History, Hypergraphs, and Memory: The Exact Complexity of Deviation-Rational Control

Anonymous authors
Paper under double-blind review

Keywords: reinforcement learning, big worlds, bounded-memory agents, history processes, deviation regret, hypergraph transversals, convex geometry, combinatorial optimization.

Summary

We study the minimum internal memory required for ε-deviation-rational control in arbitrary history processes. Our starting point is a past-compression view of bounded agency: after each attainable history, the agent compresses the entire past into one of finitely many memory labels, and that label must decode to a continuation controller whose deviation regret is at most ε. This formulation is entirely agent-side. It does not assume that the environment is Markov, fully observable, deterministic, finitely enumerable, or representable by a small latent model.

Our main result is an exact structural law: retentive complexity equals the transversal number of the family of history-indexed low-regret continuation sets. In the finite-controller mixed-continuation model, this yields a precise hypergraph law at exact rationality ($\varepsilon = 0$): the required memory is exactly the transversal number of the family of rational-support sets. We further show that every finite hypergraph can be realized as such a support family in a one-step fully observable instance, establishing combinatorial universality. Finally, we prove that pairwise compatibility is insufficient in general—even the complete compatibility graph can conceal large memory gaps—while convexity yields a Helly-type local certificate for one-state memory. Together, these results recast bounded-agent memory in big worlds as a problem in combinatorial optimization and convex geometry rather than environment-state abstraction.

Contribution(s)

1. We prove an exact structural law: for every deviation-stable continuation system and every $\varepsilon \geq 0$, retentive complexity equals the transversal number of the family of history-indexed ε-rational continuation sets.
   Context: This is an exact characterization of memory complexity, not an algorithmic tractability result.
2. In the finite-controller mixed-continuation model, we show that low-regret sets are convex polytopes and that, at exact rationality, retentive complexity is exactly the hypergraph transversal number of the family of rational-support sets; moreover, every finite hypergraph is realizable as such a support family in a one-step fully observable instance.
   Context: The hypergraph law and universality statement are proved for finite controller dictionaries with continuation values affine under mixing.
3. We prove a higher-order impossibility result: pairwise compatibility does not determine exact retentive complexity, and for every order q admitting a finite projective plane there exist two instances with identical complete compatibility graph K_{q^2+q+1} but exact retentive complexities 1 and $q+1$.
   Context: This is a negative structural result; the projective-plane construction provides an explicit obstruction, not a classification of all obstructions.
4. We identify the convex regime in which local structure suffices: one-state memory admits a Helly-type certificate, and in the finite-controller simplex model any deviation class containing the external deviations yields exactly the same low-regret sets and retentive complexities as the external-deviation model for all $\varepsilon \geq 0$.
   Context: The Helly and robustness statements are finite-controller simplex results; they are not claimed for arbitrary continuation spaces.

History, Hypergraphs, and Memory: The Exact Complexity of Deviation-Rational Control

Anonymous authors
Paper under double-blind review

Abstract

We study the minimum internal memory required for ε-deviation-rational control in arbitrary history processes. We formalize this as the minimum number of labels in a past-compression architecture that maps each attainable history to a finite memory state and decodes that state to a continuation controller. Our first result is an exact law: this retentive complexity equals the transversal number of the family of history-indexed ε-rational continuation sets. We then specialize to finite controller dictionaries with mixed continuations on a simplex and show that the low-regret sets are convex polytopes; at exact rationality ($\varepsilon = 0$), the problem collapses to a hypergraph law, in which the required memory is exactly the transversal number of the family of rational-support sets. We further prove universality by showing that every finite hypergraph arises as such a support family in a one-step fully observable instance. Finally, we prove that pairwise compatibility is insufficient in general—even the complete compatibility graph can conceal large memory gaps—while convexity yields a Helly-type local certificate for one-state memory. Together, these results recast bounded-agent memory in big worlds as a problem in combinatorial optimization and convex geometry rather than environment-state abstraction.

1 Introduction

Reinforcement-learning agents act in histories, not in states supplied by the analyst. In large or partially observed worlds, the agent must compress the past into a bounded internal memory and act from that compressed state thereafter. This raises a basic question: *how much memory is required to summarize the past while preserving rational future behavior?* We study that question in an arbitrary history-based interface, under a deviation-regret criterion, and ask for the minimum memory cardinality sufficient to remain ε-deviation-rational from every attainable history.

This question is naturally history-first. Elelimy et al. (2025) advocate a continual-RL viewpoint based on histories, deviations, and behavioral consequences rather than on optimality in a fixed MDP, while hindsight-rationality provides rich deviation classes for sequential decision problems (Morrill et al., 2021). Related recent work studies what capable agents must internally represent, including predictive/world-model necessity results (Richens et al., 2025; Nayebi, 2026), bounded-interaction quotients (Nixon, 2026), and decision-centric memory compression (Zou et al., 2026). Those results are closely related in spirit, but they do not provide an exact cardinal law for the minimum memory required for ε-deviation-rational control in arbitrary history processes.

Our formulation separates the *past* from the *future*. For each attainable history h, let K_h^ε be the set of continuation controllers that are ε-deviation-rational from h. Memory then becomes a compression problem: histories are mapped to finitely many labels, and each label must decode to a continuation controller lying in the appropriate low-regret set for every history assigned to that label. This formulation is entirely agent-side: the relevant objects are the continuation-controller class, the de-

viation class, and the number of internal memory labels, not any latent state-space or observability assumption.

Our contribution is an exact theoremic characterization of this memory problem. First, retentive complexity equals the transversal number of the family of history-indexed low-regret continuation sets. Second, in the finite-controller mixed-continuation model, exact rationality ($\varepsilon = 0$) collapses to a hypergraph law: the required memory is exactly the transversal number of the family of rational-support sets. Third, every finite hypergraph arises as such a support family in a one-step, fully observable instance. Fourth, pairwise compatibility is insufficient in general: even the complete compatibility graph can conceal large memory gaps. Finally, convexity yields a Helly-type local certificate for one-state memory.

These results yield a different conceptual picture of bounded-agent memory than the one suggested by quotient or metric viewpoints. The governing object is the higher-order intersection geometry of the family of low-regret continuation sets. In the finite-controller model this geometry becomes a hypergraph transversal problem; under convexity it admits local-to-global certificates; and, in general, it is strictly richer than anything visible from pairwise compatibility alone. Finally, the theory interfaces directly with ordinary finite-state agents: any deterministic recurrent agent whose current internal state determines a continuation controller must have at least as many reachable internal states as the retentive complexity defined here.

2 Model and Retentive Complexity

We work directly with histories and deviations, following the recent history-process and deviation-regret perspective in continual reinforcement learning and hindsight-rationality (Elelimy et al., 2025; Morrill et al., 2021). The formulation below is intentionally agent-side: the environment enters only through the attainable-history set and the continuation values induced after each history, while the complexity notion itself concerns only the family of continuation controllers available to the agent and the deviations against which those controllers are evaluated.

Definition 2.1 (Deviation-stable continuation system)**.** A *deviation-stable continuation system* is a quadruple

$$\mathfrak{M} = (\mathcal{H}, \mathcal{C}, \Phi, J),$$

where $\mathcal{H} \neq \varnothing$ is a set of attainable histories, $\mathcal{C} \neq \varnothing$ is a class of admissible continuation controllers, $\Phi \neq \varnothing$ is a family of self-maps $\phi : \mathcal{C} \to \mathcal{C}$, and $J : \mathcal{H} \times \mathcal{C} \to \mathbb{R}$ is a bounded continuation-value functional. We write ϕc for $\phi(c)$. The deviation regret of $c \in \mathcal{C}$ after $h \in \mathcal{H}$ is

$$\mathrm{Reg}_\Phi(h, c) := \sup_{\phi \in \Phi} \Big(J(h, \phi c) - J(h, c) \Big). \tag{2.1}$$

Definition 2.2 (ε-rational continuation set)**.** For $\varepsilon \geq 0$ and $h \in \mathcal{H}$, the ε-rational continuation set is

$$K_h^\varepsilon := \{c \in \mathcal{C} : \mathrm{Reg}_\Phi(h, c) \leq \varepsilon\}. \tag{2.2}$$

An element $c \in \mathcal{C}$ should be read as the entire closed-loop behavior the agent can execute *from the current moment onward* once its internal memory label has been fixed. Thus K_h^ε is the set of continuations that remain ε-deviation-rational from history h.

Definition 2.3 (Past-compression architecture)**.** For $m \in \mathbb{N}$, write $[m] := \{1, \ldots, m\}$. A *past-compression architecture of size* m is a pair

$$(\sigma, d), \qquad \sigma : \mathcal{H} \to [m], \quad d : [m] \to \mathcal{C}.$$

It is ε*-rational* if

$$d(\sigma(h)) \in K_h^\varepsilon \qquad \forall h \in \mathcal{H}. \tag{2.3}$$

Definition 2.4 (Retentive complexity)**.** The *retentive complexity* of $\mathfrak{M}$ at tolerance ε is

$$m_{\Phi}^{\mathrm{pc}}(\varepsilon) := \inf\{m \in \mathbb{N} : \exists\, \varepsilon\text{-rational past-compression architecture of size } m\}. \tag{2.4}$$

If no finite ε-rational architecture exists, we interpret $m_{\Phi}^{\mathrm{pc}}(\varepsilon) = \infty$.

The quantity $m_{\Phi}^{\mathrm{pc}}(\varepsilon)$ is the minimum number of memory labels needed to compress the *entire past* while still ensuring that, from every attainable history, at least one continuation chosen solely from the compressed label is ε-deviation-rational.

Proposition 2.5 (Finite-state compilation lower bound)**.** *Let Z be a finite internal state set, let $\zeta : \mathcal{H} \to Z$ be a deterministic history-to-state map, and let $d : Z \to \mathcal{C}$ assign a continuation controller to each internal state. If*

$$d(\zeta(h)) \in K_h^{\varepsilon} \qquad \forall h \in \mathcal{H}, \tag{2.5}$$

then

$$|Z| \geq m_{\Phi}^{\mathrm{pc}}(\varepsilon).$$

Proof. Let $Z_{\mathrm{reach}} := \zeta(\mathcal{H}) \subseteq Z$, and choose a bijection $b : Z_{\mathrm{reach}} \to [|Z_{\mathrm{reach}}|]$. Define

$$\sigma := b \circ \zeta, \qquad \tilde{d}(i) := d\big(b^{-1}(i)\big) \quad \text{for } i \in [|Z_{\mathrm{reach}}|].$$

Then, for every $h \in \mathcal{H}$,

$$\tilde{d}(\sigma(h)) = d(\zeta(h)) \in K_h^{\varepsilon}$$

by (2.5). Hence $(\sigma, \tilde{d})$ is an ε-rational past-compression architecture of size $|Z_{\mathrm{reach}}|$. By Definition 2.4,

$$m_{\Phi}^{\mathrm{pc}}(\varepsilon) \leq |Z_{\mathrm{reach}}| \leq |Z|,$$

which proves the claim. □

3 Exact Reduction to a Transversal Problem

The central structural fact of the paper is that retentive complexity is exactly a hitting-set problem over the family of low-regret continuation sets. This reduction is purely set-theoretic: it uses only the definitions of K_h^{ε} and $m_{\Phi}^{\mathrm{pc}}(\varepsilon)$, and therefore does not rely on any simplifying assumption about the environment or the deviation class.

For $\varepsilon \geq 0$, write

$$\mathcal{K}_{\varepsilon} := \{K_h^{\varepsilon} : h \in \mathcal{H}\}.$$

A set $P \subseteq \mathcal{C}$ is a *transversal* for $\mathcal{K}_{\varepsilon}$ if

$$P \cap K_h^{\varepsilon} \neq \varnothing \qquad \forall h \in \mathcal{H}.$$

Its minimum possible cardinality is the transversal number

$$\tau(\mathcal{K}_{\varepsilon}) := \inf\{m \in \mathbb{N} : \exists P \subseteq \mathcal{C},\ |P| \leq m,\ P \cap K_h^{\varepsilon} \neq \varnothing\ \forall h \in \mathcal{H}\}, \tag{3.1}$$

with the convention $\inf \varnothing = \infty$.

Theorem 3.1 (Exact transversal law)**.** *For every deviation-stable continuation system and every $\varepsilon \geq 0$,*

$$m_{\Phi}^{\mathrm{pc}}(\varepsilon) = \tau(\mathcal{K}_{\varepsilon}). \tag{3.2}$$

Proof. We show that, for each memory budget $m \in \mathbb{N}$, there exists an ε-rational past-compression architecture of size m if and only if there exists a transversal of $\mathcal{K}_{\varepsilon}$ of cardinality at most m.

First, suppose (σ, d) is an ε-rational past-compression architecture of size m, so

$$\sigma : \mathcal{H} \to [m], \qquad d : [m] \to \mathcal{C},$$

and

$$d(\sigma(h)) \in K_h^\varepsilon \qquad \forall h \in \mathcal{H}.$$

Let

$$P := d([m]) \subseteq \mathcal{C}.$$

Then $|P| \leq m$. Moreover, for every $h \in \mathcal{H}$,

$$d(\sigma(h)) \in P \cap K_h^\varepsilon,$$

so $P \cap K_h^\varepsilon \neq \varnothing$. Hence P is a transversal for $\mathcal{K}_\varepsilon$ of size at most m.

Conversely, suppose $P \subseteq \mathcal{C}$ is a transversal for $\mathcal{K}_\varepsilon$ with $|P| \leq m$. Enumerate its distinct elements as

$$P = \{c_1, \ldots, c_r\}, \qquad r \leq m.$$

Define $d : [m] \to \mathcal{C}$ by $d(i) = c_i$ for $1 \leq i \leq r$, and pad the remaining labels arbitrarily, say $d(i) = c_1$ for $r < i \leq m$. Since P intersects every K_h^ε, the set

$$I(h) := \{i \in [r] : c_i \in K_h^\varepsilon\}$$

is nonempty for every $h \in \mathcal{H}$. Define

$$\sigma(h) := \min I(h).$$

Then $\sigma : \mathcal{H} \to [m]$ is well defined, and by construction

$$d(\sigma(h)) = c_{\sigma(h)} \in K_h^\varepsilon \qquad \forall h \in \mathcal{H}.$$

Thus (σ, d) is an ε-rational past-compression architecture of size m.

We have proved the equivalence

$$\exists\ \varepsilon\text{-rational architecture of size } m \iff \exists \text{ transversal of } \mathcal{K}_\varepsilon \text{ of size } \leq m.$$

Taking the infimum over $m \in \mathbb{N}$ on both sides yields (3.2). □

Theorem 3.1 identifies the exact combinatorial object governing memory: not a quotient of histories, not a pairwise metric, and not a descriptive compression criterion, but the piercing number of the family of history-indexed low-regret continuation sets.

4 Finite Controller Dictionaries: Hypergraph Law and Universality

To expose the combinatorial content of Theorem 3.1, we specialize to a finite controller dictionary $\{e_1, \ldots, e_k\}$ and its simplex of mixed continuations

$$\Delta_k := \left\{ x \in \mathbb{R}_+^k : \sum_{i=1}^k x_i = 1 \right\}.$$

We assume that continuation values are affine under mixing, so that for every history $h \in \mathcal{H}$ and $x \in \Delta_k$,

$$J(h, x) = \sum_{i=1}^k x_i J(h, e_i).$$

Let

$$v(h) := (v_1(h), \dots, v_k(h)) \in \mathbb{R}^k, \qquad v_i(h) := J(h, e_i),$$

and take the deviation class to be the external deviations

$$\Phi_{\text{ext}} := \{\phi_i : i \in [k]\}, \qquad \phi_i(x) := e_i.$$

Thus, after any history, a deviation may replace the mixed continuation by any pure library controller.

Theorem 4.1 (Polyhedral low-regret sets)**.** *For every history $h \in \mathcal{H}$ and every $x \in \Delta_k$,*

$$\mathrm{Reg}_{\Phi_{\text{ext}}}(h, x) = \max_{i \in [k]} v_i(h) - \langle v(h), x \rangle. \tag{4.1}$$

Consequently, for every $\varepsilon \geq 0$,

$$K_h^\varepsilon = \left\{ x \in \Delta_k : \langle v(h), x \rangle \geq \max_{i \in [k]} v_i(h) - \varepsilon \right\}. \tag{4.2}$$

In particular, K_h^ε is the intersection of the simplex with a closed half-space, and hence a convex polytope.

