ONE EDGE IS ENOUGH: ERROR-PRESERVING REPAIR FOR STOCHASTIC FIXED-POINT ITERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Accurate estimates of a nonexpansive map need not produce stable iterates, because estimation can destroy consistency between successive queries. We show that repairing one adjacent query/reply constraint suffices to preserve estimation accuracy throughout Halpern iteration. On a cube, our rule clips each fresh absolute map-value estimate against the previous reply and preserves the largest input error with constant one. This invariant yields residual $4D/(k+2)+M$ with $O(d)$ recurrence state, where D is the domain diameter and M the largest input error. Conditional robust estimation controls M, giving fixed-budget guarantees without knowledge of the noise scale or tail exponent. The coordinate geometry is essential: we prove that universal lossless repair on norm balls is possible exactly in maximum-norm spaces, even with randomization. For contractions, the same clipping argument links inherited-error decay to query travel, yielding a minimum-motion extension with geometric recovery and no additional oracle calls. Matched-cost experiments isolate stabilization across 36 settings with two estimators, while Bellman studies demonstrate recovery from corrupted replies and value-accuracy gains under variance bursts. One stored edge thus provides a low-state interface through which statistical accuracy supports stable iteration and recovery.

1 INTRODUCTION

Fixed-point algorithms underpin methods in monotone optimization and reinforcement learning. Exact Halpern iteration attains an $O(1/k)$ residual rate and motivates accelerated algorithms for these settings (Halpern, 1967; Sabach & Shtern, 2017; Contreras & Cominetti, 2023; Diakonikolas, 2020; Yoon & Ryu, 2021; Lee & Ryu, 2023; 2025). Stochastic observations can break the relation underlying this stability: a nonexpansive map sends nearby queries to nearby replies, but separately estimated values may violate that relation even when each estimate is accurate.

This discrepancy creates a gap between estimation accuracy and iterative accuracy. On a cube rotation, value errors bounded by M can produce a Halpern residual of $kM/3$. Taking $M = 1/(k+3)$ makes every input error vanish as the horizon grows, yet the output residual approaches $1/3$ (Appendix A.1). The algorithm must therefore control how its estimates fit together along the query path.

Minibatching and variance reduction address this challenge through statistical error control (Cai et al., 2022; 2024; Bravo & Contreras, 2026). Clipping stochastic differences can also enforce pathwise consistency (Diakonikolas, 2026). We study the corresponding structural question for fresh absolute map-value estimates: *what is the least history needed to restore stability while preserving the accuracy already supplied by the estimator?* Enforcing consistency against a noisy earlier reply appears to create a conflict, because the true current value can lie outside the admissible region.

Our key observation is that one adjacent edge resolves this conflict in the maximum norm. Clip the fresh estimate to the ball centered at the preceding reply, with radius equal to the distance between the two queries, and retract onto the domain. Coordinate endpoints preserve the historical error bound even when the true value lies outside that ball. Halpern's residual becomes $4D/(k+2)+M$: the largest estimation error enters once, regardless of the horizon. Because this guarantee holds pathwise, statistical estimation bounds transfer directly to the repaired iteration (Figure 1).

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

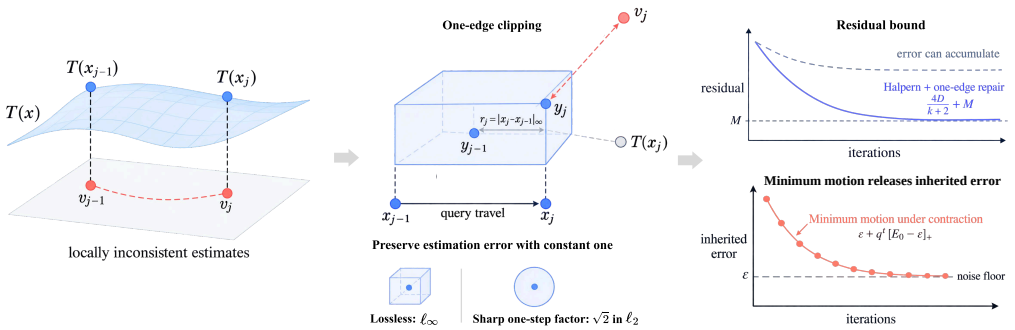


Figure 1: **One edge preserves estimation accuracy and stabilizes iteration.** Left: noisy absolute estimates can violate local consistency. Center: clipping against the previous repaired reply, with radius $r_j = \|x_j - x_{j-1}\|_\infty$, followed by domain retraction preserves the largest input error with constant one. Universal lossless repair on norm balls characterizes maximum-norm geometry; the sharp Euclidean one-step factor is $\sqrt{2}$. Right: schematic residual and recovery bounds. Halpern's residual is bounded by $4D/(k+2) + M$, where M is the largest input error. Under contraction, minimum motion gives geometric recovery toward ε with $q = (1+\gamma)/2$; recovery time starts at a stage after which fresh errors remain at most ε.

The invariant reveals how geometry and query selection shape stability. Its coordinate structure leads to a characterization of the norms that permit lossless repair. Its historical maximum can retain an early error even after estimates improve; under contraction, we show that query travel makes this error decay. These consequences lead to three contributions:

- **Error-preserving stabilization.** One adjacent query/reply constraint gives a constant-one maximum-error guarantee and a low-state Halpern residual bound. Conditional median-of-means estimation yields fixed-budget guarantees without tail parameters, matching the class-aware high-dimensional sample complexity up to a logarithmic factor for each fixed $1 < p \leq 2$.
- **A characterization of lossless repair.** We reduce universal causal repair to ball intersection and characterize its possible ambient norms, including randomized rules. In Euclidean geometry, the sharp one-step amplification coefficient is $\sqrt{2}$.
- **Recovery by controlling query travel.** We derive a travel-dependent error bound and a minimum-motion rule that releases inherited error at a geometric rate while retaining constant-one repair. A sharp pulse bound shows how a dense disturbance becomes a one-coordinate query displacement, yielding inverse-dimension attenuation under doubly stochastic dynamics.

Matched-estimator experiments show stabilization across all 36 settings with both ordinary and robust estimates. Bellman studies then separate recovery from disturbance rejection and measure their effect on value accuracy under variance bursts.

2 Related Work

Anchoring and stochastic fixed-point methods. Halpern's residual rates (Leuştean, 2007; Sabach & Shtern, 2017; Cheval & Leuştean, 2024; Lieder, 2021; Contreras & Cominetti, 2023) support proximal, monotone, minimax and Bellman methods (Kim, 2021; Diakonikolas, 2020; Yoon & Ryu, 2021; Park & Ryu, 2022; Lee & Ryu, 2023; 2025), alongside averaging approaches (Cominetti et al., 2014; Bravo & Cominetti, 2024). Stochastic Halpern methods control accumulated noise through minibatching (Bravo & Contreras, 2026) or variance reduction (Cai et al., 2022; 2024); related estimators reuse sampled information in variational inequalities and optimization (Alacaoglu & Malitsky, 2022; Nguyen et al., 2017; Fang et al., 2018; Li et al., 2021). Inexact Halpern analysis prescribes weighted error tolerances under cocoercivity (Liang et al., 2025, Theorem 3.1). Dual anchoring obtains stochastic guarantees from expected cocoercivity or square-nonexpansivity (Yoon & Loizou, 2026). Our repair uses maximum-norm nonexpansivity of the mean map to make the largest

absolute estimation error enter the residual with coefficient one. This separates structural stability from the estimator's statistical analysis.

Repairing observations along the query path. Consistency can also be restored by modifying noisy observations. VR-GHAL enforces pathwise Lipschitzness by adding averaged radially clipped stochastic differences to a recursive estimator (Diakonikolas, 2026, Eqs. (9)–(10), Lemma 5). Its differences sample both adjacent queries, with clipping bias and martingale fluctuations controlled under finite variance in quadratically smoothable spaces. COCO jointly denoises gradient observations under pairwise cocoercivity constraints (Madeira et al., 2021). One-edge repair acts causally on a fresh absolute map-value estimate: it holds the emitted reply fixed and uses a coordinate endpoint invariant to preserve the historical maximum error. Thus one adjacent constraint supports Halpern stability without resampling the stored endpoint.

Geometry of consistent corrections. Property-preserving and Lipschitz reconstruction produce globally valid corrections (Ailon et al., 2008; Jha & Raskhodnikova, 2013), with geometry-dependent costs (Awasthi et al., 2015; Lange et al., 2024). Extension and noisy interpolation likewise concern consistent functions (Ciosmak, 2021; Huang et al., 2023). For iteration, adjacent consistency can suffice even when the full transcript has no nonexpansive interpolant. Our two-query reduction characterizes the geometry of this local guarantee by linking universal causal repair to compatible ball intersection; Nachbin's classical theorem then identifies the ambient norm (Nachbin, 1950; Pavón, 2014). Nonexpansive retractions connect this geometry to domains beyond boxes (Aronszajn & Panitchpakdi, 1956; Descombes & Pavón, 2017).

Robust estimation and oracle complexity. Mean estimators control heavy-tailed observations under weak moments (Catoni, 2012; Devroye et al., 2016; Lugosi & Mendelson, 2019a). Geometric and multivariate medians address other geometries (Minsker, 2015; Lugosi & Mendelson, 2019b; Hopkins, 2020; Hsu & Sabato, 2016), while martingale truncation handles dependent noise with infinite variance (Whitehouse et al., 2024). Robust estimation also supports risk minimization, bandits and stochastic gradients (Brownlees et al., 2015; Bubeck et al., 2013; Cutkosky & Mehta, 2021; Hübler et al., 2025; Liu & Zhou, 2025). Here the estimator supplies the accuracy that repair preserves. Coordinate median-of-means fits the maximum norm, and successive conditioning controls its blocks after adaptive histories. Recursive anchoring also gives a near-optimal finite-variance upper bound: on unit maximum-norm balls, Diakonikolas et al. (2026, Theorem 5, Corollary 6) attain $\widetilde{O}(\varepsilon^{-1} + \nu^2\varepsilon^{-3})$ calls, suppressing logarithmic factors. Their construction uses independent mini-batches, a known variance bound and a Lipschitz retraction, and extends to other norm geometries. Our cube algorithm attains the same precision exponent at $p = 2$ through a constant-one, pathwise interface for absolute estimates. Its fixed-budget allocation takes neither the noise scale nor p, and successive conditioning covers martingale noise with $1 < p \leq 2$ moments. Extending their sign-chain lower-bound construction to weak moments then establishes the class-wise comparison in Corollary 4.

3 One-edge repair: mechanism and guarantees

One-edge repair preserves estimation accuracy while restoring consistency. This invariant controls Halpern residuals; under contraction, query motion further releases inherited error. We develop these mechanisms before characterizing the geometry that makes lossless repair possible.

Setting. Let $C \subset \mathbb{R}^d$ be nonempty and bounded, with maximum-norm diameter D. Assume a specified nonexpansive retraction $R : \mathbb{R}^d \to C$: it fixes each point of C and satisfies $\|R(u) - R(v)\| \leq \|u - v\|$. Every unmarked norm is the maximum norm. We study a nonexpansive self-map $T : C \to C$ and the residual $\mathrm{res}_T(x) = \|x - T(x)\|$. Queries $x_j \in C$ and fresh estimates $v_j \in \mathbb{R}^d$ may be adaptive. Retraction allows nonconvex C; on $[0, 1]^d$, take coordinate clipping $R = P_C$ and $D = 1$.

3.1 ONE-EDGE REPAIR AND ERROR PRESERVATION

Nonexpansiveness limits reply changes by query travel. We enforce this relation by clipping each fresh estimate around the previous emitted reply, which stays fixed. Let $Q_{a,r}$ clip coordinate i to $[a_i - r, a_i + r]$. Initialize $y_0 = R(v_0)$ and set

$$r_j = \|x_j - x_{j-1}\|, \qquad y_j = R\big(Q_{y_{j-1},r_j}(v_j)\big), \quad j \geq 1. \tag{1}$$

Retraction restores domain membership, so on the cube the composition clips to $C \cap B_\infty(y_{j-1}, r_j)$ (Figure 1, center). Although the box is centered at a noisy reply, enforcing this constraint preserves the largest input error with constant one:

Lemma 1 (Error-preserving repair on a retract)**.** *For every nonexpansive $T : C \to C$ and every finite query/estimate sequence, the replies in* (1) *belong to C and satisfy*

$$\|y_j - y_{j-1}\| \leq \|x_j - x_{j-1}\|, \qquad \max_{j \leq k} \|y_j - T(x_j)\| \leq \max_{j \leq k} \|v_j - T(x_j)\|.$$

The statement holds pathwise for adaptive sequences and uses no numerical error bound as input.

The reason is coordinatewise: if the previous reply and fresh estimate have errors at most M, nonexpansiveness ensures that each clipping interval intersects the current M-error band. Clipping an estimate already in that band keeps it there, even when the true value lies outside the box. Retraction preserves both this bound and the adjacent constraint (Appendix A).

3.2 HALPERN STABILITY AND STATISTICAL EFFICIENCY

The two parts of Lemma 1 control different components of the residual: adjacent consistency controls the gap between a query and its reply, while error preservation controls the gap from that reply to the true map value. To exploit this decomposition, fix an anchor $x_0 \in C$. At stages $j = 0, \ldots, k$, obtain v_j and its local repair y_j. For $j < k$, use the retracted anchor update

$$x_{j+1} = R\left(\frac{2}{j+3} x_0 + \frac{j+1}{j+3} y_j\right). \tag{2}$$

Return x_k after obtaining its reply, using exactly $k + 1$ approximate-value queries. On a convex domain, R fixes the anchor mixture, recovering the classical recurrence.

Theorem 2 (Halpern stability on a retract)**.** *For every $k \geq 1$ and every realized finite history,*

$$\|x_k - y_k\| \leq \frac{4D}{k+2}, \qquad \mathrm{res}_T(x_k) \leq \frac{4D}{k+2} + M, \quad M = \max_{0 \leq j \leq k} \|v_j - T(x_j)\|.$$

If evaluating R costs c_R operations and s_R auxiliary space, the recurrence uses $O(d + s_R)$ state and $O(k(d + c_R))$ work, excluding estimator workspace and the fixed representation of R.

