_{10}(\Pi_0) \ge \rho_\theta(1 - W_{n_\star}). \tag{C.13}$$

This proves the nonuniform frontier in **Corollary 2**. Under the fair-prior and uniform-target choices $\theta = 1/2$ and $w_j = 1/M$, the coverage mass reduces to $W_{n_\star} = \min\{1, 2^b B/M\}$, yielding **Theorem 2**. The relaxed-risk frontier is then given by

$$\bar{R}_{10}(b, B) = \frac{1}{2}\left(1 - \min\left\{1, \frac{2^b B}{M}\right\}\right). \tag{C.14}$$

For $0 \le \eta < 1/2$, the relaxed risk is at most η exactly when the resource condition $2^b B \ge M(1-2\eta)$ holds. This is the frontier of the full summary class with terminal relaxation. Feasibility in an actual subclass additionally requires an allowed terminal rule and acquisition construction.

C.5 Joint necessity in localization tasks

Proof of the localization necessity regime in Section 3.2. For i.i.d. Bernoulli records, the optimal constant decision attains the passive risks, giving $\rho_0(\tau) = R_{00} = R_{01} = \rho_\theta$. The enhanced class can compute and retain the target address, query that address, and output the return. With $B \ge 1$, this gives $R_{11} = 0$. For every actual projection $\Pi_0 \subseteq \Pi^{(b)}$, the resulting risk bound is expressed as

$$R_{10} \ge \bar{R}_{10}(\Pi_0) \ge \rho_\theta(1 - W_{n_\star}).$$

For a risk target satisfying $0 \le \eta < \rho_\theta(1-W_{n_\star})$, the restricted risk obeys $R_{10} > \eta$. The nonnegative coverage mass $W_{n_\star} \ge 0$ also gives $R_{01} = \rho_\theta > \eta$. Together, these bounds establish the separation $R_{11} \le \eta < \min\{R_{10}, R_{01}\}$, so **Corollary 1** gives $\mathcal{M}_\eta = \{(1, 1)\}$.

For a uniform target and fair records, the initial-information risk and coverage mass are given by $\rho_0 = 1/2$ and $W_{n_\star} = \min\{1, 2^b B/M\}$. At $0 \le \eta < 1/2$, the strict risk inequality becomes $2^b B/M < 1 - 2\eta$, as stated in **Section 3.2**. The budget condition $B \ge 1$ and the allowed enhanced-class construction complete the certificate. □

A numerical certificate with nonuniform targets and biased records. Choose the number of addresses as $M = 4$ and the initial observation as $O_0 = J$, with address distribution $(0.4, 0.3, 0.2, 0.1)$. The four contents are independent $\mathrm{Bernoulli}(1/4)$ variables, independent of J. Each query reads one address. Fix the query budget at $B = 1$ and the target risk at $\eta = 0.1$.

The two computational classes share their encoding, query, and terminal architecture and allow arbitrary measurable encoding and subsequent computation. The restricted summary has one possible value, compared with four for the enhanced summary. After encoding, all actual query, stopping, and terminal rules use only the summary, labelled returns, and independent randomness. The classes

are nested, with the restricted acquisition projection contained in $\Pi^{(0)}$. Constant outputs attain the passive risks, which are given by

$$R_{00} = R_{01} = 0.25.$$

Substituting the coverage mass $W_{\min\{4,2^0\cdot 1\}} = W_1 = 0.4$ gives the numerical risk bound

$$R_{10} \geq \bar{R}_{10} = 0.25(1 - 0.4) = 0.15 > 0.1.$$

Here the restricted acquisition projection is the full zero-bit summary class, giving an optimal terminally relaxed risk of exactly 0.15. This lower bound suffices to exclude the actual restricted class. The enhanced class encodes J in a four-valued summary and answers correctly after one query, giving $R_{11} = 0$. Together these facts certify $\mathcal{M}_{0.1} = \{(1, 1)\}$. The certificate applies to the declared interface and computational classes.

D Model interpretation and structural extensions

D.1 Location clues, answer contents, and control constraints

The addressable environment in **Section 3** represents checks on objects whose states remain fixed during interaction. A routing rule maps the object identified by the request to a storage location where reading the stored state supplies the answer. The objects O_0, h, and V_J play these three roles. For binary states such as presence or activation, an exact read of the target record provides the answer. This correspondence specifies the retained task structure. The state distribution must still be established for the application being modeled.

Independence has two distinct roles. Mutual independence of the V_i, together with the condition $V \perp O_0$, prevents other observations from predicting an unread target and yields a uniform residual risk under target omission. The bandwidth coverage bound, by contrast, requires only the independence condition $V \perp O_0$. Returns can change later actions, but carry no additional information about the address designated by the request. The summary therefore controls how acquisition distinguishes targets. Copies of the answer at other locations can lower the residual risk under target omission, while pointers in query returns can improve localization. These changes affect the two factors in the general reduction separately.

Context-retention bandwidth is a separately specified controller resource. **Definition 4** routes every influence of O_0 on subsequent acquisition through a summary of at most b bits. This models programs that organize queries through a bounded control message, a separation between message capacity and local computation familiar from communication complexity (Kushilevitz & Nisan, 1997). After the encoder forms the summary using any measurable computation, the query stage can retain newly acquired information. The coverage bound therefore isolates limited retained localization information. Further restrictions on encoding or program form yield subclasses $\Pi_0 \subseteq \Pi^{(b)}$ to which the same bound applies.

Although the full history H_t records O_0 and every interaction, the class $\mathcal{C}_c$ determines which parts the controller can actually access. Thus the history can contain O_0 while the acquisition stage accesses the initial observation only through a summary. In the matched instance, the actual terminal rule also sees only the summary and query history. The auxiliary rule $\bar{\psi}$ receives full-history access and arbitrary measurable computation. Acquisition and stopping remain in Π_0, making the resulting risk a lower bound on the original class.

Identical Bernoulli priors across independent records fix unread-target risk at ρ_θ, whereas the address distribution determines the probability mass that queries can cover. The following extensions examine heterogeneous priors, actual terminal decisions, multiple evidence records, and correlated contents.

Active feature acquisition selects observations according to their expected predictive value (Li & Oliva, 2021), and adaptive submodularity provides conditions for efficient sequential acquisition policies (Golovin & Krause, 2011). Our coverage bound quantifies the best performance of all policies satisfying the stated summary constraint. For language-model readers, evidence use can also depend on where relevant passages occur in the input (Liu et al., 2024a). This empirical effect motivates separating coverage from actual readout performance, without identifying positional effects with a fixed summary-bit budget.

Formal task representation. Let $\mathcal{O}$ be the initial-observation space, P_0 the distribution of O_0, and $h : \mathcal{O} \to [M]$ a measurable localization function. The record model in **Corollary 2** is a fixed-target task whose environment space and distribution are given by

$$\Omega = \mathcal{O} \times \{0,1\}^M, \qquad P = P_0 \otimes \mathrm{Bernoulli}(\theta)^{\otimes M}, \qquad 0 < \theta < 1. \tag{D.1}$$

For an environment $\omega = (o, v)$, define the initial observation by $O_0(\omega) = o$ and the record values by $V_i(\omega) = v_i$. The target address and answer are defined by $J(\omega) = h(o)$ and $y^\star(\omega) = v_{h(o)}$, respectively. Querying address i returns v_i. The decision space is $\mathcal{D} = \{0,1\}$ and the loss is $\ell(d, \omega) = \mathbf{1}\{d \neq v_{h(o)}\}$. The model allows arbitrary initial distributions and localization functions for records fixed before interaction.

D.2 Heterogeneous content priors

For independent and identically distributed records $V_i \sim \mathrm{Bernoulli}(\theta)$ and any nonempty acquisition and stopping class Π_0, **Appendix C.3** establishes the risk relations

$$R_{10} \geq \bar{R}_{10} = \rho_\theta\big(1 - \mathrm{Addr}_B(\Pi_0)\big), \qquad R_{00} = R_{01} = \rho_\theta. \tag{D.2}$$

The target address can have any distribution. Identical record priors give the same optimal constant-decision risk for every unacquired target record. When priors differ across addresses, we can instead use a uniform lower bound.

Address-dependent priors. Let the V_i be mutually independent with a joint distribution satisfying $V \perp O_0$. Use the record priors $V_i \sim \mathrm{Bernoulli}(\theta_i)$ with parameters $\theta_i \in [\varepsilon, 1-\varepsilon]$ and $0 < \varepsilon \leq 1/2$. Under the same localization map, exact address interface, and remaining task conditions, the risks are bounded as

$$R_{10} \geq \bar{R}_{10} \geq \varepsilon\big(1 - \mathrm{Addr}_B(\Pi_0)\big), \qquad R_{00}, R_{01} \geq \varepsilon. \tag{D.3}$$

In the adaptive-execution comparison, fix j. The event $\{J = j\} \cap E_j$ and the terminal output on this event are determined by (O_0, U, V_{-j}). Conditional on these variables, V_j retains the $\mathrm{Bernoulli}(\theta_j)$ prior, so the conditional error under target omission is at least $\min\{\theta_j, 1-\theta_j\} \geq \varepsilon$. Summing over addresses gives a bound uniform over acquisition rules and full-history terminal rules. **Proposition 4** then bounds the relaxed risk. Passive solvers never read the target, so the same argument gives their bounds. Address contents can therefore differ in predictability while preserving a localization-based risk certificate.

D.3 Target concentration and resource substitution

For a single target record, this separation gives an explicit extension to identically distributed biased records and a nonuniform target. Retain the exact address interface, the localization rule $J = h(O_0)$, the independence condition $V \perp O_0$, and the full summary class. Let the records be i.i.d. $\mathrm{Bernoulli}(\theta)$, with $0 < \theta < 1$. The prior determines the optimal constant-decision risk $\rho_\theta = \min\{\theta, 1-\theta\}$. The address probabilities $w_j = \Pr(J = j)$ determine the cumulative mass $W_n = \sum_{j=1}^{n} w_{(j)}$ of the n most probable addresses, with $W_0 = 0$.

Corollary 2 (Nonuniform targets and biased records)**.** *Under these conditions, the initial-information and passive risks satisfy $\rho_0 = R_{00} = R_{01} = \rho_\theta$, while the terminally relaxed risks are given by*

$$\bar{R}_{10} = \rho_\theta\big(1 - \mathrm{Addr}_B(\Pi_0)\big), \qquad \bar{R}_{10}(b, B) = \rho_\theta\big(1 - W_{\min\{M, 2^b B\}}\big). \tag{D.4}$$

The second expression lower-bounds R_{10} for every actual subclass $\Pi_0 \subseteq \Pi^{(b)}$.