Proof. By affineness of J under mixing,

$$J(h, x) = \sum_{i=1}^{k} x_i J(h, e_i) = \langle v(h), x \rangle.$$

Moreover, for each external deviation ϕ_i,

$$J(h, \phi_i x) = J(h, e_i) = v_i(h).$$

Substituting into Definition 2.1 gives

$$\mathrm{Reg}_{\Phi_{\text{ext}}}(h, x) = \sup_{i \in [k]} (v_i(h) - \langle v(h), x \rangle) = \max_{i \in [k]} v_i(h) - \langle v(h), x \rangle,$$

which proves (4.1). The description (4.2) then follows by Definition 2.2. □

At exact rationality, the geometry collapses to simplex faces. Define the maximizer set

$$S_h := \arg\max_{i \in [k]} v_i(h) \subseteq [k],$$

and for any nonempty $S \subseteq [k]$, write

$$\Delta(S) := \{x \in \Delta_k : \mathrm{supp}(x) \subseteq S\}.$$

Corollary 4.2 (Exact hypergraph law)**.** *For every $h \in \mathcal{H}$,*

$$K_h^0 = \Delta(S_h).$$

Hence, if

$$\mathcal{S} := \{S_h : h \in \mathcal{H}\},$$

then

$$m^{\text{pc}}_{\Phi_{\text{ext}}}(0) = \tau(\mathcal{S}), \tag{4.3}$$

where

$$\tau(\mathcal{S}) := \inf\{|T| : T \subseteq [k],\ T \cap S \neq \varnothing\ \forall S \in \mathcal{S}\}.$$

Proof. By Theorem 4.1,

$$K_h^0 = \left\{ x \in \Delta_k : \langle v(h), x \rangle = \max_{i \in [k]} v_i(h) \right\}.$$

A convex combination attains the maximum of the coordinates of $v(h)$ iff it places mass only on maximizing indices. Therefore $K_h^0 = \Delta(S_h)$.

Now apply Theorem 3.1 to the family $\{\Delta(S_h) : h \in \mathcal{H}\}$. If $T \subseteq [k]$ hits every support set in $\mathcal{S}$, then the vertex set $\{e_i : i \in T\}$ pierces every face $\Delta(S_h)$. Conversely, if $P \subseteq \Delta_k$ pierces every $\Delta(S_h)$, choose for each $x \in P$ an arbitrary index $i(x) \in \operatorname{supp}(x)$, and set

$$T := \{i(x) : x \in P\} \subseteq [k].$$

Whenever $x \in P \cap \Delta(S_h)$, one has $\operatorname{supp}(x) \subseteq S_h$, hence $i(x) \in S_h$; thus T hits every set in $\mathcal{S}$. The minimum piercing number of the faces $\Delta(S_h)$ is therefore exactly $\tau(\mathcal{S})$, proving (4.3). □

Corollary 4.2 identifies exact retentive complexity with a classical hypergraph transversal number. The next theorem shows that this is not merely an analogy: every finite hypergraph can arise exactly from a history-dependent decision problem.

Theorem 4.3 (Universality)**.** *For every finite hypergraph*

$$\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\},$$

there exists a one-step, fully observable history process with controller dictionary $\{e_1, \ldots, e_k\}$ *such that*

$$\{S_h : h \in \mathcal{H}\} = \mathcal{S}.$$

Equivalently, every finite hypergraph can be realized as the family of exact rational-support sets of some mixed-controller instance.

Proof sketch. For each hyperedge $S \in \mathcal{S}$, create a distinct observable history h_S. After observing h_S, the agent chooses an action $A \in [k]$ and receives reward

$$R = \mathbf{1}\{A \in S\},$$

after which the episode terminates. Then

$$J(h_S, e_i) = \mathbf{1}\{i \in S\},$$

so the maximizing indices are exactly those in S, i.e.,

$$S_{h_S} = \arg\max_{i \in [k]} J(h_S, e_i) = S.$$

Hence the family $\{S_h : h \in \mathcal{H}\}$ is exactly $\mathcal{S}$. The full proof in Supplement E gives the stronger statement that, for every $\varepsilon \geq 0$,

$$K_{h_S}^{\varepsilon} = \left\{ x \in \Delta_k : \sum_{i \in S} x_i \geq 1 - \varepsilon \right\}.$$

□

Combining Theorem 4.3 with Corollary 4.2, every finite hitting-set instance already appears as an exact retentive-complexity problem in a one-step history process. Thus the memory geometry of exact deviation-rational control is combinatorially universal even in the simplest fully observable setting.

5 Pairwise Compatibility is Not Enough

The exact hypergraph law identifies retentive complexity with a transversal number of the family $\mathcal{S} = \{S_h : h \in \mathcal{H}\}$. A natural question is whether this family can be reduced to pairwise intersection data alone.

For the mixed-controller model at exact rationality, define the *pairwise compatibility graph*

$$G_{\text{pair}} = (\mathcal{H}, E_{\text{pair}})$$

by declaring distinct histories $h, h' \in \mathcal{H}$ adjacent iff

$$K_h^0 \cap K_{h'}^0 \neq \varnothing, \tag{5.1}$$

which, by Corollary 4.2, is equivalent to $S_h \cap S_{h'} \neq \varnothing$. Thus G_{pair} retains only pairwise intersections and forgets all higher-order intersection structure.

Theorem 5.1 (Projective-plane separation)**.** *Let $\Pi_q = (P_q, \mathcal{L}_q, \in)$ be a finite projective plane of order q, and set $n := q^2 + q + 1$. Then there exist two one-step, fully observable mixed-controller instances, each with controller dictionary $\{e_1, \dots, e_n\}$ and external deviation class Φ_{ext}, such that both have pairwise compatibility graph K_n, but their exact retentive complexities are 1 and $q + 1$, respectively.*

The only numerical input is the following classical counting fact.

Lemma 5.2. *If $\mathscr{L}_q = \{L \subseteq P_q : L \in \mathcal{L}_q\}$ is the line hypergraph of a projective plane of order q, then*

$$\tau(\mathscr{L}_q) = q + 1.$$

Proof sketch of Theorem 5.1. Construct two instances on the same controller simplex Δ_n. In Instance A, reward is always $\mathbf{1}\{A = 1\}$, so $K_{h_j}^0 = \{e_1\}$ for all histories and the pairwise graph is K_n; hence $m_{\Phi_{\text{ext}}}^{\text{pc}}(0) = 1$. In Instance B, histories are indexed by lines $L \in \mathcal{L}_q$, reward at history h_L is $\mathbf{1}\{A \in L\}$, and Theorem 4.3 yields $K_{h_L}^0 = \Delta(L)$. Distinct lines intersect, so the pairwise graph is again K_n. But Corollary 4.2 gives $m_{\Phi_{\text{ext}}}^{\text{pc}}(0) = \tau(\mathscr{L}_q) = q + 1$. □

Proof sketch of Lemma 5.2. Any line $L^\star \in \mathcal{L}_q$ contains $q + 1$ points and intersects every other line, so $\tau(\mathscr{L}_q) \leq q + 1$. Conversely, any set of at most q points lies on at most $q(q + 1) = q^2 + q$ lines, while the plane contains $q^2 + q + 1$ lines in total. Hence some line is missed, so no q-point set can be a transversal. □

Corollary 5.3 (No pairwise characterization)**.** *No function of the pairwise compatibility graph alone can determine $m_{\Phi_{\text{ext}}}^{\text{pc}}(0)$ for all finite-controller mixed-continuation instances.*

Proof. Theorem 5.1 gives two instances with the same graph $G_{\text{pair}} = K_n$ but different exact retentive complexities. □

In particular, exact deviation-rational memory cannot in general be reduced to any pairwise graph-theoretic abstraction.

Thus even the complete compatibility graph can be maximally uninformative: what governs memory is the higher-order intersection geometry of the family $\{K_h^0 : h \in \mathcal{H}\}$, or equivalently of the support hypergraph $\mathcal{S}$. This is precisely the structural gap between pairwise compatibility and exact memory.

6 Local Convex Certificate and Discussion

The preceding sections show that exact retentive complexity is governed by a higher-order set system and, in general, cannot be read from pairwise compatibility alone. In the mixed-controller model, however, the geometry of the low-regret sets yields a sharp local certificate for the special case of one-state memory.

Corollary 6.1 (Helly certificate for one-state memory)**.** *Consider the finite-controller mixed-continuation model of Section 4, with controller simplex Δ_k and external deviations Φ_{ext}. For every $\varepsilon \geq 0$, the following are equivalent:*

1. $m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(\varepsilon) = 1$;
2. $\bigcap_{h \in \mathcal{H}} K^{\varepsilon}_h \neq \varnothing$;
3. *every subfamily of $\{K^{\varepsilon}_h : h \in \mathcal{H}\}$ of cardinality at most k has nonempty intersection.*

Proof. The equivalence of (1) and (2) is immediate from Theorem 3.1: a one-state architecture exists if and only if the family $\{K^{\varepsilon}_h : h \in \mathcal{H}\}$ admits a transversal of size 1, and this is equivalent to the existence of a single point $x \in \Delta_k$ belonging to every K^{ε}_h.

The implication (2) $\Rightarrow$ (3) is trivial. It remains to prove (3) $\Rightarrow$ (2). By Theorem 4.1, each K^{ε}_h is a convex polytope in Δ_k. Since Δ_k lies in the affine hyperplane

$$\mathrm{aff}(\Delta_k) = \left\{ x \in \mathbb{R}^k : \sum_{i=1}^{k} x_i = 1 \right\},$$

the relevant affine dimension is $k-1$. Hence, by Helly's theorem in affine dimension $k-1$ (Matoušek, 2002), every finite subfamily $\mathcal{F} \subseteq \{K^{\varepsilon}_h : h \in \mathcal{H}\}$ has nonempty intersection whenever every k or fewer members of $\mathcal{F}$ intersect. Assumption (3) therefore implies that *every finite* subfamily has nonempty intersection.

Finally, each K^{ε}_h is closed and contained in the compact set Δ_k. Since the family has the finite intersection property, compactness of Δ_k implies

$$\bigcap_{h \in \mathcal{H}} K^{\varepsilon}_h \neq \varnothing.$$

Thus (2) holds. □

Corollary 6.1 reveals a striking contrast with Section 5. For general memory budgets, exact retentive complexity is governed by the full higher-order intersection geometry of the family $\{K^{\varepsilon}_h : h \in \mathcal{H}\}$, and pairwise data can be maximally uninformative. For the special case of one-state memory, by contrast, convexity collapses the problem to a local certificate of order k, where k is the size of the controller dictionary rather than any measure of environmental complexity. This threshold is exactly the Helly number of the affine hull of Δ_k, whose dimension is $k-1$. In this sense, one-state certification depends only on controller complexity, whereas general retentive complexity remains a genuinely higher-order set-system invariant.

Taken together, Theorems 3.1 and 5.1, Corollary 4.2, and Corollary 6.1 yield a unified picture: deviation-rational memory is neither a quotient of histories nor a pairwise metric invariant. Its exact complexity is a piercing problem over history-indexed low-regret continuation sets; in the finite-controller model this becomes a hypergraph transversal law, while convexity supplies additional local-to-global structure only in special regimes. This places bounded-agent memory in arbitrary worlds squarely at the interface of reinforcement learning, combinatorial optimization, and convex geometry.

References

Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, Heidelberg, 3 edition, 2006.

Charles Castaing and Michel Valadier. *Convex Analysis and Measurable Multifunctions*, volume 580 of *Lecture Notes in Mathematics*. Springer, Berlin, Heidelberg, 1977.

Esraa Elelimy, David Szepesvari, Martha White, and Michael Bowling. Rethinking the foundations for continual reinforcement learning. *Reinforcement Learning Journal*, 6:2174–2194, 2025.

Jiří Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer, 2002.

Dustin Morrill, Ryan D'Orazio, Marc Lanctot, James R Wright, Michael Bowling, and Amy R Greenwald. Efficient deviation types and learning for hindsight rationality in extensive-form games. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7818–7828. PMLR, 2021.

Aran Nayebi. What capable agents must know: Selection theorems for robust decision-making under uncertainty. *arXiv preprint arXiv:2603.02491*, 2026. URL https://arxiv.org/abs/2603.02491.

Anthony T. Nixon. The myhill–nerode theorem for bounded interaction: Canonical abstractions via agent-bounded indistinguishability. *arXiv preprint arXiv:2603.21399*, 2026. URL https://arxiv.org/abs/2603.21399.

Jonathan Richens, Tom Everitt, and David Abel. General agents need world models. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 51659–51687. PMLR, 2025. URL https://proceedings.mlr.press/v267/richens25a.html.

Mingxi Zou, Zhihan Guo, Langzhang Liang, Zhuo Wang, Qifan Wang, Qingsong Wen, Irwin King, Lizhen Qu, and Zenglin Xu. Remember the decision, not the description: A rate-distortion framework for agent memory. *arXiv preprint arXiv:2605.10870*, 2026. URL https://arxiv.org/abs/2605.10870.

Supplementary Materials

The following content was not necessarily subject to peer review.

A Formal Interface and Measurable Refinements

This section serves two purposes. First, it records a concrete interface from a standard history-based reinforcement-learning model to the abstract continuation-system formalism used in the main text. Second, it gives a measurable refinement of Theorem 3.1, showing that under standard regularity conditions the finite-memory compression map can be chosen measurably without changing the exact memory law.

A.1 From measurable history processes to continuation systems

We formalize a general history-based decision problem in the spirit of the history-process viewpoint used in recent continual RL work (Elelimy et al., 2025).

Definition A.1 (Discounted measurable history process)**.** A *discounted measurable history process* consists of the following data:

1. a standard Borel action space $\mathcal{A}$;
2. a standard Borel signal space X, whose elements may encode observations, rewards, terminal indicators, or any other percept available to the agent;
3. the set of pre-action histories
$$\mathcal{H} := \bigsqcup_{t\geq 0} (\mathsf{X}\times\mathcal{A})^t \times \mathsf{X};$$
4. a stochastic kernel
$$Q(\cdot \mid h, a)$$
from $\mathcal{H}\times\mathcal{A}$ to X, interpreted as the law of the next signal after history h and action a;
5. a bounded measurable reward extractor
$$\rho : \mathsf{X} \to [-R, R]$$
for some finite $R > 0$;
6. a discount factor $\gamma \in [0, 1)$;
7. a nonempty class Π of admissible continuation controllers, where each $\pi \in \Pi$ is a universally measurable stochastic kernel from $\mathcal{H}$ to $\mathcal{A}$;
8. a nonempty deviation class Φ, where each $\phi \in \Phi$ is a self-map of Π.

For $h \in \mathcal{H}$ and $\pi \in \Pi$, let
$$(H_t, A_t, X_{t+1})_{t\geq 0}$$
denote the controlled process started from initial history $H_0 = h$, with
$$A_t \sim \pi(\cdot \mid H_t), \qquad X_{t+1} \sim Q(\cdot \mid H_t, A_t),$$
and history recursion
$$H_{t+1} := H_t \oplus (A_t, X_{t+1}),$$
where $\oplus$ denotes concatenation of one action-signal pair. Standard kernel-composition arguments yield a unique law for this process.

Proposition A.2 (Induced continuation system)**.** *Let* $(\mathcal{A}, \mathsf{X}, \mathcal{H}, Q, \rho, \gamma, \Pi, \Phi)$ *be a discounted measurable history process. Define*

$$J(h, \pi) := \mathbb{E}_h^\pi \left[\sum_{t=0}^{\infty} \gamma^t \rho(X_{t+1}) \right]. \tag{A.1}$$

Then J *is well defined, bounded by* $R/(1-\gamma)$ *in absolute value, and*

$$(\mathcal{H}, \Pi, \Phi, J)$$

is a deviation-stable continuation system in the sense of Definition 2.1.

Proof. Fix $h \in \mathcal{H}$ and $\pi \in \Pi$. Since $|\rho(X_{t+1})| \le R$ almost surely for every t,

$$\sum_{t=0}^{\infty} \gamma^t |\rho(X_{t+1})| \le \sum_{t=0}^{\infty} \gamma^t R = \frac{R}{1-\gamma} \qquad \text{a.s.}$$

Hence the discounted sum in (A.1) converges absolutely almost surely and is integrable. Therefore $J(h, \pi)$ is well defined and satisfies

$$|J(h, \pi)| \le \frac{R}{1-\gamma}.$$

All ingredients of Definition 2.1 are thus satisfied, with continuation class $\mathcal{C} = \Pi$ and value functional J given by (A.1). □

Remark A.3 (Why the abstraction is lossless for our purposes). The memory theory developed in the main text depends only on the attainable history set, the continuation-controller class, the deviation class, and the continuation values $J(h, c)$. Once these objects are fixed, no additional representation of the environment is needed. In particular, none of the main results requires that histories arise from a Markov state process, from a finite latent model, or even from an explicitly parameterized environment.