Under (2), adjacent consistency bounds the query increments by $\|x_j - x_{j-1}\| \leq 2Dj/((j+1)(j+2))$. Comparing the next anchor mixture to y_k then gives $\|x_k - y_k\| \leq 4D/(k+2)$. Adding the reply error from Lemma 1 gives the residual bound, so estimation error enters once, regardless of the horizon (Appendix A).

Statistical efficiency without tail inputs. Theorem 2 reduces stochastic residual control to a single statistical task: bounding the largest input error along an adaptive trajectory. Conditional median-of-means estimation supplies this bound using a fixed sampling budget, without taking the noise scale or tail exponent as inputs.

Raw calls return $Y_t = T(x_t) + \zeta_t$ at predictable queries, with integrable noise and $\mathbb{E}[\zeta_t \mid \mathcal{F}_{t-1}] = 0$; the filtration includes prior calls and algorithmic randomness. After every full starting history, each coordinate of a B-sample mean at a frozen query has conditional expected absolute error at most a common deterministic $b_B < \infty$. Appendix B.2 derives this premise from conditional moments; the algorithm does not receive b_B.

Specialize to $C = [0,1]^d$, $R = P_C$ and $D = 1$, and split a raw-call budget N into equal blocks. For $d, N \geq 1$ and $a, \eta \in (0,1)$, choose

$$k = \lceil 4/a \rceil, \quad K = k+1, \quad L = \min\{l \geq 3 : l \text{ odd}, \, l \geq 2\log_2(dK/\eta)\}, \quad B = \lfloor N/(KL) \rfloor. \tag{3}$$

For $B \geq 1$, each v_j is the coordinate median of L successive B-sample means at x_j (Lugosi & Mendelson, 2019a), followed by repair and (2). Here a allocates accuracy to the iteration term: reducing it increases the number of queries and leaves fewer samples per estimate. Algorithm 1 includes the zero-sample branch, which returns the center with residual at most $1/2$.

Theorem 3 (Conditional mean-error transfer)**.** *On the branch $B \geq 1$, the full-history block-mean premise gives*

$$\Pr\{\mathrm{res}_T(x_k) \leq a + 8b_B\} \geq 1 - \eta, \qquad \mathbb{E}\,\mathrm{res}_T(x_k) \leq a + 16b_B.$$

The method uses exactly $KLB \leq N$ raw vectors. Streaming and deterministic median selection require $O(dL)$ workspace and $O(dN)$ response-processing work.

Successive conditioning controls block failures after adaptive histories. A union bound over coordinates and queries controls the largest input error, and integrating its tail gives the expectation bound (Appendix B.3).

To quantify the sample cost, suppose $\mathbb{E}[\|\zeta_t\|^p \mid \mathcal{F}_{t-1}] \leq \nu^p$, $1 < p \leq 2$. Then $b_B \leq 3\nu B^{-(p-1)/p}$ (Appendix B.2), and taking $a = \varepsilon/2$ converts the transfer bound into a sufficient budget. Comparison with a class-aware lower bound then quantifies the cost of unknown moment parameters: for fixed p in the regime below, a logarithmic factor separates the budgets.

Let $N^*_{d,p,\nu}(\varepsilon, \eta)$ be the least hard raw-call cap attaining residual ε with failure probability at most η uniformly over nonexpansive cube maps with conditionally centered noise satisfying $\mathbb{E}[\|\zeta_t\|^p \mid \mathcal{F}_{t-1}] \leq \nu^p$. Algorithms access the map through the oracle, query and output in the cube, and may know p, ν.

Corollary 4 (No polynomial cost for unknown moment order)**.** *Fix $1 < p \leq 2$, $0 < \nu \leq 1$, $0 < \varepsilon < 1/128$ and $0 < \eta < 1/4$. Set $s = p/(p-1)$ and $m = \lfloor 1/(32\varepsilon) \rfloor$. If $d \geq m$, then*

$$\frac{m}{64} \max\{1, (\nu/(8\varepsilon))^s\} < N^*_{d,p,\nu}(\varepsilon, \eta) \leq KL \left\lceil \max\{1, (48\nu/\varepsilon)^s\} \right\rceil,$$

where $K = \lceil 8/\varepsilon \rceil + 1$ and L is as in (3)*. The same local-repair algorithm, receiving only d, N, ε, η, attains the upper bound for every p, ν. For each fixed p, the two bounds differ by at most $O_p(\log(ed/(\varepsilon\eta)))$.*

The lower bound embeds the sign-chain geometry of Diakonikolas et al. (2026) in a weak-moment oracle (Appendix B.4). Constants can grow as $p \downarrow 1$; the sufficient cap depends analytically on the unknown moment parameters, while the algorithm runs to its fixed budget.

3.3 Recovery and pulse attenuation

The preceding guarantees protect the trajectory against error accumulation. When fresh estimates improve, recovery further requires the influence of an early error to decay. A repeated query has zero repair radius and can retain its old reply even with exact new estimates (Appendix A.1). For a γ-contractive map, query travel provides the missing decay. Write $e_j = \|v_j - T(x_j)\|$, $E_j = \|y_j - T(x_j)\|$ and $s_j = \|x_j - x_{j-1}\|$. The endpoint argument behind Lemma 1 now gives

$$E_j \leq \max\{e_j, [E_{j-1} - (1-\gamma)s_j]_+\}, \qquad E_k \leq \max_{0 \leq i \leq k} \left[e_i - (1-\gamma) \sum_{j=i+1}^{k} s_j \right]_+. \tag{4}$$

The distance from $T(x_j)$ to the clipping box is at most $[E_{j-1} - (1-\gamma)s_j]_+$; clipping and retraction give the first inequality, and induction gives the second. Travel therefore discounts inherited error.

Minimum-motion update. This bound suggests choosing the next query far enough away to discount inherited error. Since that error is unknown, we use the observable discrepancy between the fresh estimate and its repaired reply to set the required travel.

On $C = [0, D]^d$, let $b_t = P_C v_t$, repair it by (1) to obtain y_t, and measure $d_t = \|b_t - y_t\|$. Given a base proposal $h_t \in C$, choose

$$x_{t+1} \in \operatorname*{argmin}_{x \in C:\, \|x - x_t\| \geq d_t/2} \|x - h_t\|. \tag{5}$$

This rule minimally perturbs a base proposal while enforcing the movement floor $d_t/2$. Keep a feasible h_t; otherwise change one coordinate to a feasible boundary $x_{t,i} \pm d_t/2$, minimizing the change in $O(d)$ work (Appendix C.2). We use either the repaired proposal $h_t = (y_t + y_{t-1})/2$ or the fresh-estimate proposal $h_t = (b_t + b_{t-1})/2$, initializing each with its current value. After $K \geq 2$ queries, both policies return the last repaired average $(y_{K-1} + y_{K-2})/2$. Fresh proposals guide queries with new estimates; repaired proposals retain the consistency correction. Both policies retain $O(d)$ recurrence state and use one fresh estimate per stage; movement adds no oracle call.

Theorem 5 (Recovery on the original query stream)**.** *For a γ-contractive cube self-map, $0 \leq \gamma < 1$, set $e_t = \|b_t - T(x_t)\|$, $E_t = \|y_t - T(x_t)\|$ and $q = (1+\gamma)/2$. The unchanged one-edge repair with* (5) *satisfies*

$$E_{t+1} \leq \max\{e_{t+1}, qE_t + (1-q)e_t\}. \tag{6}$$

All physical query/reply edges remain consistent and the historical maximum error retains coefficient one. If $e_i \leq \epsilon$ for $i \geq m$, then $E_{m+k} \leq \epsilon + q^k[E_m - \epsilon]_+$. Consequently $e_t \to 0$ implies convergence to the fixed point for either query policy above. The movement rule requires no extra query or contraction parameter.

The link between the update and the theorem is $d_t \geq [E_t - e_t]_+$: the observable correction bounds the inherited error exceeding the fresh error. Substituting travel of at least $d_t/2$ into (4) yields the geometric factor q. The perturbation to either base proposal vanishes with the estimation errors, giving convergence of its contraction recursion. Appendix C gives the recovery proof and finite-budget propagation.

From recovery to pulse attenuation. Minimum motion also controls how disturbances enter future queries: repair can reject a dense error while movement changes only one coordinate. For affine contractions, this sparse displacement yields a dimension-dependent gain over raw adjacent averaging.

Corollary 6 (Dimension-dependent pulse attenuation)**.** *Let $T(x) = v^* + \gamma P(x - v^*)$ be a cube self-map with doubly stochastic P. Start raw adjacent averaging and repaired-proposal minimum motion at $x_0 = y_0 = v^*$. Give both a feasible completed-estimate pulse $b_1 = v^* + a\xi$, $a > 0$ and $\xi \in \{-1, 1\}^d$, followed by an exact estimate at each method's actual query. After these three estimates,*

$$\frac{\|h_2^{\mathrm{rep}} - v^*\|_2^2}{\|h_2^{\mathrm{raw}} - v^*\|_2^2} \leq \frac{\gamma^2}{d(2-\gamma)^2}.$$

The constant is sharp for every even d. For independent symmetric signs, the ratio of expected squared errors is at most $\gamma^2/(4d)$.

In this construction, the repeated first query has zero repair radius, so $y_1 = v^*$. Minimum motion sends x_2 to $v^* + sae_j/2$ for one feasible coordinate and sign. Consequently the two output errors are $sa\gamma Pe_j/4$ and $a(2I+\gamma P)\xi/4$: repair transports one column of P, whereas raw averaging transports the entire pulse. Appendix C.4 proves the bound and extends it to nonuniform state occupancy and background errors; Appendix C.5 treats approximate initialization and later outputs.

3.4 The geometry required by lossless repair

Both mechanisms rely on preserving compatible coordinate error bands. To determine whether lossless correction requires this geometry, we now allow any causal repair rule. The following characterization shows that maximum-norm geometry is essential to the universal guarantee.

A causal rule uses only C and the observed query/estimate/reply history, receives neither T nor an error radius, and returns $y_j \in C$. Losslessness means adjacent consistency, $\|y_j - y_{j-1}\| \leq \|x_j - x_{j-1}\|$, and the prefix bound

$$\max_{h \leq j} \|y_h - T(x_h)\| \leq \max_{h \leq j} \|v_h - T(x_h)\|. \tag{7}$$

Randomized rules use fixed measurable functions of this history and a global seed independent of T; expectations cover seed and input randomness.

Theorem 7 (Lossless repair characterizes maximum-norm geometry)**.** *A real normed space of dimension $d \geq 1$ admits universal deterministic causal lossless repair on every closed norm ball if and*

only if it is linearly isometric to ℓ_∞^d. The guarantee must hold for every nonexpansive self-map and estimate sequence. Necessity uses only prefixes $j = 0, 1$ with arbitrary x_0 and $x_1 = (2x_0 + y_0)/3$. It also holds for randomized rules with almost-sure adjacent consistency and expectations on both sides of (7)*. In maximum-norm geometry, the deterministic one-edge rule is lossless pathwise at every finite horizon.*

The necessity argument reduces this geometric question to two queries. An exact constant-map transcript forces the initial reply. Different nonexpansive maps can then produce the same next observed estimate but different true targets. Preserving every such target lets the rule add one ball constraint without losing earlier compatible constraints. Repeating this operation gives the binary intersection property, which classically characterizes finite-dimensional maximum-norm spaces (Nachbin, 1950, Theorem 3); see also Pavón (2014). Appendix D.1 supplies the witness maps and full reduction.

Within maximum-norm geometry, the retraction in (1) extends the rule to coupled and nonconvex domains (Aronszajn & Panitchpakdi, 1956; Descombes & Pavón, 2017); Appendix D.2 constructs ordered-vector and branching examples.

The cost of changing geometry. In Euclidean norm, two maps with identical observed estimates can force an error amplification approaching $\sqrt{2}$ at the second reply (Figure 1, geometric comparison). Projection attains this coefficient for one repair when the previous repaired error and current raw error are each at most M. Theorem 12 and Lemma 13 give the sharp statements in Appendix D.3; a full-trajectory coefficient requires a separate invariant.

4 Experiments

We test whether one edge stabilizes a shared estimator, whether query motion releases an inherited error, and when these mechanisms improve value accuracy. Comparisons match actual sample cost. Halpern runs return x_k; contraction runs return the last repaired average.

4.1 Does one edge stabilize iteration?

Setup. Three repairs share Halpern iteration and the 81-query allocation ($a = \eta = 0.05$): range-only clipping, one edge and full history. We cross shift/cycle maps, dimensions $8, 32, 128$, budgets 20k/100k and three noise laws, with 32 seeds per cell. Bounded-skew, Pareto-1.3 and log-only noise span bounded observations through finite first moments with no higher moment. Coordinate median-of-means (MoM) and ordinary means use the same samples (Appendix E.1).

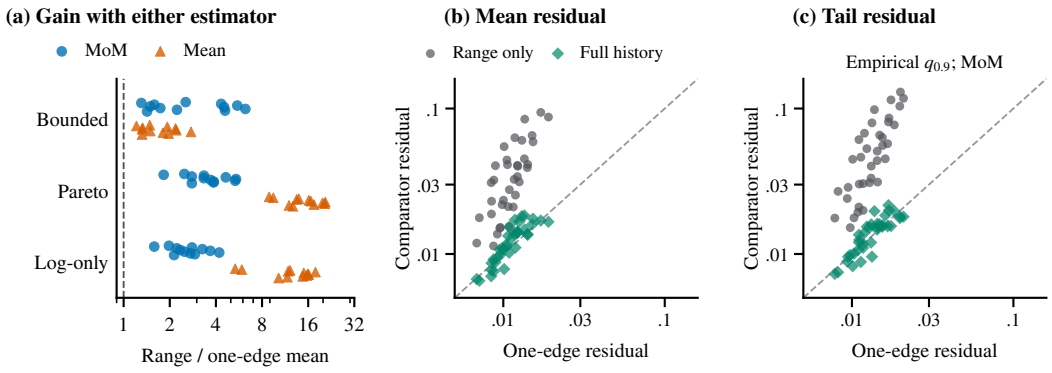


Figure 2: **One edge stabilizes iteration without retaining the full history.** (a) Each dot is a range-only/one-edge mean-residual ratio; each noise law has 12 settings per estimator. (b)–(c) Absolute mean and empirical 90th-percentile residuals with MoM: each comparator has 36 points, with points above the diagonal favoring one edge. Every cell uses 32 paired seeds and matched calls. Full setting identities and statistics are in Appendix E.1.

