PRIVATE-RESPONSE COORDINATION: INFORMATION LIMITS AND FEASIBLE DECISIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

A coordinator can query each party's allocation at a price without observing its objective. We study when these responses suffice to produce and certify a feasible joint decision. We construct coupled strongly convex objectives with different optima that admit the same error-bounded vector response at every price. Their two-world minimax regret equals the conditional-error budget, even with arbitrary adaptive queries and unlimited memory. A comparison family separately gives sharp query, state, and ambiguity bounds. We then certify feasible candidates from response positions, public smoothness, and local error bounds without fitting their objectives. On identical transcripts and recovered averages, the position certificate verifies target accuracy after a mean of 8.6 and 20.3 high-fidelity calls for quadratic and log-cosh objectives. Neither recovery-based certificate does so within 512 calls. Coupled nonlinear studies show how response fitting and validation complement each other. The results distinguish irreducible response ambiguity from avoidable certificate conservatism.

1 INTRODUCTION

Several regions must agree on an electricity schedule. Each knows its costs and operating limits, while a coordinator can ask what it would supply at a given price. If the proposed supplies do not meet demand, the coordinator changes the price and asks again. What does another local solve contribute to the final allocation?

Decomposition coordinates local optimization through prices (Dantzig & Wolfe, 1960; Palomar & Chiang, 2006; Boyd et al., 2011), but a response budget alone does not answer this question. Distinct preferences may admit the same responses. A price update may overlook information those responses reveal. A sufficient certificate may remain inconclusive after the decision has become accurate. Additional queries need not remove all three obstacles. We ask *when responses suffice for a good feasible decision, and when apparent limits reflect how those responses are used*.

We study *private-response coordination*: parties retain their local objectives and respond to public queries, while their decisions satisfy a shared constraint. Query depth counts local response calls, persistent state measures retained information, and fidelity describes response accuracy. Private denotes restricted access to the local problems, not protection against inference from their responses.

1. Information limits. Coupled strongly convex objectives with different optima admit the same error-bounded response at every price (Figure 1(a)). Their exact two-world regret floor holds for arbitrary adaptive queries and decoders, even with unlimited memory. A comparison family separately gives sharp query–state–ambiguity bounds.

2. Response-position certificates. We certify feasible candidates from complete responses, public smoothness, and error bounds, without objective values, gradients, or a fitted model (Figure 1(b)). Worst-case regret over transcript-consistent problems separates information uncertainty from certificate looseness. Both bounds coincide at the construction's common zero-price response.

3. Candidate quality and certification. On identical histories and candidates, the position certificate verifies accuracy earlier than two recovery-based bounds. Stable probes and coupled nonlinear tests compare fitting with feedback. An initialization-based error bound increases acceptance without changing candidates or responses.

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.

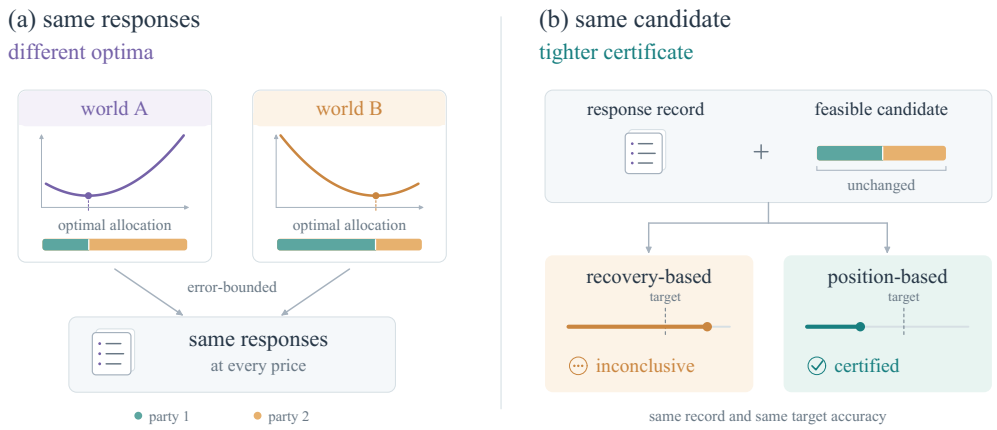


Figure 1: Response ambiguity and certificate tightness.

2 Related Work

Coordination and certification. Prices coordinate separable objectives in resource allocation and energy networks (Kelly et al., 1998; Palomar & Chiang, 2006; Kraning et al., 2014; Boyd et al., 2011). Averaged primal recovery bounds feasibility and objective error under exact (Nedić & Ozdaglar, 2009) or inexact (Necoara & Nedelcu, 2014) dual responses. We use this analysis as a reference. Inexact-oracle accuracy certificates already support stopping and primal recovery (Gladin et al., 2023). Computable duality gaps offer another route (Walwil & Fercoq, 2025). Our certificate instead checks an already feasible candidate from complete conditional responses and public curvature and error bounds, without observed objective values or gradients.