A.2 Measurable low-regret correspondences

The main text deliberately states the transversal law in a purely set-theoretic form. For measurable decision problems, one can strengthen it by requiring the history-to-memory encoder to be measurable.

We isolate a convenient regularity hypothesis. It is stronger than necessary, but it cleanly covers all finite-controller constructions used in the paper.

Assumption A.4 (Regular measurable setting)**.** Fix $\varepsilon \ge 0$. In addition to the hypotheses of Definition 2.1, assume:

1. $\mathcal{H}$ is a standard Borel space;
2. $\mathcal{C}$ is a compact metric space;
3. the deviation class is countable,
$$\Phi = \{\phi_1, \phi_2, \ldots\},$$
and each $\phi_n : \mathcal{C} \to \mathcal{C}$ is continuous;
4. the continuation-value functional $J : \mathcal{H} \times \mathcal{C} \to \mathbb{R}$ is bounded and Carathéodory, meaning:
   (a) for each fixed $c \in \mathcal{C}$, the map $h \mapsto J(h, c)$ is Borel measurable;
   (b) for each fixed $h \in \mathcal{H}$, the map $c \mapsto J(h, c)$ is continuous.

Proposition A.5 (Borel graph and closed values)**.** *Under Assumption A.4, the regret map*

$$(h, c) \longmapsto \mathrm{Reg}_\Phi(h, c)$$

is Borel measurable on $\mathcal{H} \times \mathcal{C}$. Moreover, for each fixed $h \in \mathcal{H}$, the function $c \mapsto \mathrm{Reg}_\Phi(h, c)$ is lower semicontinuous. Consequently, for every $h \in \mathcal{H}$, the set

$$K_h^\varepsilon = \{c \in \mathcal{C} : \mathrm{Reg}_\Phi(h, c) \leq \varepsilon\}$$

is closed, and the graph

$$\mathrm{Gr}(K^\varepsilon) := \{(h, c) \in \mathcal{H} \times \mathcal{C} : c \in K_h^\varepsilon\}$$

is a Borel subset of $\mathcal{H} \times \mathcal{C}$.

Proof. For each $n \in \mathbb{N}$, define

$$f_n(h, c) := J(h, \phi_n(c)) - J(h, c).$$

Because ϕ_n is continuous and J is Carathéodory, each $f_n : \mathcal{H} \times \mathcal{C} \to \mathbb{R}$ is Borel measurable. Indeed, for fixed c,

$$h \longmapsto J(h, \phi_n(c)) - J(h, c)$$

is Borel measurable, and for fixed h,

$$c \longmapsto J(h, \phi_n(c)) - J(h, c)$$

is continuous. Therefore f_n is jointly Borel measurable.

By definition of deviation regret,

$$\mathrm{Reg}_\Phi(h, c) = \sup_{n \geq 1} f_n(h, c).$$

A countable supremum of Borel functions is Borel, so Reg_Φ is Borel measurable on $\mathcal{H} \times \mathcal{C}$. For each fixed h, the map $c \mapsto f_n(h, c)$ is continuous, hence

$$c \longmapsto \sup_{n \geq 1} f_n(h, c)$$

is lower semicontinuous as a pointwise supremum of continuous functions. This proves the lower semicontinuity of $c \mapsto \mathrm{Reg}_\Phi(h, c)$.

Now fix $h \in \mathcal{H}$. Since K_h^ε is the ε-sublevel set of the lower semicontinuous function $c \mapsto \mathrm{Reg}_\Phi(h, c)$, it is closed in $\mathcal{C}$. Finally,

$$\mathrm{Gr}(K^\varepsilon) = \{(h, c) \in \mathcal{H} \times \mathcal{C} : \mathrm{Reg}_\Phi(h, c) \leq \varepsilon\} = \bigcap_{n \geq 1} \{(h, c) : f_n(h, c) \leq \varepsilon\},$$

which is Borel because each f_n is Borel. □

Corollary A.6 (Borel membership sections)**.** *Under Assumption A.4, for every fixed $c \in \mathcal{C}$, the membership set*

$$A_c^\varepsilon := \{h \in \mathcal{H} : c \in K_h^\varepsilon\}$$

is Borel measurable.

Proof. For fixed $c \in \mathcal{C}$,

$$A_c^\varepsilon = \{h \in \mathcal{H} : (h, c) \in \mathrm{Gr}(K^\varepsilon)\}.$$

Since $\mathrm{Gr}(K^\varepsilon)$ is Borel in $\mathcal{H} \times \mathcal{C}$, its section at the fixed point c is Borel in $\mathcal{H}$. □

A.3 A measurable refinement of the transversal law

We now define the measurable analogue of retentive complexity.

Definition A.7 (Measurable past-compression complexity)**.** Assume $\mathcal{H}$ is equipped with a σ-algebra. For $\varepsilon \geq 0$, the *measurable retentive complexity* is

$$m_{\Phi}^{\mathrm{pc,meas}}(\varepsilon) := \inf \left\{ m \in \mathbb{N} : \begin{array}{l} \exists\, \varepsilon\text{-rational past-compression architecture } (\sigma, d) \\ \text{of size } m \text{ with measurable } \sigma \end{array} \right\},$$

with the convention that the infimum of the empty set is ∞.

Theorem A.8 (Measurable finite-transversal law)**.** *Under Assumption A.4, for every $\varepsilon \geq 0$,*

$$m_{\Phi}^{\mathrm{pc,meas}}(\varepsilon) = m_{\Phi}^{\mathrm{pc}}(\varepsilon) = \tau(\mathcal{K}_{\varepsilon}). \tag{A.2}$$

Proof. Any measurable past-compression architecture is, in particular, a past-compression architecture, so

$$m_{\Phi}^{\mathrm{pc}}(\varepsilon) \leq m_{\Phi}^{\mathrm{pc,meas}}(\varepsilon).$$

By Theorem 3.1, it therefore suffices to prove

$$m_{\Phi}^{\mathrm{pc,meas}}(\varepsilon) \leq \tau(\mathcal{K}_{\varepsilon}).$$

If $\tau(\mathcal{K}_{\varepsilon}) = \infty$, there is nothing to prove. Assume $\tau(\mathcal{K}_{\varepsilon}) = m < \infty$. Then there exists a transversal

$$P = \{c_1, \ldots, c_m\} \subseteq \mathcal{C}$$

for $\mathcal{K}_{\varepsilon}$. Define $d : [m] \to \mathcal{C}$ by $d(i) = c_i$. For each $i \in [m]$, let

$$A_i := \{h \in \mathcal{H} : c_i \in K_h^{\varepsilon}\}.$$

By Corollary A.6, each A_i is Borel. Since P is a transversal, for every $h \in \mathcal{H}$,

$$\{i \in [m] : h \in A_i\} \neq \varnothing.$$

Define

$$\sigma(h) := \min\{i \in [m] : h \in A_i\}. \tag{A.3}$$

This is well defined. Moreover, for each $i \in [m]$,

$$\{h \in \mathcal{H} : \sigma(h) = i\} = A_i \cap \bigcap_{j<i} A_j^c,$$

which is Borel. Hence $\sigma : \mathcal{H} \to [m]$ is measurable.

Finally, by construction of σ,

$$c_{\sigma(h)} \in K_h^{\varepsilon} \qquad \forall h \in \mathcal{H}.$$

Equivalently,

$$d(\sigma(h)) \in K_h^{\varepsilon} \qquad \forall h \in \mathcal{H}.$$

Thus (σ, d) is a measurable ε-rational past-compression architecture of size m. Therefore

$$m_{\Phi}^{\mathrm{pc,meas}}(\varepsilon) \leq m = \tau(\mathcal{K}_{\varepsilon}).$$

Combining this with Theorem 3.1 yields (A.2). □

Remark A.9 (Why finite transversals avoid measurable-selection subtleties). The proof of Theorem A.8 does not invoke any general measurable-selection theorem. Once a finite transversal $P = \{c_1, \ldots, c_m\}$ is fixed, the encoder is obtained canonically by the minimum-index rule (A.3). Measurability then reduces to the Borel measurability of the membership sections A_i.

A.4 A selection corollary

Although not needed for the finite-memory results in the main text, it is worth recording the corresponding one-selector statement. Standard measurable-selection theory implies that, under suitable hypotheses, one can choose a measurable continuation $s(h) \in K_h^\varepsilon$ whenever every low-regret set is nonempty.

Remark A.10 (Classical measurable-selection consequence). Suppose, in addition to Assumption A.4, that the correspondence $h \mapsto K_h^\varepsilon$ is weakly measurable and has nonempty closed values. Then the Kuratowski–Ryll-Nardzewski theorem yields a Borel selector

$$s_\varepsilon : \mathcal{H} \to \mathcal{C} \qquad \text{such that} \qquad s_\varepsilon(h) \in K_h^\varepsilon \ \ \forall h \in \mathcal{H};$$

see, for example, Aliprantis & Border (2006, Theorem 18.13) or Castaing & Valadier (1977, Chapter III). We do not use this result in the main text, because Theorem A.8 already gives the finite-memory statement needed in the paper.

Remark A.11 (Coverage of the main constructions). All finite-controller constructions in Sections 4 and 5 satisfy Assumption A.4. Indeed, $\mathcal{C} = \Delta_k$ is compact metrizable, the external deviation class is finite and continuous, and $J(h, \cdot)$ is affine (hence continuous) for each fixed history h. Thus no additional measurability burden is hidden by the abstract formulation used in the main paper.

B Finite-State Realizations and Proof of Proposition 2.5

This section proves the lower bound that links our retentive-complexity notion to ordinary bounded-memory agents. The argument is entirely internal to the framework: it uses only the definitions of the low-regret continuation sets K_h^ε, the past-compression architectures, and the induced complexity $m_\Phi^{\mathrm{pc}}(\varepsilon)$. In particular, the proof does *not* assume any Markov structure, latent-state enumerability, observability condition, or simplicity property of the environment.

We begin by isolating the exact abstraction of a finite-memory agent that is needed for the proposition.

Definition B.1 (Finite-state realization)**.** Fix a deviation-stable continuation system

$$\mathfrak{M} = (\mathcal{H}, \mathcal{C}, \Phi, J)$$

in the sense of Definition 2.1. A *finite-state realization* of $\mathfrak{M}$ is a triple

$$(Z, \zeta, d),$$

where:

1. Z is a finite set of internal states;
2. $\zeta : \mathcal{H} \to Z$ is a deterministic history-to-state map;
3. $d : Z \to \mathcal{C}$ is a decoder assigning to each internal state an admissible continuation controller.

We interpret $\zeta(h)$ as the internal state reached by the agent after history h, and $d(z)$ as the continuation controller induced by starting the agent from internal state z. The realization is said to be *ε-deviation-rational* if

$$\mathrm{Reg}_\Phi(h,\, d(\zeta(h))) \le \varepsilon \qquad \forall h \in \mathcal{H}. \tag{B.1}$$

The next proposition is stronger than Proposition 2.5 in the main text because it lower bounds the number of *reachable* internal states rather than the ambient cardinality of the declared state space.

Proposition B.2 (Reachable-state strengthening)**.** *Let (Z, ζ, d) be a finite-state realization of $\mathfrak{M}$ that is ε-deviation-rational in the sense of* (B.1)*. Define the reachable set*

$$Z_{\mathrm{reach}} := \zeta(\mathcal{H}) \subseteq Z.$$

Then

$$|Z_{\mathrm{reach}}| \geq m_{\Phi}^{\mathrm{pc}}(\varepsilon). \tag{B.2}$$

Consequently,

$$|Z| \geq m_{\Phi}^{\mathrm{pc}}(\varepsilon). \tag{B.3}$$

Proof. Fix an ε-deviation-rational finite-state realization (Z, ζ, d), and let

$$m := |Z_{\mathrm{reach}}|.$$

Since Z_{reach} is finite, there exists a bijection

$$b : Z_{\mathrm{reach}} \to [m].$$

Using this bijection, define

$$\sigma := b \circ \zeta : \mathcal{H} \to [m]$$

and

$$\tilde{d} : [m] \to \mathcal{C}, \qquad \tilde{d}(i) := d\big(b^{-1}(i)\big).$$

We claim that $(\sigma, \tilde{d})$ is an ε-rational past-compression architecture of size m in the sense of Definition 2.3.

To verify the claim, fix any history $h \in \mathcal{H}$. Because the realization is ε-deviation-rational, (B.1) yields

$$\mathrm{Reg}_{\Phi}(h,\, d(\zeta(h))) \leq \varepsilon.$$

By Definition 2.2, this is equivalent to the membership statement

$$d(\zeta(h)) \in K_h^{\varepsilon}.$$

On the other hand, since b is a bijection on Z_{reach} and $\zeta(h) \in Z_{\mathrm{reach}}$, we have

$$\tilde{d}(\sigma(h)) = \tilde{d}(b(\zeta(h))) = d\big(b^{-1}(b(\zeta(h)))\big) = d(\zeta(h)).$$

Combining the two displays gives

$$\tilde{d}(\sigma(h)) \in K_h^{\varepsilon} \qquad \forall h \in \mathcal{H}. \tag{B.4}$$

Thus $(\sigma, \tilde{d})$ satisfies the defining feasibility condition for an ε-rational past-compression architecture.

By Definition 2.4 of $m_{\Phi}^{\mathrm{pc}}(\varepsilon)$, the existence of an ε-rational past-compression architecture of size m implies

$$m_{\Phi}^{\mathrm{pc}}(\varepsilon) \leq m = |Z_{\mathrm{reach}}|.$$

This proves (B.2). Since $Z_{\mathrm{reach}} \subseteq Z$, we immediately obtain

$$m_{\Phi}^{\mathrm{pc}}(\varepsilon) \leq |Z_{\mathrm{reach}}| \leq |Z|,$$

which is (B.3). □

Proof of Proposition 2.5. Proposition 2.5 is the weaker ambient-cardinality statement (B.3), which follows immediately from Proposition B.2. □

Remark B.3 (Why this proposition is the correct bridge). The proof uses only the following two facts: first, every history h is mapped to a current internal state $\zeta(h)$; second, every internal state z determines an admissible continuation controller $d(z)$. No further structure is needed. In particular, the proposition applies verbatim to any deterministic recurrent architecture whose hidden state is a deterministic function of the realized history and whose future closed-loop behavior is completely determined by the current hidden state.

The proposition is intentionally stated for deterministic history-to-state realizations. If the internal state after a fixed history were itself randomized, then the natural comparison object would be a *distributional* or *mixed* version of past-compression complexity. That richer formalism is orthogonal to the present paper and is not needed for any of our main results.

C Full Proof of Theorem 3.1

This section proves the exact reduction of retentive complexity to a transversal problem. The proof is purely set-theoretic and requires no topological, measurable, or compactness assumptions. In particular, it does not use any simplifying assumption on the environment, such as latent-state finiteness, Markovity, observability, or determinism.

Throughout this section, we write

$$\mathcal{K}_\varepsilon := \{K_h^\varepsilon : h \in \mathcal{H}\},$$

where K_h^ε is the ε-rational continuation set from Definition 2.2. Since $\mathcal{H}$ is the set of attainable histories, it is nonempty in every sequential decision problem of interest. If one wishes to allow the degenerate case $\mathcal{H} = \varnothing$, all statements below remain valid after the harmless convention that both sides of Theorem 3.1 are interpreted as 0.

We begin by making the hitting-set notion explicit.

Definition C.1 (Transversal number)**.** Let $\mathcal{F}$ be a family of subsets of $\mathcal{C}$. A set $P \subseteq \mathcal{C}$ is a *transversal* (or *piercing set*) for $\mathcal{F}$ if

$$P \cap F \neq \varnothing \qquad \forall F \in \mathcal{F}.$$

The *transversal number* of $\mathcal{F}$ is

$$\tau(\mathcal{F}) := \inf\{m \in \mathbb{N} : \exists P \subseteq \mathcal{C} \text{ with } |P| \leq m \text{ and } P \cap F \neq \varnothing \ \forall F \in \mathcal{F}\},$$

with the convention $\inf \varnothing = \infty$.