The Bernoulli record prior gives a residual risk of ρ_θ under target omission, while target nonuniformity changes the greatest mass covered by $2^b B$ addresses to $W_{\min\{M, 2^b B\}}$. Grouping the most probable addresses attains the frontier (**Appendices C.3** and **C.4**). Concentrated targets increase coverage, while predictable contents reduce the residual risk under target omission. The uniform fair case recovers **(13)**. In the matched summary classes, the enhanced class still succeeds with one query, so both permissions are necessary under the budget and risk conditions $B \geq 1$ and $\eta < \rho_\theta(1 - W_{\min\{M, 2^b B\}})$.

Appendix C.4 establishes the exact nonuniform coverage frontier:

$$\mathrm{Addr}_B^{(b)} = W_{\min\{M,2^bB\}}, \qquad W_n = \sum_{j=1}^{n} w_{(j)}, \quad W_0 = 0. \tag{D.5}$$

The coverage proof allows correlations among content coordinates and requires independence of the content vector from O_0. Independence and identical priors across records are used separately to convert the probability of target omission into an exact risk.

With nonuniform targets, one additional bandwidth bit and twice the query budget still produce the same value of 2^bB and the same optimal coverage. The marginal gain is the mass of newly coverable addresses. Although the two resource changes give the same optimal coverage, their costs and effects on a fixed implementation require separate evaluation. Target concentration thus affects resource requirements in addition to the number of addresses.

For i.i.d. Bernoulli records, an allowed zero-risk enhanced solver, and any actual projection $\Pi_0 \subseteq \Pi^{(b)}$, a sufficient condition for agent irreplaceability is given by

$$W_{\min\{M,2^bB\}} < 1 - \frac{\eta}{\rho_\theta} \quad \Longrightarrow \quad \mathcal{M}_\eta = \{(1,1)\}, \qquad B \geq 1, \quad 0 \leq \eta < \rho_\theta. \tag{D.6}$$

Appendix C.5 provides the proof and a numerical certificate. Crossing the displayed boundary makes this localization bound inconclusive. Establishing actual feasibility then requires a successful solver in the restricted class.

D.4 Covering the target and identifying the answer

Complete coverage alone may leave the actual terminal rule unable to identify the target. To examine this distinction, use the matched summary classes with $M = 2$, $b = 0$, and $B = 2$. Let $J = h(O_0)$ be uniform, with independent fair bits V_1, V_2 satisfying $(V_1, V_2) \perp O_0$. Actual acquisition, stopping, and terminal rules access only the summary and query history. The relaxed terminal rule also accesses the original O_0. The actual and relaxed risks are given by

$$R_{10} = \frac{1}{4}, \qquad \bar{R}_{10} = 0. \tag{D.7}$$

Proof of the two-address readout gap. Under the zero-bit constraint $b = 0$, the summary is constant. With all internal randomness collected in an independent seed U, every actual history and output is determined by (V_1, V_2, U), which is independent of J. Even after reading both records, the conditional error of the actual terminal rule is bounded as

$$\begin{aligned} \Pr\{D_G \neq V_J \mid V_1, V_2, U\} &= \frac{1}{2}\mathbf{1}\{D_G \neq V_1\} + \frac{1}{2}\mathbf{1}\{D_G \neq V_2\} \\ &\geq \frac{1}{2}\mathbf{1}\{V_1 \neq V_2\}. \end{aligned} \tag{D.8}$$

The records differ with probability $1/2$, giving risk at least $1/4$. Querying both addresses and outputting V_1 attains this bound, since an error occurs exactly on the intersection of the events $J = 2$ and $V_1 \neq V_2$. With the same acquisition procedure, the relaxed rule recovers J from O_0 and selects the labelled return V_J, achieving zero risk with no additional query. □

The gap comes from terminal access to localization information. Both records are in the history, but when they disagree, the actual terminal rule still cannot identify which one the request designates.

At $b = 0, B = 1$, the class is a subclass of the two-query class and still has risk at least $1/4$. Querying address 1 and outputting V_1 attains this bound. At $b = 1, B = 1$, encoding and querying address J gives zero actual risk. The zero-bit classes have optimal coverage $1/2$ and 1 at budgets 1 and 2, respectively. The one-bit, one-query class also has coverage 1. All three comparisons use the same task and exact interface, with the stated summary capacities and query budgets.

D.5 COMPLETE COVERAGE OF MULTIPLE EVIDENCE RECORDS

To analyze acquisition when an answer depends on several records, replace the target address with a target evidence set. For initial observation O_0 and address space $[M]$, define the nonempty target evidence set by

$$\mathcal{K} = K(O_0) \subseteq [M], \qquad \mathcal{K} \neq \varnothing, \tag{D.9}$$

where K is a task-specified measurable map. Target sets can vary in size and overlap across instances. For an acquisition and stopping class Π, the *complete-coverage probability* is the supremal probability of acquiring every target record, given by

$$\mathrm{Cov}_B(\Pi) := \sup_{(\pi,T)\in\Pi} \Pr\{\mathcal{K} \subseteq S_{\pi,T}\}. \tag{D.10}$$

We first characterize optimal complete coverage under the summary constraint, then relate this coverage to a task requiring every target record. Queries return exact coordinate values fixed before interaction. The class $\Pi^{(b)}$ permits any encoder and subsequent rules satisfying **Definition 4**.

Theorem 3 (Coverage frontier for target evidence sets)**.** *Let the address count satisfy $M \geq 1$ and let b, B be nonnegative integers. Assume the independence condition $V \perp O_0$ and use internal randomness independent of the environment. With the number of summary values given by $q = 2^b$, the optimal complete-coverage probability is given by*

$$\mathrm{Cov}_B(\Pi^{(b)}) = \max_{L_1,\ldots,L_q\subseteq[M],|L_c|\leq B\ (c\in[q])} \Pr\{\exists c \in [q] : \mathcal{K} \subseteq L_c\}. \tag{D.11}$$

The identity allows dependent coordinates of V and an arbitrary target-set distribution.

Proof of Theorem 3. Under the zero-query budget $B = 0$, every query set is empty while $\mathcal{K}$ is nonempty, so both sides are zero. For the remaining argument, take a positive budget $B \geq 1$.

For the upper bound, fix an encoding $C = f(O_0) \in [q]$ and the associated acquisition and stopping rules. Collect all internal randomness in U. Fixing (V, U) makes the visited-address set from summary c a deterministic set $L_c(V, U)$. This is an analytical representation of the execution, which reads only queried values. Because every influence of O_0 on acquisition passes through C, the visited-address set is expressed as

$$S_{\pi,T} = L_{f(O_0)}(V, U). \tag{D.12}$$

For each c with positive probability $\Pr(C = c) > 0$, independence of (V, U) and C transfers the conditional budget bound to $|L_c(V, U)| \leq B$ almost surely. Finitely many such values give a common probability-one event on which all bounds hold. For unused summary values, assign the empty visited set $L_c(V, U) = \varnothing$.

Complete coverage implies that one of these sets contains $\mathcal{K}$. Independence of the target set $\mathcal{K} = K(O_0)$ from (V, U) then bounds the complete-coverage probability as

$$\Pr\{\mathcal{K} \subseteq S_{\pi,T}\} \leq \mathbb{E}_{V,U}\left[\Pr_{O_0}\{\exists c \in [q] : K(O_0) \subseteq L_c(V, U)\}\right]. \tag{D.13}$$

For every (V, U) in the common probability-one event, the inner probability uses fixed address sets of size at most B and is therefore bounded by the maximum in the theorem. Taking the supremum over allowed policies proves the upper bound.

For attainability, finiteness of the address space ensures that some $(L_1^\star, \ldots, L_q^\star)$ attains the maximum. After reading O_0, the encoder stores the smallest index of a set containing $K(O_0)$, or index 1 if none exists. This is a measurable encoding into $[q]$. The query stage reads the indicated set in a fixed order and stops after at most B queries. The event of complete coverage is given by

$$\{\exists c \in [q] : \mathcal{K} \subseteq L_c^\star\}.$$

This procedure belongs to $\Pi^{(b)}$ and attains the maximum. □

The frontier in **(D.11)** gives the largest probability of fully covering the target evidence set with a collection of query lists. For $\mathcal{K} = \{J\}$, the frontier reduces to the largest probability mass on at

most $2^b B$ addresses, giving $\min\{1, 2^b B/M\}$ for a uniform target. Multiple-record coverage also depends on set size and overlap. A target set larger than B cannot be fully covered, while one list may cover several target sets.

To connect coverage to risk, let all record values be independent fair bits, independent of O_0, and let the target be their parity over the evidence set, i.e.,

$$V_1, \dots, V_M \overset{\text{iid}}{\sim} \text{Bernoulli}(1/2), \qquad V \perp O_0, \qquad y^\star(\omega) = \bigoplus_{i \in \mathcal{K}} V_i. \tag{D.14}$$

Here $\bigoplus$ denotes addition modulo 2. The decision space is $\{0, 1\}$ and loss is zero-one. Every unread target bit can change the answer, so partial coverage leaves decision uncertainty. Retain $(\pi, T) \in \Pi_0$ and give the terminal rule the complete stopped history and arbitrary measurable computation. The terminally relaxed risk is the infimum of prediction error over these acquisition procedures and terminal rules, given by

$$\bar{R}_{10}^{\mathcal{K}} := \inf_{(\pi,T)\in\Pi_0, \bar{\psi} \text{ measurable}} \Pr\{\bar{\psi}(H_T) \neq y^\star(\omega)\}. \tag{D.15}$$

The terminal rule may use independent internal randomness and makes zero additional queries.

Theorem 4 (Exact risk for multiple-record parity)**.** *For any nonempty acquisition and stopping projection Π_0, the actual and relaxed risks are related by*

$$R_{10} \geq \bar{R}_{10}^{\mathcal{K}} = \frac{1}{2}\big(1 - \text{Cov}_B(\Pi_0)\big), \tag{D.16}$$

where R_{10} is the class-optimal risk under the original computational restrictions.

Proof of Theorem 4. Fix $(\pi, T) \in \Pi_0$ and collect all acquisition, stopping, and terminal randomness in U. The incomplete-coverage event is defined by

$$F = \{\mathcal{K} \not\subseteq S_{\pi,T}\}.$$

Conditional on (O_0, U), all rules and $\mathcal{K}$ are fixed. The vector V is uniform on the finite set $\{0, 1\}^M$ and independent of (O_0, U). We can therefore choose a probability-one set of (O_0, U) on which the budget holds for every v. Fix such a pair.

For a vector v producing event F, select the first unread target coordinate by

$$j(v) = \min\big(\mathcal{K} \setminus S_{\pi,T}(v)\big),$$

and obtain v' by complementing only coordinate $j(v)$. Both executions start from the same initial observation and seed. If their histories agree, their query and stopping decisions agree. The original execution never queries $j(v)$, so every queried value remains unchanged. Induction gives identical stopping times, stopped histories, and visited-address sets.

Thus v' also produces F with the same target set, so the selected coordinate satisfies $j(v') = j(v)$. This transformation is a fixed-point-free involution on F. Uniformity of V makes this involution measure-preserving. Each paired execution produces the same terminal output and opposite target parities. For every binary decision D based on the stopped history and seed, the error contribution from incomplete coverage is given by

$$\Pr\{D \neq y^\star(\omega), F\} = \frac{1}{2}\Pr(F). \tag{D.17}$$

Although the selected unread coordinate depends on the execution, complementing the selected coordinate preserves that execution and the selection event. The argument includes adaptive and repeated queries and early stopping.