Information in responses. Allocation communication complexity (Nisan & Segal, 2006; Dobzinski et al., 2019), revealed preferences (Ji et al., 2018), and inverse optimization (Chen & Kılınç-Karzan, 2021) study information recovered from choices. Our scalar analysis uses comparison search, quantization, and distributional lower bounds (Nowak, 2011; Gray & Neuhoff, 1998; Yao, 1977). Memory-constrained optimization (Blanchard et al., 2023a;b) motivates retained-state limits. The exact law here concerns terminal resolution in the specified comparison families. Complete vector responses motivate the separate all-price construction and identification baselines.

Imperfect feedback. Imprecise comparisons (Ajtai et al., 2015) and inexact first-order oracles (Devolder et al., 2014; Schmidt et al., 2011) distinguish persistent ambiguity from numerical approximation. Tunable-oracle schedules address work allocation (Van Dessel & Glineur, 2023). We examine what the resulting records can identify and certify. Smooth strongly convex interpolation (Taylor et al., 2016) characterizes consistency from values and gradients. Our gradient-compatible set imposes tractable necessary conditions when neither is observed. Response-only access does not supply the formal confidentiality of differentially private allocation (Chen & Chua, 2024).

3 Response Access and Decision Quality

Party $m \in \{1, \dots, M\}$ chooses $x_m \in \mathcal{X}_m \subseteq \mathcal{V}_m$ with objective f_m. Public linear maps $A_m : \mathcal{V}_m \to \mathcal{Y}$ define the shared interface. All spaces are finite dimensional. For a public coupling set $\mathcal{C}$,

$$F^\star = \min_{x \in \mathcal{F}} F(x), \qquad F(x) = \sum_m f_m(x_m), \qquad Ax = \sum_m A_m x_m,$$
$$\mathcal{F} = \{x : x_m \in \mathcal{X}_m,\ Ax \in \mathcal{C}\}. \tag{1}$$

We assume an attained finite optimum. The public description I_{pub} contains the spaces, maps, coupling set, query domain, and stated bounds. Each complete pair $(f_m, \mathcal{X}_m)$ remains local unless explicitly made public.

Queries and fidelity. At price $\lambda \in \Lambda \subseteq \mathcal{Y}$, the conditional objective is $\phi_m(x_m; \lambda) = f_m(x_m) + \langle \lambda, A_m x_m \rangle$, and the exact response is

$$R_m(\lambda) \in \arg\min_{x_m \in \mathcal{X}_m} \phi_m(x_m; \lambda), \qquad y_m(\lambda) = A_m R_m(\lambda). \tag{2}$$

Queried minima are attained, with a fixed measurable tie rule. Normally only y_m is returned. The constructive certificate below explicitly uses complete local responses. An approximate locally feasible response $\widetilde{x}_m$ has

$$\delta_m = \phi_m(\widetilde{x}_m; \lambda) - \min_{\mathcal{X}_m} \phi_m(\cdot; \lambda), \qquad e_m = A_m \widetilde{x}_m - y_m(\lambda). \tag{3}$$

A declared solver and public bounds may supply $\widehat{\delta}_m \geq \delta_m$. Local work J is not itself an error tolerance. Every round queries all parties once. Q rounds use Q calls per party and MQ responses. Any further private-oracle invocation, including validation, is counted. Computation from public information and received responses requires no additional private call, but its computational cost is reported separately.

Retention and output. A full-transcript protocol chooses prices from $T_{q-1} = ((\lambda^j, y^j))_{j<q}$ and instance-independent randomness. A strict (Q, S) protocol retains a state with at most 2^S values, through which every response-dependent record used later must pass. Its feasible terminal output is decoded from this state. The public round index and program are free. This restriction measures retained information, not temporary workspace. Real-valued price and local averaging protocols form a separate class. Appendix A specifies both conventions. For a feasible decision, static coordination regret is $\mathrm{Reg}(x) = F(x) - F^\star$, not cumulative online loss. An infeasible response has a separate coupling violation and signed objective gap. Any recovery P must be executable from its declared information. Projecting an interface need not recover a complete feasible local decision. We report raw violation and recovery distance alongside $\mathrm{Reg}(Px)$. Below, $\mathcal{C} = \{b\}$, $\mathcal{X} = \prod_m \mathcal{X}_m$, and the dual lower bound is

$$d(\lambda) = \inf_{x \in \mathcal{X}} \{F(x) + \langle \lambda, Ax - b \rangle\}, \qquad h = -d. \tag{4}$$

4 Information Limits

4.1 An error floor for every price strategy

Approximate vector responses can be identical even when the optimal feasible allocations differ. The next construction permits arbitrary adaptive prices, unlimited transcript retention, and any feasible terminal decoder.