Remark C.2 (Budgeted versus exact size). A past-compression architecture of size m is, by Definition 2.3, a pair (σ, d) with $\sigma : \mathcal{H} \to [m]$ and $d : [m] \to \mathcal{C}$. Unused labels are permitted. Consequently, for each fixed budget m, the existence of an ε-rational architecture of size m is equivalent to the existence of one using at most m labels. We use this monotonicity implicitly below.

The heart of the argument is the following budget-by-budget equivalence.

Proposition C.3 (Budget equivalence between memory and transversals)**.** *Fix $m \in \mathbb{N}$. The following are equivalent.*

1. *There exists an ε-rational past-compression architecture of size m.*
2. *There exists a transversal $P \subseteq \mathcal{C}$ for $\mathcal{K}_\varepsilon$ with $|P| \leq m$.*

Proof. We prove the two implications separately.

(1) $\Rightarrow$ (2). Assume there exists an ε-rational past-compression architecture (σ, d) of size m, so

$$\sigma : \mathcal{H} \to [m], \qquad d : [m] \to \mathcal{C},$$

and

$$d(\sigma(h)) \in K_h^\varepsilon \qquad \forall h \in \mathcal{H}. \tag{C.1}$$

Define

$$P := d([m]) \subseteq \mathcal{C}.$$

Then $|P| \leq m$. Fix any history $h \in \mathcal{H}$. By (C.1),

$$d(\sigma(h)) \in K_h^\varepsilon.$$

Since $d(\sigma(h)) \in P$ by construction, we obtain

$$P \cap K_h^\varepsilon \neq \varnothing.$$

Because h was arbitrary, P is a transversal for $\mathcal{K}_\varepsilon$ of cardinality at most m.

(2) ⇒ (1). Assume there exists a transversal $P \subseteq \mathcal{C}$ for $\mathcal{K}_\varepsilon$ with $|P| \le m$. Since $\mathcal{H} \neq \varnothing$ and P intersects every K_h^ε, the set P is nonempty. Write

$$P = \{c_1, \dots, c_r\} \qquad \text{with} \qquad 1 \le r \le m,$$

where the c_i are pairwise distinct.

Define a decoder $d : [m] \to \mathcal{C}$ by

$$d(i) := \begin{cases} c_i, & 1 \le i \le r, \\ c_1, & r < i \le m. \end{cases} \tag{C.2}$$

Thus the first r labels encode the transversal points, and the remaining labels (if any) are padded arbitrarily using c_1.

Now fix $h \in \mathcal{H}$. Since P is a transversal for $\mathcal{K}_\varepsilon$, we have

$$P \cap K_h^\varepsilon \neq \varnothing.$$

Equivalently,

$$I(h) := \{i \in [r] : c_i \in K_h^\varepsilon\}$$

is nonempty. Define

$$\sigma(h) := \min I(h). \tag{C.3}$$

This is well defined because $I(h) \neq \varnothing$ and $I(h) \subseteq [r] \subseteq [m]$. By the definition of $I(h)$,

$$c_{\sigma(h)} \in K_h^\varepsilon.$$

Using (C.2), we have $d(\sigma(h)) = c_{\sigma(h)}$, hence

$$d(\sigma(h)) \in K_h^\varepsilon \qquad \forall h \in \mathcal{H}.$$

Therefore (σ, d) is an ε-rational past-compression architecture of size m. □

We now derive the main theorem.

Proof of Theorem 3.1. Recall that

$$m_\Phi^{\mathrm{pc}}(\varepsilon) = \inf\{m \in \mathbb{N} : \exists\, \varepsilon\text{-rational past-compression architecture of size } m\},$$

with the convention that the infimum of the empty set is ∞; see Definition 2.4. By Proposition C.3, for each $m \in \mathbb{N}$,

$$\exists\, \varepsilon\text{-rational architecture of size } m \quad \Longleftrightarrow \quad \exists\, P \subseteq \mathcal{C},\ |P| \le m,\ P \cap K_h^\varepsilon \neq \varnothing\ \forall h \in \mathcal{H}.$$

Therefore the set of feasible memory budgets coincides exactly with the set of feasible transversal budgets:

$$\{m \in \mathbb{N} : \exists\, \varepsilon\text{-rational architecture of size } m\} = \left\{ m \in \mathbb{N} : \begin{array}{l} \exists\, P \subseteq \mathcal{C},\ |P| \le m, \\ P \cap K_h^\varepsilon \neq \varnothing\ \forall h \in \mathcal{H} \end{array} \right\}.$$

Taking infima of these two identical subsets of $\mathbb{N}$ yields

$$m_\Phi^{\mathrm{pc}}(\varepsilon) = \tau(\mathcal{K}_\varepsilon).$$

This is exactly the claimed identity

$$m_\Phi^{\mathrm{pc}}(\varepsilon) = \tau(\{K_h^\varepsilon : h \in \mathcal{H}\}).$$

□

Remark C.4 (No choice-theoretic subtlety). The reverse implication in Proposition C.3 does not invoke any form of the axiom of choice. Once a finite transversal $P = \{c_1, \dots, c_r\}$ is enumerated, the encoder σ is defined canonically by the minimum-index rule (C.3). This finitary feature is important for later measurable refinements: whenever the membership events $\{h : c_i \in K_h^\varepsilon\}$ carry additional structure, the minimum-index construction inherits that structure automatically.

Remark C.5 (What the theorem does *not* use). Theorem 3.1 is purely structural. It does not rely on convexity of the continuation class, finiteness of the history set, compactness of $\mathcal{C}$, or any algebraic property of the deviation class Φ. All such structure becomes relevant only when we attempt to compute, bound, or geometrically analyze the transversal number in subsequent sections.

D Full Proof of Theorem 4.1 and Corollary 4.2

This section gives the full derivation of the simplex geometry used in Section 4. Throughout, we fix:

- a controller dictionary $\{e_1, \dots, e_k\}$,
- the mixed-continuation class

$$\Delta_k := \left\{ x \in \mathbb{R}_+^k : \sum_{i=1}^k x_i = 1 \right\},$$

- the external deviation class

$$\Phi_{\text{ext}} := \{\phi_i : i \in [k]\}, \qquad \phi_i(x) := e_i,$$

- and a continuation-value functional that is affine under mixing:

$$J(h, x) = \sum_{i=1}^k x_i J(h, e_i) \qquad \forall h \in \mathcal{H},\ \forall x \in \Delta_k.$$

For each history $h \in \mathcal{H}$, write

$$v(h) := (v_1(h), \dots, v_k(h)) \in \mathbb{R}^k, \qquad v_i(h) := J(h, e_i).$$

We begin with a simple geometric lemma that will be used repeatedly.

Lemma D.1 (Simplex maximization identity)**.** *Let $v \in \mathbb{R}^k$, and let*

$$M(v) := \max_{i \in [k]} v_i, \qquad A(v) := \arg\max_{i \in [k]} v_i.$$

Then, for every $x \in \Delta_k$,

$$\langle v, x \rangle \le M(v). \tag{D.1}$$

Moreover, equality holds if and only if

$$\operatorname{supp}(x) \subseteq A(v). \tag{D.2}$$

Proof. Fix $x \in \Delta_k$. Since $x_i \ge 0$ for all i and $\sum_{i=1}^k x_i = 1$,

$$\langle v, x \rangle = \sum_{i=1}^k x_i v_i \le \sum_{i=1}^k x_i M(v) = M(v),$$

proving (D.1).

For the equality characterization, observe that

$$M(v) - \langle v, x\rangle = \sum_{i=1}^{k} x_i\big(M(v) - v_i\big). \tag{D.3}$$

Each summand on the right-hand side is nonnegative, because $x_i \geq 0$ and $M(v) - v_i \geq 0$. Therefore the sum in (D.3) vanishes if and only if every summand vanishes. Since $x_i \geq 0$, this is equivalent to

$$x_i > 0 \implies M(v) - v_i = 0,$$

or, equivalently,

$$x_i > 0 \implies i \in A(v).$$

This is exactly (D.2). □

We now prove the polyhedral description of the low-regret continuation sets.

Proof of Theorem 4.1. Fix a history $h \in \mathcal{H}$ and a mixed continuation $x \in \Delta_k$. By affineness of J under mixing,

$$J(h, x) = \sum_{i=1}^{k} x_i J(h, e_i) = \sum_{i=1}^{k} x_i v_i(h) = \langle v(h), x\rangle. \tag{D.4}$$

Moreover, for each $i \in [k]$,

$$J(h, \phi_i x) = J(h, e_i) = v_i(h).$$

Substituting these identities into the definition of deviation regret yields

$$\begin{aligned}
\mathrm{Reg}_{\Phi_{\mathrm{ext}}}(h, x) &= \sup_{i\in[k]} \Big(J(h, \phi_i x) - J(h, x)\Big) \\
&= \sup_{i\in[k]} \Big(v_i(h) - \langle v(h), x\rangle\Big) \\
&= \max_{i\in[k]} v_i(h) - \langle v(h), x\rangle.
\end{aligned}$$

This proves (4.1).

Now apply Definition 2.2. A point $x \in \Delta_k$ belongs to K_h^ε if and only if

$$\mathrm{Reg}_{\Phi_{\mathrm{ext}}}(h, x) \leq \varepsilon,$$

which, by the displayed identity above, is equivalent to

$$\max_{i\in[k]} v_i(h) - \langle v(h), x\rangle \leq \varepsilon.$$

Rearranging gives

$$\langle v(h), x\rangle \geq \max_{i\in[k]} v_i(h) - \varepsilon,$$

and therefore

$$K_h^\varepsilon = \left\{x \in \Delta_k : \langle v(h), x\rangle \geq \max_{i\in[k]} v_i(h) - \varepsilon\right\}.$$

This is exactly (4.2).

Finally, Δ_k is a simplex, hence a compact polytope, and

$$\left\{x \in \mathbb{R}^k : \langle v(h), x\rangle \geq \max_{i\in[k]} v_i(h) - \varepsilon\right\}$$

is a closed half-space. Their intersection is therefore a convex polytope. □

To derive the exact hypergraph law, it is convenient to isolate the conversion between hitting subsets of $[k]$ and piercing simplex faces.

Proposition D.2 (Budgeted face-piercing equivalence)**.** *Let $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$ be any family of nonempty subsets, and for each $S \in \mathcal{S}$ define the simplex face*

$$\Delta(S) := \{x \in \Delta_k : \operatorname{supp}(x) \subseteq S\}.$$

Fix $m \in \mathbb{N}$. Then the following are equivalent:

1. *there exists a set $P \subseteq \Delta_k$ with $|P| \leq m$ such that*

$$P \cap \Delta(S) \neq \varnothing \qquad \forall S \in \mathcal{S};$$

2. *there exists a set $T \subseteq [k]$ with $|T| \leq m$ such that*

$$T \cap S \neq \varnothing \qquad \forall S \in \mathcal{S}.$$

Consequently,

$$\tau(\{\Delta(S) : S \in \mathcal{S}\}) = \tau(\mathcal{S}), \tag{D.5}$$

where the left-hand side is the transversal number in the ambient space Δ_k, and the right-hand side is the usual hypergraph transversal number on the vertex set $[k]$.

Proof. We prove the two implications separately.

(2) ⇒ (1). Assume $T \subseteq [k]$ hits every set in $\mathcal{S}$ and $|T| \leq m$. Define

$$P_T := \{e_i : i \in T\} \subseteq \Delta_k.$$

Then $|P_T| = |T| \leq m$. Let $S \in \mathcal{S}$. Since $T \cap S \neq \varnothing$, choose $i \in T \cap S$. Then $\operatorname{supp}(e_i) = \{i\} \subseteq S$, so $e_i \in \Delta(S)$. Because $e_i \in P_T$, we obtain

$$P_T \cap \Delta(S) \neq \varnothing.$$

Since $S \in \mathcal{S}$ was arbitrary, P_T pierces every face $\Delta(S)$.

(1) ⇒ (2). Assume $P \subseteq \Delta_k$ pierces every face $\Delta(S)$ and $|P| \leq m$. Because P is finite, we may write

$$P = \{x^{(1)}, \ldots, x^{(r)}\} \qquad \text{with} \qquad r \leq m.$$

Each $x^{(j)} \in \Delta_k$ has nonempty support, since its coordinates are nonnegative and sum to 1. For each $j \in [r]$, choose an index

$$i_j \in \operatorname{supp}\Big(x^{(j)}\Big).$$

Define

$$T := \{i_1, \ldots, i_r\} \subseteq [k].$$

Then $|T| \leq r \leq m$. We claim that T hits every $S \in \mathcal{S}$.

Fix $S \in \mathcal{S}$. Since P pierces $\Delta(S)$, there exists some $x^{(j)} \in P \cap \Delta(S)$. By definition of $\Delta(S)$,

$$\operatorname{supp}\Big(x^{(j)}\Big) \subseteq S.$$

Because $i_j \in \operatorname{supp}(x^{(j)})$, it follows that $i_j \in S$. Since $i_j \in T$, we conclude that

$$T \cap S \neq \varnothing.$$

Thus T is a hitting set for $\mathcal{S}$.

The identity (D.5) follows immediately by taking the infimum over $m \in \mathbb{N}$. □

Proof of Corollary 4.2. Fix $h \in \mathcal{H}$, and recall the definition

$$S_h := \arg\max_{i\in[k]} v_i(h).$$

Since $[k]$ is finite, $S_h \neq \varnothing$.

We first identify K_h^0. By Theorem 4.1,

$$K_h^0 = \left\{ x \in \Delta_k : \langle v(h), x\rangle \geq \max_{i\in[k]} v_i(h) \right\}.$$

By Lemma D.1, every $x \in \Delta_k$ satisfies

$$\langle v(h), x\rangle \leq \max_{i\in[k]} v_i(h),$$

with equality if and only if

$$\operatorname{supp}(x) \subseteq \arg\max_{i\in[k]} v_i(h) = S_h.$$

Hence

$$K_h^0 = \{x \in \Delta_k : \operatorname{supp}(x) \subseteq S_h\} = \Delta(S_h).$$

Now let

$$\mathcal{S} := \{S_h : h \in \mathcal{H}\}.$$

By Theorem 3.1,

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = \tau\big(\{K_h^0 : h \in \mathcal{H}\}\big).$$

Using the identity $K_h^0 = \Delta(S_h)$, this becomes

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = \tau(\{\Delta(S_h) : h \in \mathcal{H}\}).$$

Finally, Proposition D.2 implies

$$\tau(\{\Delta(S_h) : h \in \mathcal{H}\}) = \tau(\mathcal{S}).$$

Therefore

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = \tau(\mathcal{S}),$$

which is exactly (4.3). □

Remark D.3 (Why exact rationality collapses to support geometry). At tolerance $\varepsilon = 0$, the only mixed continuations that remain deviation-rational after history h are those that place mass exclusively on maximizing pure controllers. Thus the exact low-regret set retains only the *support pattern* S_h, and all numerical information in the values $v_i(h)$ beyond the maximizer set becomes irrelevant. This is precisely what allows exact retentive complexity to collapse to a hypergraph transversal problem.

E Full Proof of Theorem 4.3

This section proves a stronger version of Theorem 4.3. Rather than realizing only the family of exact rational-support sets, we will construct, for an arbitrary finite hypergraph, a one-step history process whose entire family of ε-rational continuation sets is available in closed form for every $\varepsilon \geq 0$.

The proof is intentionally constructive and minimal. The realized instance is a fully observable one-step decision problem with binary rewards. This is important conceptually: it shows that the universality phenomenon established by Theorem 4.3 is not an artifact of a rich latent-state model, long-horizon dependence, or partial observability. It already appears in the simplest possible sequential setting.

E.1 Construction

Fix $k \in \mathbb{N}$, and let

$$\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$$

be a finite hypergraph on vertex set $[k]$. We first dispose of the degenerate case $\mathcal{S} = \varnothing$: if the hypergraph is empty, the theorem is trivial, since one may take a decision process with no attainable pre-action histories, in which case the realized family of rational-support sets is empty. Accordingly, throughout the remainder of this section we assume

$$\mathcal{S} \neq \varnothing.$$

For each hyperedge $S \in \mathcal{S}$, introduce a distinct observation symbol o_S. Consider the following one-step history process:

1. Nature samples a hyperedge $S \in \mathcal{S}$ from any distribution μ on $\mathcal{S}$ with full support, e.g.

$$\mu(S) = \frac{1}{|\mathcal{S}|} \qquad \forall S \in \mathcal{S}.$$

2. The agent observes o_S.
3. The agent selects an action $A \in [k]$.
4. The reward is

$$R := \mathbf{1}\{A \in S\}.$$

5. The episode terminates immediately.

Let h_S denote the pre-action history consisting of the single observation o_S. Since μ has full support, the attainable pre-action histories are exactly

$$\mathcal{H}_{\mathcal{S}} := \{h_S : S \in \mathcal{S}\}.$$

We now specify the continuation-controller class and the deviation class. Let

$$\mathcal{C}_{\mathcal{S}} := \Delta_k$$

be the simplex over the canonical basis $e_1, \ldots, e_k$, where e_i is the pure controller that deterministically chooses action i. Thus the controller dictionary has size k, namely $\{e_1, \ldots, e_k\}$, and $\mathcal{C}_{\mathcal{S}}$ is the class of all mixed continuations over this dictionary.