The error on complete coverage is nonnegative, so total risk is at least $\frac{1}{2}\Pr(F)$. A relaxed terminal rule recovers $\mathcal{K}$ from O_0. On complete coverage, the rule selects the labelled target values and computes their parity. Otherwise, the output is 0, attaining conditional error $1/2$ under incomplete coverage by **(D.17)**. The optimal terminal risk for the fixed acquisition procedure is therefore given by

$$\inf_{\bar{\psi}} \Pr\{\bar{\psi}(H_T) \neq y^\star(\omega)\} = \frac{1}{2}\left(1 - \Pr\{\mathcal{K} \subseteq S_{\pi,T}\}\right). \tag{D.18}$$

Taking the infimum over $(\pi, T) \in \Pi_0$ gives the relaxed class risk as

$$\bar{R}^{\mathcal{K}}_{10} = \frac{1}{2}\big(1 - \mathrm{Cov}_B(\Pi_0)\big).$$

The equality follows from the infimum and supremum even when an optimal acquisition policy is unattained. Every actual terminal rule is allowed by the relaxation, giving $R_{10} \geq \bar{R}^{\mathcal{K}}_{10}$. □

Because $\mathcal{K}$ is nonempty, the parity over $\mathcal{K}$ is fair conditional on O_0. Both passive classes therefore have risk $1/2$, attained by constants. Suppose the enhanced class can compute and retain $\mathcal{K}$, query every address in $\mathcal{K}$, and return the parity, with $|\mathcal{K}| \leq B$ almost surely. This construction gives $R_{11} = 0$. For risk targets $0 \leq \eta < 1/2$, if the class inclusion $\Pi_0 \subseteq \Pi^{(b)}$ holds and the coverage frontier in **Theorem 3** is below $1 - 2\eta$, the matched-risk criterion gives $\mathcal{M}_\eta = \{(1, 1)\}$. This certifies joint necessity for targets distributed over several records.

D.6 Residual risk under target omission with correlated contents

Other records can partially predict the target, while target omission still leaves irreducible risk. Retain the localization rule $J = h(O_0)$, target V_J, binary decisions, and the exact address interface. For an address count $M \geq 1$, define a shared state and private noise model by

$$\begin{aligned} W \sim \mathrm{Bernoulli}(1/2), \qquad Z_1, \ldots, Z_M \overset{\mathrm{iid}}{\sim} \mathrm{Bernoulli}(\varepsilon), \\ V_i = W \oplus Z_i \quad (i \in [M]), \qquad 0 < \varepsilon < \frac{1}{2}. \end{aligned} \tag{D.19}$$

Take the variable W and all Z_i to be mutually independent, with the vector $(W, Z_1, \ldots, Z_M)$ also independent of O_0. The symbol $\oplus$ denotes XOR. Each V_i is fair, while distinct records share W and satisfy $\Pr\{V_i = V_j\} = (1 - \varepsilon)^2 + \varepsilon^2 > 1/2$ for $i \neq j$. Thus the contents are correlated and still satisfy $V \perp O_0$.

Proposition 5 (Residual risk for correlated records)**.** *For this task, any finite query budget, any allowed (π, T), and any relaxed full-history terminal rule, define the decision by $D = \bar{\psi}(H_T)$. For a positive probability of target omission $\Pr(E_{\pi,T}) > 0$, the conditional risk is bounded as*

$$\mathbb{E}[\ell(D, \omega) \mid E_{\pi,T}] \geq \varepsilon. \tag{D.20}$$

Consequently, for every nonempty restricted acquisition projection Π_0, the class risks satisfy the following relations:

$$R_{10} \geq \bar{R}_{10} \geq \varepsilon\big(1 - \mathrm{Addr}_B(\Pi_0)\big), \qquad R_{00} = R_{01} = \frac{1}{2}. \tag{D.21}$$

With a budget $B \geq 1$, an enhanced class that can compute J, query that address, and output the return achieves zero risk, giving $R_{11} = 0$.

Proof of Proposition 5. Fix the acquisition, stopping, and relaxed terminal rules. Collect their randomness in an independent seed U. For each j, let Z_{-j} contain the remaining private variables and define the omission event by $E_j = \{j \notin S_{\pi,T}\}$. Compare executions sharing (O_0, U, W, Z_{-j}) with $Z_j = 0$ and $Z_j = 1$. Before the first query to j, all queried contents agree, so query and stopping decisions agree. Induction shows that both first query j at the same step or both stop before querying that address. The finite budget terminates this comparison.

The event E_j and the terminal output on that event consequently depend only on (O_0, U, W, Z_{-j}). The localization rule $J = h(O_0)$ gives the same dependence for the event $\{J = j\} \cap E_j$. Conditional on these variables, Z_j retains the $\mathrm{Bernoulli}(\varepsilon)$ prior and W is fixed. Either output bit therefore has error at least ε when predicting $V_j = W \oplus Z_j$. The error contribution from omitting target j is bounded as

$$\Pr\{D \neq V_j,\ J = j,\ E_j\} \geq \varepsilon \Pr\{J = j,\ E_j\}. \tag{D.22}$$

Sum over j and divide by the positive probability of target omission to obtain the uniform conditional-risk bound. **Proposition 4** bounds $\bar{R}_{10}$, thereby giving a lower bound for the actual class through $R_{10} \geq \bar{R}_{10}$. The proof includes adaptive and repeated queries and early stopping.

Because the independence condition $V \perp O_0$ holds and every V_j is fair, the target remains fair conditional on O_0 and internal randomness. Both passive classes have risk $1/2$, attained by constants. The enhanced class reads the target once and has zero risk. □

The model retains the independence condition $V \perp O_0$, so the coverage analysis in **Appendix D.3** applies. For a uniform target and any subclass $\Pi_0 \subseteq \Pi^{(b)}$, addressability is bounded as

$$\mathrm{Addr}_B(\Pi_0) \le \mathrm{Addr}_B^{(b)} = \min\left\{1, \frac{2^b B}{M}\right\}. \tag{D.23}$$

Use a budget $B \ge 1$ and suppose the enhanced class contains the zero-risk construction above. For risk targets $0 \le \eta < \varepsilon$, a sufficient condition for agent irreplaceability is expressed as

$$\min\left\{1, \frac{2^b B}{M}\right\} < 1 - \frac{\eta}{\varepsilon}. \tag{D.24}$$

This condition establishes the risk relations $R_{10} > \eta$, $R_{01} = 1/2 > \eta$, and $R_{11} = 0$. **Proposition 1** yields $\mathcal{M}_\eta(\tau) = \{(1,1)\}$.

The shared factor W gives other records predictive value. Even exact knowledge of W leaves minimum error ε from the target's private variable Z_J. The actual risk conditional on target omission also depends on how well the observations identify W, so ε is a uniform lower bound. Unobserved private information can therefore preserve localization risk in a correlated environment.

E Computational models and bounds for history processing

We first transfer risk bounds on complete raw data to interactive solver classes, then instantiate the transfer in an online fan-in-two circuit model. A two-block parity task gives the separation in **(16)**. We then extend the argument to general targets and derive the bound for anomaly aggregation.

E.1 Risk transfer through output-preserving simulation

Let $X \sim P_X$ be the raw data and $F(X) \in \{0,1\}$ the target under zero-one loss. Let $\mathcal{A}(w)$ contain programs with access to the complete raw data and cost at most w. The optimal prediction risk within this class is the infimum of program error probabilities, given by

$$r_F(w) := \inf_{A \in \mathcal{A}(w)} \Pr\{A(X) \ne F(X)\}. \tag{E.1}$$

Full-data access follows the program model's reading, selection, and processing costs. Program descriptions, fixed preprocessing, and the distribution of auxiliary randomness are chosen independently of the current unknown data. Complete access supplies the raw data. Aggregate statistics and the answer are computed within the program's resource bound.

Proposition 6 (Simulation transfer of computational risk)**.** *Let $\mathcal{G}_{\mathrm{res}}$ be a nonempty class of restricted interactive implementations with computation budget W_0, instance size n, and query budget B. Suppose every $G \in \mathcal{G}_{\mathrm{res}}$ has a program A_G, chosen independently of the current input, that satisfies the following uniform cost bound and joint-distribution identity:*

$$A_G \in \mathcal{A}\big(\Phi(W_0, n, B)\big), \qquad (X, A_G(X)) \overset{\mathrm{d}}{=} (X, D_G), \tag{E.2}$$

where the joint distributions include each procedure's internal randomness. The optimal interactive risk is then bounded as

$$\inf_{G \in \mathcal{G}_{\mathrm{res}}} \Pr\{D_G \ne F(X)\} \ge r_F\big(\Phi(W_0, n, B)\big). \tag{E.3}$$

Proof of Proposition 6. For each $G \in \mathcal{G}_{\mathrm{res}}$, the simulation belongs to the program class defining the right-hand infimum. Equality of the joint input-output distributions gives the risk comparison

$$r_F\big(\Phi(W_0, n, B)\big) \le \Pr\{A_G(X) \ne F(X)\} = \Pr\{D_G \ne F(X)\}.$$

Taking the infimum over G proves the result. □

The simulator reconstructs the initial observation, conditional interface returns, and required interaction randomness. The simulation cost includes dynamic addressing, stopping, and final output

selection. Preserving the joint distribution with X ensures preservation of risk. Existence of a sufficient acquisition procedure is a separate condition used below to establish zero terminally relaxed risk.

For $\mathcal{G}_{\mathrm{res}} = \Gamma_{10}$, abbreviate the uniform simulation cost as $w_{\mathrm{sim}} = \Phi(W_0, n, B)$. The resulting lower bound on restricted acquisition risk is expressed as

$$R_{10} \geq r_F(w_{\mathrm{sim}}). \tag{E.4}$$

The proposition permits any computational cost model. To instantiate the transfer, we next give a pointwise simulation for raw block access under a depth bound. Applying the transfer to runtime or another resource requires the corresponding simulation cost and distributional risk bound.

E.2 A TWO-SLOT ONLINE MODEL AND MATCHED COMPUTATIONAL CLASSES

Fix a positive integer L and define the total input size by $N = 2L$. The external data comprise two raw bit blocks, represented as

$$X = (X^{(0)}, X^{(1)}) \in \{0,1\}^L \times \{0,1\}^L. \tag{E.5}$$

Index the concatenated coordinates by $X_1, \ldots, X_N$. The initial observation is a fixed description of the task, block length, and interface, independent of the current data. An action $a \in \{0,1\}$ returns $X^{(a)}$ exactly under the query budget $B = 2$. The target is a specified binary function $F(X)$. The following subsections specify the data distribution and target function.