Theorem 4.1 (All-price indistinguishability). *Let $\mathcal{X} = \prod_m \mathcal{X}_m$ be nonempty, closed, and convex. Suppose $Ac = b$, $Av = 0$, $v \neq 0$, and $c \pm v \in \mathcal{X}$. Choose $B = \mathrm{diag}(B_1, \ldots, B_M) \succ 0$ and define*

$$F_\sigma(x) = \tfrac{1}{2} \|x - c - \sigma v\|_B^2, \qquad \delta_m = \tfrac{1}{2} \|v_m\|_{B_m}^2, \quad \delta = \sum_m \delta_m, \quad \sigma \in \{-1, +1\}. \tag{5}$$

The complete response function

$$R_0(\lambda) = \arg\min_{x \in \mathcal{X}} \{\tfrac{1}{2} \|x - c\|_B^2 + \langle \lambda, Ax - b \rangle\} \tag{6}$$

has local conditional error at most δ_m in both worlds at every price. Hence every adaptive, possibly randomized response-only protocol with a transcript-measurable feasible output satisfies

$$\max_{\sigma \in \{-1,+1\}} \mathbb{E}\, \mathrm{Reg}_\sigma(\widehat{x}) \geq \delta. \tag{7}$$

For the two-world experiment, this minimax value is exactly δ, attained by outputting c.

The matrices B_m may be nondiagonal, so local coordinates can be coupled. For $y = R_0(\lambda)$, strong convexity and optimality give, with $\phi_\sigma(x; \lambda) = F_\sigma(x) + \langle \lambda, Ax - b \rangle$,

$$\phi_\sigma(y; \lambda) - \phi_\sigma(x; \lambda) \leq -\tfrac{1}{2} \|x - y\|_B^2 + \sigma \langle Bv, x - y \rangle \leq \delta. \tag{8}$$

The same argument holds locally with δ_m. Since a common oracle is fixed for every price, coupling the protocol randomness produces identical transcripts and output distributions. Their average regret is $\frac{1}{2}\mathbb{E} \|\widehat{x} - c\|_B^2 + \delta$. Appendix H.1 gives the full proof, including data-dependent stopping and a coupled box example. The obstruction concerns the class of oracles with fixed conditional-error caps, not every numerical solver within it. Requesting a smaller cap changes the interface. Under the fixed caps, neither another price nor additional memory removes the ambiguity.

4.2 Queries, retained state, and comparison ambiguity

The all-price construction isolates irreducible response error. A comparison model separately quantifies the effects of limited querying and retention. Consider

$$\min_{x_1,x_2\in[0,1]} \beta x_1 + \tfrac{1}{2}x_2^2 \quad \text{subject to } x_1 + x_2 = 1, \qquad \beta \in [0,1]. \tag{9}$$

At price $\lambda = -p$, $p \in [0,1]$, select

$$y_1(p;\beta) = \mathbf{1}\{p > \beta\}, \qquad y_2(p) = p. \tag{10}$$

The optimum is $(1-\beta, \beta)$, and every feasible output $(1-\widehat{\beta}, \widehat{\beta})$ has

$$\mathrm{Reg}_\beta(\widehat{x}) = \tfrac{1}{2}(\widehat{\beta} - \beta)^2. \tag{11}$$

For public $\tau \in [0, 1/2]$, replace the first response by the set-valued comparison

$$\mathcal{Y}_\tau(p;\beta) = \begin{cases} \{0\}, & p < \beta - \tau, \\ \{1\}, & p > \beta + \tau, \\ \{0,1\}, & |p-\beta| \le \tau. \end{cases} \tag{12}$$

A nonanticipatory selector may choose any admissible bit. Let $\mathfrak{D}^{\mathrm{dz}}_{Q,S,\tau}$ and $\mathfrak{R}^{\mathrm{dz}}_{Q,S,\tau}$ denote deterministic and randomized minimax regret over coefficients and selectors for strict (Q,S) protocols.

Theorem 4.2 (Joint query, state, and ambiguity limit)**.** *For integers $Q, S \ge 0$, set $n = \min\{Q,S\}$ and $U_n(\tau) = \frac{1}{2}[\tau + (\frac{1}{2} - \tau)2^{-n}]^2$. Then*

$$\mathfrak{D}^{\mathrm{dz}}_{Q,S,\tau} = U_n(\tau), \tag{13}$$

$$\max\left\{\frac{1}{24\,4^n},\ \frac{\tau^2}{2} + \frac{(1-2\tau)^2}{24\,4^n}\right\} \le \mathfrak{R}^{\mathrm{dz}}_{Q,S,\tau} \le U_n(\tau). \tag{14}$$

Thus $\mathfrak{R}^{\mathrm{dz}}_{Q,S,\tau} = \Theta(4^{-n} + \tau^2)$ with universal constants, tending to $\tau^2/2$ as $n \to \infty$.

Robust bisection attains radius $\tau + (1/2 - \tau)2^{-n}$. Conversely, the response $\mathbf{1}\{p > c\}$ is admissible for both $\beta = c - \tau$ and $c + \tau$, for every $c \in [\tau, 1-\tau]$. A protocol must locate the center while unable to distinguish the endpoints. Quantizing these pairs proves the lower bounds (Appendix B.2). Exact-response Bayes and deterministic constants, finite-candidate recovery, and vector/work extensions are given in Appendices B.1–C. The exact state law is specific to these comparison families. Theorem 4.1 isolates a separate fidelity obstruction for coupled vector responses.