Structural benefit. One-edge repair lowers mean residual relative to range-only clipping in all 36 settings with either estimator. Figure 2a shows range-only/one-edge ratios of 1.31–6.23 with MoM and 1.21–20.68 with ordinary means. Simultaneous 95% intervals place the probability of a lower residual with one edge above $1/2$ in every setting (Appendix E.1.3).

How much history is useful? Figure 2b–c compares absolute residuals. Against full-history repair with MoM, one edge improves 23/36 means and 19/36 upper quantiles. The recurrence state falls from $O(dk)$ to $O(d)$ and total repair work from $O(dk^2)$ to $O(dk)$. One edge therefore supplies the structural guarantee at lower storage and repair cost, with competitive residual accuracy (comparator: Appendix E.1.1).

Repair and estimator choice. With one edge, ordinary/MoM mean-residual ratios are 0.61–0.99 under bounded noise, 1.13–1.43 under Pareto noise and 1.14–1.29 under log-only noise (12 settings each; Table 3). Both estimators benefit from the edge constraint; robust aggregation adds accuracy under heavier tails, while ordinary means use bounded-noise samples more efficiently. Appendix E.1 also compares the complete-batch n^4 schedule of Bravo & Contreras (2026) under IID sampling.

4.2 Does query motion release a corrupted reply?

Setup and controls. Navigation and queue Bellman maps use $T(v) = (1-\gamma)r + \gamma Pv$, $\gamma = 0.95$, and maximum-norm value error. The normal grid crosses both tasks, 64/256 states, bounded/Pareto rewards and caps 20k/100k. All methods use MoM; 16 held-out seed units aggregate paired log errors across settings (Appendix E.3).

Table 1: **Query proposals separate precision from robustness.** Value-error ratios to raw adjacent averaging; smaller is better. The first two rows share strict one-edge repair and minimum motion. Normal noise uses 16 settings, variance bursts eight, and each corruption test four. All comparisons match actual calls and use 16 seed units; intervals are pointwise approximate paired-log Student intervals.

Query policy / method	Baseline noise ratio [95% interval]	Variance bursts ratio [95% interval]	Initial corruption	Terminal corruption
Fresh proposals	0.986 [0.979, 0.993]	0.990 [0.981, 0.998]	0.998	0.022
Repaired proposals	1.054 [1.036, 1.073]	0.788 [0.722, 0.860]	1.107	0.022
Stationary repair	1.081 [1.063, 1.100]	0.727 [0.639, 0.828]	69.943	0.022
Calibrated damping	1.479 [1.437, 1.523]	1.085 [0.965, 1.220]	2.122	1.016
Tail average	1.230 [1.192, 1.270]	1.086 [1.030, 1.145]	1.411	0.065
Fresh-edge Picard	1.089 [1.068, 1.110]	0.819 [0.733, 0.915]	1.107	0.030

Recovery and rejection are different capabilities. We replace the first completed estimate with the current query, or the last with the all-ones vector, retaining all sample charges. Stationary repair retains the first error, reaching a ratio of 69.943 relative to raw averaging. Fresh- and repaired-proposal movement reduce this to 0.998 and 1.107, respectively, while both suppress terminal corruption to 0.022 (Table 1). Query motion releases the inherited error while repair rejects the new disturbance.

Fresh proposals give ordinary-noise precision (ratio 0.986); repaired proposals give stronger burst rejection (0.788), at an ordinary-noise ratio of 1.054. Both preserve the reply invariant and recovery guarantee.

4.3 When does repair improve value accuracy?

Independent comparison. We next compare value accuracy against controls that either iterate on estimates or reuse observations to fit the Bellman map. An independent 1,280-run study measures MSE $d^{-1}\|\widehat{v} - v^*\|_2^2$ for these methods on both Bellman tasks, with 64 states, caps 20k/100k and stochastic transitions. Bounded reward noise is supplemented by centered bursts of amplitude 10 and mean duration 20 vector rounds. Model-fitting controls reuse observations with $O(d^2)$ count storage; repair uses $O(d)$ recurrence state (Appendix E.4).

Under bursts (Figure 3a), the primary repaired/raw ratio is 0.695 (95% interval $[0.489, 0.988]$); ratios to calibrated damping, tail averaging and mean-reward model fitting are 0.433, 0.543 and 0.443. Under baseline noise, the repaired/raw ratio is 1.251 and model reuse is more accurate. The task/budget breakdown (Figure 3b) gives burst ratios of 0.638–0.804 across all four cells, showing a consistent direction of the estimated effect.

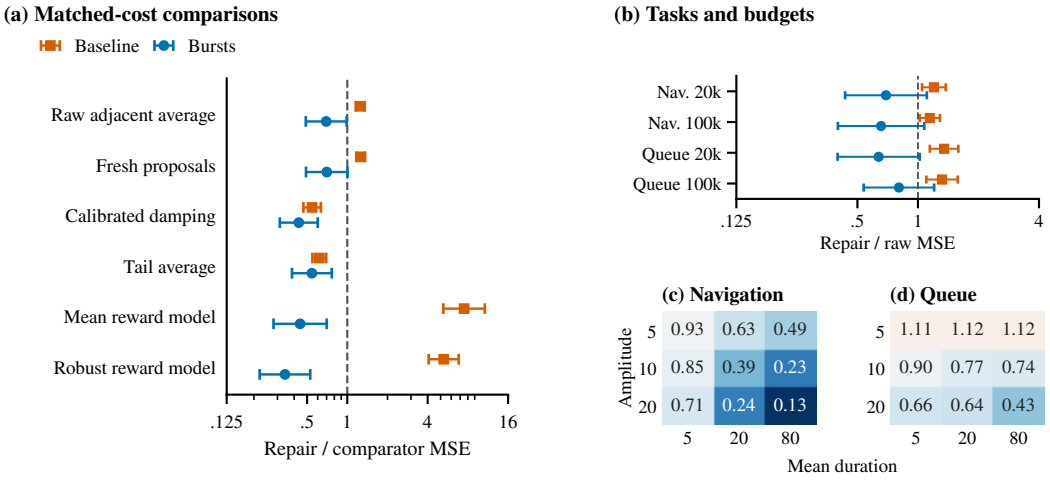


Figure 3: **Value accuracy across comparators, tasks and disturbance regimes.** Ratios below one favor repaired-proposal motion. (a) Paired geometric MSE ratios pooled over tasks and budgets; (b) the same repair/raw comparison split by task and budget. Intervals are 95% paired-log Student intervals over 16 seeds; burst/raw in (a) is primary, others pointwise. (c)–(d) All 18 burst settings from a separate study at cap 100k; duration counts raw vector rounds, with hot probability $1/11$. No-burst ratios are 1.04 (navigation) and 1.73 (queue).

Disturbance strength and persistence. Figure 3c–d varies amplitude $5, 10, 20$ and duration $5, 20, 80$, with hot probability $1/11$, 16 independent seeds and cap 100k in all 18 task/settings cells. Larger and longer bursts generally favor repair. At amplitude 5 the queue task favors raw averaging in all three durations, whereas the navigation task already benefits.

The pooled movement/raw ratio is 0.583 with interval $[0.422, 0.806]$. Stationary repair rejects these bursts more strongly, giving 0.447 $[0.323, 0.618]$; the movement/stationary ratio is 1.305 $[1.089, 1.563]$. The three intervals control nominal 95% family coverage (Appendix E.5), showing that movement pays a measurable precision cost for the ability to escape an inherited error.

5 Discussion and Conclusion

One-edge repair separates the statistical task of estimating a map value from the structural task of using that estimate in an iteration. The constant-one invariant lets a solver improve its estimator without redesigning the Halpern stability argument: a bound on input accuracy transfers directly to the residual. The norm characterization explains why: maximum-norm intervals preserve all compatible coordinate error bands, restoring consistency without knowing the true map value or an error radius.

For contractions, the travel bound turns this protection into a mechanism for recovery. A rejected correction reveals how much the fresh estimate violates the edge constraint; minimum motion converts that observable discrepancy into travel that releases inherited error. The experiments separate the resulting capabilities. Repair stabilizes both ordinary and robust estimates and improves value accuracy under variance bursts, while movement enables recovery after initial corruption.

A direction for further work is an adaptive movement policy that balances the travel needed to discount old errors against perturbations of the base update. One stored edge preserves estimation accuracy; query travel controls how long inherited error persists.

Reproducibility statement

The appendix gives complete proofs, algorithm specifications and experimental protocols. The anonymous code supplement contains the final solvers, fixed experimental grids and seeds, and commands to regenerate the reported experiments, statistical summaries and main quantitative figures. Its analysis checks grid completeness, sample costs and output errors. A small execution check is provided separately from the full experimental runs. True maps are available only to generators and evaluators.

AI use statement

We used generative AI tools to aid or polish writing, for retrieval and discovery, and for research ideation or execution. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

References

Nir Ailon, Bernard Chazelle, Seshadhri Comandur, and Ding Liu. Property-preserving data reconstruction. *Algorithmica*, 51(2):160–182, 2008.

Ahmet Alacaoglu and Yura Malitsky. Stochastic variance reduction for variational inequality methods. In *Proceedings of the Conference on Learning Theory*, volume 178, pp. 778–816. PMLR, 2022.

Nachman Aronszajn and Prom Panitchpakdi. Extension of uniformly continuous transformations and hyperconvex metric spaces. *Pacific Journal of Mathematics*, 6(3):405–439, 1956. doi: 10.2140/pjm.1956.6.405.

Pranjal Awasthi, Madhav Jha, Marco Molinaro, and Sofya Raskhodnikova. Limitations of local filters of Lipschitz and monotone functions. *ACM Transactions on Computation Theory*, 7(1):2:1–2:16, 2015. doi: 10.1145/2692372.2692373.

Manuel Blum, Robert W. Floyd, Vaughan Pratt, Ronald L. Rivest, and Robert E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973. doi: 10.1016/S0022-0000(73)80033-9.

Mario Bravo and Roberto Cominetti. Stochastic fixed-point iterations for nonexpansive maps: Convergence and error bounds. *SIAM Journal on Control and Optimization*, 62(1):191–219, 2024. doi: 10.1137/22M1515550.

Mario Bravo and Juan Pablo Contreras. Stochastic Halpern iteration in normed spaces and applications to reinforcement learning. *Mathematical Programming*, 2026. URL https://doi.org/10.1007/s10107-026-02345-1.

Christian Brownlees, Emilien Joly, and Gábor Lugosi. Empirical risk minimization for heavy-tailed losses. *The Annals of Statistics*, 43(6):2507–2536, 2015. doi: 10.1214/15-AOS1350.

Sébastien Bubeck, Nicolò Cesa-Bianchi, and Gábor Lugosi. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013. URL https://arxiv.org/abs/1209.1727.

Xufeng Cai, Chaobing Song, Cristóbal Guzmán, and Jelena Diakonikolas. Stochastic Halpern iteration with variance reduction for stochastic monotone inclusions. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Xufeng Cai, Ahmet Alacaoglu, and Jelena Diakonikolas. Variance reduced Halpern iteration for finite-sum monotone inclusions. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=0i6Z9N5MLY.

Olivier Catoni. Challenging the empirical mean and empirical variance: A deviation study. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 48(4):1148–1185, 2012. URL https://arxiv.org/abs/1009.2048.

Horaţiu Cheval and Laurenţiu Leuştean. Linear rates of asymptotic regularity for Halpern-type iterations. *Mathematics of Computation*, 2024. doi: 10.1090/mcom/3991.

Krzysztof J. Ciosmak. Continuity of extensions of Lipschitz maps, 2021. URL https://arxiv.org/abs/1904.02993v4. Version 4; first posted 2019.

Roberto Cominetti, José A. Soto, and José Vaisman. On the rate of convergence of Krasnosel'skiĭ–Mann iterations and their connection with sums of Bernoullis. *Israel Journal of Mathematics*, 199:757–772, 2014. doi: 10.1007/s11856-013-0045-4.

Juan Pablo Contreras and Roberto Cominetti. Optimal error bounds for non-expansive fixed-point iterations in normed spaces. *Mathematical Programming*, 199:343–374, 2023. doi: 10.1007/s10107-022-01830-7.

Ashok Cutkosky and Harsh Mehta. High-probability bounds for non-convex stochastic optimization with heavy tails. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Dominic Descombes and Maël Pavón. Injective subsets of $l_\infty(i)$. *Advances in Mathematics*, 317: 91–107, 2017. doi: 10.1016/j.aim.2017.06.025. URL https://arxiv.org/abs/1511.04761.

Luc Devroye, Matthieu Lerasle, Gábor Lugosi, and Roberto I. Oliveira. Sub-gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016. doi: 10.1214/16-AOS1440.

Jelena Diakonikolas. Halpern iteration for near-optimal and parameter-free monotone inclusion and strong solutions to variational inequalities. In *Proceedings of the Conference on Learning Theory*, volume 125, pp. 1428–1451. PMLR, 2020.

Jelena Diakonikolas. Solving Stochastic Fixed-Point Equations with High Probability. *arXiv preprint*, 2026. URL https://arxiv.org/abs/2607.09097v1.

Jelena Diakonikolas, Cristóbal Guzmán, and David Martínez-Rubio. Oracle complexity of stochastic fixed-point equations with nonexpansive maps, 2026. URL https://arxiv.org/abs/2609.09524v1. Version 1, September 8, 2026.

Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Benjamin Halpern. Fixed points of nonexpanding maps. *Bulletin of the American Mathematical Society*, 73(6):957–961, 1967. doi: 10.1090/S0002-9904-1967-11864-0.

Samuel B. Hopkins. Mean estimation with sub-gaussian rates in polynomial time. *The Annals of Statistics*, 48(2):1193–1213, 2020. doi: 10.1214/19-AOS1843.

Daniel Hsu and Sivan Sabato. Loss minimization and parameter estimation with heavy tails. *Journal of Machine Learning Research*, 17(18):1–40, 2016. URL https://www.jmlr.org/papers/v17/14-273.html.