Both computational classes use a two-slot causal Boolean circuit. Each implementation has a finite fixed circuit description and random input U independent of X. The circuit may depend on N, the task distribution, and resource parameters and may have any finite size for each N. The circuit description and random-input distribution are fixed independently of the current unknown X. Ordinary gates have fan-in at most two and may implement any Boolean truth table, including negation (NOT), conjunction (AND), OR, and binary XOR, using the standard bounded-fan-in circuit convention (Arora & Barak, 2009). Fixed wires, constants, and fan-out have zero depth cost.

At slot $s \in \{1,2\}$, the circuit produces an enable bit $e_s \in \{0,1\}$ and address $a_s \in \{0,1\}$. These signals are computed from public constants, U, and returns from earlier slots, before the current return is available. Every execution satisfies $e_2 \leq e_1$. The enabled state $e_s = 1$ triggers an interface call, whereas the disabled state $e_s = 0$ makes a local wrapper return the placeholder 0^L. Return coordinate i is therefore given by

$$R_{s,i} = e_s \wedge \big((\neg a_s \wedge X_i^{(0)}) \vee (a_s \wedge X_i^{(1)})\big), \qquad i \in [L]. \tag{E.6}$$

With $T = e_1 + e_2$ actual acquisitions, H_T records only actual calls and returns. Disabled-slot placeholders are generated from the known enable state and supply no external evidence. The circuit implements prefix enabling, stopping, and final answer selection within the declared cost model.

Depth and state access. The circuit designates one final output bit D. The ordinary depth d is the maximum number of ordinary gates on a directed path in the complete causal graph. Enabled block reads are separate interface primitives with zero internal depth weight, while ordinary depth accumulated before a read remains on the path afterward. Gates for enabling, address selection, state updates, stopping, and final answer selection all count toward d. Data-dependent register selection, array access, and candidate-answer selection must be implemented with ordinary gates.

The representation permits storage of all returns and any number of intermediate wires. Computation can occur between acquisitions and reuse intermediate results across rounds, but depth accumulates across those rounds. This model bounds the causal depth of fan-in-two operations, while B separately bounds external calls.

Matched classes. Let $\mathcal{C}(d)$ contain all solvers with a representation of ordinary depth at most d. Each solver chooses the query, stopping, online-processing, and output rules. For nonnegative integers $d_0 \leq d_1$, define the matched computational classes by

$$\mathcal{C}_0 = \mathcal{C}(d_0), \qquad \mathcal{C}_1 = \mathcal{C}(d_1). \tag{E.7}$$

The classes share the task, interface, two-slot representation, history storage, and query budget. Only the ordinary-depth limit changes. The passive class Γ_{0c} requires $T = 0$ on every execution, while Γ_{1c} allows at most two calls. A constant zero-query circuit has depth zero and belongs to both classes. Nested depth bounds imply $\mathcal{C}_0 \subseteq \mathcal{C}_1$.

These causal circuits specify the implementation class for this instance. Another program representation can use the depth result after a cost-preserving representation or simulation has been established.

E.3 Interface simulation and output support

For the simulation cost, extend the representation to B causal slots with the same two-block interface and prefix condition $e_{s+1} \le e_s$. All other rules remain unchanged. This extension expresses the overhead as a function of slot count. The matched instance continues to use $B = 2$.

Proposition 7 (Interface simulation on complete raw data)**.** *Every B-slot implementation of ordinary depth at most d has a pointwise simulation by an ordinary fan-in-two circuit on (X, U). With a uniform overhead $\Delta_B \le 4B$, the simulation depth is bounded as*

$$d_{\mathrm{sim}} \le d + \Delta_B. \tag{E.8}$$

The simulation preserves actual queries, stopping, and the final output for each (x, u) and therefore preserves the joint output distribution with X.

Proof of Proposition 7. Retain the original ordinary gates and wires. For slot s and return coordinate i, replace the read primitive in **(E.6)** by four successive gate layers:

$$\begin{gathered} n_s = \neg a_s, p_{s,i} = n_s \wedge X_i^{(0)}, q_{s,i} = a_s \wedge X_i^{(1)}, \\ m_{s,i} = p_{s,i} \vee q_{s,i}, r_{s,i} = e_s \wedge m_{s,i}. \end{gathered} \tag{E.9}$$

The second layer computes $p_{s,i}$ and $q_{s,i}$ in parallel. The signal n_s fans out to every coordinate, and all coordinates are expanded in parallel. The longest path in **(E.9)** passes through NOT, AND, OR, and the enabling AND, adding four gates.

Fix (x, u). Because the first slot uses the same constants and seed for the enable bit and address, the expansion returns the same block or zero placeholder. If earlier returns agree, the next slot's control signals and intermediate gate values also agree, preserving the expanded return. Induction preserves every actual query and return. Prefix enabling preserves early stopping and subsequent zero padding. The designated output wire and all dynamic selection gates remain unchanged, so the final output agrees pointwise.

Consider any directed path in the simulation. For a path beginning at an original constant or random input, contracting each read expansion gives a path in the original causal graph with at most d ordinary gates. A path beginning at a raw data coordinate first enters a read expansion. Contracting later expansions gives a suffix of an original path, again with at most d ordinary gates. Causal order allows at most B read expansions along a path, each adding at most four gates. The depth bound $d + 4B$ therefore permits $\Delta_B = 4B$. Pointwise agreement also establishes the required joint-distribution equality for random inputs. □

The original interactive circuit accesses X only through interface returns. The simulation uses raw input wires to implement those returns while retaining ordinary gates for read selection, enabling, and answer selection. For two slots, the choice $\Delta_2 = 8$ gives a conservative overhead bound.

Output support. To relate depth to data dependence, we use the support of a Boolean function $f : \{0,1\}^N \to \{0,1\}$, the set of input coordinates on which the output depends (O'Donnell, 2014), expressed as

$$\mathrm{Supp}(f) := \{j \in [N] : \exists x, x' \text{ differing only at } j,\ f(x) \ne f(x')\}. \tag{E.10}$$

After fixing independent random inputs, a node of ordinary depth h in a fan-in-two circuit has support size at most 2^h. The proof is inductive. A data input has singleton support, whereas a

constant has empty support. A gate can depend only on coordinates that affect an input to the gate. Thus, for a node v with predecessors v_1, v_2, the support inclusion is expressed as

$$\mathrm{Supp}(v) \subseteq \mathrm{Supp}(v_1) \cup \mathrm{Supp}(v_2). \tag{E.11}$$

The induction hypothesis and the identity $h(v) = 1 + \max\{h(v_1), h(v_2)\}$ bound the support size as

$$|\,\mathrm{Supp}(v)| \le |\,\mathrm{Supp}(v_1)| + |\,\mathrm{Supp}(v_2)| \le 2^{h(v_1)} + 2^{h(v_2)} \le 2^{1+\max\{h(v_1),h(v_2)\}} = 2^{h(v)}. \tag{E.12}$$

Unary gates satisfy the same bound. Induction over the causal graph extends the bound to the output. The simulation gives the output-support bound for every depth-d interactive implementation and fixed seed u:

$$\big|\mathrm{Supp}(D_G(\cdot, u))\big| \le \min\{N, 2^{d+\Delta_B}\}. \tag{E.13}$$

This bounds the support of the entire output function. The support can vary with u but is fixed across current inputs x. A coordinate may already be present in the history while lying outside this support.

E.4 The parity separation and terminal relaxation

Let $X_1, \ldots, X_N$ be independent $\mathrm{Bernoulli}(1/2)$ variables. The parity target and zero-one loss are defined by

$$F(X) = \bigoplus_{j=1}^{N} X_j, \qquad \ell(D, X) = \mathbf{1}\{D \neq F(X)\}. \tag{E.14}$$

Here $\bigoplus$ denotes addition modulo 2. Use the two-block interface, $B = 2$, and matched computational classes from **Appendix E.2**.

Proposition 8 (Processing separation for two blocks)**.** *For a total input size $N = 2L$, let nonnegative integers $d_0 \le d_1$ satisfy the depth conditions*

$$2^{d_0+8} < N, \qquad d_1 \ge \lceil \log_2 N \rceil. \tag{E.15}$$

The resulting matched risks are given by

$$R_{00} = R_{01} = R_{10} = \frac{1}{2}, \qquad R_{11} = 0. \tag{E.16}$$

Consequently, the minimal sufficient set is given by $\mathcal{M}_\eta = \{(1,1)\}$ for every risk target $0 \le \eta < 1/2$.

Proof of Proposition 8. Take any allowed implementation of $G \in \Gamma_{10}$ and fix the implementation seed u. The simulation and **(E.13)** bound the output support as

$$\big|\mathrm{Supp}(D_G(\cdot, u))\big| \le 2^{d_0+8} < N.$$

Choose a coordinate $j(u)$ outside this support. Fixing all other coordinates, complementing $X_{j(u)}$ preserves the output and changes the parity target. The bit is independent and fair, so exactly one input in each pair is misclassified. The fixed-seed error is given by

$$\Pr_X\{D_G(X, u) \neq F(X)\} = \frac{1}{2}. \tag{E.17}$$

Independence of U and X preserves error $1/2$ after averaging over u. This argument covers all query, stopping, and online-processing rules in the restricted class, proving $R_{10} = 1/2$.

Passive implementations use a fixed O_0 and independent random inputs. The fair target is independent of both, so passive risk is $1/2$, attained by constant outputs.

The enhanced implementation fixes $e_1 = e_2 = 1$, $a_1 = 0$, and $a_2 = 1$, acquiring both raw blocks. Define the required tree depth by $h = \lceil \log_2 N \rceil$. Feed the N returned bits and $2^h - N$ zero constants into a balanced binary XOR tree of depth h with the root designated as the output. Query control and stopping are constant. Under the common cost convention, the depth bound $h \le d_1$ and the two acquisitions make this construction admissible. The output is correct on every input, establishing $R_{11} = 0$. The matched-risk criterion gives the stated minimal configuration. □

The regime is nonempty. For example, the parameter choices $N = 1024$, $d_0 = 1$, and $d_1 = 10$ satisfy **(E.15)**. This condition certifies a separation regime, while the exact depth frontier remains unspecified.

Terminal relaxation with fixed acquisition rules. Let Π_0 be the projection of Γ_{10} onto acquisition and stopping rules. As in the main text, the terminally relaxed risk is the infimum over this acquisition class and all measurable terminal decision rules, given by

$$\bar{R}_{10}(\tau_{\mathrm{proc}}) := \inf_{(\pi,T)\in\Pi_0} \inf_{\bar{\psi}\ \text{measurable}} \Pr\{\bar{\psi}(H_T) \neq F(X)\}. \tag{E.18}$$

The terminal rule accesses the complete stopped history and internal randomness, with zero additional queries. Risk uses the same target $F(X)$, while acquisition and stopping retain their original restrictions.