Let the deviation class be the external-deviation family

$$\Phi_{\text{ext}} := \{\phi_i : i \in [k]\}, \qquad \phi_i(x) := e_i \quad \forall x \in \Delta_k.$$

Finally, define the continuation value functional

$$J_{\mathcal{S}}(h_S, x) := \mathbb{E}[R \mid h_S,\ x] \qquad \forall S \in \mathcal{S},\ \forall x \in \Delta_k.$$

Since the process terminates after one action, this expectation is completely explicit.

Lemma E.1 (Continuation values in the realization instance)**.** *For every $S \in \mathcal{S}$ and every $x \in \Delta_k$,*

$$J_{\mathcal{S}}(h_S, x) = \sum_{i \in S} x_i. \tag{E.1}$$

In particular, for each pure library controller e_i,

$$J_{\mathcal{S}}(h_S, e_i) = \mathbf{1}\{i \in S\}. \tag{E.2}$$

Proof. Fix $S \in \mathcal{S}$ and $x \in \Delta_k$. Under controller x, the action A is sampled from $[k]$ according to x, so

$$\mathbb{P}(A = i \mid h_S, x) = x_i \qquad \forall i \in [k].$$

By definition of the reward,

$$R = \mathbf{1}\{A \in S\}.$$

Therefore

$$\begin{aligned} J_{\mathcal{S}}(h_S, x) &= \mathbb{E}[\mathbf{1}\{A \in S\} \mid h_S, x] \\ &= \sum_{i=1}^{k} \mathbb{P}(A = i \mid h_S, x)\, \mathbf{1}\{i \in S\} \\ &= \sum_{i=1}^{k} x_i\, \mathbf{1}\{i \in S\} = \sum_{i \in S} x_i, \end{aligned}$$

which proves (E.1). Evaluating this identity at $x = e_i$ gives

$$J_{\mathcal{S}}(h_S, e_i) = \sum_{j \in S} (e_i)_j = \mathbf{1}\{i \in S\},$$

which is (E.2). □

The next proposition is the stronger realization statement alluded to above.

Proposition E.2 (Strong realization theorem)**.** *For the one-step history process constructed above, the following hold for every hyperedge $S \in \mathcal{S}$.*

1. *The pure-controller value vector satisfies*

$$v(h_S) = (\mathbf{1}\{1 \in S\}, \ldots, \mathbf{1}\{k \in S\}) \in \{0,1\}^k.$$

2. *The exact rational-support set at history h_S is*

$$\arg\max_{i \in [k]} J_{\mathcal{S}}(h_S, e_i) = S. \tag{E.3}$$

3. *For every $\varepsilon \geq 0$, the ε-rational continuation set is*

$$K^{\varepsilon}_{h_S} = \left\{ x \in \Delta_k : \sum_{i \in S} x_i \geq 1 - \varepsilon \right\}. \tag{E.4}$$

4. *In particular, at exact rationality $\varepsilon = 0$,*

$$K^{0}_{h_S} = \Delta(S) := \{x \in \Delta_k : x_i = 0 \textit{ for all } i \notin S\}. \tag{E.5}$$

Proof. Fix $S \in \mathcal{S}$.

Part 1. By Lemma E.1,

$$v_i(h_S) := J_{\mathcal{S}}(h_S, e_i) = \mathbf{1}\{i \in S\} \qquad \forall i \in [k].$$

This is exactly the stated formula for $v(h_S)$.

Part 2. Because $S \neq \varnothing$, there exists at least one $i \in S$, and hence

$$\max_{j \in [k]} J_{\mathcal{S}}(h_S, e_j) = 1.$$

By Part 1,

$$J_{\mathcal{S}}(h_S, e_i) = 1 \iff i \in S.$$

Therefore the set of maximizers is precisely S, proving (E.3).

Part 3. Let $x \in \Delta_k$. By definition of external deviations,

$$\phi_i(x) = e_i \qquad \forall i \in [k].$$

Hence the deviation regret at history h_S is

$$\begin{aligned}\mathrm{Reg}_{\Phi_{\mathrm{ext}}}(h_S, x) &= \sup_{i\in[k]} (J_{\mathcal{S}}(h_S, \phi_i(x)) - J_{\mathcal{S}}(h_S, x)) \\ &= \sup_{i\in[k]} (J_{\mathcal{S}}(h_S, e_i) - J_{\mathcal{S}}(h_S, x)).\end{aligned}$$

By Part 1 and Lemma E.1,

$$\sup_{i\in[k]} J_{\mathcal{S}}(h_S, e_i) = 1, \qquad J_{\mathcal{S}}(h_S, x) = \sum_{i\in S} x_i.$$

Therefore

$$\mathrm{Reg}_{\Phi_{\mathrm{ext}}}(h_S, x) = 1 - \sum_{i\in S} x_i. \tag{E.6}$$

By Definition 2.2,

$$x \in K^{\varepsilon}_{h_S} \iff \mathrm{Reg}_{\Phi_{\mathrm{ext}}}(h_S, x) \le \varepsilon.$$

Using (E.6), this is equivalent to

$$1 - \sum_{i\in S} x_i \le \varepsilon,$$

or, equivalently,

$$\sum_{i\in S} x_i \ge 1 - \varepsilon.$$

This proves (E.4).

Part 4. Setting $\varepsilon = 0$ in (E.4) yields

$$K^0_{h_S} = \left\{ x \in \Delta_k : \sum_{i\in S} x_i \ge 1 \right\}.$$

Since $x \in \Delta_k$ implies $x_i \ge 0$ for all i and $\sum_{i=1}^k x_i = 1$, the inequality

$$\sum_{i\in S} x_i \ge 1$$

is equivalent to the equality

$$\sum_{i\in S} x_i = 1,$$

which in turn is equivalent to

$$\sum_{i\notin S} x_i = 0.$$

Because all coordinates of x are nonnegative, this happens if and only if

$$x_i = 0 \qquad \forall i \notin S.$$

Thus

$$K^0_{h_S} = \{x \in \Delta_k : x_i = 0 \; \forall i \notin S\} = \Delta(S),$$

proving (E.5). □

We can now derive the main-text realization theorem immediately.

Proof of Theorem 4.3. If $\mathcal{S} = \varnothing$, the result is trivial, as noted above. Assume henceforth that $\mathcal{S} \neq \varnothing$, and consider the one-step history process constructed in this section.

By Proposition E.2, for every $S \in \mathcal{S}$,

$$\arg\max_{i\in[k]} J_{\mathcal{S}}(h_S, e_i) = S.$$

Therefore the family of exact rational-support sets of the constructed instance is

$$\left\{ \arg\max_{i\in[k]} J_{\mathcal{S}}(h, e_i) : h \in \mathcal{H}_{\mathcal{S}} \right\} = \left\{ \arg\max_{i\in[k]} J_{\mathcal{S}}(h_S, e_i) : S \in \mathcal{S} \right\} = \mathcal{S}.$$

Equivalently, if the main text denotes

$$S_h := \arg\max_{i\in[k]} v_i(h), \qquad v_i(h) = J(h, e_i),$$

then the realized family satisfies

$$\{S_h : h \in \mathcal{H}_{\mathcal{S}}\} = \mathcal{S}.$$

This is exactly the claim of Theorem 4.3. □

Corollary E.3 (Memory universality at exact rationality)**.** *For every finite hypergraph $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$, there exists a one-step history process such that*

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = \tau(\mathcal{S}).$$

Proof. Apply Theorem 4.3 to realize $\mathcal{S}$ as the family of exact rational-support sets. Then Corollary 4.2 implies that the exact retentive complexity of the realized instance is the transversal number of $\mathcal{S}$. □

Remark E.4 (Why nonempty hyperedges are necessary). The restriction $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$ is not cosmetic. If $S = \varnothing$, then in the one-step binary-reward model one has

$$J_{\mathcal{S}}(h_S, e_i) = 0 \qquad \forall i \in [k],$$

so the maximizer set would be all of $[k]$, not $\varnothing$. Thus empty hyperedges cannot be realized as exact rational-support sets in this model. This is precisely why the theorem is formulated for hypergraphs with nonempty edges.

Remark E.5 (Minimality of the construction). The realizing instance is a contextual bandit with a single decision, binary rewards, and full observability. In particular, the universality statement does not rely on recurrence, hidden state, long temporal dependencies, or inaccurate models. These richer phenomena are central elsewhere in the paper, but they are not needed to show that exact deviation-rational memory already has the full combinatorial expressivity of arbitrary finite hypergraphs.

F Proof of Theorem 5.1 and Lemma 5.2

This section proves the quantitative separation showing that pairwise compatibility information can be maximally uninformative about exact retentive complexity. The proof combines a finite-geometric construction with the exact hypergraph law from Section 4.

F.1 Projective planes and their line hypergraphs

We first record the finite-geometric object that drives the separation.

Definition F.1 (Finite projective plane of order q)**.** Let $q \in \mathbb{N}$. A *finite projective plane of order* q is an incidence structure

$$\Pi_q = (P_q, \mathcal{L}_q, \in),$$

where P_q is a finite set of points and $\mathcal{L}_q$ is a finite set of lines, satisfying:

1. every line $L \in \mathcal{L}_q$ contains exactly $q+1$ points;
2. every point $p \in P_q$ lies on exactly $q+1$ lines;
3. any two distinct points lie on a unique common line;
4. any two distinct lines intersect in a unique common point;
5. $$|P_q| = |\mathcal{L}_q| = q^2 + q + 1.$$

Associated to Π_q is the *line hypergraph*

$$\mathscr{L}_q := \{L \subseteq P_q : L \in \mathcal{L}_q\},$$

viewed as a hypergraph on the vertex set P_q.

Remark F.2. Theorem 5.1 is conditional on the existence of a projective plane of order q. We do not need any classification result about which orders admit such planes. Once Π_q is given, the proof below is entirely elementary.

The next lemma is the key finite-geometric counting fact.

Lemma F.3. *Let $\Pi_q = (P_q, \mathcal{L}_q, \in)$ be a finite projective plane of order q, and let $\mathscr{L}_q$ be its line hypergraph. Then*

$$\tau(\mathscr{L}_q) = q + 1.$$

Proof. We prove the upper and lower bounds separately.

Upper bound. Fix any line $L^\star \in \mathcal{L}_q$. Since $L^\star$ contains exactly $q+1$ points, it is a subset of P_q of cardinality $q+1$. We claim that $L^\star$ is a transversal for $\mathscr{L}_q$. Indeed, let $L \in \mathcal{L}_q$ be any line. If $L = L^\star$, then trivially $L^\star \cap L = L^\star \neq \varnothing$. If $L \neq L^\star$, then by the projective-plane axiom that any two distinct lines intersect in a unique point,

$$L^\star \cap L \neq \varnothing.$$

Thus $L^\star$ intersects every hyperedge of $\mathscr{L}_q$, so

$$\tau(\mathscr{L}_q) \leq |L^\star| = q + 1.$$

Lower bound. Let $T \subseteq P_q$ be any set with $|T| \leq q$. We will show that T is *not* a transversal for $\mathscr{L}_q$.

For each point $p \in T$, let

$$\mathcal{L}(p) := \{L \in \mathcal{L}_q : p \in L\}$$

denote the set of lines through p. By Definition F.1,

$$|\mathcal{L}(p)| = q + 1 \qquad \forall p \in T.$$

Hence the set of lines incident to at least one point of T satisfies

$$\begin{aligned}\left|\bigcup_{p \in T} \mathcal{L}(p)\right| &\leq \sum_{p \in T} |\mathcal{L}(p)| \\ &= |T|(q+1) \\ &\leq q(q+1) \\ &= q^2 + q.\end{aligned}$$

On the other hand, the total number of lines in Π_q is

$$|\mathcal{L}_q| = q^2 + q + 1.$$

Therefore

$$\left| \bigcup_{p \in T} \mathcal{L}(p) \right| < |\mathcal{L}_q|,$$

so there exists at least one line $L \in \mathcal{L}_q$ such that

$$L \notin \bigcup_{p \in T} \mathcal{L}(p).$$

By definition of the union, this means $p \notin L$ for every $p \in T$, i.e.

$$T \cap L = \varnothing.$$

Thus T fails to intersect the hyperedge L, and hence T is not a transversal.

We have shown that no subset of P_q of cardinality at most q can hit all lines. Therefore

$$\tau(\mathscr{L}_q) \geq q + 1.$$

Combining the upper and lower bounds yields

$$\tau(\mathscr{L}_q) = q + 1.$$

□

F.2 Pairwise compatibility graph

For clarity, we restate the graph notion used in the main theorem.

Definition F.4 (Pairwise compatibility graph at exact rationality)**.** Let $\mathfrak{M}$ be an instance of the mixed-controller model with attainable history set $\mathcal{H}$, controller class Δ_n, and exact low-regret sets $\{K_h^0 : h \in \mathcal{H}\}$. The *pairwise compatibility graph* of $\mathfrak{M}$ is the simple graph

$$G_{\text{pair}}(\mathfrak{M}) = (\mathcal{H}, E),$$

where two distinct histories $h, h' \in \mathcal{H}$ are adjacent iff

$$K_h^0 \cap K_{h'}^0 \neq \varnothing.$$

Remark F.5. The graph $G_{\text{pair}}(\mathfrak{M})$ records only whether pairs of exact low-regret sets intersect. It forgets all higher-order intersection geometry. Theorem 5.1 shows that this loss of information can be severe: even the complete graph does not determine exact retentive complexity.

We now prove the main separation theorem.

Proof of Theorem 5.1. Fix a projective plane $\Pi_q = (P_q, \mathcal{L}_q, \in)$ of order q, and let

$$n := q^2 + q + 1 = |P_q| = |\mathcal{L}_q|.$$

Enumerate the points of P_q as $[n]$. We construct two one-step instances, denoted $\mathfrak{M}_q^{(A)}$ and $\mathfrak{M}_q^{(B)}$, each with:

- controller dictionary $\{e_1, \ldots, e_n\}$,
- mixed continuation class Δ_n,

- external deviation class
$$\Phi_{\text{ext}} = \{\phi_i : i \in [n]\}, \qquad \phi_i(x) = e_i.$$

Both instances have exactly n attainable pre-action histories.

We will show:

1. $G_{\text{pair}}(\mathfrak{M}_q^{(A)})$ is the complete graph on n vertices;
2. $G_{\text{pair}}(\mathfrak{M}_q^{(B)})$ is also the complete graph on n vertices;
3. $$m^{\text{pc}}_{\Phi_{\text{ext}}}(0; \mathfrak{M}_q^{(A)}) = 1, \qquad m^{\text{pc}}_{\Phi_{\text{ext}}}(0; \mathfrak{M}_q^{(B)}) = q + 1.$$

This will prove the theorem.

Instance A: complete pairwise graph with memory 1. Let the attainable histories of $\mathfrak{M}_q^{(A)}$ be

$$\mathcal{H}^{(A)} := \{h_1, \dots, h_n\},$$

represented concretely by n distinct observations. The agent then chooses an action $A \in [n]$, and the reward is

$$R^{(A)} := \mathbf{1}\{A = 1\}.$$

The episode terminates immediately.

For any history h_j and any mixed controller $x \in \Delta_n$,

$$J^{(A)}(h_j, x) = x_1.$$

Consequently, for each pure controller e_i,

$$J^{(A)}(h_j, e_i) = \mathbf{1}\{i = 1\}.$$

Hence the unique maximizing controller is e_1, and therefore

$$K^{0,(A)}_{h_j} = \Delta(\{1\}) = \{e_1\} \qquad \forall j \in [n].$$

It follows immediately that, for any two distinct histories $h_j, h_{j'}$,

$$K^{0,(A)}_{h_j} \cap K^{0,(A)}_{h_{j'}} = \{e_1\} \neq \varnothing.$$

Thus

$$G_{\text{pair}}(\mathfrak{M}_q^{(A)})$$

is the complete graph on n vertices.