Julien Walden Huang, Stephen Roberts, and Jan-Peter Calliess. Lipschitz interpolation: Non-parametric convergence under bounded stochastic noise. *arXiv preprint arXiv:2310.07039*, 2023. URL https://arxiv.org/abs/2310.07039v1.

Florian Hübler, Ilyas Fatkhullin, and Niao He. From gradient clipping to normalization for heavy tailed SGD. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, volume 258, pp. 2413–2421. PMLR, 2025.

Madhav Jha and Sofya Raskhodnikova. Testing and reconstruction of Lipschitz functions with applications to data privacy. *SIAM Journal on Computing*, 42(2):700–731, 2013. doi: 10.1137/110840741.

Donghwan Kim. Accelerated proximal point method for maximally monotone operators. *Mathematical Programming*, 190:57–87, 2021. doi: 10.1007/s10107-021-01643-0.

Jane Lange, Ephraim Linder, Sofya Raskhodnikova, and Arsen Vasilyan. Local Lipschitz filters for bounded-range functions with applications to arbitrary real-valued functions. *arXiv preprint arXiv:2308.14716v3*, 2024. URL https://arxiv.org/abs/2308.14716v3.

Jongmin Lee and Ernest K. Ryu. Accelerating value iteration with anchoring. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Jongmin Lee and Ernest K. Ryu. Optimal non-asymptotic rates of value iteration for average-reward Markov decision processes. In *International Conference on Learning Representations*, 2025.

Laurenţiu Leuştean. Rates of asymptotic regularity for Halpern iterations of nonexpansive mappings. *Journal of Universal Computer Science*, 13(11):1680–1691, 2007. doi: 10.3217/JUCS-013-11-1680.

Zhize Li, Hongyan Bao, Xiangliang Zhang, and Peter Richtárik. PAGE: A simple and optimal probabilistic gradient estimator for nonconvex optimization. In *Proceedings of the International Conference on Machine Learning*, 2021. URL https://arxiv.org/abs/2008.10898v3.

Ling Liang, Zusen Xu, Kim-Chuan Toh, and Jia-Jie Zhu. An inexact Halpern iteration with application to distributionally robust optimization. *arXiv preprint arXiv:2402.06033v3*, 2025. URL https://arxiv.org/abs/2402.06033v3.

Felix Lieder. On the convergence rate of the Halpern-iteration. *Optimization Letters*, 15:405–418, 2021. doi: 10.1007/s11590-020-01617-9.

Zijian Liu and Zhengyuan Zhou. Nonconvex stochastic optimization under heavy-tailed noises: Optimal convergence without gradient clipping. In *International Conference on Learning Representations*, 2025.

Gabor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions—a survey. *arXiv preprint arXiv:1906.04280*, 2019a. URL https://arxiv.org/abs/1906.04280.

Gábor Lugosi and Shahar Mendelson. Sub-gaussian estimators of the mean of a random vector. *The Annals of Statistics*, 47(2), 2019b. doi: 10.1214/17-AOS1639.

Manuel Madeira, Renato Negrinho, Joao Xavier, and Pedro M. Q. Aguiar. COCO denoiser: Using co-coercivity for variance reduction in stochastic convex optimization. *arXiv preprint arXiv:2109.03207*, 2021. URL https://arxiv.org/abs/2109.03207.

Stanislav Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4):2308–2335, 2015. doi: 10.3150/14-BEJ645.

Leopoldo Nachbin. A theorem of the Hahn–Banach type for linear transformations. *Transactions of the American Mathematical Society*, 68(1):28–46, 1950. doi: 10.1090/S0002-9947-1950-0032932-3.

Lam M. Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In *Proceedings of the International Conference on Machine Learning*, volume 70, pp. 2613–2621. PMLR, 2017.

Jisun Park and Ernest K Ryu. Exact optimal accelerated complexity for fixed-point iterations. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 17420–17457, 2022. URL https://proceedings.mlr.press/v162/park22c.html.

Maël Pavón. Injective convex polyhedra, 2014. URL https://arxiv.org/abs/1410.7306v1.

Shoham Sabach and Shimrit Shtern. A first order method for solving convex bilevel optimization problems. *SIAM Journal on Optimization*, 27(2):640–660, 2017. doi: 10.1137/16M105592X.

Justin Whitehouse, Ben Chugg, Diego Martinez-Taboada, and Aaditya Ramdas. Mean Estimation in Banach Spaces Under Infinite Variance and Martingale Dependence. *arXiv preprint arXiv:2411.11271v2*, 2024. URL https://arxiv.org/abs/2411.11271v2. Version 2, March 2025.

TaeHo Yoon and Nicolas Loizou. Direct acceleration of stochastic root-finding without variance reduction and regularization. *arXiv preprint arXiv:2608.12043v1*, 2026. URL https://arxiv.org/abs/2608.12043v1.

Taeho Yoon and Ernest K. Ryu. Accelerated algorithms for smooth convex-concave minimax problems with $O(1/k^2)$ rate on squared gradient norm. In *Proceedings of the International Conference on Machine Learning*, volume 139, pp. 12098–12109. PMLR, 2021.

Guide to the appendix. Appendix A proves error preservation and Halpern stability. Appendix B gives the sampling algorithm and statistical guarantees. Appendix C develops minimum-motion recovery and pulse attenuation. Appendix D establishes the norm classification and geometric constructions. Appendix E collects experimental protocols, baselines and supporting results.

A Error preservation and Halpern stability

Proof of Lemma 1. Fix k and set $M = \max_{j\le k} \|v_j - T(x_j)\|$. Since R fixes $T(x_0)$, its nonexpansiveness proves the initial bound. If the bound holds at $j-1$, then for $z = T(x_j)$ and $a = y_{j-1}$,

$$\|a - z\| \le M + r_j, \qquad a_i - r_j \le z_i + M, \quad a_i + r_j \ge z_i - M.$$

Clipping $v_{j,i} \in [z_i - M, z_i + M]$ to $[a_i - r_j, a_i + r_j]$ keeps it in the same error band: a replaced coordinate lands at an endpoint inside that band. Thus $\|Q_{a,r_j}(v_j) - z\| \le M$. Applying R preserves this bound because $R(z) = z$. Also $R(a) = a$, so $\|R(Q_{a,r_j}(v_j)) - a\| \le r_j$. Induction proves both claims, including zero radii. □

Proof of Theorem 2. Retraction keeps all queries in C. Put $\Delta_j = \|x_j - x_{j-1}\|$. Nonexpansiveness of R gives $\Delta_1 \le D/3$. Subtracting the pre-retraction anchor mixtures and using adjacent consistency yields

$$\Delta_{j+1} \le \frac{j+1}{j+3}\Delta_j + \frac{2D}{(j+2)(j+3)}, \qquad \Delta_j \le \frac{2Dj}{(j+1)(j+2)}$$

by induction. Define x_{k+1} by (2) for analysis, using the existing y_k. Since $R(y_k) = y_k$, its distance to y_k is at most $2D/(k+3)$. Therefore

$$\|x_k - y_k\| \le \Delta_{k+1} + \frac{2D}{k+3} \le \frac{4D}{k+2}.$$

Lemma 1 supplies the comparator error. Each stage stores a constant number of vectors and evaluates R at most twice. The singleton case $D = 0$ is immediate. □

A.1 Diagnostic examples

An exact example of error accumulation. Let $c = (1/2, 1/2)$, $u = (1, 0)$ and $R(z_1, z_2) = (-z_2, z_1)$. The map $T(x) = c + R(x - c)$ is a maximum-norm isometry of the cube. For a fixed horizon k, supply the deterministic stage-indexed error $e_j = MR^{j+1}u$, where $0 < M \le 3/[2(k+3)]$. Direct Halpern iteration anchored at c and using $v_j = T(x_j) + e_j$ satisfies

$$x_j = c + \frac{jM}{3}R^j u, \qquad v_j = c + \frac{(j+3)M}{3}R^{j+1}u, \qquad \mathrm{res}_T(x_k) = \frac{kM}{3}.$$

Substitution into (2) proves the first two formulas by induction; the last follows from $\|u - Ru\| = 1$. Every v_j lies in the cube, so range-only clipping leaves this run unchanged. For $M = 1/(k+3)$, the maximum raw error tends to zero while the residual tends to $1/3$. With local repair, the same map and stage-indexed error rule, evaluated at the repaired method's own queries, instead give residual at most $4/(k+2) + M \to 0$.

Global interpolation can fail. With scalar anchor zero, replies $1/4, 5/16, 1/4, 39/160$ produce queries $0, 1/12, 5/32, 3/20$ under (2). Every adjacent reply difference is at most the corresponding query difference, so taking these replies as raw estimates realizes the transcript through local repair. But replies 1 and 3 differ by $11/160$, whereas their query distance is $1/15$. A nonexpansive map cannot interpolate both. Theorem 2 remains valid because its proof uses adjacent constraints directly.

The maximum retains early error. Consider $T \equiv 0$ in one dimension, anchor $x_0 = 1/2$, first estimate $v_0 = 1/2$, and every later estimate equal to zero. The first reply equals the anchor, so the next query is unchanged. The local radius is zero and every repaired reply remains $1/2$. The residual stays $1/2$ although later estimates are exact. Thus the theorem controls the historical maximum error: exact later estimates do not remove an earlier error from the guarantee.

B STATISTICAL GUARANTEES AND SAMPLE COMPLEXITY

B.1 FIXED-BUDGET SAMPLING ALGORITHM

Algorithm 1 One-edge Halpern iteration with a fixed sampling budget

Input: sampler at any $x \in [0,1]^d$, anchor x_0, budget N, accuracy allocation a, confidence η.

1. Set k, K, L, B by (3). If $B = 0$, return $\frac{1}{2}\mathbf{1}$.
2. For $j = 0, \dots, k$:
   (a) Collect L successive means of B responses at x_j; let v_j be their coordinate median.
   (b) Set $y_0 = P_C v_0$ if $j = 0$; otherwise compute y_j by the one-edge clipping rule (1).
   (c) If $j < k$, set $x_{j+1} = \frac{2}{j+3} x_0 + \frac{j+1}{j+3} y_j$.
3. Return the queried point x_k.

Calls: $KLB \le N$. **Precision target** ε**:** take $a = \varepsilon/2$.

B.2 CONDITIONAL ORDINARY-MEAN BOUND

Assume $\mathbb{E}[\|\zeta_i\|^p \mid \mathcal{F}_{i-1}] \le \nu^p$ and conditional centering, for $1 < p \le 2$. Freeze a query for B calls and fix a coordinate error ξ_i. Set $A = \nu B^{1/p}$, $c_i = \mathrm{clip}_{[-A,A]} \xi_i$, $e_i = \xi_i - c_i$ and $m_i = \mathbb{E}[c_i \mid \mathcal{F}_{i-1}] = -\mathbb{E}[e_i \mid \mathcal{F}_{i-1}]$. For $\nu > 0$,

$$\mathbb{E}[c_i^2 \mid \mathcal{F}_{i-1}] \le A^{2-p}\nu^p, \qquad \mathbb{E}[|e_i| \mid \mathcal{F}_{i-1}] \le A^{1-p}\nu^p.$$

Conditional on the block's starting history, martingale orthogonality bounds the expected absolute mean of $c_i - m_i$ by $\sqrt{A^{2-p}\nu^p/B} = \nu B^{-(p-1)/p}$. Each of the means of m_i and e_i contributes at most the same amount. Their sum is the original sample mean, proving $b_B \le 3\nu B^{-(p-1)/p}$ after every starting history. When $\nu = 0$, all errors vanish almost surely. Truncation appears only in this analysis; the estimator averages the original observations.

B.3 PROOF OF THE CONDITIONAL MEAN-ERROR TRANSFER

Proof. For $b_B > 0$, conditional Markov bounds a block-coordinate failure at threshold $8b_B u$ by $(8u)^{-1}$ for $u \ge 1$. A bad median requires at least $q = (L+1)/2$ bad blocks. Successive conditioning bounds the joint failure of any chosen q blocks by $(8u)^{-q}$, even under dependence. A subset union gives a median-coordinate tail at most $2^L(8u)^{-q} \le 2^{-L/2}u^{-q}$. Tower conditioning and the finite coordinate/query union imply

$$\Pr\{M > 8b_B u\} \le dK2^{-L/2}u^{-q} \le \eta u^{-q}.$$

The case $u = 1$ and Theorem 2 prove the probability bound. Tonelli's layer-cake formula gives $\mathbb{E}M \le 8b_B(1 + 1/(q-1)) \le 16b_B$, proving the expectation bound. If $b_B = 0$, all finitely many block errors vanish almost surely by conditional nonnegativity. Streaming retains L means and block sums. Sorting coordinates takes $O(dKL\log(eL))$ comparisons; deterministic selection (Blum et al., 1973) reduces this to $O(dKL)$. Adding sample processing and local repair yields $O(dN)$ exact-real operations, excluding oracle generation. □

B.4 PROOF OF THE HIGH-DIMENSIONAL PRECISION CONSEQUENCE

The upper bound in Corollary 4 is the power-tail consequence of Theorem 3. We prove the lower bound using the sign-chain of Diakonikolas et al. (2026, Lemma 10 and the proof of Theorem 11). Their parameter is set to $\delta = 4\varepsilon$, their chain length is exactly $m = \lfloor 1/(8\delta) \rfloor$, and $d \ge m$. For signs $\sigma \in \{-1,1\}^m$, the resulting map S_σ on $[-1,1]^d$ is nonexpansive, has displacement at most 2δ, and residual at most δ requires the output to encode every sign. The stronger prefix oracle in their information argument leaves the unexposed suffix independent and uniform, and reveals at most two new signs in conditional expectation per query. We use these established geometric and information properties to extend the oracle lower bound to conditional pth moments.

Normalize the map by $T_\sigma(x) = (S_\sigma(2x - \mathbf{1}) + \mathbf{1})/2$. Then T_σ maps C to itself, is nonexpansive and has $\|T_\sigma(x) - x\| \le 4\varepsilon$. Moreover, $\mathrm{res}_{T_\sigma}(x) \le \varepsilon$ implies $\mathrm{res}_{S_\sigma}(2x - \mathbf{1}) \le 2\varepsilon < \delta$, and hence recovery of every sign.