5 Certifying a Feasible Decision from Responses

A conservative certificate may miss information even when no further query is needed. A response is not an arbitrary local point: at its queried price, no feasible local move improves its conditional objective by more than the declared error. We use that fact to check a feasible candidate without reconstructing its objective.

5.1 The response-position certificate

Suppose the concatenated local decisions $y \in \mathcal{X}$ are returned at price λ, rather than only their interface images. Assume the public bound

$$\psi_\lambda(y) - \min_{x\in\mathcal{X}} \psi_\lambda(x) \le \widehat{\delta}, \qquad \psi_\lambda(x) = F(x) + \langle \lambda, Ax - b\rangle. \tag{15}$$

Let F be convex with L-Lipschitz gradient, $L > 0$, on $\mathcal{X}$. Fix any $z \in \mathcal{X}$ satisfying $Az = b$, and write $d_z = z - y$. The response rules out conditional gradients that would allow too large a decrease within the feasible set. For a public box $\mathcal{X} = [0,u]$ and public bounds $|\nabla_i F(x)| \le G_i$, these restrictions define

$$h_y(g) = \max_{w\in\mathcal{X}-y} \{-\langle g, w\rangle - \tfrac{L}{2}\|w\|^2\},$$
$$\mathcal{G}_y = \{g : |g_i - (A^\top\lambda)_i| \le G_i,\ h_y(g) \le \widehat{\delta}\}. \tag{16}$$

Proposition 5.1 (Response-position certificate)**.** *Under the preceding assumptions,*

$$\mathrm{Reg}(z) \le H_{\mathrm{pos}}(y, \lambda, z) := \widehat{\delta} + \sup_{g \in \mathcal{G}_y} \langle g, d_z \rangle + \tfrac{L}{2} \|d_z\|^2 . \tag{17}$$

For $\mathcal{X} = \mathbb{R}^n$ under the same global smoothness and error conditions,

$$\mathrm{Reg}(z) \le \left(\sqrt{\widehat{\delta}} + \sqrt{L/2}\,\|z - y\|\right)^2. \tag{18}$$

Why the certificate is valid. Set $g = \nabla\psi_\lambda(y)$. For every $w \in \mathcal{X} - y$, near-optimality and smoothness imply

$$-\widehat{\delta} \le \psi_\lambda(y + w) - \psi_\lambda(y) \le \langle g, w \rangle + \tfrac{L}{2} \|w\|^2 . \tag{19}$$

Hence $g \in \mathcal{G}_y$. Weak duality bounds feasible regret by $\psi_\lambda(z) - \min \psi_\lambda$, and the smooth upper model gives (17). On $\mathbb{R}^n$, choosing $w = -g/L$ gives $\|g\|^2 \le 2L\widehat{\delta}$, proving (18). That step may be infeasible on a box. The constrained set retains boundary information (Appendix H.2).

Boundary information. For an exact box response and $u_i > 0$, $g_i = 0$ at an interior coordinate, $g_i \ge 0$ at its lower bound, and $g_i \le 0$ at its upper bound. An exact boundary response need not have a zero conditional gradient. The position bound retains these one-sided restrictions when evaluating a feasible displacement.

Computing the certificate. Separate local error and smoothness bounds give a sum of local certificates even when each f_m couples its coordinates. For one local block, suppress its index and set $a = A_m^\top \lambda$, $d_z = z - y$, $\ell_i = a_i - G_i$, and $r_i = a_i + G_i$.

1. Bound the support. For any $t \ge 0$, compute

$$h_i(s) = \max_{-y_i \le w \le u_i - y_i} \{-sw - Lw^2/2\},$$
$$U_y(t) = t\widehat{\delta} + \sum_i \max_{\ell_i \le s \le r_i} \{(d_z)_i s - t h_i(s)\}, \qquad t \ge 0.$$

Weak Lagrange duality gives $\sup_{g \in \mathcal{G}_y} \langle g, d_z \rangle \le U_y(t)$. Each tested multiplier therefore supplies a valid upper bound, without requiring the multiplier search to converge.

2. Evaluate scalar maxima. For a trial s, set

$$w_i(s) = \mathrm{clip}(-s/L, -y_i, u_i - y_i), \qquad h_i(s) = -s w_i(s) - L w_i(s)^2/2.$$

For $t > 0$, evaluate $(d_z)_i s - t h_i(s)$ at ℓ_i, r_i, and $\mathrm{clip}(L(d_z)_i/t, \ell_i, r_i)$, and take the largest value. At $t = 0$, the two endpoints suffice.

3. Combine valid bounds. For a nonempty finite multiplier set $\mathcal{T} \subseteq [0, \infty)$, the local gap is at most $\widehat{\delta} + \min_{t \in \mathcal{T}} U_y(t) + L\|d_z\|^2/2$. Summing local bounds at a common price certifies the joint candidate. Appendix H.2 derives the calculation and specifies the tested grid. Central calculation uses complete response positions, but the feasible candidate may come from feedback, averaging, or fitting. No prescribed recovery map or constant K is required.