Moreover, since the singleton set $\{e_1\}$ intersects every exact low-regret set, Theorem 3.1 implies

$$m^{\text{pc}}_{\Phi_{\text{ext}}}(0; \mathfrak{M}_q^{(A)}) = 1.$$

Instance B: projective-plane geometry. Let the attainable histories of $\mathfrak{M}_q^{(B)}$ be indexed by the lines of Π_q:

$$\mathcal{H}^{(B)} := \{h_L : L \in \mathcal{L}_q\}.$$

Since $|\mathcal{L}_q| = n$, this instance also has n histories.

At history h_L, the agent chooses an action $A \in [n]$, interpreted as a point of P_q, and receives reward

$$R^{(B)} := \mathbf{1}\{A \in L\}.$$

The episode terminates immediately.

By Proposition E.2 applied to the hypergraph

$$\mathscr{L}_q = \{L \subseteq P_q : L \in \mathcal{L}_q\},$$

the exact low-regret set at history h_L is

$$K^{0,(B)}_{h_L} = \Delta(L) \qquad \forall L \in \mathcal{L}_q.$$

Now let $L, L' \in \mathcal{L}_q$ be distinct. Since Π_q is a projective plane, the two lines intersect in a unique point $p \in P_q$. Therefore

$$e_p \in \Delta(L) \cap \Delta(L') = K^{0,(B)}_{h_L} \cap K^{0,(B)}_{h_{L'}}.$$

Hence every pair of distinct histories in $\mathfrak{M}_q^{(B)}$ is adjacent, so

$$G_{\mathrm{pair}}(\mathfrak{M}_q^{(B)})$$

is also the complete graph on n vertices.

Finally, by Corollary 4.2, the exact retentive complexity of $\mathfrak{M}_q^{(B)}$ is the transversal number of the family of support sets $\mathscr{L}_q$. Lemma 5.2 therefore gives

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0; \mathfrak{M}_q^{(B)}) = \tau(\mathscr{L}_q) = q + 1.$$

Conclusion. The two instances $\mathfrak{M}_q^{(A)}$ and $\mathfrak{M}_q^{(B)}$ have the same pairwise compatibility graph—namely the complete graph on n vertices, where $n = q^2 + q + 1$—but their exact retentive complexities are 1 and $q + 1$, respectively. This proves Theorem 5.1. □

Corollary F.6 (Quantitative failure of pairwise determination)**.** *For every projective plane of order q, there exist two exact-rationality instances with identical pairwise compatibility graph whose retentive complexities differ by a factor of $q + 1$.*

Proof. This is immediate from Theorem 5.1, since the two realized complexities are 1 and $q+1$. □

Remark F.7 (Why the complete graph is the strongest possible host graph). Theorem 5.1 uses the complete compatibility graph, which is the graph containing the *most* pairwise information one could hope for at the graph level. Thus the obstruction is not that the graph is too sparse to reveal the right answer. Rather, the obstruction is genuinely higher-order: pairwise intersections do not determine the transversal number of the family of exact low-regret sets.

Remark F.8 (Same controller class, same deviation class). The two instances in the proof are matched in a stronger sense than claimed in the theorem statement. They use the same controller simplex Δ_n, the same external deviation family, and the same number n of attainable histories. The only difference lies in the geometry of the exact low-regret continuation sets. This isolates the source of the separation precisely.

G Full Proof of Corollary 6.1

This section gives a complete proof of Corollary 6.1. The argument has two ingredients. The first is a finite-dimensional Helly reduction inside the affine hull of the simplex. The second is a compactness argument upgrading finite intersection to intersection of the entire history-indexed family.

G.1 Finite-family Helly reduction inside the simplex

We begin with the finite-family statement.

Lemma G.1 (Helly reduction on the simplex)**.** *Let $\mathcal{F}$ be a finite family of convex subsets of Δ_k. If every subfamily of $\mathcal{F}$ of cardinality at most k has nonempty intersection, then*

$$\bigcap_{F\in\mathcal{F}} F \neq \varnothing.$$

Proof. Let

$$H := \left\{ x \in \mathbb{R}^k : \sum_{i=1}^{k} x_i = 1 \right\}.$$

Then $\Delta_k \subseteq H$, and H is an affine hyperplane of dimension $k-1$. Consider the affine bijection

$$T : \mathbb{R}^{k-1} \to H, \qquad T(z_1, \dots, z_{k-1}) = \left(z_1, \dots, z_{k-1},\, 1 - \sum_{i=1}^{k-1} z_i \right).$$

For each $F \in \mathcal{F}$, define

$$\widetilde{F} := T^{-1}(F) \subseteq \mathbb{R}^{k-1}.$$

Since T is affine and F is convex, each $\widetilde{F}$ is convex in $\mathbb{R}^{k-1}$. Moreover, because T is bijective,

$$\bigcap_{j=1}^{m} F_j \neq \varnothing \quad \Longleftrightarrow \quad \bigcap_{j=1}^{m} \widetilde{F}_j \neq \varnothing$$

for every finite subfamily $\{F_1, \dots, F_m\} \subseteq \mathcal{F}$.

Hence every subfamily of $\left\{ \widetilde{F} : F \in \mathcal{F} \right\}$ of cardinality at most $k = (k-1)+1$ has nonempty intersection. By Helly's theorem in $\mathbb{R}^{k-1}$ (Matoušek, 2002, Chapter 2),

$$\bigcap_{F\in\mathcal{F}} \widetilde{F} \neq \varnothing.$$

Applying T gives

$$\bigcap_{F\in\mathcal{F}} F \neq \varnothing,$$

as claimed. □

G.2 Closedness and compactness

The next lemma upgrades the finite-family conclusion to the full history-indexed family.

Lemma G.2 (Finite intersection property on the simplex)**.** *Let $\{F_\alpha : \alpha \in I\}$ be a family of closed subsets of the simplex Δ_k. If every finite subfamily has nonempty intersection, then*

$$\bigcap_{\alpha\in I} F_\alpha \neq \varnothing.$$

Proof. The simplex Δ_k is compact as a closed and bounded subset of $\mathbb{R}^k$. Suppose, for contradiction, that

$$\bigcap_{\alpha\in I} F_\alpha = \varnothing.$$

Then

$$\Delta_k \subseteq \bigcup_{\alpha\in I} (\Delta_k \setminus F_\alpha).$$

Because each F_α is closed in Δ_k, the relative complement $\Delta_k \setminus F_\alpha$ is open in Δ_k. Thus $\{\Delta_k \setminus F_\alpha : \alpha \in I\}$ is an open cover of the compact space Δ_k, and so admits a finite subcover:

$$\Delta_k \subseteq \bigcup_{j=1}^{m} (\Delta_k \setminus F_{\alpha_j})$$

for some $\alpha_1, \ldots, \alpha_m \in I$. Equivalently,

$$\bigcap_{j=1}^{m} F_{\alpha_j} = \varnothing,$$

contradicting the finite intersection hypothesis. □

We now verify the closedness property needed to apply Lemma G.2 to the low-regret continuation sets.

Lemma G.3 (Closedness of low-regret sets in the mixed-controller model)**.** *In the finite-controller mixed-continuation model of Section 4, for every history $h \in \mathcal{H}$ and every $\varepsilon \geq 0$, the set K_h^ε is a closed convex subset of Δ_k.*

Proof. By Theorem 4.1,

$$K_h^\varepsilon = \left\{ x \in \Delta_k : \langle v(h), x \rangle \geq \max_{i \in [k]} v_i(h) - \varepsilon \right\}.$$

The map $x \mapsto \langle v(h), x \rangle$ is affine, hence continuous, on $\mathbb{R}^k$. It follows that

$$\left\{ x \in \mathbb{R}^k : \langle v(h), x \rangle \geq \max_{i \in [k]} v_i(h) - \varepsilon \right\}$$

is a closed half-space. Therefore K_h^ε, being the intersection of this closed half-space with the simplex Δ_k, is closed in Δ_k. It is also convex because both sets in the intersection are convex. □

G.3 Proof of the corollary

Full proof of Corollary 6.1. We prove the equivalence of the three statements in the corollary.

(1) $\iff$ (2). By Theorem 3.1,

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(\varepsilon) = 1 \quad \iff \quad \tau(\{K_h^\varepsilon : h \in \mathcal{H}\}) = 1.$$

Since $\mathcal{H} \neq \varnothing$ by Definition 2.1, a transversal of cardinality at most 1 must be a singleton $\{x\}$ for some $x \in \Delta_k$. Such a singleton pierces every K_h^ε if and only if

$$x \in K_h^\varepsilon \qquad \forall h \in \mathcal{H},$$

which is equivalent to

$$\bigcap_{h \in \mathcal{H}} K_h^\varepsilon \neq \varnothing.$$

This proves (1) $\iff$ (2).

(2) $\implies$ (3). This implication is immediate: if the total intersection is nonempty, then every subfamily, in particular every subfamily of cardinality at most k, also has nonempty intersection.

(3) $\Longrightarrow$ (2). Assume that every subfamily of

$$\mathcal{K}_\varepsilon := \{K_h^\varepsilon : h \in \mathcal{H}\}$$

of cardinality at most k has nonempty intersection.

We first show that every *finite* subfamily of $\mathcal{K}_\varepsilon$ has nonempty intersection. Let

$$\mathcal{F} \subseteq \mathcal{K}_\varepsilon$$

be finite. By Lemma G.3, every member of $\mathcal{F}$ is convex and contained in Δ_k. Moreover, every subfamily of $\mathcal{F}$ of cardinality at most k has nonempty intersection by the assumption of (3). Therefore Lemma G.1 applies and yields

$$\bigcap_{F \in \mathcal{F}} F \neq \varnothing.$$

Since $\mathcal{F}$ was arbitrary, the family $\mathcal{K}_\varepsilon$ has the finite intersection property.

Next, Lemma G.3 shows that every K_h^ε is closed in the compact simplex Δ_k. Hence Lemma G.2 implies

$$\bigcap_{h \in \mathcal{H}} K_h^\varepsilon \neq \varnothing.$$

This is exactly statement (2).

Combining the three implications proves the equivalence of (1), (2), and (3). □

Remark G.4 (Why the threshold is k). The threshold k is not arbitrary. The family $\{K_h^\varepsilon : h \in \mathcal{H}\}$ lives inside the affine hull of the simplex,

$$\operatorname{aff}(\Delta_k) = \left\{ x \in \mathbb{R}^k : \sum_{i=1}^{k} x_i = 1 \right\},$$

which has dimension $k-1$. Helly's theorem therefore yields a local intersection certificate of order $(k-1)+1 = k$. The certificate depends only on the size of the controller dictionary and not on any measure of environmental complexity.

Remark G.5 (Finite-history special case). If $\mathcal{H}$ is finite, the compactness step is unnecessary. In that case statement (3) implies statement (2) directly from Lemma G.1 by applying it to the finite family $\{K_h^\varepsilon : h \in \mathcal{H}\}$.

H Richer Deviation Classes

The main text focuses on external deviations because they already yield a sharp, exact memory law. This section shows that the geometric picture is robust to substantially richer deviation families. In particular:

1. finite affine deviation classes still induce polyhedral low-regret sets;
2. the one-state Helly certificate extends unchanged to such classes;
3. enriching the deviation class can only increase memory complexity in general;
4. once the deviation class contains the external deviations, the entire regret geometry is already saturated: for every $\varepsilon \geq 0$, the low-regret sets and retentive complexities coincide exactly with those of the external-deviation model.

This places the external-deviation analysis in the main text on a firmer foundation: the support geometry is not an artifact of using only the simplest deviation family, but rather a stable consequence of allowing at least the pure-action replacement deviations emphasized in the hindsight-rationality literature (Morrill et al., 2021).

H.1 Finite affine deviation classes

Fix the finite-controller mixed-continuation model of Section 4. Thus the continuation class is the simplex

$$\Delta_k := \left\{ x \in \mathbb{R}^k_+ : \sum_{i=1}^{k} x_i = 1 \right\},$$

and continuation values are affine under mixing:

$$J(h, x) = \langle v(h), x \rangle, \qquad v(h) := (v_1(h), \dots, v_k(h)) \in \mathbb{R}^k.$$

Definition H.1 (Finite affine deviation class)**.** A *finite affine deviation class* on Δ_k is a finite family

$$\Phi = \{\phi^1, \dots, \phi^m\}$$

of affine self-maps $\phi^\ell : \Delta_k \to \Delta_k$, each of the form

$$\phi^\ell(x) = M_\ell x + b_\ell,$$

where $M_\ell \in \mathbb{R}^{k \times k}$ and $b_\ell \in \mathbb{R}^k$ satisfy $\phi^\ell(\Delta_k) \subseteq \Delta_k$.

Remark H.2 (Examples). This class includes:

1. *external deviations*: $\phi_i(x) = e_i$, obtained by taking $M_i = 0$ and $b_i = e_i$;
2. *permutation deviations*: $\phi(x) = Px$, where P is a permutation matrix;
3. *stochastic mixing deviations*: $\phi(x) = Mx$, where M is a column-stochastic matrix;
4. *reset-and-mix deviations*: $\phi(x) = \alpha e_i + (1 - \alpha)Mx$ with $\alpha \in [0, 1]$, $i \in [k]$, and M column-stochastic.

Thus the affine formalism captures a broad family of structured deviations while preserving exact tractability.

Proposition H.3 (Polyhedral low-regret sets for finite affine deviations)**.** *Let $\Phi = \{\phi^1, \dots, \phi^m\}$ be a finite affine deviation class on Δ_k, with*

$$\phi^\ell(x) = M_\ell x + b_\ell.$$

Then, for every history $h \in \mathcal{H}$ and every $x \in \Delta_k$,

$$\mathrm{Reg}_\Phi(h, x) = \max_{1 \le \ell \le m} \Big(\big\langle (M_\ell^\top - I) v(h), x \big\rangle + \langle v(h), b_\ell \rangle \Big). \tag{H.1}$$

Consequently, for every $\varepsilon \ge 0$,

$$K^\varepsilon_{h,\Phi} = \Delta_k \cap \bigcap_{\ell=1}^{m} \big\{ x \in \mathbb{R}^k : \big\langle (M_\ell^\top - I) v(h), x \big\rangle \le \varepsilon - \langle v(h), b_\ell \rangle \big\}. \tag{H.2}$$

In particular, $K^\varepsilon_{h,\Phi}$ is a closed convex polytope.

Proof. Fix $h \in \mathcal{H}$, $x \in \Delta_k$, and $\ell \in [m]$. Since $J(h, \cdot)$ is affine under mixing and $\phi^\ell(x) = M_\ell x + b_\ell$,

$$J(h, \phi^\ell(x)) = \langle v(h), M_\ell x + b_\ell \rangle = \big\langle M_\ell^\top v(h), x \big\rangle + \langle v(h), b_\ell \rangle.$$

Therefore

$$J(h, \phi^\ell(x)) - J(h, x) = \big\langle (M_\ell^\top - I) v(h), x \big\rangle + \langle v(h), b_\ell \rangle.$$

Taking the maximum over $\ell \in [m]$ gives (H.1).

Now $x \in K^\varepsilon_{h,\Phi}$ iff $\mathrm{Reg}_\Phi(h, x) \le \varepsilon$. By (H.1), this is equivalent to the simultaneous inequalities

$$\big\langle (M_\ell^\top - I) v(h), x \big\rangle + \langle v(h), b_\ell \rangle \le \varepsilon \qquad \forall \ell \in [m],$$

or equivalently

$$\left\langle (M_\ell^\top - I)v(h), x \right\rangle \le \varepsilon - \langle v(h), b_\ell \rangle \qquad \forall \ell \in [m].$$

This proves (H.2). Since $K^\varepsilon_{h,\Phi}$ is the intersection of the simplex Δ_k with finitely many closed half-spaces, it is a closed convex polytope. □

Corollary H.4 (Helly certificate for finite affine deviations)**.** *Let Φ be any finite affine deviation class on Δ_k. Then, for every $\varepsilon \ge 0$, the following are equivalent:*

1. $m^{\mathrm{pc}}_\Phi(\varepsilon) = 1$;
2. $\bigcap_{h \in \mathcal{H}} K^\varepsilon_{h,\Phi} \neq \varnothing$;
3. *every subfamily of* $\left\{ K^\varepsilon_{h,\Phi} : h \in \mathcal{H} \right\}$ *of cardinality at most k has nonempty intersection.*

Proof. By Proposition H.3, every $K^\varepsilon_{h,\Phi}$ is a closed convex subset of Δ_k. The proof of Corollary 6.1 in Supplement G applies verbatim with K^ε_h replaced by $K^\varepsilon_{h,\Phi}$. □

H.2 Monotonicity under richer deviation classes

We next record a monotonicity principle that will be useful for interpreting the main results as lower bounds for more demanding deviation notions.