Put $w = \max\{1, (\nu/(8\varepsilon))^{p/(p-1)}\}$ and $\pi = 1/w$. At query t, draw $A \sim \mathrm{Bernoulli}(\pi)$ independently of $\mathcal{F}_{t-1}$ and return

$$Y = x + \frac{A}{\pi}\Delta_\sigma(x), \qquad \Delta_\sigma = T_\sigma - \mathrm{Id}\,.$$

Conditional on the full interaction history, its mean is $T_\sigma(x)$. If $\pi = 1$, the noise is zero. Otherwise its conditional pth moment is bounded by

$$\begin{aligned}\mathbb{E}[\|Y - T_\sigma(x)\|^p \mid \mathcal{F}_{t-1}] &= \|\Delta_\sigma(x)\|^p\big[(1-\pi) + \pi^{1-p}(1-\pi)^p\big] \\ &\le 2(4\varepsilon)^p\pi^{1-p} = 2^{1-p}\nu^p \le \nu^p.\end{aligned}$$

The bound holds at every predictable query, discharging the full-history premise and integrability. The oracle is allowed to return values outside C.

For an algorithm with a hard cap N, reveal each Bernoulli indicator and the stronger exact prefix response whenever $A = 1$. This augmented transcript simulates the actual responses. Under uniformly random signs, the indicators convey no sign information; at each call the expected new prefix length is at most 2π, conditionally on the augmented history and the independent algorithm seed. Padding stopped algorithms to N calls and writing J_N for the revealed prefix length gives $\mathbb{E}J_N \le 2N\pi$. If $N \le mw/64$, then $\Pr(J_N = m) \le 1/32$. On $J_N < m$, at least one independent fair suffix sign remains. Any output encodes every sign with conditional probability at most $1/2$. Therefore success probability is at most $1/32 + 1/2 = 17/32$, and averaging over the finite sign family supplies a fixed instance with failure probability at least $1/4$. This excludes every integer cap at or below $mw/64$ even for class-aware randomized algorithms, proving the strict lower inequality.

Finally, $m \ge 1/(64\varepsilon)$, $K = O(\varepsilon^{-1})$ and $L = O(\log(ed/(\varepsilon\eta)))$. With $s = p/(p-1)$, the ratio between the two moment factors is at most 384^s, and the ceiling costs at most a factor two. This proves the fixed-p logarithmic comparison uniformly in the displayed $d, \nu, \varepsilon, \eta$ range. The stated dimension condition supplies the full sign chain.

C Minimum-motion recovery and pulse attenuation

C.1 Geometric recovery

Proof of Theorem 5. Since $d_t \le D$, a feasible coordinate boundary exists. Its distance from h_t is at most d_t, hence $\|x_{t+1} - h_t\| \le d_t$. Also $d_t \ge [E_t - e_t]_+$. Substituting the guaranteed travel $d_t/2$ into (4) proves (6). Induction gives the tail bound. Finally $d_t \le E_t + e_t \to 0$; the query perturbations vanish, and the corresponding averaged contraction recursion converges. Appendix C.3 gives the finite-time propagation and sample-cost bound for both query policies. □

C.2 Exact minimum-motion update

The repair is (1) on $C = [0, D]^d$. At each query project the fresh estimate to $b_t \in C$, emit y_t, and put $d_t = \|b_t - y_t\|$. These quantities are available before selecting the next query. For $r_t = d_t/2$, the feasible set in (5) is the union of the coordinate half-spaces $x_i \le x_{t,i} - r_t$ and $x_i \ge x_{t,i} + r_t$, intersected with C. If h_t is outside the open ball, it is an exact minimizer. Otherwise projection onto any feasible half-space changes just that coordinate to its boundary. Compare these at most $2d$ scalar distances and use the smallest, with ties broken by coordinate order and then the negative sign. This proves optimality of the stated $O(d)$ operation.

Because $r_t \le D/2$, at least one sign is feasible for each coordinate. Inside the open ball every candidate displacement is at most $r_t + \|h_t - x_t\| < 2r_t = d_t$. The case $r_t = 0$ returns h_t directly. Thus

$$\|x_{t+1} - x_t\| \ge d_t/2, \qquad \|x_{t+1} - h_t\| \le d_t \le E_t + e_t.$$

C.3 Fixed-point error and sample cost

Write $a_t = \|x_t - x^*\|$. For repaired adjacent-average proposals, $t \geq 1$, the bound is

$$a_{t+1} \leq \frac{\gamma}{2}(a_t + a_{t-1}) + \frac{3E_t + E_{t-1}}{2} + e_t.$$

For $0 < \gamma < 1$, its nonnegative comparison matrix has spectral radius $\rho = (\gamma + \sqrt{\gamma^2 + 8\gamma})/4 < 1$ and positive eigenvector $(\rho, 1)$. The case $\gamma = 0$ follows directly from the forcing terms. For fresh proposals, the forcing is $(e_t + e_{t-1})/2 + \delta_t$, where $\delta_t = \|x_{t+1} - h_t\| \leq E_t + e_t$. Both forcings vanish when $e_t \to 0$, proving convergence of the queries and final repaired readout. If $e_t \leq \epsilon$ eventually, these recursions give asymptotic fixed-point error at most $3\epsilon/(1-\gamma)$.

For any deterministic block allocation $B_0, \ldots, B_{K-1} \geq 1$, the original conditional MoM argument, with odd $L \geq 2\log_2(dK/\eta)$, gives simultaneous

$$e_t \leq A_t := 24\nu B_t^{-(p-1)/p}, \qquad 0 \leq t < K,$$

with probability at least $1 - \eta$ under the moment premise of Section 3.2. Set $U_0 = A_0$ and $U_t = \max\{A_t, qU_{t-1} + (1-q)A_{t-1}\}$. Induction gives $E_t \leq U_t$. Since $U_t \geq A_t$, the fresh-proposal forcing $U_t + 3A_t/2 + A_{t-1}/2$ is at most $(3U_t + U_{t-1})/2 + A_t$, the same envelope as for repaired proposals. The final repaired average $\widehat{x} = (y_{K-1} + y_{K-2})/2$, for $K \geq 2$, consequently satisfies

$$\|\widehat{x} - x^*\| \leq D\rho^{K-1} + \sum_{t=1}^{K-1} \rho^{K-1-t}\left(\frac{3U_t + U_{t-1}}{2} + A_t\right). \tag{8}$$

For $\gamma > 0$, the weighted maximum norm of the initial pair is at most D/ρ. Iterating the comparison matrix proves the expression for the next query. For the final repaired average, the forcing is $(U_{K-1} + U_{K-2})/2$ with no movement term, also below this envelope. At $\gamma = 0$, only the final summand remains, with $0^0 = 1$. This is a finite-budget bound for the implemented allocation, counting $L\sum_t B_t$ raw vectors. The event applies directly to conditionally centered oracle samples.

C.4 Dense errors become sparse query displacements

Corollary 6 extends to nonuniform state occupancy and bounded background errors. The repaired output transports one column of the transition matrix, so its cost depends on the weight assigned to that column.

Corollary 8 (Pulse attenuation and state occupancy)**.** *Let $T(x) = v^* + \gamma P(x - v^*)$ be a cube self-map, where P is nonnegative row-stochastic and $0 \leq \gamma < 1$. Start both methods at $x_0 = y_0 = v^*$, supply $b_1 = v^* + a\xi \in C$ with arbitrary $\xi \in \{-1, 1\}^d$, and suppose the next completed estimates have errors at most ϵ in maximum norm. Let π be any stationary probability vector of P, set $m = \max_i \pi_i$, and write $\|u\|_\pi^2 = \sum_i \pi_i u_i^2$ (a seminorm when weights vanish). For raw adjacent averaging and repaired averaging with the radius-$d_t/2$ rule* (5)*,*

$$\frac{\|h_2^{\mathrm{rep}} - v^*\|_\pi^2}{\|h_2^{\mathrm{raw}} - v^*\|_\pi^2} \leq \left(\frac{\gamma a\sqrt{m} + 2\epsilon}{(2-\gamma)a - 2\epsilon}\right)^2 \tag{9}$$

whenever the denominator on the right is positive. Thus $\epsilon < a(2 - \gamma - \gamma\sqrt{m})/4$ ensures strict improvement on every pulse realization. If P is doubly stochastic, take $\pi_i = 1/d$; with $\epsilon = 0$ the uniform-MSE ratio is at most $\gamma^2/[d(2-\gamma)^2]$, and this constant is attained for every even d.

Proof. The first repair returns $y_1 = v^*$. Minimum motion produces $x_2 = v^* + sae_j/2$ for a feasible coordinate j and sign s. With exact next estimates, the repaired and raw output errors are respectively

$$sa\gamma Pe_j/4, \qquad a(2I + \gamma P)\xi/4.$$

Since $0 \leq P_{ij} \leq 1$ and $\pi^\top P = \pi^\top$, we have $\sum_i \pi_i P_{ij}^2 \leq \pi_j \leq m$. Each raw coordinate has magnitude at least $a(2-\gamma)/4$, because $|(P\xi)_i| \leq 1$. The next background error changes either average by at most $\epsilon/2$: the preceding repaired reply is exact, so coefficient-one repair preserves the current error bound. Triangle and reverse triangle inequalities prove (9). The improvement threshold follows by comparing numerator and denominator. For sharpness, let P swap disjoint coordinate pairs, choose opposite pulse signs within each pair, and take $v^* = \frac{1}{2}\mathbf{1}$. Then $P\xi = -\xi$ and every column of P has Euclidean norm one, attaining the displayed clean uniform-MSE ratio. □

The result concerns h_2, after three completed estimates b_0, b_1, b_2, and includes the displacement needed for recovery. It holds for dependent signs and biased background. Stationary weighting measures error under the chain's state occupancy; it equals uniform MSE when π is uniform.

For an approximate initial state satisfying Proposition 9's activation and feasibility assumptions, the same proof gives

$$\|h_2^{\mathrm{raw}} - v^*\|_\pi \geq \left(a(2-\gamma)/4 - B_2^{\mathrm{raw}}\right)_+, \qquad \|h_2^{\mathrm{rep}} - v^*\|_\pi \leq a\gamma\sqrt{m}/4 + B_2^{\mathrm{rep}}.$$

This extension also needs no independence among pulse signs. A useful regime is determined jointly by the pulse size, initial accuracy, background error and concentration of state occupancy.

Independent symmetric signs give a stronger expected uniform-MSE bound. With $c = \max_j \|Pe_j\|_2$, $s_P = \|2I + \gamma P\|_{\mathrm{F}}/\sqrt{d}$, and exact initialization,

$$\frac{R_2^{\mathrm{rep}}}{R_2^{\mathrm{raw}}} \leq \left(\frac{\gamma ac/\sqrt{d} + 2\epsilon}{as_P - 2\epsilon}\right)^2, \qquad as_P > 2\epsilon.$$

For doubly stochastic P, $c \leq 1$ and $s_P \geq 2$. Hence the clean expected ratio is at most $\gamma^2/(4d)$, and $\epsilon < a(2 - \gamma/\sqrt{d})/4$ certifies expected improvement. For a cyclic permutation with zero diagonal, the exact clean expected ratio is $\gamma^2/[d(4+\gamma^2)]$, so the inverse-dimension order is attained. These local bounds complement the full-horizon response condition in Proposition 9; repeated pulses and later outputs require their own transported-error terms.

C.5 Approximate initialization and later outputs

Let $T(x) = v^* + A(x - v^*)$ be a cube self-map with $A = \gamma P$ and nonnegative row-stochastic P. Assume

$$\|x_0 - v^*\|_\infty \leq \beta, \qquad y_0 = b_0 = T(x_0) + e_0 \in C, \qquad \|e_0\|_\infty \leq \epsilon_0.$$

Both methods query $x_1 = y_0$. At stage 1, let the completed estimate be $b_1 = T(x_1) + a\xi \in C$, where the coordinates of ξ are independent symmetric signs and $a > 0$ is fixed. For every $t \geq 2$, each method's completed estimate at its actual query has error at most ϵ almost surely. These background errors may be biased, dependent and adaptive. Compare raw adjacent averaging with one-edge repair and minimum motion. All error bounds concern completed estimates, including any estimation or domain clipping.

Define the clean pulse-response matrices

$$U_2 = (A + 2I)/4, \quad U_3 = (A^2 + 4A)/8, \qquad V_2 = A/4, \quad V_3 = (A^2 + 2A)/8,$$

with $U_t = A(U_{t-1} + U_{t-2})/2$ and the same recurrence for V_t, $t \geq 4$. Set $F_t = \|U_t\|_{\mathrm{F}}$ and $\kappa_t = \max_j \|V_t e_j\|_2$. The maximum over columns accommodates every coordinate selected by minimum motion.

The initial uncertainty enters through scalar quantities

$$\delta = \gamma\beta + \epsilon_0, \quad L = \beta + 2\delta, \quad D_0 = L + \gamma\delta, \quad q = (1+\gamma)/2,$$
$$X = \max\{\delta + D_0/2, (\delta + L)/2\}, \qquad E = \max\{D_0, \epsilon\}.$$

For raw averaging start $X_2^{\mathrm{raw}} = q\delta$, $Y_1^{\mathrm{raw}} = \gamma\delta$ and iterate

$$Y_t^{\mathrm{raw}} = \gamma X_t^{\mathrm{raw}} + \epsilon, \quad B_t^{\mathrm{raw}} = (Y_t^{\mathrm{raw}} + Y_{t-1}^{\mathrm{raw}})/2, \quad X_{t+1}^{\mathrm{raw}} = B_t^{\mathrm{raw}}.$$

For repair start $X_2^{\mathrm{rep}} = X$, $Y_1^{\mathrm{rep}} = L$ and iterate

$$Y_t^{\mathrm{rep}} = \gamma X_t^{\mathrm{rep}} + E, \quad B_t^{\mathrm{rep}} = (Y_t^{\mathrm{rep}} + Y_{t-1}^{\mathrm{rep}})/2, \quad X_{t+1}^{\mathrm{rep}} = B_t^{\mathrm{rep}} + 2E.$$

These deterministic recurrences separate the transported primary error from all subsequent uncertainty.