5.2 Information uncertainty and certificate looseness

For a transcript T and fixed feasible candidate z, let $\mathcal{W}(T)$ be the worlds consistent with the public assumptions and observed responses. Define

$$C(T, z) = \sup_{\omega \in \mathcal{W}(T)} \mathrm{Reg}_\omega(z). \tag{20}$$

Every uniformly valid bound $H(T, z)$ is at least $C(T, z)$. For the true world ω_0,

$$H(T, z) - \mathrm{Reg}_{\omega_0}(z) = \underbrace{H(T, z) - C(T, z)}_{\text{certificate looseness}} + \underbrace{C(T, z) - \mathrm{Reg}_{\omega_0}(z)}_{\text{remaining information uncertainty}} . \tag{21}$$

In Theorem 4.1, a transcript containing the zero-price response c and error cap δ gives $C(T, c) = \delta$: the cap upper-bounds regret and the two worlds attain it. The position certificate has zero displacement and also equals δ. In general $\mathcal{G}_y$ is a necessary-condition relaxation, not an exact computation of C or full multi-query function consistency (Taylor et al., 2016).

Comparison with recovery-based certificates. Averaged inexact-dual recovery (Nedić & Ozdaglar, 2009; Necoara & Nedelcu, 2014) gives $O(Q^{-1} + \bar{\delta}_Q)$ regret under strong primal convexity and an executable controlled recovery, where $\bar{\delta}_Q$ is the average aggregate conditional-error bound. Its stopping bound H_Q pays a global recovery penalty $K \|A\bar{x}_Q - b\|$. The coordinate version H_Q^{coord} uses the actual displacement. Appendix D proves both bounds and $H_Q^{\mathrm{coord}} \le H_Q$ for the tested constants. All three certificates can check the same recovered average, isolating differences in certification from differences in candidate quality.

Using more of the same record. For any retained indices $\mathcal{I}_Q$, each computed position bound $\widehat{H}_q$ can check the same current candidate z_Q:

$$\mathrm{Reg}(z_Q) \le \widehat{H}_{\mathrm{hist}}(z_Q) = \min_{q \in \mathcal{I}_Q} \widehat{H}_q(z_Q). \tag{22}$$

Each term must use a common price across parties. Historical bounds for different candidates must be recomputed. Expanded multiplier search and this reuse require computation and memory, but no new response. The experiments measure whether the tighter bounds justify this additional work.

Local computation with additional messages. At a common price, party m can retain its response and send a scalar h_m bounding its candidate's local conditional gap. For a jointly feasible candidate, $\mathrm{Reg}(z) \le \sum_m h_m$. With convex local sets, retained feasible responses can form $z_m = \sum_q \alpha_{mq} r_m^q$. The coordinator uses only interface images to choose $\alpha_{mq} \ge 0$, $\sum_q \alpha_{mq} = 1$, and $\sum_{m,q} \alpha_{mq} A_m r_m^q = b$. Convexity ensures local feasibility. The last equality ensures joint feasibility. Such weights exist only if b lies in the sum of the observed interface convex hulls. Appendix H.3 proves both constructions. Feasible recovery and quality certification remain separate. This extension adds certificate messages, local computation, retained responses, and any validation calls to the accounting. The experiments below use the central full-response interface.

6 Experiments

We first keep responses and outputs fixed to compare certificates, then vary how candidates are proposed. All reported targets concern feasible regret, separately from raw imbalance and recovery distance.

Setup and accounting. The separable studies use four parties, each with 24 coordinates, public boxes $0 \le x_{mt} \le u_{mt}$, and balance $\sum_m x_{mt} = d_t$. Private quadratic objectives have curvatures a_{mt} and centers p_{mt}. Log-cosh objectives add $\gamma_{mt} \log\cosh(x_{mt} - p_{mt})$. Each study has 200 instances in ten seed blocks of 20, paired across methods. Calls return complete local vectors. J derivative bisections plus two boundary checks cost $J + 2$ comparisons per coordinate. Certificates use public $V = \|u\|_F^2$ and target $\mathrm{Reg}/V \le 10^{-3}$. Identification scores using $F(0)$ are stated separately. Curvature, solver, and recovery specifications are in Appendix F. Standard deviations refer to seed means, not confidence intervals.

6.1 Same responses and decisions, earlier certification

The recovery, coordinate, and position certificates are evaluated on the same price trajectories and recovered running averages $z_Q = P\bar{x}_Q$, with step $\eta = 1/16$. Public Euclidean box–balance projection uses no private completion. All three receive identical responses and public constants. The position rule uses the current complete response. No extra solve or candidate change is made for this comparison.