Proposition H.5 (Monotonicity under deviation enrichment)**.** *Let Φ_1 and Φ_2 be two deviation classes on the same continuation system, with*

$$\Phi_1 \subseteq \Phi_2.$$

Then, for every history $h \in \mathcal{H}$ and every $\varepsilon \ge 0$,

$$K^\varepsilon_{h,\Phi_2} \subseteq K^\varepsilon_{h,\Phi_1}. \tag{H.3}$$

Consequently,

$$m^{\mathrm{pc}}_{\Phi_2}(\varepsilon) \ge m^{\mathrm{pc}}_{\Phi_1}(\varepsilon). \tag{H.4}$$

Proof. Fix $h \in \mathcal{H}$ and $c \in \mathcal{C}$. Since $\Phi_1 \subseteq \Phi_2$,

$$\mathrm{Reg}_{\Phi_1}(h,c) = \sup_{\phi \in \Phi_1} \big(J(h,\phi c) - J(h,c) \big) \le \sup_{\phi \in \Phi_2} \big(J(h,\phi c) - J(h,c) \big) = \mathrm{Reg}_{\Phi_2}(h,c).$$

Hence $\mathrm{Reg}_{\Phi_2}(h,c) \le \varepsilon$ implies $\mathrm{Reg}_{\Phi_1}(h,c) \le \varepsilon$, which is exactly (H.3).

Now let

$$\mathcal{K}_{\Phi_r,\varepsilon} := \left\{ K^\varepsilon_{h,\Phi_r} : h \in \mathcal{H} \right\}, \qquad r \in \{1,2\}.$$

By (H.3), every transversal of $\mathcal{K}_{\Phi_2,\varepsilon}$ is also a transversal of $\mathcal{K}_{\Phi_1,\varepsilon}$. Therefore

$$\tau(\mathcal{K}_{\Phi_2,\varepsilon}) \ge \tau(\mathcal{K}_{\Phi_1,\varepsilon}).$$

Applying Theorem 3.1 on both sides yields (H.4). □

H.3 Saturation by external deviations

The next theorem is the corrected robustness statement. In the finite-controller simplex model, once external deviations are available, no additional simplex-valued deviations can produce larger regret. Thus the external deviations already determine the *entire* low-regret geometry, not merely its exact-rationality slice.

Theorem H.6 (Saturation by external deviations)**.** *Consider the finite-controller mixed-continuation model on Δ_k, and let Φ be any family of self-maps $\phi : \Delta_k \to \Delta_k$ such that*

$$\Phi_{\text{ext}} \subseteq \Phi.$$

Then, for every history $h \in \mathcal{H}$, every $x \in \Delta_k$, and every $\varepsilon \geq 0$,

$$\mathrm{Reg}_{\Phi}(h, x) = \mathrm{Reg}_{\Phi_{\text{ext}}}(h, x) = \max_{i \in [k]} v_i(h) - \langle v(h), x \rangle. \tag{H.5}$$

Consequently,

$$K^{\varepsilon}_{h,\Phi} = K^{\varepsilon}_{h,\Phi_{\text{ext}}} = \left\{ x \in \Delta_k : \langle v(h), x \rangle \geq \max_{i \in [k]} v_i(h) - \varepsilon \right\}, \tag{H.6}$$

and

$$m^{\text{pc}}_{\Phi}(\varepsilon) = m^{\text{pc}}_{\Phi_{\text{ext}}}(\varepsilon) \qquad \forall \varepsilon \geq 0. \tag{H.7}$$

Proof. Fix $h \in \mathcal{H}$ and $x \in \Delta_k$, and write

$$M_h := \max_{i \in [k]} v_i(h).$$

Because $\Phi_{\text{ext}} \subseteq \Phi$, we immediately have

$$\mathrm{Reg}_{\Phi}(h, x) \geq \mathrm{Reg}_{\Phi_{\text{ext}}}(h, x).$$

By Theorem 4.1,

$$\mathrm{Reg}_{\Phi_{\text{ext}}}(h, x) = M_h - \langle v(h), x \rangle.$$

Hence

$$\mathrm{Reg}_{\Phi}(h, x) \geq M_h - \langle v(h), x \rangle. \tag{H.8}$$

For the reverse inequality, let $\phi \in \Phi$. Since $\phi(x) \in \Delta_k$, Lemma D.1 implies

$$J(h, \phi(x)) = \langle v(h), \phi(x) \rangle \leq M_h.$$

Therefore

$$J(h, \phi(x)) - J(h, x) \leq M_h - \langle v(h), x \rangle.$$

Taking the supremum over $\phi \in \Phi$ yields

$$\mathrm{Reg}_{\Phi}(h, x) \leq M_h - \langle v(h), x \rangle.$$

Combining this with (H.8) proves (H.5).

The set identity (H.6) follows immediately from the definition of low-regret sets:

$$x \in K^{\varepsilon}_{h,\Phi} \iff \mathrm{Reg}_{\Phi}(h, x) \leq \varepsilon \iff M_h - \langle v(h), x \rangle \leq \varepsilon.$$

Finally, applying Theorem 3.1 to the two equal families $\left\{ K^{\varepsilon}_{h,\Phi} : h \in \mathcal{H} \right\}$ and $\left\{ K^{\varepsilon}_{h,\Phi_{\text{ext}}} : h \in \mathcal{H} \right\}$ gives (H.7). □

Corollary H.7 (Exact hypergraph law under arbitrary external enrichment)**.** *In the finite-controller mixed-continuation model, for every deviation class Φ satisfying*

$$\Phi_{\text{ext}} \subseteq \Phi \subseteq \{\phi : \Delta_k \to \Delta_k\},$$

one has

$$K^{0}_{h,\Phi} = \Delta(S_h) \qquad \forall h \in \mathcal{H},$$

and therefore

$$m^{\text{pc}}_{\Phi}(0) = \tau(\{S_h : h \in \mathcal{H}\}).$$

Proof. Set $\varepsilon = 0$ in (H.6), then invoke Lemma D.1 exactly as in the proof of Corollary 4.2. □

Corollary H.8 (Helly certificate under arbitrary external enrichment)**.** *Under the same hypotheses as Corollary H.7, for every $\varepsilon \geq 0$, the following are equivalent:*

1. $m^{\mathrm{pc}}_{\Phi}(\varepsilon) = 1$;
2. $\bigcap_{h \in \mathcal{H}} K^{\varepsilon}_{h,\Phi} \neq \varnothing$;
3. *every subfamily of* $\left\{K^{\varepsilon}_{h,\Phi} : h \in \mathcal{H}\right\}$ *of cardinality at most k has nonempty intersection.*

Proof. By Theorem H.6,

$$K^{\varepsilon}_{h,\Phi} = K^{\varepsilon}_{h,\Phi_{\mathrm{ext}}} \qquad \forall h \in \mathcal{H},\ \forall \varepsilon \geq 0.$$

Thus the statement is exactly Corollary 6.1. □

Remark H.9 (Where richer deviations still matter). If Φ is an arbitrary finite affine deviation class, then Proposition H.3 shows that the geometry of $K^{\varepsilon}_{h,\Phi}$ generally depends on the chosen deviations. The saturation theorem identifies the precise threshold at which that dependence disappears: once the external deviations are included, the finite-controller simplex model is already maximally adversarial, and no further simplex-valued deviation can enlarge regret or increase memory complexity at any tolerance $\varepsilon \geq 0$.

I Fractional Memory Relaxation and LP Duality

This section develops a linear-programming relaxation of exact retentive complexity in the finite-controller model. The resulting object is a fractional hitting-set problem, or equivalently a fractional face-piercing problem on the simplex. This viewpoint is not needed for any claim in the main paper, but it sharpens the optimization structure underlying the exact hypergraph law.

Throughout this section, fix a finite support family

$$\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\},$$

typically $\mathcal{S} = \{S_h : h \in \mathcal{H}\}$ from Corollary 4.2.

I.1 Fractional exact retentive complexity

Definition I.1 (Fractional transversal number)**.** The *fractional transversal number* of $\mathcal{S}$ is

$$\tau^*(\mathcal{S}) := \min\left\{\sum_{i=1}^{k} w_i : w \in \mathbb{R}^k_+,\ \sum_{i \in S} w_i \geq 1\ \ \forall S \in \mathcal{S}\right\}. \tag{I.1}$$

The vector w can be interpreted as assigning nonnegative *memory mass* to the pure library controllers $e_1, \ldots, e_k$, with the requirement that each support set $S \in \mathcal{S}$ receive total mass at least 1. This is the natural LP relaxation of the integral hitting-set problem.

Proposition I.2 (Fractional relaxation of exact retentive complexity)**.** *For every finite support family $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$,*

$$\tau^*(\mathcal{S}) \leq \tau(\mathcal{S}). \tag{I.2}$$

Consequently, in the exact-rationality mixed-controller model,

$$\tau^*(\{S_h : h \in \mathcal{H}\}) \leq m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0). \tag{I.3}$$

Proof. Let $T \subseteq [k]$ be any transversal for $\mathcal{S}$. Define

$$w_i := \mathbf{1}\{i \in T\} \qquad \forall i \in [k].$$

Then $w \in \mathbb{R}_+^k$, $\sum_{i=1}^k w_i = |T|$, and for every $S \in \mathcal{S}$,

$$\sum_{i \in S} w_i = |S \cap T| \geq 1$$

because T hits S. Thus w is feasible for (I.1), so

$$\tau^*(\mathcal{S}) \leq |T|.$$

Taking the minimum over all integral transversals T proves (I.2). The inequality (I.3) now follows from Corollary 4.2. □

I.2 Fractional face-piercing on the simplex

The main text identifies exact memory with the transversal number of simplex faces $\Delta(S)$. The next proposition shows that the *fractional* relaxation is equally insensitive to mixing: allowing weighted mixed continuations yields exactly the same optimum as weighting only the vertices $e_1, \ldots, e_k$.

Definition I.3 (Fractional face-piercing number)**.** For a family $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$, define

$$\Delta(S) := \{x \in \Delta_k : \operatorname{supp}(x) \subseteq S\}.$$

Its *fractional face-piercing number* is

$$\tau^*_{\text{face}}(\mathcal{S}) := \inf \left\{ \sum_{x \in P} \lambda_x : \begin{array}{l} P \subseteq \Delta_k \text{ finite}, \ \lambda_x \geq 0 \ \forall x \in P, \\ \sum_{x \in P \cap \Delta(S)} \lambda_x \geq 1 \ \ \forall S \in \mathcal{S} \end{array} \right\}. \tag{I.4}$$

Proposition I.4 (Fractional face-piercing equals fractional hitting)**.** *For every finite support family* $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$,

$$\tau^*_{\text{face}}(\mathcal{S}) = \tau^*(\mathcal{S}). \tag{I.5}$$

Proof. We prove the two inequalities separately.

Step 1: $\tau^*_{\text{face}}(\mathcal{S}) \leq \tau^*(\mathcal{S})$**.** Let $w \in \mathbb{R}_+^k$ be feasible for (I.1). Set

$$P := \{e_i : w_i > 0\} \subseteq \Delta_k, \qquad \lambda_{e_i} := w_i.$$

Fix $S \in \mathcal{S}$. Since $e_i \in \Delta(S)$ whenever $i \in S$,

$$\sum_{x \in P \cap \Delta(S)} \lambda_x = \sum_{i \in S} w_i \geq 1.$$

Hence (P, λ) is feasible for (I.4) with objective value $\sum_{i=1}^k w_i$. Taking the infimum over feasible w gives

$$\tau^*_{\text{face}}(\mathcal{S}) \leq \tau^*(\mathcal{S}).$$

Step 2: $\tau^*(\mathcal{S}) \leq \tau^*_{\text{face}}(\mathcal{S})$**.** Let $P \subseteq \Delta_k$ be finite, and let $\lambda_x \geq 0$ for $x \in P$, be feasible for (I.4). Since every $x \in \Delta_k$ has nonempty support, choose once and for all an index

$$i(x) \in \operatorname{supp}(x) \qquad \forall x \in P.$$

Define $w \in \mathbb{R}_+^k$ by

$$w_i := \sum_{x \in P:\, i(x) = i} \lambda_x \qquad \forall i \in [k].$$

Then

$$\sum_{i=1}^k w_i = \sum_{x \in P} \lambda_x.$$

Now fix $S \in \mathcal{S}$. If $x \in P \cap \Delta(S)$, then $\mathrm{supp}(x) \subseteq S$, hence $i(x) \in S$. Therefore

$$\sum_{i \in S} w_i = \sum_{x \in P:\, i(x) \in S} \lambda_x \geq \sum_{x \in P \cap \Delta(S)} \lambda_x \geq 1.$$

Thus w is feasible for (I.1) with the same objective value as (P, λ). Taking the infimum over feasible (P, λ) yields

$$\tau^*(\mathcal{S}) \leq \tau^*_{\mathrm{face}}(\mathcal{S}).$$

Combining the two inequalities proves (I.5). □

Remark I.5 (No mixed-controller advantage in the fractional relaxation). Proposition I.4 shows that the LP relaxation already lives on the vertices of the controller simplex. Allowing weighted mixed continuations does not decrease the optimum beyond what can already be achieved by weighting pure controllers.

I.3 LP duality and fractional packings

Let $A \in \{0,1\}^{\mathcal{S} \times [k]}$ denote the incidence matrix of the hypergraph $\mathcal{S}$:

$$A_{S,i} := \mathbf{1}\{i \in S\} \qquad \forall S \in \mathcal{S},\ \forall i \in [k].$$

Then (I.1) can be written compactly as

$$\tau^*(\mathcal{S}) = \min\big\{\mathbf{1}^\top w : Aw \geq \mathbf{1},\ w \geq 0\big\}. \tag{I.6}$$

Its dual is the fractional packing problem.

Definition I.6 (Fractional packing number)**.** The *fractional packing number* of $\mathcal{S}$ is

$$\nu^*(\mathcal{S}) := \max\left\{\sum_{S \in \mathcal{S}} y_S : y \in \mathbb{R}^{\mathcal{S}}_+,\ \sum_{S \ni i} y_S \leq 1\ \ \forall i \in [k]\right\}. \tag{I.7}$$

Theorem I.7 (Strong duality for fractional memory)**.** *For every finite support family $\mathcal{S} \subseteq 2^{[k]} \setminus \{\varnothing\}$,*

$$\tau^*(\mathcal{S}) = \nu^*(\mathcal{S}). \tag{I.8}$$

Moreover, both optima are attained.

Proof. The primal problem (I.6) is feasible, since the vector

$$w = \mathbf{1} \in \mathbb{R}^k$$

satisfies $Aw \geq \mathbf{1}$ because every hyperedge $S \in \mathcal{S}$ is nonempty. It is also bounded below by 0.

We next show that an optimal primal solution exists. Let w be any primal-feasible vector. Define $\widetilde{w} \in \mathbb{R}^k_+$ coordinatewise by

$$\widetilde{w}_i := \min\{w_i, 1\}.$$

Then $\widetilde{w}_i \leq w_i$ for all i, so

$$\mathbf{1}^\top \widetilde{w} \leq \mathbf{1}^\top w.$$

Moreover, for every $S \in \mathcal{S}$,

$$\sum_{i \in S} \widetilde{w}_i \geq 1.$$

Indeed, if some $w_i \geq 1$ for $i \in S$, then $\widetilde{w}_i = 1$, and the sum is already at least 1; otherwise all coordinates on S satisfy $w_i < 1$, so $\widetilde{w}_i = w_i$ on S, and feasibility is preserved. Thus the primal optimum is unchanged if one restricts to the compact set

$$\big\{w \in [0,1]^k : Aw \geq \mathbf{1}\big\}.$$

By continuity of the objective, the minimum is attained.

The dual problem (I.7) is feasible, since $y = 0$ is feasible. It is bounded above because, for any dual-feasible y,

$$\sum_{S \in \mathcal{S}} y_S \leq \sum_{i=1}^{k} \sum_{S \ni i} y_S \leq k.$$

Thus the dual optimum is finite; by standard finite-dimensional LP theory it is also attained.