***Proof of the terminal-relaxation identity in* (16).** Consider the fixed acquisition procedure $(\pi^\star, T^\star)$ that reads the two distinct blocks and stops at $T^\star = 2$. Combining this acquisition procedure with a constant output gives an allowed restricted implementation of ordinary depth zero. This establishes membership $(\pi^\star, T^\star) \in \Pi_0$. The stopped history contains both raw blocks and their labels. A measurable terminal rule $\bar{\psi}^\star$ computes their parity. The relaxed risk is consequently bounded as

$$0 \leq \bar{R}_{10}(\tau_{\mathrm{proc}}) \leq \Pr\{\bar{\psi}^\star(H_{T^\star}) \neq F(X)\} = 0. \tag{E.19}$$

□

One admissible acquisition procedure and the associated relaxed terminal rule establish this upper bound, while **Proposition 8** excludes every actual restricted implementation. Localization instead retains positive risk after terminal relaxation. The localization and processing tasks therefore exhibit different failure mechanisms under the same minimal configuration. Each risk comparison uses the classes and budget fixed for the task being analyzed.

Although the successful construction uses a fixed query order, the lower bound covers adaptive queries and online processing.

E.5 General targets and anomaly aggregation

Equation (17) defines the minimum prediction risk from at most k fixed coordinates. We prove **Proposition 3** for arbitrary input distributions and binary targets, with coordinate sets and prediction rules chosen independently of the current input.

Proof of Proposition 3. For fixed G and seed u, denote the output support by $S_u = \mathrm{Supp}(D_G(\cdot, u))$. The output is some $g_u(X_{S_u})$. Indeed, two inputs agreeing on S_u can be connected by complementing coordinates outside S_u, each of which preserves the output. The support bound $|S_u| \leq k$ makes this prediction eligible for the defining infimum. The fixed-seed prediction risk is bounded as

$$\Pr_X\{D_G(X, u) \neq F(X)\} \geq \varepsilon_F(k). \tag{E.20}$$

Averaging over the independent seed gives the same lower bound for every implementation, so taking the infimum proves the claim. The two-slot simulation gives $|S_u| \leq \min\{N, 2^{d_0+8}\} = k_0$, yielding **(18)**. The dependence bound applies to every target and also to the full-input case $k_0 = N$, where the prediction risk is given by $\varepsilon_F(N) = 0$. □

For fair parity and $k < N$, an independent fair coordinate remains outside the support, giving $\varepsilon_F(k) = 1/2$. At $k = N$, the rule $g = F$ gives zero risk. The next target yields a risk that decreases as more coordinates become available.

Anomaly aggregation. Retain the two-block raw-record interface and computational model. Interpret each coordinate as an anomaly flag and ask whether any flag is present. Define the target and record distribution by

$$F(X) = \bigvee_{j=1}^{N} X_j, \qquad X_j \overset{\text{iid}}{\sim} \mathrm{Bernoulli}(p_N), \qquad p_N = 1 - 2^{-1/N}. \tag{E.21}$$

This prior gives $\Pr(F(X) = 0) = (1 - p_N)^N = 1/2$, preserving a balanced binary target.

Corollary 3 (Prediction risk for anomaly aggregation)**.** *The optimal fixed-coordinate prediction risk for the anomaly task is given by*

$$\varepsilon_F(k) = 2^{-k/N} - \frac{1}{2}, \qquad 0 \le k \le N. \tag{E.22}$$

For the matched two-slot classes, define the support bound by $k_0 = \min\{N, 2^{d_0+8}\}$*. Under the enhanced-depth condition* $d_1 \ge \lceil \log_2 N \rceil$*, the class risks satisfy the following relations:*

$$R_{00} = R_{01} = \frac{1}{2}, \qquad R_{10} \ge 2^{-k_0/N} - \frac{1}{2}, \qquad R_{11} = \bar{R}_{10} = 0. \tag{E.23}$$

Hence the minimal sufficient set is given by $\mathcal{M}_\eta = \{(1,1)\}$ *for risk targets* $0 \le \eta < 2^{-k_0/N} - 1/2$*.*

Proof of Corollary 3. Define the zero-flag probability by $q = 1 - p_N = 2^{-1/N}$ and fix S of size m. If any coordinate in X_S equals 1, the target is certainly 1. Otherwise, an event of probability q^m, independence gives conditional probability $q^{N-m} \ge q^N = 1/2$ for target 0. The optimal prediction is therefore 0. The resulting optimal risk is given by

$$q^m(1 - q^{N-m}) = q^m - q^N = 2^{-m/N} - \frac{1}{2}. \tag{E.24}$$

The rule that outputs 1 after observing a flag and 0 otherwise attains this risk. The risk decreases with m. Since all coordinates have the same distribution, the choice $m = k$ proves **(E.22)**, including $k = 0$ and $k = N$. Randomization preserves the conditional minimum-risk bound.

Proposition 3 transfers this bound to the actual restricted class. The fixed initial observation and balanced target give passive risk $1/2$. The enhanced solver reads both blocks and evaluates a balanced binary OR tree, padding to a power of two with zeros. The tree has depth $\lceil \log_2 N \rceil$, giving $R_{11} = 0$. The same fixed acquisition procedure with a constant output belongs to the restricted class. A relaxed terminal rule computes OR from the complete history, giving $\bar{R}_{10} = 0$. The matched-risk criterion proves joint necessity. □

Equation (E.22) is exact for the auxiliary prediction class and lower-bounds the actual interactive class. The result extends the same output-dependence analysis to aggregation over many records, where complementing a single bit changes the answer only on some inputs.

F Experimental setup for natural-language tasks

F.1 Data and evaluation metrics

FinQA (Chen et al., 2021) and HotpotQA (Yang et al., 2018) evaluate numerical reasoning and multi-hop question answering, respectively. The main experiments use 1,000 questions from each dataset, corresponding to 275 report clusters and 893 clusters of questions sharing candidate articles. FinQA accuracy follows the original numerical-answer scoring rules. HotpotQA uses official answer exact match (EM) and the token-level F1 score, the harmonic mean of precision and recall, supplemented by supporting-fact and joint metrics. Complete annotated-support coverage is the fraction of questions for which all annotated supporting records are present in the acquired history, and is computed separately from answer accuracy.

To examine overlap with development material, we also report results on 896 FinQA questions with nonoverlapping reports and 937 HotpotQA questions with nonoverlapping articles. The main analysis retains the full test cohorts, and the subset analyses use the same configurations and scoring procedures. The evidence-replacement experiment retains 961 FinQA questions in 271 report clusters that satisfy the fragment-matching conditions. **Appendix I** describes this construction.

F.2 Acquisition procedures and processing methods

The main model is Qwen2.5-14B-Instruct (Qwen Team, 2024). Configurations without acquisition use only the initial question. Fixed BM25 retrieval (Robertson & Zaragoza, 2009) selects 2, 4, or 8 evidence blocks in one step according to relevance. Adaptive acquisition uses acquired records to

Table G.1. Class risks (%) and measured errors of specified programs in parentheses. Minimal sufficient configurations correspond to target risks $0 \le \eta < 1/2$. The answer summary is available to both acquisition classes.

Task or interface	R_{00}	R_{10}	R_{01}	R_{11}	Minimal
Initial coordinate	0 (0.00)	0 (0.00)	0 (0.00)	0 (0.00)	$(0, 0)$
External coordinate	50 (49.60)	0 (0.00)	50 (49.60)	0 (0.00)	$(1, 0)$
Initial parity	50 (49.92)	50 (49.92)	0 (0.00)	0 (0.00)	$(0, 1)$
External parity	50 (49.79)	50 (49.93)	50 (49.79)	0 (0.00)	$(1, 1)$
External parity + summary	50 (50.15)	0 (0.00)	50 (50.15)	0 (0.00)	$(1, 0)$

select subsequent queries and returns at most 8 blocks. All-evidence configurations receive every candidate record. Different readouts reuse the original acquisition history of the same procedure.

The FinQA direct readout produces a final numerical answer. The program readout generates and executes operations in a fixed domain-specific language (DSL). Annotated programs are used only for evaluation, and generated programs access only visible records. Draft enhancement first generates intermediate analysis and then applies the corresponding direct or program readout. The HotpotQA enhanced readout organizes existing evidence before answering and citing supporting facts. Acquisition control, intermediate drafts, and final answers are logged separately.

F.3 Paired statistical comparisons

The four primary FinQA comparisons are fixed B8 direct versus program readout, all-evidence program readout with versus without a draft, fixed B8 versus adaptive program readout, and the difference in draft gains between all-evidence and B8 inputs. Each reports a 98.75% interval. The two primary HotpotQA comparisons are fixed B8 versus adaptive direct answering and all-evidence direct versus organized readout, each with a 97.5% interval. Bonferroni correction (Dunn, 1961) is applied separately to these two families, each targeting 95% nominal simultaneous coverage.

Confidence intervals use 10,000 cluster bootstrap replicates (Field & Welsh, 2007). Each replicate resamples complete report or article clusters, retains their questions and paired answers, and computes the overall score difference with question-count weighting. The evidence-replacement experiment reports 95% cluster intervals. The four primary comparisons for the second model report 98.75% intervals.

F.4 Resource accounting

Acquisition volume is measured by the number of returned record blocks. Model calls include acquisition control, drafts, and final answers, and token counts include their inputs and outputs. Drafts and control records reused across configurations are charged to the deployment cost of every configuration that uses them. Summing deployment costs across configurations therefore differs from the total cost of executing the experiment. Program execution time is not included in the token metric. Quality–token curves compare the mean costs of evaluated configurations, while acquisition counts and sequential calls are reported separately.

G Controlled-task setup and complete results

G.1 Comparison of the four configurations

We separate acquisition from computation by varying evidence location and the target function under the same record representation. Single-coordinate tasks read a specified bit, while parity tasks compute the parity of $N = 1024$ bits. Raw records are placed either in the initial observation or behind the external interface. The query budget is $B = 2$, and restricted and enhanced programs have ordinary gate depths of 1 and 10. We measure each specified program's error on 20,000 inputs and compare the measured error with the risk certificate for the corresponding solver class.

Table G.2. Complete localization results (%), with 20,000 inputs per setting. The restricted-terminal reference is the Bayes risk under the fixed encoding and acquisition procedure.

b	B	Analytical coverage	Measured coverage	Relaxed reference	Relaxed measured	Restricted reference	Actual error
4	8	12.500	12.185	43.750	43.415	48.291	48.480
4	16	25.000	24.765	37.500	37.070	47.545	47.515
4	48	75.000	75.110	12.500	12.275	45.704	45.780
4	51	79.688	79.645	10.156	10.095	45.527	45.630
4	52	81.250	81.270	9.375	9.370	45.527	45.515
4	64	100.000	100.000	0.000	0.000	45.033	44.760
5	8	25.000	24.845	37.500	36.940	46.582	46.340
5	24	75.000	74.770	12.500	12.420	43.956	44.170
5	25	78.125	77.970	10.938	10.930	43.704	43.590
5	26	81.250	81.260	9.375	9.355	43.704	43.570
5	32	100.000	100.000	0.000	0.000	43.003	42.970
10	1	100.000	100.000	0.000	0.000	0.000	0.000

Table G.1 shows that moving the raw records changes acquisition requirements, whereas changing the target function changes processing requirements. External parity requires both capabilities and thus the full agent configuration. When the interface provides an answer summary to both acquisition classes, a restricted acquisition program can also solve the task, and the minimal sufficient configuration becomes $(1, 0)$. Necessity therefore depends on the alternative solution procedures permitted by the task interface.