At $J = 8$, all instances are certified within 512 calls (Figure 2, Table 1). Mean first depths are 8.585 and 20.315, compared with offline target attainment at 2.070 and 3.185. Neither recovery-based certificate verifies the target within this cap. Their much longer crossing times appear in Appendix G.6. Thus much of their delay is avoidable looseness, not a limitation shared by every certificate using the same information.

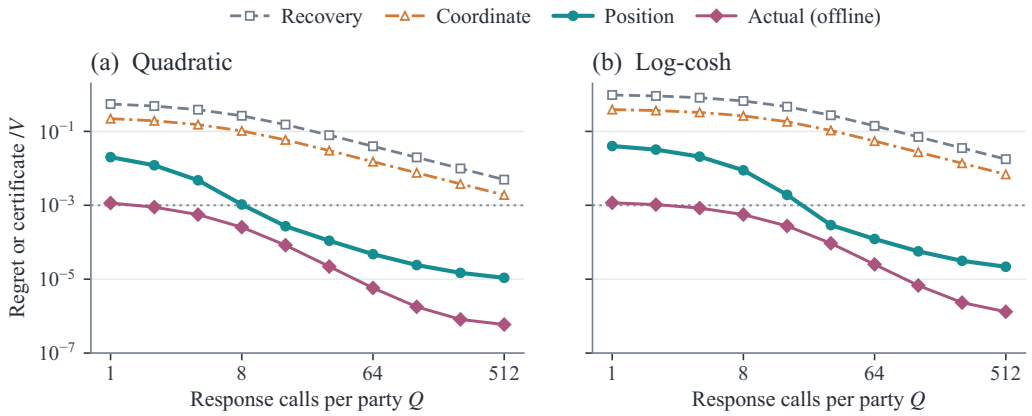


Figure 2: Same transcript and recovered average, different certificates ($J = 8$). Means over 200 paired instances. The dotted line is the 10^{-3} target. First-crossing statistics are computed per instance, not from these mean curves. Actual regret is evaluated offline.

Table 1: Certification on the same outputs, with a shared cap of 512 calls. Depth is mean $\pm$ sample standard deviation across ten seed means. Neither recovery-based certificate crosses within this cap. A dash denotes censoring, not depth 512.

Family	J	Offline first target	Position first crossing	Certified
Quadratic	4	2.35 ± 0.16	—	0/200
Quadratic	8	2.07 ± 0.18	8.59 ± 0.20	200/200
Log-cosh	4	4.11 ± 0.47	—	0/200
Log-cosh	8	3.18 ± 0.48	20.31 ± 0.46	200/200

At $J = 4$, every instance attains the offline target, but none of the position bounds certifies the target accuracy, even with expanded search or retained-prefix reuse. Reference calculation costs 52–56 microseconds per instance in batch, excluding local solves and recovery. Appendix J.1 reports censoring and the small additional history gains at selected high-fidelity prefixes.

6.2 Few-query identification requires stable probes

For the declared quadratic family, two interior probes recover every coordinate:

$$p_{mt} = y_{mt}(0), \qquad a_{mt} = \frac{0.02}{y_{mt}(0) - y_{mt}(0.02)}. \tag{23}$$

All baselines know the same family and parameter ranges, but not realized coefficients. At local work $C = Q(J + 2) = 64$, feedback with $(Q, J) = (8, 6)$ has mean regret 6.217×10^{-5} on fresh instances, while two-probe identification with $J = 30$ reaches numerical-reference precision. Both use offline normalization by $F(0)$. This separates a feedback rule's iteration needs from the interface's information content (Appendix G.3).

For log-cosh objectives, exact identifiability from three probes does not guarantee stable recovery from approximate responses. The identification system is $T\theta = (-\ell_1, -\ell_2)^\top$, where $\theta = (a, \gamma)^\top$ and row j of T is $(d_j, \tanh d_j)$, with $d_j = y(\ell_j) - y(0)$. For small displacements these columns are nearly dependent. If each response error is at most ϵ and $4\epsilon < \sigma_{\min}(T)$, then

$$\left\| \widehat{\theta} - \theta \right\| \le \frac{4\epsilon \, \|\theta\|}{\sigma_{\min}(T) - 4\epsilon}. \tag{24}$$

Appendix I derives this bound, gives observable constants, and relates gradient estimation error to feasible regret, including inexact fitted solutions.

The paired probe study uses 200 fresh log-cosh instances and the same three calls and local work for both policies. The adaptive policy changes only probe geometry, using the zero-price response and public error intervals to stay away from saturation. At $J = 19$, mean $\mathrm{Reg}\,/V$ falls from 2.25×10^{-4} to 2.91×10^{-6}. Both policies meet the 10^{-3} target on all 200 instances. These are offline decision-quality measurements, not certificate acceptance rates. At $J = 30$, both reach numerical-reference accuracy (Appendix J.2). A fitted candidate can also be checked directly by the position certificate, without a global guarantee on the fitted objective.