Proposition 9 (Prospective attenuation with bounded background)**.** *Under the preceding conditions, suppose $a > D_0 + \beta + \delta$ and $a \leq \min_i \min(v_i^*, 1 - v_i^*)$. Write $R_t = \mathbb{E}\|h_t - v^*\|_2^2/d$. For every $t \geq 2$,*

$$\sqrt{R_t^{\mathrm{raw}}} \geq \left(aF_t/\sqrt{d} - B_t^{\mathrm{raw}}\right)_+, \qquad \sqrt{R_t^{\mathrm{rep}}} \leq a\kappa_t/\sqrt{d} + B_t^{\mathrm{rep}}.$$

In particular, the precomputable condition

$$a(F_t - \kappa_t)/\sqrt{d} > B_t^{\mathrm{raw}} + B_t^{\mathrm{rep}} \tag{10}$$

ensures $R_t^{\mathrm{rep}} < R_t^{\mathrm{raw}}$. *When the raw lower bound is positive, the squared ratio of the displayed upper and lower square-root bounds bounds* $R_t^{\mathrm{rep}}/R_t^{\mathrm{raw}}$.

Proof. First $\|y_0 - v^*\|_\infty \le \delta$ and the first query distance is at most $\beta + \delta$. Adjacent consistency therefore gives $\|y_1 - v^*\|_\infty \le L$. If $d_1 = \|b_1 - y_1\|_\infty$, the pulse form implies $|d_1 - a| \le D_0$. Also $\|h_1 - x_1\|_\infty = \|y_1 - y_0\|_\infty/2 \le (\beta + \delta)/2$. The activation condition forces minimum motion to choose a single coordinate j and sign s, setting that coordinate of x_2 to $x_{1,j} + sd_1/2$ and leaving the others at h_1. Consequently its distance from $v^* + sae_j/2$ is at most X.

For each outcome take an exact-value reference trajectory starting from this feasible single-coordinate displacement. Its first reply is v^* and its output error at stage t is saV_te_j. The raw exact-value reference has output error $aU_t\xi$. The two exceptional stage-3 expressions account for the first moved query. Later stages satisfy the homogeneous matrix recurrences above. Independence of the pulse signs gives $\mathbb{E}\|U_t\xi\|_2^2 = F_t^2$, whereas $\|V_te_j\|_2 \le \kappa_t$ for any selected coordinate.

At the actual stage-1 repaired reply, error relative to $T(x_1)$ is at most D_0. Starting from this existing reply, the coefficient-one induction bounds every later repaired error by E. The fresh-to-repaired displacement is at most $\epsilon + E \le 2E$, and nearest feasible movement shifts the base proposal by at most that displacement. Subtracting the exact-value reference gives the repaired scalar recurrence, starting with query discrepancy X and previous-reply discrepancy L. The raw discrepancies start at $q\delta$ and $\gamma\delta$, and obey the raw recurrence. Thus the corresponding output discrepancies are bounded by B_t^{rep} and B_t^{raw} in maximum norm. Triangle and reverse triangle inequalities in the joint probability-coordinate L^2 norm prove the risk bounds and (10). □

For $\beta = \epsilon_0 = 0$, write $B_t^{\mathrm{raw}} = \epsilon J_t$ and $B_t^{\mathrm{rep}} = \epsilon H_t$. Here $J_2 = H_2 = 1/2$, $J_3 = 1 + \gamma/4$ and $H_3 = 1 + 5\gamma/4$. Condition (10) becomes

$$a(F_t - \kappa_t)/\sqrt{d} > \epsilon(J_t + H_t).$$

This defines a finite attenuation window: the primary-error response decays while persistent background leaves a nonzero allowance. A failed condition gives no performance ordering. The bounds are analytical assumptions, not extra oracle access; high-probability estimator events need an explicit failure contribution to bound unconditional risk. Actual pulse feasibility follows, for example, from $a + \gamma\delta \le \min_i \min(v_i^*, 1 - v_i^*)$.

D GEOMETRY OF LOSSLESS REPAIR

D.1 PROOF OF THE NORM CLASSIFICATION

Proof of Theorem 7. Write $\overline{B}(a, r) = \{x : \|x - a\| \le r\}$. The rule is fixed before the unknown map and estimate sequence. It receives no error radius.

We prove necessity by deriving the binary intersection property from two-query repair. Fix finitely many compatible balls $\overline{B}(a_i, r_i)$, $1 \le i \le n$. Choose a unit vector $e \in E$ and a norm-one linear functional ℓ with $\ell(e) = 1$, whose existence follows from the Hahn–Banach theorem applied to $\mathrm{span}\{e\}$. Choose one positive-radius closed ball C containing all a_i and all $a_i - 3r_ie$. Convexity then also puts $a_i - 2r_ie$ and every segment between the centers in C. Fix the rule in Theorem 7 for this C.

For a center $a = a_i$ and radius $r = r_i > 0$, restart the rule with

$$x_0 = a - 3re, \qquad v_0 = a.$$

The same input is exact for the constant map $T \equiv a$. The maximum-error invariant at $j = 0$ therefore forces $y_0 = a$, irrespective of the actual hidden map. Consequently $x_1 = a - 2re$ and $\|x_1 - x_0\| = r$. Given any $v \in C$ as v_1, denote its reply by q. Adjacent consistency gives $q \in C \cap \overline{B}(a, r)$.

For each $z \in C$, define

$$d_z = \|z - a\|, \quad u_z = \begin{cases} a, & d_z \le r, \\ a + (1 - r/d_z)(z - a), & d_z > r. \end{cases}$$

Then $u_z \in [a, z] \subset C$, $\|z - u_z\| \le r$, and $\|u_z - a\| = (d_z - r)_+$. The map

$$T_z(x) = u_z + \mathrm{clip}_{[0,1]}\left(\frac{\ell(x - x_0)}{r}\right)(z - u_z) \tag{11}$$

is a self-map of C, since its image is in $[u_z, z] \subset C$. It is nonexpansive because scalar clipping is 1-Lipschitz and $\|z - u_z\|/r \le 1$. Moreover $T_z(x_0) = u_z$ and $T_z(x_1) = z$. Every T_z gives the same observed transcript (x_0, a, x_1, v) and hence the same reply q. Applying the maximum-error invariant at $j = 1$ gives the simultaneous bound

$$\|q - z\| \le \max\{\|v - z\|, (\|a - z\| - r)_+\} \qquad (z \in C). \tag{12}$$

When $r = 0$, define $q = a$ directly; (12) still holds and $q \in C \cap \overline{B}(a, 0)$.

Start with $v = a_1$, which lies in the first ball. At each $i = 2, \dots, n$, use the preceding construction with a_i, r_i and current point v, and replace v by q. If v already lies in each previous ball $\overline{B}(a_h, r_h)$, $h < i$, then (12) at $z = a_h$, together with pairwise compatibility, yields

$$\|q - a_h\| \le \max\{r_h, (\|a_i - a_h\| - r_i)_+\} \le r_h.$$

The new point also lies in $\overline{B}(a_i, r_i)$ and in C. Induction gives a point in all n balls. No future center was revealed to the repair rule; universality over T_z is what gives simultaneous preservation.

For an arbitrary nonempty compatible family, fix one ball B_0 in it. Every finite subfamily together with B_0 intersects by the finite argument. The intersections with B_0 are closed subsets of the compact set B_0, since E is finite dimensional. The finite intersection property gives an intersection of the whole family. The empty family has intersection E. Thus every compatible family of balls intersects.

Classical ball-intersection classification. The finite-dimensional specialization of Nachbin (1950, Theorem 3) states: *a real normed space E of finite dimension $d \ge 1$ is linearly isometric to ℓ_∞^d if and only if every family of pairwise-intersecting closed norm balls in E has nonempty total intersection.* Finite dimensionality supplies completeness. The same characterization is stated in the introduction of Pavón (2014). Pairwise compatibility and pairwise intersection are equivalent for norm balls: a suitable point on the segment between two centers belongs to both. The property proved above therefore verifies the classification's hypothesis and establishes necessity.

For randomized rules, write replies as measurable functions of the observed prefix and one internal seed whose law is fixed independently of T. The exact constant-map input forces $\mathbb{E}\|y_0 - a\| = 0$, hence $y_0 = a$ almost surely. The anchored x_1 is therefore deterministic, and every T_z induces the same law of the internal state and second reply Q. Almost-sure domain membership and consistency put Q in $C \cap \overline{B}(a, r)$. The raw prefix maximum for T_z is the deterministic right side of (12), so the expected-prefix guarantee implies

$$\mathbb{E}\|Q - z\| \le \max\{\|v - z\|, (\|a - z\| - r)_+\} \quad (z \in C).$$

The barycenter $q = \mathbb{E}Q$ exists because C is bounded and finite dimensional. Convexity puts q in $C \cap \overline{B}(a, r)$, and Jensen's inequality yields (12) for every z. Each bound concerns the same distribution of Q; no common probability-one event over all z is needed. The compatible-ball induction and classification then apply unchanged. This argument already uses admissible anchored adaptive queries; allowing arbitrary predictable query histories only strengthens the premise.

For sufficiency, use a linear isometry from E to ℓ_∞^d. A norm ball becomes a coordinate box, whose coordinate-clipping retraction is nonexpansive. Lemma 1 then gives the all-horizon lossless guarantee, including the two anchored prefixes required here. Transporting the replies back through the isometry proves sufficiency. □

D.2 Retraction constructions

The following constructions supply the retractions used in Section 3.4. Their formulas give explicit computational costs in addition to existence.

Lemma 10 (Ordered domain with linear-time retraction)**.** *For* $C_{\mathrm{ord}} = \{x \in [0,1]^d : x_1 \le \cdots \le x_d\}$,

$$R_{\mathrm{ord}}(v)_i = \max_{1 \le h \le i} \min\{1, \max\{0, v_h\}\}$$

is a nonexpansive retraction computable in $O(d)$ *time and* $O(d)$ *output space. For* $d \ge 2$ *the domain is not a Cartesian box. Theorems 1–2 apply with* $D = 1$ *and* $O(d)$ *recurrence state and work per stage.*

Proof. Clipping each coordinate is nonexpansive in the maximum norm. The difference between two maxima over the same finite index set is bounded by the maximum coordinate difference. The output is bounded and nondecreasing, and a bounded nondecreasing vector is fixed. A running maximum computes the output in one pass. Both constant-zero and constant-one vectors belong to the domain, so its maximum-norm diameter is one. For $d \ge 2$, each coordinate ranges over $[0,1]$ but a vector with first coordinate one and last coordinate zero is excluded, proving the non-box claim. □

Lemma 11 (A nonconvex branching domain)**.** *Let* $d \ge 3$*, and let* $s_i \in \{-1,1\}^d$ *have coordinate* i *equal to* 1 *and all others equal to* -1*. The set*

$$C_{\mathrm{star}} = \{ts_i : 0 \le t \le 1,\ 1 \le i \le d\}$$

has maximum-norm diameter two and is nonconvex. For $v \in \mathbb{R}^d$*, choose any index* i *attaining its largest coordinate; let* $m_1 = v_i$ *and* $m_2 = \max_{h \ne i} v_h$*. Set*

$$R_{\mathrm{star}}(v) = \min\{1, (m_1 - m_2)/2\}\, s_i.$$

This formula is independent of tie-breaking and is a nonexpansive retraction onto C_{star}*, computable in* $O(d)$ *time. Thus retracted anchoring gives the residual bound* $8/(k+2) + M$ *with* $O(d)$ *state and work per stage on this branching domain.*

Proof. A tie gives $m_1 = m_2$ and hence the origin. Each point ts_i with $t > 0$ has largest coordinate t and second largest $-t$, so it is fixed. Let u, v choose branches i, j, with $a = (u_i - \max_{h \ne i} u_h)/2$, $b = (v_j - \max_{h \ne j} v_h)/2$ and $\delta = \|u - v\|_\infty$. If $i = j$, both the largest coordinate and the maximum of the remaining coordinates change by at most δ, so $|a - b| \le \delta$; clipping amplitudes at one preserves this inequality. If $i \ne j$, then

$$2(a+b) = u_i - \max_{h \ne i} u_h + v_j - \max_{h \ne j} v_h \le (u_i - v_i) + (v_j - u_j) \le 2\delta.$$

The distance between output points on different branches is the sum of their capped amplitudes, which is at most $a + b \le \delta$. This proves nonexpansiveness, including ties. Different tips have distance two, and all distances are at most two. The midpoint of s_1 and s_2 has first two coordinates zero and all remaining coordinates -1, so it is outside the set. The largest and second-largest coordinates require one pass; forming the output requires another. Apply Theorem 2 with $D = 2$. □

D.3 Sharp Euclidean repair coefficient

Theorem 12 (Sharp Euclidean one-step obstruction)**.** *On* $C = [-4,4]^2$ *with Euclidean norm, no deterministic causal rule independent of* T *can enforce adjacent consistency and*

$$\max_{j \le k} \|y_j - T(x_j)\|_2 \le c \max_{j \le k} \|v_j - T(x_j)\|_2$$

for every finite prefix, nonexpansive self-map T*, and estimate sequence with a fixed* $c < \sqrt{2}$*. The obstruction holds at* $x_1 = (2x_0 + y_0)/3$*, including randomized rules when both maxima are replaced by their expectations and adjacent consistency holds almost surely.*

Proof of Theorem 12. The rule maps $H_j = (x_0, v_0, y_0, \dots, x_{j-1}, v_{j-1}, y_{j-1}, x_j, v_j)$ to $y_j \in C$, independently of T, and adjacent consistency means $\|y_j - y_{j-1}\|_2 \le \|x_j - x_{j-1}\|_2$. For randomized rules, the internal seed law is fixed independently of T.