G.2 Retaining localization information and query budgets

The localization experiments use $M = 1024$ addresses and a b-bit summary of initial target information, while varying the query budget B. **Table G.2** reports coverage, terminally relaxed risk, and the actual restricted terminal's error. The relaxed terminal accesses the initial information in the complete history. The actual terminal accesses only the summary and query records.

At $b = 4, B = 64$, every target is covered, yet actual error is 44.76%. The setting $b = 5, B = 32$ also attains complete coverage. The query records contain the target state but do not restore the full target identity lost in the summary. Retaining the complete address allows the specified program at $b = 10, B = 1$ to achieve zero error. This risk gap reflects a limit on information retention, distinct from the depth limit when complete evidence is available.

G.3 Structural extensions and analytical calibration

We further vary the target distribution, required evidence sets, and record correlation to examine how the risk characterizations depend on task structure. The nonuniform localization grid uses $M \in \{64, 256\}$ and varies state bias, the Zipf exponent, retention bandwidth, and query budget. Multiple-evidence tasks vary the size and overlap of evidence sets. Correlated-record tasks vary private noise. **Table G.3** summarizes these grids and the discrepancies between program measurements and analytical references.

For nonuniform tasks, the probability mass covered by a fixed address capacity changes with distribution concentration. More concentrated target locations yield lower relaxed risk under the same bandwidth and query budget, as shown in **Figure 4(a)**. Multiple-evidence tasks extend single-target coverage to acquisition of the required record set. Overlapping evidence sets allow different targets to share query returns, so the required budget also depends on set structure.

For fixed-support predictors with complete raw evidence, fair parity retains risk $1/2$ for every $k < N$ and drops to zero at $k = N$. Balanced OR risk decreases gradually with k, while a single-coordinate target can be predicted exactly once that coordinate is accessible. **Figure 2(b)** compares these analytical risks with 33 measurements of specified programs. Applying this prediction risk to interactive solvers requires the full-data simulation relation in **Section 3.3**.

Table G.3. Experimental scope and analytical calibration for structural extensions. Errors are in percentage points.

Structural variation	Settings	Mean abs. difference	Max. abs. difference	Reference
Nonuniform locations and biased states	288	0.134	0.783	Terminally relaxed risk
Evidence sets and overlap	216	0.207	0.725	Fixed-route relaxed risk
Correlated records and private noise	120	0.120	0.480	Fixed-route relaxed risk

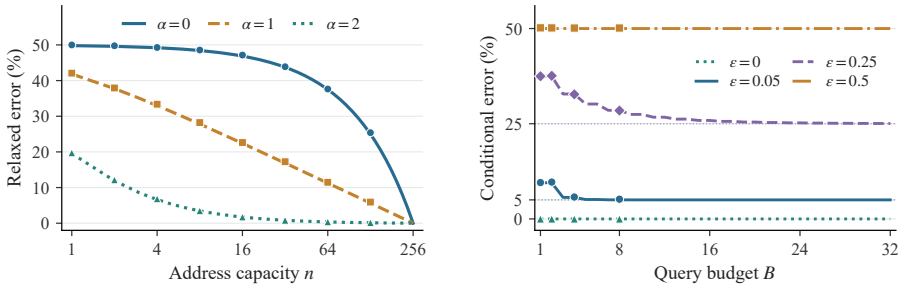


Figure 4. Localization under nonuniform targets and correlated records. (a) Terminally relaxed risk for $M = 256$ independent fair records with target probabilities proportional to $j^{-\alpha}$. Address capacity is $n = \min\{M, 2^b B\}$. (b) Conditional risk under target omission for $M = 256, b = 0$ with one fair latent bit affecting every record and independent record-specific flips of probability ε. Curves give analytical values and points give measurements of specified programs.

Record correlation allows observed records to inform an unobserved target. **Figure 4(b)** fixes $M = 256, b = 0$ and varies the number of queries and private noise. As queries increase, conditional risk under target omission approaches the corresponding noise floor. The acquisition lower bound thus depends on both omission probability and conditional risk after omission. The more predictive other records are, the less value there is in reading the target record directly.

H Complete natural-task configurations and resource use

This appendix reports every natural-task configuration compared in the main text. FinQA has 20 configurations and HotpotQA has 12, each evaluated on a fixed denominator of 1,000 questions. **Tables H.1 and H.2** report answer quality, complete support coverage, acquisition volume, model calls, and total input and output tokens. Token counts are given in thousands (k). FinQA also reports valid-output counts.

The 32 configurations produce 32,000 final answers. Including drafts and acquisition control, the experiments execute 26,714 model generations for FinQA and 20,600 for HotpotQA. Quality differences on the same acquisition history correspond to changes in processing. Comparisons across acquisition procedures also include changes in control and acquired evidence.

H.1 Primary paired results

To distinguish net gains from changes on individual questions, we report confidence intervals and answer-transition counts for the primary comparisons. Wrong-to-correct transitions indicate an incorrect baseline answer corrected by the intervention. Correct-to-wrong transitions indicate a previously correct answer made incorrect. All comparisons use the same questions and the cluster resampling procedure in **Appendix F.3**.

On FinQA with fixed B8 retrieval, program readout corrects 358 answers and loses 26, giving a net gain of 33.2 percentage points. Adding a draft to the all-evidence program readout corrects 136

Table H.1. Complete FinQA results. Accuracy and coverage are percentages, and resources are means per question. Valid counts are totals of valid outputs. Bold marks the highest accuracy for each readout.

Acquisition	Readout	Accuracy	Coverage	Blocks	Calls	Tokens (k)	Valid
No acquisition	Program	0.8	0.0	0.00	1.00	0.698	485
No acquisition	Direct	0.2	0.0	0.00	1.00	0.319	30
No acquisition	Draft + direct	0.2	0.0	0.00	2.00	0.760	25
No acquisition	Draft + program	0.5	0.0	0.00	2.00	1.113	40
B2	Program	32.1	40.0	2.00	1.00	0.854	838
B2	Direct	12.6	40.0	2.00	1.00	0.484	551
B4	Program	39.7	56.1	4.00	1.00	1.017	864
B4	Direct	15.3	56.1	4.00	1.00	0.646	748
B8	Program	50.3	73.4	8.00	1.00	1.325	914
B8	Direct	17.1	73.4	8.00	1.00	0.958	866
B8	Draft + direct	32.9	73.4	8.00	2.00	2.328	869
B8	Draft + program	54.2	73.4	8.00	2.00	2.691	872
All evidence	Program	**57.3**	100.0	31.60	1.00	2.665	927
All evidence	Direct	17.5	100.0	31.60	1.00	2.297	955
All evidence	Draft + direct	**37.9**	100.0	31.60	2.00	5.023	978
All evidence	Draft + program	**64.9**	100.0	31.60	2.00	5.387	962
Adaptive	Program	51.6	77.9	7.09	3.71	2.593	912
Adaptive	Direct	**18.8**	77.9	7.09	3.71	2.225	906
Adaptive	Draft + direct	34.0	77.9	7.09	4.71	3.538	923
Adaptive	Draft + program	55.2	77.9	7.09	4.71	3.902	911

Table H.2. Complete HotpotQA results. EM, F1, and coverage are percentages, and resources are means per question. Bold marks the highest EM for each readout.

Acquisition	Readout	EM	F1	Coverage	Blocks	Calls	Tokens (k)
No acquisition	Direct	14.6	19.00	0.0	0.00	1.00	0.181
No acquisition	Organized	24.0	31.90	0.0	0.00	2.00	0.612
B2	Direct	33.9	43.89	13.4	2.00	1.00	0.338
B2	Organized	33.7	44.90	13.4	2.00	2.00	0.894
B4	Direct	42.5	55.22	30.6	4.00	1.00	0.501
B4	Organized	42.4	55.29	30.6	4.00	2.00	1.247
B8	Direct	48.0	62.21	51.8	7.99	1.00	0.829
B8	Organized	48.8	63.74	51.8	7.99	2.00	1.899
All evidence	Direct	**52.3**	66.63	100.0	41.15	1.00	3.375
All evidence	Organized	**56.4**	73.26	100.0	41.15	2.00	7.003
Adaptive	Direct	52.2	66.88	63.9	6.66	3.60	1.886
Adaptive	Organized	52.8	68.73	63.9	6.66	4.60	2.827

answers and loses 60, for a net gain of 7.6 points. Adaptive program readout gains 1.3 points, with an interval containing zero. The intervals for HotpotQA adaptive direct answering and all-evidence organization are both positive. The four FinQA comparisons retain the same interval conclusions in the subset with nonoverlapping reports. The two HotpotQA gains in the subset with nonoverlapping articles are 4.269 and 3.842 points.

H.2 Evidence coverage and answer transitions

Table H.4 groups questions by complete annotated-support coverage under fixed B8 and adaptive acquisition and compares the resulting answer transitions. FinQA uses program readout and Hot-

Table H.3. Primary paired comparisons. The four FinQA comparisons use 98.75% intervals and the two HotpotQA comparisons use 97.5% intervals. Score differences are in percentage points (pp). Bold marks gains whose intervals exclude zero.

Prespecified comparison	Gain (pp)	Adjusted interval	Wrong → correct / correct → wrong
FinQA, B8 direct → program	**+33.2**	[28.30,37.93]	358/26
FinQA, all-evidence program → draft + program	**+7.6**	[3.93,11.34]	136/60
FinQA, B8 program → adaptive program	+1.3	[-1.44,4.02]	66/53
FinQA, draft gain all evidence minus B8	+3.7	[-0.50,7.95]	–
HotpotQA, B8 → adaptive direct	**+4.2**	[1.51,6.93]	93/51
HotpotQA, all-evidence direct → organized	**+4.1**	[1.19,7.01]	102/61

Table H.4. Answer transitions under different support-coverage conditions. Net change is the number of wrong-to-correct transitions minus correct-to-wrong transitions.

Task	Support coverage	Questions	Fixed correct	Adaptive correct	Wrong → correct	Correct → wrong	Net
FinQA	Both complete	699	435	425	24	34	-10
FinQA	Adaptive only	80	5	29	27	3	+24
FinQA	Fixed only	35	25	13	2	14	-12
FinQA	Neither complete	186	38	49	13	2	+11
HotpotQA	Both complete	408	254	257	15	12	+3
HotpotQA	Adaptive only	231	76	133	63	6	+57
HotpotQA	Fixed only	110	53	38	1	16	-15
HotpotQA	Neither complete	251	97	94	14	17	-3

Table H.5. Supporting-fact and joint EM on HotpotQA (%). Supporting-fact matching evaluates model outputs and is computed separately from complete support coverage in the input history.