Since (I.6) and (I.7) are a primal-dual pair of feasible finite linear programs with finite optimal values, strong duality applies, yielding

$$\tau^*(\mathcal{S}) = \nu^*(\mathcal{S}).$$

□

The dual variables y_S assign load to histories/support sets, subject to the constraint that the total load incident to any pure controller i is at most 1. In this sense, $\nu^*(\mathcal{S})$ measures the maximum normalized amount of *history obstruction* that can be packed into the support family.

Corollary I.8 (Complementary slackness interpretation)**.** *Let $w^\star$ and $y^\star$ be primal and dual optimal solutions. Then:*

1. *if $w_i^\star > 0$, then*

$$\sum_{S \ni i} y_S^\star = 1;$$

2. *if $y_S^\star > 0$, then*

$$\sum_{i \in S} w_i^\star = 1.$$

Proof. This is the standard complementary-slackness condition for the primal-dual pair (I.6)–(I.7). □

I.4 Projective-plane example

The projective-plane family used in Section 5 also admits an exact fractional analysis.

Proposition I.9 (Fractional value for projective-plane line hypergraphs)**.** *Let $\Pi_q = (P_q, \mathcal{L}_q, \in)$ be a projective plane of order q, and let*

$$\mathscr{L}_q = \{L \subseteq P_q : L \in \mathcal{L}_q\}$$

be its line hypergraph. Then

$$\tau^*(\mathscr{L}_q) = \nu^*(\mathscr{L}_q) = \frac{q^2 + q + 1}{q + 1}. \tag{I.9}$$

Proof. For the primal problem, assign uniform weight

$$w_p := \frac{1}{q+1} \qquad \forall p \in P_q.$$

Since each line $L \in \mathscr{L}_q$ contains exactly $q + 1$ points,

$$\sum_{p \in L} w_p = 1,$$

so w is primal-feasible. Its objective value is

$$\sum_{p \in P_q} w_p = \frac{|P_q|}{q+1} = \frac{q^2+q+1}{q+1}.$$

Hence

$$\tau^*(\mathscr{L}_q) \leq \frac{q^2+q+1}{q+1}.$$

For the dual problem, assign uniform weight

$$y_L := \frac{1}{q+1} \qquad \forall L \in \mathcal{L}_q.$$

Each point $p \in P_q$ lies on exactly $q+1$ lines, so

$$\sum_{L \ni p} y_L = 1.$$

Thus y is dual-feasible. Its objective value is

$$\sum_{L \in \mathcal{L}_q} y_L = \frac{|\mathcal{L}_q|}{q+1} = \frac{q^2+q+1}{q+1}.$$

Hence

$$\nu^*(\mathscr{L}_q) \geq \frac{q^2+q+1}{q+1}.$$

Now apply Theorem I.7:

$$\tau^*(\mathscr{L}_q) = \nu^*(\mathscr{L}_q).$$

The upper and lower bounds above therefore coincide, proving (I.9). □

Corollary I.10 (Fractional value of the projective-plane realization)**.** *In the exact-rationality instance of Section 5 generated from the line hypergraph of a projective plane of order q,*

$$\tau^*(\{S_h : h \in \mathcal{H}\}) = \frac{q^2+q+1}{q+1}, \qquad m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = q+1.$$

Proof. The first identity is Proposition I.9; the second is Lemma 5.2 together with Corollary 4.2. □

Remark I.11 (Integrality gap). For the projective-plane family,

$$\frac{m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0)}{\tau^*(\mathscr{L}_q)} = \frac{q+1}{(q^2+q+1)/(q+1)} = \frac{(q+1)^2}{q^2+q+1} = 1 + \frac{q}{q^2+q+1}.$$

Thus the projective-plane obstruction separating exact memory from pairwise compatibility has only a modest fractional gap: the LP relaxation is strictly smaller than the exact complexity, but asymptotically close to it.

Remark I.12 (Total unimodularity and exactness). The gap between $\tau^*(\mathcal{S})$ and $\tau(\mathcal{S})$ is precisely an integrality phenomenon of the incidence matrix A. In particular, if the system

$$Aw \geq \mathbf{1}, \qquad w \geq 0$$

defines an integral polyhedron—for example, if the relevant incidence matrix is totally unimodular after sign normalization—then the fractional and exact retentive complexities coincide. We do not pursue such structural conditions here, but they suggest a natural bridge from bounded-memory RL to polyhedral combinatorics.

J Small Examples and Counterexamples

This section collects small finite examples that isolate the structural phenomena proved in the main text. The goal is not to add new theory, but to make the geometry of exact and approximate retentive complexity completely explicit in low-dimensional instances.

J.1 A minimal pairwise obstruction

We begin with the smallest nontrivial support family exhibiting failure of pairwise determination.

Definition J.1 (Triangle support family)**.** Let

$$\mathcal{S}_\triangle := \{\{1,2\},\{1,3\},\{2,3\}\} \subseteq 2^{[3]} \setminus \{\varnothing\}.$$

Proposition J.2 (Minimal complete-graph obstruction)**.** *There exists a one-step, fully observable mixed-controller instance with controller dictionary $\{e_1, e_2, e_3\}$ such that:*

1. *the family of exact rational-support sets is*

$$\mathcal{S}_\triangle = \{\{1,2\},\{1,3\},\{2,3\}\};$$

2. *the pairwise compatibility graph is the complete graph K_3;*
3. *the exact retentive complexity is*

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = 2.$$

Proof. By Theorem 4.3, there exists a one-step instance realizing the support family $\mathcal{S}_\triangle$. Let the attainable histories be h_{12}, h_{13}, h_{23}, with

$$S_{h_{12}} = \{1,2\}, \qquad S_{h_{13}} = \{1,3\}, \qquad S_{h_{23}} = \{2,3\}.$$

Every two sets in $\mathcal{S}_\triangle$ intersect, so

$$S_{h_{12}} \cap S_{h_{13}} = \{1\} \neq \varnothing, \quad S_{h_{12}} \cap S_{h_{23}} = \{2\} \neq \varnothing, \quad S_{h_{13}} \cap S_{h_{23}} = \{3\} \neq \varnothing.$$

Hence the pairwise compatibility graph is K_3.

To compute the exact retentive complexity, apply Corollary 4.2. Thus it suffices to compute $\tau(\mathcal{S}_\triangle)$. No singleton hits all three sets, because

$$\{1\} \cap \{2,3\} = \varnothing, \qquad \{2\} \cap \{1,3\} = \varnothing, \qquad \{3\} \cap \{1,2\} = \varnothing.$$

Hence $\tau(\mathcal{S}_\triangle) \geq 2$. On the other hand, $\{1,2\}$ hits all three sets, so $\tau(\mathcal{S}_\triangle) \leq 2$. Therefore

$$\tau(\mathcal{S}_\triangle) = 2,$$

and so

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(0) = 2.$$

□

Remark J.3. Proposition J.2 is the smallest complete-graph obstruction in the sense that it uses only three histories and three pure controllers. It is the finite toy version of the projective-plane separation in Section 5.

J.2 A tiny integrality gap

The same triangle family also exhibits the smallest nontrivial fractional gap.

Proposition J.4 (Fractional value of the triangle family)**.** *For the support family $\mathcal{S}_\triangle$,*

$$\tau(\mathcal{S}_\triangle) = 2, \qquad \tau^*(\mathcal{S}_\triangle) = \frac{3}{2}.$$

Proof. The exact value $\tau(\mathcal{S}_\triangle) = 2$ was proved in Proposition J.2.

For the fractional value, consider the primal LP

$$\tau^*(\mathcal{S}_\triangle) = \min\left\{ w_1 + w_2 + w_3 : \begin{array}{c} w_1 + w_2 \geq 1, \\ w_1 + w_3 \geq 1, \\ w_2 + w_3 \geq 1, \\ w_1, w_2, w_3 \geq 0 \end{array} \right\}.$$

The uniform vector

$$w_1 = w_2 = w_3 = \frac{1}{2}$$

is feasible, so

$$\tau^*(\mathcal{S}_\triangle) \leq \frac{3}{2}.$$

Conversely, summing the three covering constraints gives

$$2(w_1 + w_2 + w_3) \geq 3,$$

and therefore

$$w_1 + w_2 + w_3 \geq \frac{3}{2}$$

for every feasible w. Hence

$$\tau^*(\mathcal{S}_\triangle) = \frac{3}{2}.$$

□

Remark J.5. The triangle family therefore has integrality gap

$$\frac{\tau(\mathcal{S}_\triangle)}{\tau^*(\mathcal{S}_\triangle)} = \frac{2}{3/2} = \frac{4}{3}.$$

This is the smallest explicit example in the paper where exact memory exceeds its fractional relaxation.

J.3 A sharp threshold for one-state memory

We now turn to positive tolerance $\varepsilon > 0$. The next two-history example shows that approximate retentive complexity can exhibit a sharp phase transition.

Proposition J.6 (Sharp threshold for one-state approximate memory)**.** *Consider the mixed-controller model with $k = 2$, controller simplex Δ_2, external deviations Φ_{ext}, and two histories h_1, h_2 with value vectors*

$$v(h_1) = (1, 0), \qquad v(h_2) = (0, 1).$$

Then, for every $\varepsilon \geq 0$,

$$K^\varepsilon_{h_1} = \{x \in \Delta_2 : x_1 \geq 1 - \varepsilon\}, \qquad K^\varepsilon_{h_2} = \{x \in \Delta_2 : x_1 \leq \varepsilon\}.$$

Consequently,

$$m^{\text{pc}}_{\Phi_{\text{ext}}}(\varepsilon) = \begin{cases} 2, & 0 \leq \varepsilon < \frac{1}{2}, \\ 1, & \varepsilon \geq \frac{1}{2}. \end{cases} \tag{J.1}$$

Proof. By Theorem 4.1,

$$K^{\varepsilon}_{h_1} = \{x \in \Delta_2 : \langle (1,0), x\rangle \geq 1-\varepsilon\} = \{x \in \Delta_2 : x_1 \geq 1-\varepsilon\},$$

and similarly

$$K^{\varepsilon}_{h_2} = \{x \in \Delta_2 : \langle (0,1), x\rangle \geq 1-\varepsilon\} = \{x \in \Delta_2 : x_2 \geq 1-\varepsilon\} = \{x \in \Delta_2 : x_1 \leq \varepsilon\}.$$

Their intersection is nonempty iff there exists $x_1 \in [0,1]$ satisfying

$$1-\varepsilon \leq x_1 \leq \varepsilon,$$

which is equivalent to $\varepsilon \geq \frac{1}{2}$. By Corollary 6.1, this is equivalent to

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(\varepsilon) = 1 \quad \Longleftrightarrow \quad \varepsilon \geq \frac{1}{2}.$$

It remains to show that, when $\varepsilon < \frac{1}{2}$, the exact value is 2. Since the total intersection is empty, one-state memory is impossible. On the other hand,

$$e_1 \in K^{\varepsilon}_{h_1}, \qquad e_2 \in K^{\varepsilon}_{h_2} \qquad \forall \varepsilon \geq 0,$$

so the set $\{e_1, e_2\}$ is a transversal for $\{K^{\varepsilon}_{h_1}, K^{\varepsilon}_{h_2}\}$. By Theorem 3.1,

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(\varepsilon) \leq 2 \qquad \forall \varepsilon \geq 0.$$

Hence $m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}(\varepsilon) = 2$ for $0 \leq \varepsilon < \frac{1}{2}$, proving (J.1). □

Remark J.7. At the threshold $\varepsilon = \frac{1}{2}$, the total intersection is the singleton

$$K^{1/2}_{h_1} \cap K^{1/2}_{h_2} = \left\{ \left(\frac{1}{2}, \frac{1}{2}\right) \right\}.$$

Thus the one-state feasible continuation appears through a single barycentric point of the simplex.

J.4 Mixed controllers can strictly help at positive tolerance

The previous example also shows that mixed continuations are not merely a technical convenience: they can strictly reduce approximate memory complexity.

Definition J.8 (Pure-restricted approximate retentive complexity)**.** Fix a finite-controller mixed-continuation instance with dictionary $\{e_1, \ldots, e_k\}$. Its *pure-restricted approximate retentive complexity* is

$$m^{\mathrm{pc}}_{\mathrm{pure}}(\varepsilon) := \inf \left\{ m \in \mathbb{N} : \begin{array}{l} \exists\, (\sigma, d) \text{ of size } m \text{ with } d : [m] \to \{e_1, \ldots, e_k\}, \\ d(\sigma(h)) \in K^{\varepsilon}_h \ \forall h \in \mathcal{H} \end{array} \right\}.$$

Proposition J.9 (Strict mixed-controller advantage)**.** *In the two-history instance of Proposition J.6,*

$$m^{\mathrm{pc}}_{\Phi_{\mathrm{ext}}}\left(\frac{1}{2}\right) = 1, \qquad m^{\mathrm{pc}}_{\mathrm{pure}}\left(\frac{1}{2}\right) = 2.$$

Proof. The first identity is Proposition J.6. Indeed, the mixed controller

$$x^{\star} = \left(\frac{1}{2}, \frac{1}{2}\right)$$

belongs to both $K^{1/2}_{h_1}$ and $K^{1/2}_{h_2}$, so one-state memory is feasible.

For the pure-restricted problem, note that

$$e_1 \notin K_{h_2}^{1/2}, \qquad e_2 \notin K_{h_1}^{1/2},$$

because

$$K_{h_1}^{1/2} = \left\{x \in \Delta_2 : x_1 \geq \frac{1}{2}\right\}, \qquad K_{h_2}^{1/2} = \left\{x \in \Delta_2 : x_1 \leq \frac{1}{2}\right\}.$$

Hence no single pure controller belongs to both low-regret sets, so $m_{\text{pure}}^{\text{pc}}(1/2) \neq 1$.

On the other hand, the pure controllers e_1 and e_2 satisfy

$$e_1 \in K_{h_1}^{1/2}, \qquad e_2 \in K_{h_2}^{1/2},$$

so a two-label pure architecture exists. Therefore

$$m_{\text{pure}}^{\text{pc}}\left(\frac{1}{2}\right) = 2.$$

□

Remark J.10. At exact rationality $\varepsilon = 0$, the face-piercing reduction in Section 4 shows that mixed continuations do not reduce memory beyond what can already be achieved by selecting vertices of the simplex. The strict separation in Proposition J.9 is therefore a genuinely positive-tolerance phenomenon.

J.5 A small complete-graph example below the Helly threshold

The final example ties together the pairwise and Helly phenomena.

Proposition J.11 (Complete pairwise graph does not imply one-state memory)**.** *For the triangle family $\mathcal{S}_\triangle$, every pair of exact low-regret sets intersects, but the total intersection is empty:*

$$\Delta(\{1,2\}) \cap \Delta(\{1,3\}) \neq \varnothing, \qquad \Delta(\{1,2\}) \cap \Delta(\{2,3\}) \neq \varnothing, \qquad \Delta(\{1,3\}) \cap \Delta(\{2,3\}) \neq \varnothing,$$

whereas

$$\Delta(\{1,2\}) \cap \Delta(\{1,3\}) \cap \Delta(\{2,3\}) = \varnothing.$$

Proof. The pairwise intersections are

$$\begin{aligned}
\Delta(\{1,2\}) \cap \Delta(\{1,3\}) &= \{e_1\}, \\
\Delta(\{1,2\}) \cap \Delta(\{2,3\}) &= \{e_2\}, \\
\Delta(\{1,3\}) \cap \Delta(\{2,3\}) &= \{e_3\}.
\end{aligned}$$

all of which are nonempty.

Now let $x \in \Delta(\{1,2\}) \cap \Delta(\{1,3\}) \cap \Delta(\{2,3\})$. Then simultaneously

$$\operatorname{supp}(x) \subseteq \{1,2\}, \qquad \operatorname{supp}(x) \subseteq \{1,3\}, \qquad \operatorname{supp}(x) \subseteq \{2,3\}.$$

Hence

$$\operatorname{supp}(x) \subseteq \{1,2\} \cap \{1,3\} \cap \{2,3\} = \varnothing,$$

which is impossible for a point $x \in \Delta_3$. Therefore the triple intersection is empty. □

Remark J.12. Proposition J.11 is consistent with Corollary 6.1: the simplex Δ_3 has affine dimension 2, so one-state memory is controlled by intersections of subfamilies of size at most 3, not merely by pairwise intersections.