Proof. Any finite coefficient forces $y_0 = 0$ on input $x_0 = (-3, 0)$, $v_0 = 0$, by comparison with the zero map. Thus $x_1 = (-2, 0)$ and adjacent consistency requires $\|y_1\|_2 \le 1$. Fix $0 < \varepsilon \le 1$, let $a_\varepsilon = \sqrt{1 + 2\varepsilon}$, $w_\pm = (a_\varepsilon, \pm\varepsilon)$ and $v_1 = (a_\varepsilon, 0)$. Since $\|w_\pm\|_2 = 1 + \varepsilon$, the self-maps

$$T_\pm(u) = P_C\left(\frac{u_1 + 3 + \varepsilon}{1 + \varepsilon} w_\pm\right)$$

are nonexpansive. At the two queries, projection is inactive and the values are $\varepsilon w_\pm/(1+\varepsilon)$ and $w_\pm$. Both raw errors equal ε. The common observed transcript forces the same second reply for the two maps. For every feasible y, averaging squared distances and using the reverse triangle inequality gives

$$\max_{\pm} \|y - w_\pm\|_2^2 \ge \|y - v_1\|_2^2 + \varepsilon^2 \ge (a_\varepsilon - 1)^2 + \varepsilon^2.$$

Consequently the error amplification is at least

$$g(\varepsilon) = \sqrt{1 + \left(\frac{2}{\sqrt{1 + 2\varepsilon} + 1}\right)^2} \longrightarrow \sqrt{2} \quad (\varepsilon \downarrow 0).$$

For every $c < \sqrt{2}$, choose a positive $\varepsilon \le 1$ with $g(\varepsilon) > c$. One of the two maps violates the proposed promise. At $y = (1, 0)$ both squared distances equal $(a_\varepsilon - 1)^2 + \varepsilon^2$, attaining the geometric lower bound for each fixed ε. Fix any $c < \sqrt{2}$ and any randomized rule with the factor-c expected-prefix promise from Theorem 12, on the same square $C = [-4, 4]^2$. The zero-map input forces $y_0 = 0$ almost surely. Both witness maps thus have the same second-reply law Q, supported in the unit disk. Its barycenter $q = \mathbb{E}Q$ is feasible. If both expected errors were at most $c\varepsilon$, Jensen would give $\|q - w_\pm\|_2 \le c\varepsilon$ for both signs, contradicting the same geometric bound. Thus randomization does not lower this one-step coefficient. □

Lemma 13 (Sharp one-step Euclidean repair coefficient)**.** *Let $a, z, v \in \mathbb{R}^d$, $r, M \ge 0$, with $\|a - z\|_2 \le r + M$ and $\|v - z\|_2 \le M$. Euclidean projection $q = P_{B_2(a,r)}v$ satisfies*

$$\|q - a\|_2 \le r, \qquad \|q - z\|_2 \le \sqrt{2}M.$$

The constant $\sqrt{2}$ is the best uniform coefficient over these inputs for a repair that is not supplied z: the two-target construction in Theorem 12 gives the matching lower bound. If a, z also lie in a closed convex domain C, the reply $y = P_C q$ satisfies the same two bounds and belongs to C.

Proof. If v is in the ball, $q = v$ and the claim follows. Otherwise put $t = \|v - q\|_2 > 0$ and $n = (v - q)/t$. Radial projection onto a Euclidean ball gives $q = a + rn$, also when $r = 0$. Therefore

$$\langle n, z - q\rangle = \langle n, z - a\rangle - r \le \|z - a\|_2 - r \le M.$$

Expanding the square yields

$$\begin{aligned}\|q - z\|_2^2 &= \|v - z\|_2^2 - t^2 + 2t\langle n, z - q\rangle \\ &\le M^2 - t^2 + 2tM = 2M^2 - (t - M)^2 \le 2M^2.\end{aligned}$$

For the constrained version, Euclidean projection onto C fixes both a and z and is nonexpansive, so both distances can only decrease. The lower-bound targets $z = w_\pm$ with $a = 0$, $r = 1$, $M = \varepsilon$, $v = (\sqrt{1 + 2\varepsilon}, 0)$ satisfy the assumptions and force ratios approaching $\sqrt{2}$ for every feasible reply selected without knowing the target. □

This is a one-step bound with previous-reply error bounded by M. Iterating it alone does not preserve the historical raw-error maximum with constant $\sqrt{2}$, because its next premise involves the preceding repaired error.

E EXPERIMENTAL PROTOCOLS AND SUPPORTING RESULTS

E.1 STABILITY AND ESTIMATOR COMPARISONS

E.1.1 FULL-HISTORY BASELINE

The full-history comparator enforces all earlier query/reply constraints. At query x_j, replace the local interval by

$$\ell_j = \max\{0, \max_{h<j}(y_h - \|x_j - x_h\|\mathbf{1})\}, \qquad u_j = \min\{\mathbf{1}, \min_{h<j}(y_h + \|x_j - x_h\|\mathbf{1})\},$$

and clip v_j coordinatewise into it; use $P_C v_0$ at initialization. This enforces every earlier query/reply constraint. The interval is nonempty by induction: for any two earlier indices h, l, pairwise consistency and the triangle inequality give $y_{h,i} - \|x_j - x_h\| \le y_{l,i} + \|x_j - x_l\|$. Earlier replies lie in C, so cube endpoints are compatible as well. Taking the maximum of lower bounds and minimum of upper bounds preserves their order. For comparator error M, each lower endpoint is at most $T(x_j)_i + M$ and each upper endpoint at least $T(x_j)_i - M$, by the same calculation as Lemma 1. The intersection therefore preserves the same maximum-error bound. It stores $O(dk)$ values and queries and scans them in $O(dk^2)$ total work over $k+1$ stages.

E.1.2 MAPS, SAMPLING AND ALLOCATION

Maps and innovations. For the anchored shift, $T(x) = (0.8, x_1, \ldots, x_{d-1})$. For the critical cycle, let θ have equally spaced coordinates from 0.25 to 0.75, put $u = x - \theta$, and set $T(x) = P_C(\theta + \mathrm{roll}(u - 0.25u^3, 1))$. The shift is nonexpansive. For the cycle, $1 - 0.75u_i^2 \in [0.578125, 1]$ on the cube; the coordinate map, cyclic permutation and projection are all nonexpansive. Their fixed points are respectively the all-0.8 vector and θ. Neither map nor fixed point is exposed to the solver.

Independent scalar innovations U_t have one of three laws: (i) 0.1 with probability 0.9, otherwise -0.9; (ii) an independent symmetric sign times $(0.3/1.3)V$, where $\Pr(V > v) = v^{-1.3}$ for $v \ge 1$; (iii) an independent symmetric sign times $\exp(H)/8$, where $\Pr(H = h)$ is proportional to $e^{-h}h^{-3/2}$ on positive integers. The last law has finite first absolute moment and no finite moment of order $1+\delta$, $\delta > 0$. The vector noise at coordinate i is

$$\zeta_{t,i} = 0.08(0.6 + 0.4x_i)s_i(x)(1 + 0.3\,\mathrm{sign}(U_{t-1}))U_t,$$

where $s_i(x)$ alternates across coordinates and changes sign at $x_i = 0.5$; $U_0 = 0$. All factors before U_t are predictable, so the noise is centered given the full preceding history. Coordinates are deliberately correlated. Precomputed innovations remain private to the sampler.

Allocation, pairing and accounting. The 32 seeds are 40000–40031. Within a fixed law, budget and seed, methods, maps and dimensions share the latent innovation stream; adaptive queries yield different observations. Seeds within a cell are independent pseudorandom replications. Cells are not independent replications of a method comparison. The first three methods use exactly 18,792/18,711/17,982 vectors at cap 20,000 and dimensions 8/32/128; at cap 100,000 they use 98,658/98,901/98,901 vectors. The published schedule uses 15,333 and 89,271 vectors at the two caps.

Published schedule. For the published baseline, set $z_0 = x_0 = \mathbf{1}/2$ and use the largest J with $\sum_{n=1}^{J} n^4 \le N$. At stage n, average n^4 raw calls at $P_C z_{n-1}$ to obtain $\bar{Y}_n$, then set $z_n = x_0/(n+1) + n\bar{Y}_n/(n+1)$. Return $P_C z_J$. The raw states and means remain unprojected, every executed batch is complete, and actual cost is $\sum_{n=1}^{J} n^4$. This is the untuned schedule of Bravo & Contreras (2026) applied to $T \circ P_C$.

Summaries. We report mean residual and linearly interpolated empirical 90th percentiles over 32 seeds. Quantiles describe the observed residual distribution; paired inference is specified below. State and repair-work comparisons are analytical costs of the production algorithms.

E.1.3 PAIRED REPAIR EFFECTS

Pair residuals by seed and count strict one-edge wins, treating ties as nonwins. MoM wins 26–32 of 32 trials per setting; ordinary averaging wins 29–32.

Let W_c be the win count in cell c and ρ_c its strict-win probability. For all 72 cells together, use two-sided Clopper–Pearson intervals at per-cell error $\alpha/72$, where $\alpha = 0.05$. For $W_c > 0$, the lower endpoint is

$$\ell_c = \mathrm{Beta}^{-1}\Big(\frac{\alpha}{144}; W_c, 33 - W_c\Big).$$

Bonferroni's inequality gives simultaneous coverage of at least 95%. The smallest lower endpoint, 0.5064, exceeds $1/2$. These intervals concern win probabilities.

Secondary comparisons. For each of the 36 settings, also pair ordinary with robust averaging under each repair choice, and full history with one edge. For these 108 comparisons, let $\Delta_{(1)} \le \cdots \le \Delta_{(32)}$ be the ordered paired residual differences. The interval $[\Delta_{(6)}, \Delta_{(27)}]$ covers a population median with error at most $2\Pr\{\mathrm{Bin}(32, 1/2) \le 5\}$. A union bound gives at least 95% simultaneous coverage over all 108 intervals. Table 2 counts intervals above, below, or containing zero. Ordinary averaging gains under bounded noise, whereas MoM gains under the heavier tails; full history and one edge each improve some settings. These median-difference intervals complement the mean summaries in the main text.

Table 2: **Paired uncertainty for secondary comparisons.** Each row covers 12 settings. Positive and negative count simultaneous median-difference intervals strictly above and below zero; zero counts intervals containing zero. The difference is the first named method's residual minus the second's.

Noise	Paired difference	Positive	Negative	Zero
Bounded	Ordinary − MoM (one edge)	0	4	8
Bounded	Ordinary − MoM (range only)	0	7	5
Bounded	Full history − one edge	1	2	9
Pareto-1.3	Ordinary − MoM (one edge)	4	0	8
Pareto-1.3	Ordinary − MoM (range only)	12	0	0
Pareto-1.3	Full history − one edge	4	1	7
Log-only	Ordinary − MoM (one edge)	3	0	9
Log-only	Ordinary − MoM (range only)	11	0	1
Log-only	Full history − one edge	4	2	6

Figure 4 displays all 144 method/settings cells; the supplementary code regenerates the runs and summaries.

E.1.4 ESTIMATOR ABLATION AND IID COMPARISON

The supplemental development protocol was fixed before its new retained outcomes, after inspecting the original development study. The pooled-mean arms change only the aggregation of the same L equal-size block means. Seeds, innovation streams, map definitions, query allocation and update coefficients are unchanged. Thus the factorial contains 4,608 unique executions, half reused from the main study, and does not constitute a new holdout.

For the IID comparison, seeds are 41000–41031 and the perturbation is $\zeta_t = 0.08U_t s$, where s is a fixed alternating-sign vector and the independent innovations are $U_t = 0.1$ with probability 0.9, otherwise -0.9. The perturbation is query independent, centered and bounded, with $\mathbb{E}\|\zeta_t\|_2^2 = 0.000576d$. The published method is applied to $S = T \circ P_C$ on $\mathbb{R}^d$: query at $P_C x$, retain the raw Halpern update, and project the final output. The extension is nonexpansive, and

$$\|P_C x - T(P_C x)\| \le \|x - S(x)\|,$$

so final projection cannot increase its residual. Its own complete n^4 batches are retained.

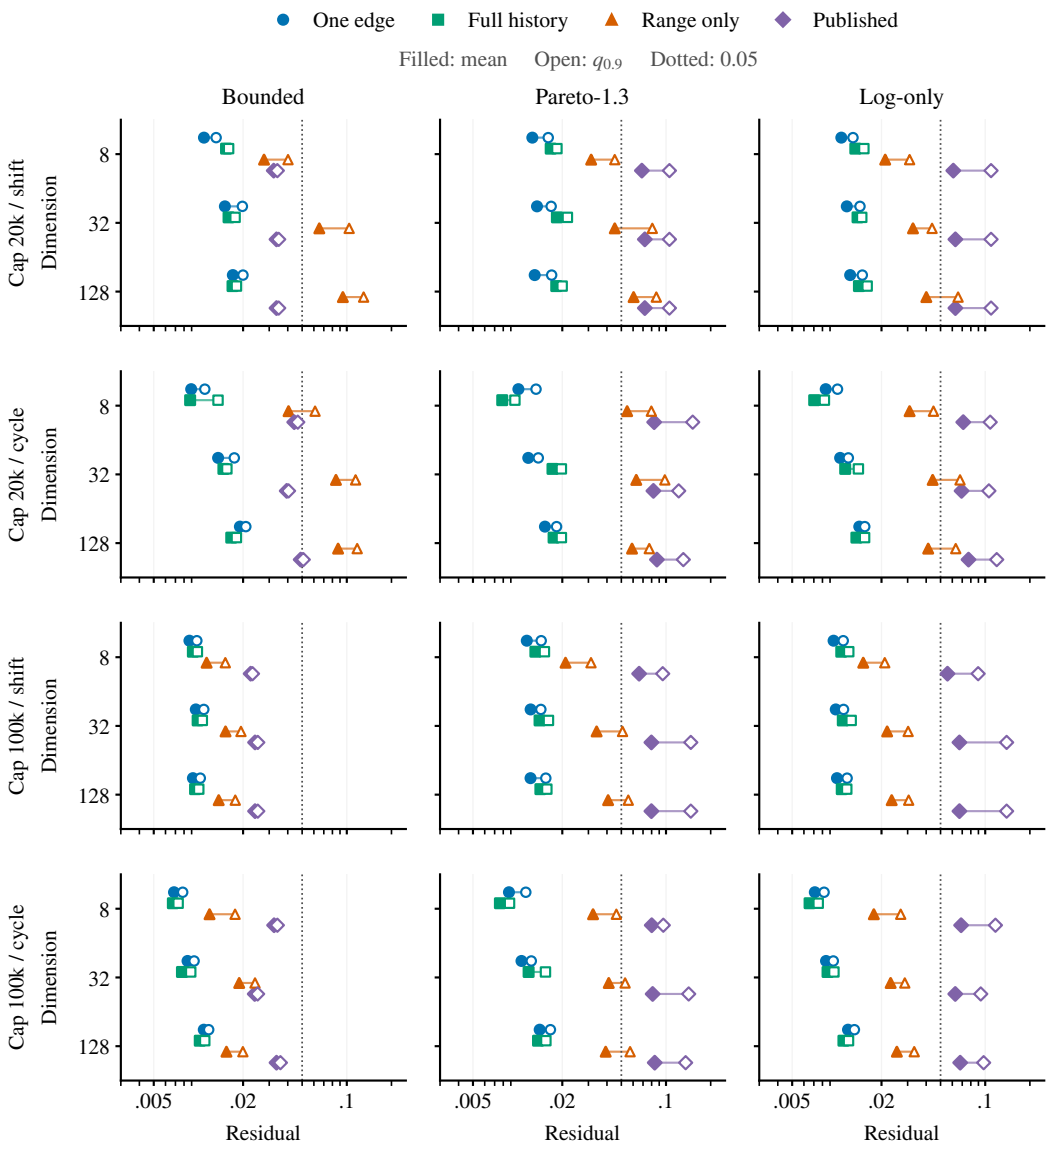


Figure 4: **Mean and tail residuals across all 144 method/settings cells.** Filled markers show means and open markers show empirical $q_{0.9}$ over 32 seeds per cell. Segments connect these two descriptive statistics. Rows identify nominal cap and map, columns identify noise law, and each group compares four methods at one dimension. All panels share a logarithmic residual axis; the dotted line marks 0.05. The three matched methods share actual costs within each setting. The allocation paragraph reports actual counts, including the published schedule's different counts; the separate IID study matches the latter's oracle premises.