Acquisition	Direct support EM	Organized support EM	Direct joint EM	Organized joint EM
No acquisition	0.0	0.7	0.0	0.4
Fixed B2	11.3	10.6	7.9	7.3
Fixed B4	18.2	18.5	12.1	13.0
Fixed B8	23.2	27.3	14.7	18.5
Adaptive $B \leq 8$	27.4	33.4	17.7	22.0
All evidence	23.2	35.4	14.5	22.4

potQA uses direct readout. Groups are formed after both histories are acquired to examine the sources of the net gain.

Among questions for which adaptive acquisition newly obtains complete support, the net gains are 24 correct answers on FinQA and 57 on HotpotQA. Answers also change when both procedures cover all support or both fail to do so, indicating that evidence organization and readout also affect performance. Of the 358 FinQA answers corrected by program readout, 335 originally had a valid direct-output format but an incorrect numerical answer. Among the 810 questions with valid outputs under both readouts, the correct count increases from 165 to 480. Program readout therefore corrects many errors in answer content.

Table H.6. HotpotQA comparisons under a common budget, with 600 questions per configuration. Three-step processing and adaptive acquisition share the stated upper limits. Bold marks the highest EM, and underlining marks the best fixed-history result.

Acquisition and readout	Correct	EM (%)
Fixed B8 direct	275/600	45.83
Fixed B8 organized	292/600	48.67
Fixed B8 three-step processing	294/600	<u>49.00</u>
Adaptive direct	300/600	**50.00**

H.3 Resource allocation under a common budget

We compare three-step processing of a fixed history with adaptive acquisition on 600 HotpotQA questions to examine alternative uses of additional calls. Both have upper limits of 8 evidence blocks, 4 model calls, and 1,088 generated tokens. Fixed-history processing produces, checks, and revises a proposal, using at most 192 tokens per step and 512 for the final answer. Each step reads the same raw history and previous notes. Adaptive acquisition starts with two BM25 blocks and permits at most three 192-token control steps, each adding at most two blocks. The adaptive procedure may stop early and uses at most 512 tokens for the final answer.

Three-step processing with fixed B8 achieves 49.0% EM, while adaptive direct answering achieves 50.0%, a difference of 6 correct answers. Spending additional calls on either fixed-history processing or subsequent evidence selection yields a higher point estimate than one direct answer. The same budget can improve a solution through different procedures, depending on the limits that remain in the current history.

On a further set of 243 FinQA questions, all-evidence direct answering achieves 18.5%, while B4 and B8 program readouts on the same questions achieve 39.92% and 52.26%. Less material combined with program processing still outperforms direct answering with all evidence, extending the resource tradeoff observed in the main text. These two supplementary comparisons use question cohorts that have already participated in development.

H.4 Configuration selection by diagnosis and development search

We compare configuration selection based on diagnosis with direct development search. Both select the same configuration on FinQA and HotpotQA. On FinQA, both choose adaptive program readout, which achieves 51.6% test accuracy. This comparison therefore shows no difference in the selected configuration or test score. The preceding paired interventions distinguish the roles of these configurations. Restoring support changes the available information, readout on a fixed history changes evidence use, and adaptive acquisition applies computation to subsequent query control.

I Evidence conditions and processing gains

I.1 Experimental setup

We hold the question, record slots, and other material fixed and replace one annotated supporting record with an unrelated fragment of similar length. This constructs paired histories with complete and missing support. In FinQA, 961 questions from 271 report clusters meet the predetermined fragment-matching conditions, while 39 do not. Replaced supporting records include 306 text fragments and 655 table fragments. Selection depends on fragment matching and does not use generated answers.

Both support conditions compare base and enhanced readouts. FinQA uses program and draft-plus-program readouts, while HotpotQA uses direct and organized readouts. The distractor sweep adds the same nested sets of distractors to paired histories, with $0, 2, 4, 8, 16, 32$ additional records. **Figure 5** reports answer scores and enhanced-readout gains.

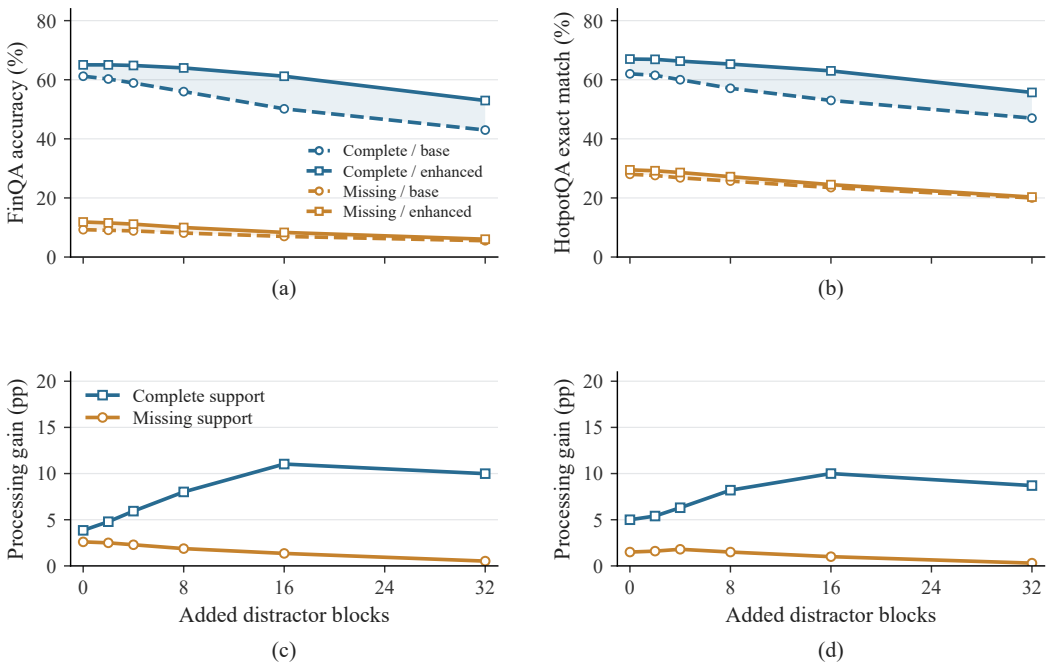


Figure 5. Distractor sweeps under different evidence conditions. Panels (a,b) show accuracy or EM, and (c,d) show the gains from enhanced readout. Blue denotes complete annotated support and orange denotes removal of one supporting record. Shaded areas show the score difference between readouts.

I.2 Restoring support and enhancing processing

In **Figure 5(a)**, the FinQA program readout without additional distractors achieves 61.19% accuracy with complete support and 9.26% with missing support. Adding a draft gives 65.04% and 11.86%, respectively. Restoring support yields gains of 51.93 and 53.17 percentage points, with 95% cluster intervals of $[48.03, 55.74]$ and $[49.35, 56.92]$. Missing support leaves the readout with insufficient information, limiting the errors that the current draft procedure can correct. Restoring support supplies the material that the solver can then organize and compute with.

I.3 Effect of distractor count

Table I.1 gives the complete values for **Figure 5**. Processing gain is the score difference between enhanced and base readouts. The interaction is the difference in processing gains between complete and missing support.

With complete support, the FinQA draft gain increases from 3.85 points without distractors to 11.03 with 16 distractors, and is 9.99 with 32. The corresponding HotpotQA organization gains are 5.0, 10.0, and 8.7 points. With missing support, processing gains generally diminish, reaching 0.52 and 0.3 points at 32 distractors. Once usable support is present, additional processing helps extract and use relevant records among distractors. When support is missing, the same processing steps have limited ability to recover the information on which the answer depends.

I.4 Paired estimates and the nonoverlapping subset

Without distractors, the FinQA draft gains with complete and missing support are 3.85 and 2.60 points, with 95% cluster intervals of $[0.80, 6.80]$ and $[0.75, 4.53]$. Their interaction is 1.25 points, with interval $[-2.14, 4.55]$. Restoring support has a clear benefit, while the difference between draft gains in these two conditions remains uncertain.

In the subset of 861 questions from 241 nonoverlapping report clusters, the interaction is 0.93 points, with a 95% interval of $[-2.69, 4.54]$. The gains from restoring support are 52.15 points for program

Table I.1. Complete distractor sweep. Scores are percentages, and gains and interactions are percentage points. Base and enhanced columns refer to the corresponding readouts.

Task	Distractors	Complete base	Complete enhanced	Missing base	Missing enhanced	Complete gain	Missing gain	Interaction
FinQA	0	61.19	65.04	9.26	11.86	3.85	2.60	1.25
FinQA	2	60.25	65.04	9.05	11.55	4.79	2.50	2.29
FinQA	4	58.90	64.83	8.84	11.13	5.93	2.29	3.64
FinQA	8	55.98	64.00	8.12	9.99	8.01	1.87	6.14
FinQA	16	50.16	61.19	6.97	8.32	11.03	1.35	9.68
FinQA	32	42.98	52.97	5.52	6.04	9.99	0.52	9.47
HotpotQA	0	62.00	67.00	28.00	29.50	5.00	1.50	3.50
HotpotQA	2	61.50	66.90	27.60	29.20	5.40	1.60	3.80
HotpotQA	4	60.00	66.30	26.80	28.60	6.30	1.80	4.50
HotpotQA	8	57.10	65.30	25.70	27.20	8.20	1.50	6.70
HotpotQA	16	53.00	63.00	23.50	24.50	10.00	1.00	9.00
HotpotQA	32	47.00	55.70	20.00	20.30	8.70	0.30	8.40

Table J.1. Paired intervention effects across models, in percentage points. The final column reports 98.75% cluster intervals for Mistral gains. Bold marks the Mistral effect whose interval excludes zero.

Paired comparison	Qwen14B (pp)	Mistral7B (pp)	Mistral 98.75% interval
FinQA, B8 direct → program	+33.2	**-3.4**	[-5.30,-1.62]
FinQA, all-evidence program + draft	+7.6	+0.6	[-0.10,1.48]
HotpotQA, B8 → adaptive direct	+4.2	+0.5	[-2.79,3.77]
HotpotQA, all-evidence direct → organized	+4.1	-2.6	[-5.87,0.79]

readout and 53.08 points for draft-plus-program readout. This subset exhibits the same pattern as the complete matched set.

J Effects of the model and output interface

J.1 Experimental setup

To examine workflow interventions across models, we repeat two primary FinQA comparisons and two primary HotpotQA comparisons with Mistral-7B-Instruct-v0.3 (Mistral AI, 2024), using 1,000 questions per task. We retain the corresponding acquisition procedures, input histories, output interfaces, and scoring rules, and compare workflow gains within each model. Each of the four comparisons uses a 98.75% paired cluster interval.

J.2 Results

Table J.1 shows that Mistral's fixed B8 program readout decreases accuracy by 3.4 points relative to direct answering. Adding a draft with all evidence yields 0.6 points. On HotpotQA, adaptive acquisition and organized readout yield 0.5 and -2.6 points, respectively. The latter three intervals all contain zero and do not reproduce the gains observed with the main model.