6.3 Candidate generation and certification under local coupling

The structural series uses four parties with eight local variables and four settings: diagonal quadratics, coupled quadratics, and coupled nonlinear objectives with perturbation strengths $\zeta \in \{0.5, 1\}$. Private banded quadratic terms and noncoordinate log-cosh directions couple local variables. All methods receive the same public boxes, balance, $\mu = 0.5$, $L = 8$, and coordinate-gradient caps $G_i = 8$. Five seed blocks of 20 instances are paired across settings and methods (Appendix I.3).

A diagonal fit uses two proposal calls. A dense symmetric positive-definite (SPD) response fit uses 19, matching feedback's proposal budget. Each candidate receives one additional validation response, so total Q is 3 or 20. Fitted candidates use their implied price. Feedback repeats its last queried price. No realized private parameters enter the certificate. Local responses use $J \in \{64, 128, 256\}$ projected-gradient steps from $u_m/2$, costing QJ vector-gradient evaluations per party, not scalar comparisons. Fitting, projection, and certification work are recorded separately.

Using the known initialization. The solver's public error bound uses its midpoint initialization:

$$\widehat{\delta}_m^{\mathrm{mid}}(J) = \frac{L}{8} \|u_m\|^2 (1 - \mu/L)^{J-1}. \tag{25}$$

The first step bounds the gap by $L \left\|x_m^0 - x_m^\star\right\|^2 /2$, and $\|u_m/2 - x_m^\star\|^2 \le \|u_m\|^2 /4$. Subsequent gaps contract by $1 - \mu/L$. This bound is valid at boundaries and for all prices. It is four times tighter than the diameter-based reference, uses no extra information, and is applied to every candidate method. Its normalized value at $J = 128$ is 2.756×10^{-4} rather than 1.103×10^{-3}.

Table 2: Acceptance (%) at identical structural candidates and responses, $J = 128$. D uses the diameter error bound. M uses (25). Both use the reference multiplier grid. Each setting has 100 paired instances. All candidates shown meet the offline regret target.

	Dense SPD fit		Feedback last	
Setting	D	M	D	M
Diagonal quadratic	0	100	0	100
Coupled quadratic	0	100	0	43
Coupled, $\zeta = 0.5$	0	100	0	30
Coupled, $\zeta = 1$	0	100	0	22

Holding every candidate, response, price, and regret fixed, the midpoint bound raises dense-fit acceptance from zero to 100% in all four settings at $J = 128$ (Table 2). The gain comes from using public solver information more sharply, not from improving the candidates. Expanded multiplier search and history reuse tighten some bounds but add no structural acceptances, so we use the less expensive midpoint-grid calculation. Appendix J.4 gives the paired acceptance changes and computation costs.

Separating candidate quality from acceptance. Dense fitting remains a strong candidate generator under nonlinear coupling (Table 3): at $\zeta = 1$, $J = 256$, mean regret is 8.19×10^{-9}, and all candidates pass. By comparison, 95% of diagonal-fit candidates meet the target but only 3% pass. Feedback last has rates 100% and 93%. Every rule remains inconclusive at $J = 64$. These results distinguish inaccurate candidates from accurate ones that the validation response and public error bound still do not certify. The full precision grid is in Appendix J.3.

Table 3: Candidate quality and certification at $J = 256$, using the midpoint error bound and reference multiplier grid. Each setting has 100 paired instances. Target means $\mathrm{Reg}\,/V \le 10^{-3}$. Q includes validation. The diagonal fit uses fewer proposal calls than the other methods. Values near zero are at numerical-reference precision.

Setting	Candidate	Q	Mean $\mathrm{Reg}\,/V$	Target (%)	Accepted (%)
Diagonal quadratic	Diagonal fit	3	8.44×10^{-19}	100	100
	Dense SPD fit	20	9.31×10^{-19}	100	100
	Feedback average	20	2.64×10^{-5}	100	100
	Feedback last	20	1.06×10^{-7}	100	100
Coupled quadratic	Diagonal fit	3	9.58×10^{-5}	100	1
	Dense SPD fit	20	9.57×10^{-19}	100	100
	Feedback average	20	4.69×10^{-5}	100	93
	Feedback last	20	3.38×10^{-6}	100	96
Coupled, $\zeta = 0.5$	Diagonal fit	3	1.92×10^{-4}	99	3
	Dense SPD fit	20	4.85×10^{-9}	100	100
	Feedback average	20	5.69×10^{-5}	100	90
	Feedback last	20	4.55×10^{-6}	100	95
Coupled, $\zeta = 1$	Diagonal fit	3	4.17×10^{-4}	95	3
	Dense SPD fit	20	8.19×10^{-9}	100	100
	Feedback average	20	8.22×10^{-5}	100	81
	Feedback last	20	7.29×10^{-6}	100	93

Validation leaves the candidate unchanged and queries the true local problems at the selected price. The certificate uses these responses and shared public bounds, not an error estimate from the fitted model. Thus fitting can supply a candidate and price without its correctness being assumed by validation. An accurate candidate may still lie far from the response at that price, leaving the position bound inconclusive (Appendix I.3).