Table 3: **Estimator preference depends on the noise law.** Each range gives the minimum and maximum of 12 matched-setting ratios (two maps, three dimensions, two budgets). Ratios below one favor ordinary averaging. Each mean and empirical $q_{0.9}$ uses 32 seeds; these are ratios of cell statistics, not pooled residuals. Event counts combine the 384 runs in each row and estimator.

		Ordinary / MoM ratio range		Events above 0.05	
Noise	Repair	Mean	$q_{0.9}$	MoM	Ordinary
Bounded	One edge	0.61–0.99	0.58–1.02	0	0
	Range only	0.24–0.90	0.25–0.87	126	0
Pareto-1.3	One edge	1.13–1.43	1.19–1.65	0	0
	Range only	4.32–6.88	4.12–13.12	110	360
Log-only	One edge	1.14–1.29	1.23–1.69	0	0
	Range only	3.16–8.01	3.98–14.60	24	303

Table 4: **IID bounded-noise comparison across all eight settings.** Each mean uses 32 seeds. Bold entries identify the lowest mean within a row. One edge, full history and range-only repair share the calls in the "Matched" column; the published schedule executes its own complete batches under the same cap. The supplementary code regenerates full cell summaries, upper quantiles and individual records.

			Mean residual $\times 10^3$				Actual calls	
Map	d	Cap (k)	One edge	History	Range	Published	Matched	Published
Shift	8	20	**9.49**	14.34	19.78	32.01	18,792	15,333
Shift	8	100	**8.57**	9.23	10.84	22.32	98,658	89,271
Shift	128	20	16.36	**16.23**	60.98	33.67	17,982	15,333
Shift	128	100	**9.86**	10.02	14.31	24.20	98,901	89,271
Cycle	8	20	8.39	**8.16**	25.44	44.74	18,792	15,333
Cycle	8	100	6.33	**6.05**	11.21	32.36	98,658	89,271
Cycle	128	20	**15.86**	16.94	46.10	47.75	17,982	15,333
Cycle	128	100	10.31	**9.92**	13.00	32.83	98,901	89,271

E.2 Bellman tasks, sampling and shared controls

Tasks and observations. We evaluate the normalized Bellman map $T(v) = (1-\gamma)r + \gamma Pv$ on $[0,1]^d$, with $\gamma = 0.95$. A navigation grid of size $\sqrt{d} \times \sqrt{d}$ has an absorbing goal with reward one and zero reward elsewhere. At other states, a fixed policy moves toward the goal with probability 0.65, slips backward with probability 0.15, moves down with probability 0.1, and waits with probability 0.1; boundary transitions are merged. The queue has states $s = 0, \ldots, d-1$, reward $1 - s/(d-1)$, and transition probabilities 0.45, 0.5, and 0.05 for increment, decrement, and stay, with boundary clipping. All transition rules are given explicitly in the accompanying generator.

At each vector call, the simulator samples one next state and reward per state. Rewards are $r_s + Z_s/2$. In the bounded case, Z_s are independent uniforms on $[-\sqrt{3}, \sqrt{3}]$. In the Pareto case, $Z_s = A\sigma_s$, where $A = (0.3/1.3)U$, $\Pr(U > u) = u^{-1.3}$ for $u \geq 1$, and independent symmetric signs σ_s share the current amplitude. Samples are fresh across calls. The returned value is $(1-\gamma)(r_s + Z_s/2) + \gamma v_{S'_s}$, whose conditional mean is $T(v)_s$.

The variance-burst law adds $10H_n\xi_{n,s}$ to the bounded reward noise, where $\xi_{n,s}$ are independent symmetric signs and the latent chain H_n starts cold, enters with probability 0.005 and exits with probability 0.05 per raw vector. Fresh signs preserve conditional centering; solvers receive map values without regime labels.

Allocation and controls. All methods start at the cube center. For nominal cap N, define

$$K(q) = 1 + \left\lceil \frac{\log(N+1)}{-\log q} \right\rceil .$$

Raw adjacent averaging, stationary repair, both minimum-motion policies, tail averaging and fresh-edge Picard use $q = \gamma$. Calibrated damping uses $q = (1 + \gamma)/2$ and update $x_{t+1} = (x_t + b_t)/2$. All methods use the smallest odd $L \geq \max\{3, 2\log_2(dK((1 + \gamma)/2)/0.05)\}$. With $S = \lfloor N/L \rfloor$, initialize each $B_t = 1$ and distribute $S - K(q)$ additional block samples proportionally to $q^{K(q)-1-t}$, using largest-remainder rounding. Each b_t is a cube-projected coordinate median of L block means of size B_t.

Raw averaging uses $x_{t+1} = (b_t + b_{t-1})/2$; stationary repair substitutes y_t, y_{t-1} and imposes no movement floor. Both initialize with the current value. Tail averaging runs projected Picard, $x_{t+1} = b_t$, and returns the mean of updates with $t \geq \lfloor 3K/4 \rfloor$. Fresh-edge Picard is a comparison rule that clips b_t around b_{t-1} with radius $\gamma\|x_t - x_{t-1}\|$ and returns its final update. The two proposed policies use (5) and return the last repaired average. Damping and averaging are generic controls; fresh-edge Picard is our raw-centered comparison, distinct from the fixed-reply repair in (1).

Each run consumes $L\sum_t B_t = L\lfloor N/L \rfloor$ vectors. The actual counts at caps 20k/100k are 19,968/99,996 for $d = 64$, and 19,995/99,975 for $d = 256$. A vector costs d transition/reward observations. The true map and fixed point are available only to the generator and evaluator.

E.3 Recovery and query-policy comparison

The normal grid crosses navigation/queue, $d = 64, 256$, bounded/Pareto rewards and caps 20k/100k. Seven methods and 16 seed units give 1,792 runs. The estimate-driven proposal was selected on 288 separate development runs; these seeds are excluded from the comparison. Initial and terminal corruption tests use both tasks and laws at $d = 64$, cap 20k, and 16 seed units, giving 896 runs. They replace the completed first estimate by the current query, or the completed final estimate by the all-ones vector. All samples are still charged. These tests isolate recovery from a corrupted estimate.

A separate 896-run study uses the same seven methods and 16 fresh seed units over the eight task/dimension/budget settings with variance bursts. It uses the variance-burst law defined in Appendix E.2. Corruption acts on completed estimates, whereas variance bursts act on raw samples.

Table 1 reports maximum-norm value-error ratios. Within each seed we average paired log ratios across settings, then exponentiate their mean over seeds. Normal and burst intervals use the paired-log Student construction with 15 degrees of freedom; corruption columns are geometric-mean ratios. Common seeds couple settings, so the seed is the independent inference unit. The table includes every method and all settings from these comparisons.

E.4 Value accuracy under variance bursts and sample reuse

Task and sampling model. The value-accuracy study uses the Bellman tasks, vector sampling model, and bounded and variance-burst reward laws in Appendix E.2, with $d = 64$, $\gamma = 0.95$ and nominal budgets 20k and 100k. The eight settings cross both tasks, budgets and laws, using 16 seeds independent of method development.

Methods and cost. Raw adjacent averaging, calibrated damping, tail averaging and both repaired query policies use the query rules, allocation and MoM estimators in Appendix E.2. The two empirical-model controls use all transition counts to estimate P, and estimate rewards using either the overall sample mean or the coordinate median of 13 block means. Both clip estimated rewards to $[0, 1]$ and solve the resulting Bellman equation. Model fitting stores $O(d^2)$ transition counts; the repair recurrence stores $O(d)$ vectors. Each non-VR run uses the same actual full-vector count, including a final actual query at its fixed returned point. That final reply does not change the returned value estimate.

The practical VR-GHAL implementation retains anchor norm-median estimation and clipped stochastic differences, using the previously selected $\beta = 0.75$ and batch multiplier 10^{-4} (Diakonikolas, 2026). These are practical constants rather than the reference's theorem-calibrated allocation. We report both a common budget cap and additional repair/averaging runs charged exactly its actual vector usage. In the latter, 99.773%–99.999% of the repair budget is used by the solver; the remaining samples form the final diagnostic reply.

Metric and uncertainty. The task loss is $d^{-1}\|\widehat{v} - v^*\|_2^2$. For each seed, we average log loss ratios across the two families and two budgets within a noise law, then exponentiate the mean over 16 seeds. Intervals are paired-log Student intervals with 15 degrees of freedom. The prespecified primary contrast is repaired proposals versus raw adjacent averaging under burst noise; the remaining intervals are pointwise secondary comparisons.

Table 5: Value-accuracy tradeoffs across all eight settings. The numerator is repaired proposals in the first seven rows; the final two rows name both methods and match actual VR oracle calls. Ratios below one favor the numerator. All 1,280 outputs and cell-level MSE values are retained in the experiment records.

Denominator	Baseline noise MSE ratio [95% interval]	Variance bursts MSE ratio [95% interval]
Raw adjacent average	1.25 [1.16, 1.35]	0.695 [0.489, 0.988]
Fresh proposals	1.26 [1.17, 1.36]	0.702 [0.491, 1]
Calibrated damping	0.545 [0.469, 0.634]	0.433 [0.312, 0.601]
Tail average	0.617 [0.548, 0.695]	0.543 [0.385, 0.765]
Robust reward model	5.27 [4.07, 6.83]	0.341 [0.221, 0.528]
Mean reward model	7.5 [5.24, 10.7]	0.443 [0.28, 0.701]
Practical VR-GHAL (cap)	0.00032 [0.000275, 0.000373]	0.000422 [0.000328, 0.000541]
Repair / VR (actual cost)	0.000346 [0.000283, 0.000424]	0.000434 [0.000348, 0.00054]
Raw average / VR (actual cost)	0.000284 [0.000236, 0.000342]	0.000638 [0.000421, 0.000968]

The actual-cost VR comparisons charge repair and raw averaging the same number of observations used by that configuration.

E.5 Burst amplitude and duration ablation

Fixed-method regime study. We vary raw reward burst amplitude over $5, 10, 20$ and mean hot duration over $5, 20, 80$ vector rounds. The exit probability is the inverse duration and the entry probability is one tenth of it, keeping stationary hot probability $1/11$ fixed. All chains start cold. A no-burst reference, both Bellman families, 16 independent evaluation seeds and the same seven methods give 2,240 outputs at dimension 64 and nominal budget 100k. The sampler agrees exactly with Appendix E.4 at amplitude 10 and duration 20. It uses independent coordinate signs and recurrent variance bursts.

The three prespecified pooled contrasts are repaired-with-movement/raw, repaired-without-movement/raw, and with/without movement. Seed-level mean log MSE ratios pool all 18 nonzero-burst family/settings. Bonferroni intervals at 98.333% per contrast control the three comparisons at nominal 95% family coverage under the Student approximation. Their ratios and intervals are respectively

$$0.583\ [0.422, 0.806], \qquad 0.447\ [0.323, 0.618], \qquad 1.305\ [1.089, 1.563].$$

Thus repair supplies burst attenuation, while movement incurs a precision cost relative to stationary repair. Stationary repair retains the early-error lock in Appendix A.1. Stronger and longer bursts generally increase repair's advantage over averaging; all three mild-amplitude queue cells favor raw averaging over the movement method. The complete grid appears in Table 6; all damping, tail-average and fitted-model outputs are retained with it.

Table 6: Complete regime grid: paired geometric MSE ratios over 16 seed units at identical actual sample cost. "Move" is repaired averaging with minimum motion; "Repair" omits movement; "raw" omits repair; "model" fits transitions and ordinary-mean rewards. Amplitude zero is the baseline reference; duration is then inactive. Individual-cell intervals and absolute MSE for all seven methods are in the accompanying results.

Model	Amplitude	Duration	Move/raw	Repair/raw	Move/repair	Move/model
Navigation	0	20	1.043	1.094	0.954	5.736
Navigation	5	5	0.926	0.913	1.014	3.366
Navigation	5	20	0.632	0.612	1.032	1.428
Navigation	5	80	0.490	0.402	1.219	0.974
Navigation	10	5	0.845	0.812	1.041	1.932
Navigation	10	20	0.386	0.368	1.049	0.488
Navigation	10	80	0.229	0.130	1.765	0.455
Navigation	20	5	0.709	0.550	1.290	0.924
Navigation	20	20	0.237	0.222	1.069	0.135
Navigation	20	80	0.128	0.058	2.196	0.158
Queue	0	20	1.725	2.000	0.862	3.758
Queue	5	5	1.114	1.199	0.929	1.257
Queue	5	20	1.120	1.322	0.847	0.767
Queue	5	80	1.121	0.993	1.129	0.886
Queue	10	5	0.904	0.901	1.003	0.644
Queue	10	20	0.770	0.802	0.960	0.255
Queue	10	80	0.744	0.337	2.208	0.477
Queue	20	5	0.664	0.455	1.460	0.339
Queue	20	20	0.640	0.472	1.357	0.095
Queue	20	80	0.432	0.100	4.302	0.238