The four FinQA configurations for Mistral produce 421, 61, 19, and 33 valid outputs for B8 direct, B8 program, all-evidence program, and all-evidence draft-plus-program readouts, respectively. Invalid programs include unevaluated expressions, nonexistent table calls, and references to undefined variables. Only 3 final outputs reach the length limit, so most failures concern interface adherence. Available computation can affect the answer only through valid intermediate represen-

Table K.1. GSM8K with necessary facts externalized, using 200 questions per configuration. The reference-evidence condition directly supplies designated supporting material. Bold marks the best score, and underlining marks the best score without designated reference evidence.

Input condition	Procedure	Correct	Accuracy (%)
Facts externalized	Direct readout	4/200	2.0
Facts externalized	Fixed retrieval + readout	<u>46/200</u>	<u>23.0</u>
Facts externalized	Plan + direct readout	2/200	1.0
Facts externalized	Planned retrieval + readout	26/200	13.0
Reference evidence	Reference-evidence readout	**161/200**	**80.5**

tations. Program-interface failures prevent the additional processing capability from translating into the corresponding task gains for this model.

K Application examples of four capability requirements

We begin with simpler alternative workflows and examine tasks with sufficient initial input, missing external facts, and limited processing of available evidence. The final example uses AgentRAN, an agent-based radio access network (RAN) controller, to compare resource requirements across degradations in the same system. Each example changes visible information or the processing procedure to examine how capability requirements depend on task conditions.

K.1 A basic workflow for self-contained questions

We use 200 Grade School Math 8K (GSM8K) questions (Cobbe et al., 2021) and Qwen2.5-32B-Instruct (Qwen Team, 2024). The initial input contains every condition in the question, and the model answers directly without external retrieval or additional planning. This workflow answers 198 questions correctly, achieving 99.0% accuracy.

Initial information and the basic readout already achieve high accuracy on these questions without external acquisition or additional planning. The direct workflow provides a simple solution corresponding to the $(0, 0)$ region. Full agent capabilities are not a prerequisite for obtaining good answers in this setting.

K.2 Acquisition requirements when facts are externalized

We next place necessary facts for the same questions outside the initial input, retaining question identifiers and target answers. Under this condition, we compare direct answering, fixed retrieval followed by answering, planning followed by direct answering, and planned retrieval. Designated reference evidence provides an additional comparison. **Table K.1** reports the correct counts for each configuration.

Direct answering falls from 198 correct answers in the self-contained condition to 4. Fixed retrieval recovers 46, and reference evidence yields 161. Moving necessary facts while preserving the target creates an acquisition requirement. Fixed retrieval exceeds planned retrieval at 26 correct answers, showing that additional planning does not automatically improve acquired evidence. This example exhibits the acquisition-dominated behavior of the $(1, 0)$ region, with the main improvement coming from material missing from the initial input.

K.3 Mathematical reasoning with complete problem statements

The American Invitational Mathematics Examination (AIME) comparison (Mathematical Association of America, n.d.) uses 60 questions with complete problem conditions in the initial input. We compare direct answering and chain-of-thought (CoT) prompting (Wei et al., 2022) with the same Qwen2.5-32B weights, under no retrieval, retrieval of solved problems, and retrieval of reference material. DeepSeek-R1-Distill-Qwen-32B (DeepSeek-AI, 2025) additionally evaluates a reasoning model with different training and decoding protocols.

Table K.2. Processing and retrieval on the American Invitational Mathematics Examination (AIME), with 60 questions per configuration. The first two rows use the same model weights, with CoT denoting chain-of-thought prompting. Bold marks the highest correct count in each row.

Model or protocol	No retrieval	Solved problems	Reference material
Qwen2.5-32B direct	**1/60**	0/60	0/60
Same weights, CoT	7/60	**8/60**	5/60
R1-Distill-Qwen-32B	**42/60**	36/60	36/60

Table K.3. Net utility J under different degradations. Regions describe the resource-response patterns of the compared configurations, with axes for additional measurements and planning depth. Bold marks the highest utility in each row.

Region	Degradation	Base	+ Meas.	+ Comp.	Joint
$(0, 0)$	Rate-limit error only	**0.81**	0.80	**0.81**	0.80
$(1, 0)$	Coverage loss only	0.42	**0.70**	0.43	0.69
$(0, 1)$	Known load and protection constraints	0.42	0.41	**0.73**	0.72
$(1, 1)$	Coverage + overload	0.31	0.48	0.48	**0.64**
$(1, 1)$	Interference + overload	0.27	0.43	0.29	**0.63**
$(1, 1)$	Triple degradation	0.16	0.25	0.20	**0.34**

With the same weights, CoT increases the correct count without retrieval from 1 to 7. The two types of retrieved material yield 8 and 5. R1 answers 42 questions correctly without retrieval and 36 with either type of material. The main difficulty is deriving the answer from conditions already given, and additional retrieval does not yield consistent improvements. This example exhibits the processing-dominated behavior of the $(0, 1)$ region. Task difficulty alone does not imply a need for external acquisition.

K.4 AGENTRAN RECOVERY TASKS ACROSS THE FOUR REGIONS

AgentRAN examines capability requirements for different degradations within the same wireless-network control environment. The OpenAirInterface platform (Kaltenberger et al., 2020) contains three cells, six stationary users, and a fifth-generation (5G) core network. Four site conditions represent co-channel interference, weak neighboring-cell coverage, concentrated traffic, and mixed carriers. Degradations include interference, coverage loss, overload, rate-limit errors, and their combinations. The controller adjusts power, hands over users, or reallocates transport capacity to restore target services while protecting neighboring users.

The two comparison axes are additional diagnostic measurements and planning depth. The base configuration requests no additional measurements and uses depth $h = 2$. Enhanced acquisition permits measurements, enhanced computation increases depth to $h = 3$, and the joint configuration permits both. Joint configurations further distinguish batch acquisition with predetermined requests from adaptive acquisition based on returned observations. Every configuration retains action-verification feedback.

Table K.3 shows a progression from base to full resource requirements. For rate-limit errors alone, the base configuration already achieves utility 0.81, and additional measurements or planning do not improve utility. For coverage loss alone, measurements increase utility from 0.42 to 0.70, while additional computation alone yields 0.43. With known load and protection constraints, deeper planning increases utility from 0.42 to 0.73, with little benefit from additional measurements. Compound degradations require both state identification and action composition under constraints, and the joint configuration exceeds either single-axis enhancement.

Figure 6 shows the same progression in verified recovery. For rate-limit errors alone, the base configuration achieves 82%. For coverage loss, measurements increase recovery from 43% to 73%, while additional planning alone yields 44%. With known load, planning increases recovery from

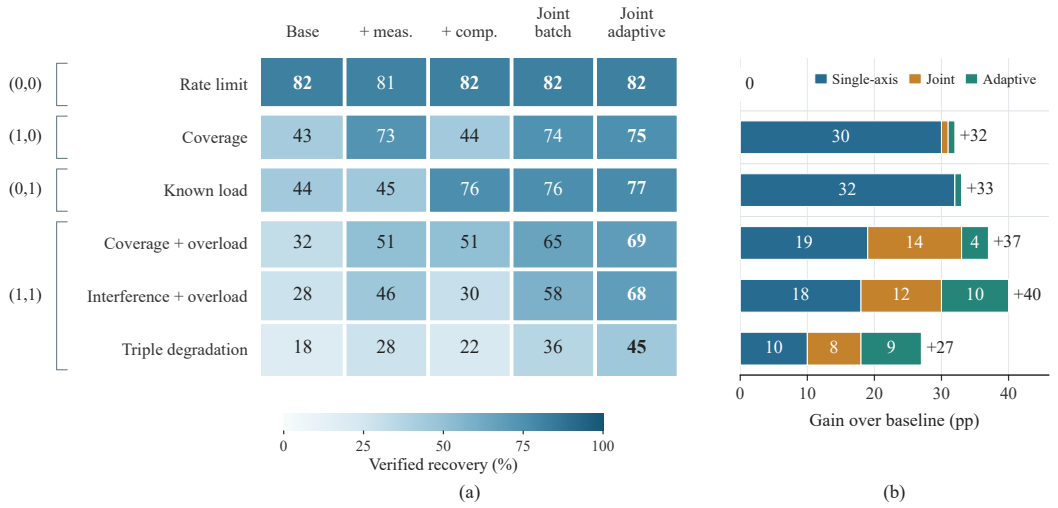


Figure 6. Resource responses in AgentRAN. (a) Verified recovery rates (%) by configuration. (b) Successive gains over the base configuration through the best single-axis reference, batch joint acquisition, and adaptive joint acquisition. The single-axis reference is the highest recovery rate among the base and two single-axis configurations.

Table K.4. Overall policy comparison in AgentRAN, with 1,920 episodes per policy. RANFix is the complete controller in this case study. Resources are means per episode. Bold marks the highest utility and recovery rate.

Policy	J	Recovery (%)	Measurements	LLM calls
Symbolic	0.624	66	3.6	0.0
LLM	0.598	63	3.2	5.8
Model + handoff	0.618	64	2.2	0.0
Model + symbolic	0.726	76	3.4	0.0
RANFix	**0.788**	**82**	3.2	1.4

44% to 76%, while additional measurements yield 45%. Even within the same control system, the required capabilities depend on the particular degradation.

When interference and overload occur together, the best single-axis, batch joint, and adaptive joint configurations achieve 46%, 58%, and 68% recovery. Under triple degradation, they achieve 28%, 36%, and 45%. Increasing only measurement access or planning depth leaves the other limitation unresolved. Joint configurations use acquired state information to organize recovery actions. Adaptive acquisition further uses early returns to guide later decisions. This fourth example thus includes cases solved by the base configuration, acquisition-dominated cases, processing-dominated cases, and cases where the two resources complement each other.

The overall policy comparison uses 1,920 test episodes, equally weighting the four sites and assigning half the weight each to single and compound degradations. The objective is $J = p_{\text{ver}} - 0.01\mathbb{E}[N_q]$, where p_{ver} is verified recovery and N_q is the measurement count. Recovery requires target and protected users to sustain their service targets. All policies share test inputs, measurement and recovery capabilities, episode budgets, and the verification procedure. The resource comparisons in **Figure 6** use a common empirical-support subset, while the overall comparison evaluates the full test cohort.

Relative to selecting actions with an LLM alone, RANFix increases recovery from 63% to 82% and reduces mean LLM calls from 5.8 to 1.4. Relative to model-plus-symbolic continuation, RANFix improves recovery by 6 points and reduces mean measurements from 3.4 to 3.2. Together with the resource comparisons, these results suggest distinguishing uncertainty about state from limits on action composition when selecting a configuration. Degradations already handled by the base work-

flow can retain a simple procedure, while compound degradations depend more on joint investment in measurement and planning.