No above-target acceptance was observed in the $4,800$ paired candidate/configuration records. An internal replay recomputed the reported certificate values from recorded candidates and response information without accessing private parameters (Appendix J.4).

7 Discussion and Conclusion

A response budget can conceal different obstacles. Under fixed conditional-error caps, coupled private objectives can remain indistinguishable to every price strategy. When useful information is present, a response's position and near-optimality can certify a feasible decision without first identifying its objective. The paired studies show that both the certificate construction and a sharper use of public solver information can change certification while leaving the candidate untouched. Fitting and validation are complementary: a flexible model proposes the decision, and actual responses check it.

The exact query–state law concerns comparison families. The coupled theorem isolates a fidelity floor, not a universal high-dimensional rate. Experiments use controlled objectives and public feasible recovery. The local-certificate extension changes the message interface and has not been tested here. The results separate ambiguity in the response interface from avoidable conservatism in certification.

AI USE STATEMENT

Generative AI tools were used solely for language editing. The authors verified the final manuscript and take full responsibility for the submission.

ETHICS STATEMENT

All experiments run entirely in simulation on synthetic navigation and a standard control benchmark, involving only generated states and trajectories.

REPRODUCIBILITY STATEMENT

The appendices in this PDF provide the protocol conventions (Appendix A), proofs and certificate calculations (Appendices B.1–E and H), and experimental specifications, including distributions, seeds, solvers, and metrics (Appendices F and I). Appendices G and J report detailed results and numerical checks. Appendix J.4 describes the internal replay. Computation times are distinguished from response calls and local work.

REFERENCES

Miklós Ajtai, Vitaly Feldman, Avinatan Hassidim, and Jelani Nelson. Sorting and selection with imprecise comparisons. *ACM Transactions on Algorithms*, 12(2):19:1–19:19, 2015. doi: 10.1145/2701427.

Moïse Blanchard, Junhui Zhang, and Patrick Jaillet. Memory-constrained algorithms for convex optimization. In *Advances in Neural Information Processing Systems*, volume 36, pp. 6156–6189, 2023a.

Moïse Blanchard, Junhui Zhang, and Patrick Jaillet. Quadratic memory is necessary for optimal query complexity in convex optimization: Center-of-mass is Pareto-optimal. In *Proceedings of the 36th Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 4696–4736. PMLR, 2023b.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011. doi: 10.1561/2200000016.

Du Chen and Geoffrey A. Chua. Noisy dual mirror descent: A near optimal algorithm for Jointly-DP convex resource allocation. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Violet Xinying Chen and Fatma Kılınç-Karzan. Online convex optimization perspective for learning from dynamically revealed preferences. *arXiv:2008.10460v3*, 2021. https://arxiv.org/abs/2008.10460.

George B. Dantzig and Philip Wolfe. Decomposition principle for linear programs. *Operations Research*, 8(1):101–111, 1960. doi: 10.1287/opre.8.1.101.

Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1–2):37–75, 2014. doi: 10.1007/s10107-013-0677-5.

Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. *Games and Economic Behavior*, 118:589–608, 2019. doi: 10.1016/j.geb.2018.02.010.

Egor Gladin, Alexander Gasnikov, and Pavel Dvurechensky. Accuracy certificates for convex minimization with inexact oracle. *arXiv:2310.00523*, 2023. https://arxiv.org/abs/2310.00523.

Robert M. Gray and David L. Neuhoff. Quantization. *IEEE Transactions on Information Theory*, 44(6):2325–2383, 1998.

Ziwei Ji, Ruta Mehta, and Matus Telgarsky. Social welfare and profit maximization from revealed preferences. In *Web and Internet Economics*, volume 11316 of *Lecture Notes in Computer Science*, pp. 264–281. Springer, 2018. doi: 10.1007/978-3-030-04612-5_18.

Frank P. Kelly, Aman K. Maulloo, and David K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, 1998. doi: 10.1057/palgrave.jors.2600523.

Michael Kraning, Eric Chu, Javad Lavaei, and Stephen Boyd. Dynamic network energy management via proximal message passing. *Foundations and Trends in Optimization*, 1(2):70–122, 2014.

Ion Necoara and Valentin Nedelcu. Rate analysis of inexact dual first-order methods: Application to dual decomposition. *IEEE Transactions on Automatic Control*, 59(5):1232–1243, 2014.

Angelia Nedić and Asuman Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, 2009. doi: 10.1137/070708111.

Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006. doi: 10.1016/j.jet.2004.10.007.

Robert D. Nowak. The geometry of generalized binary search. *IEEE Transactions on Information Theory*, 57(12):7893–7906, 2011. doi: 10.1109/TIT.2011.2169298.

Daniel P. Palomar and Mung Chiang. A tutorial on decomposition methods for network utility maximization. *IEEE Journal on Selected Areas in Communications*, 24(8):1439–1451, 2006. doi: 10.1109/JSAC.2006.879350.

Mark Schmidt, Nicolas Le Roux, and Francis Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In *Advances in Neural Information Processing Systems*, volume 24, pp. 1458–1466, 2011.

Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *arXiv:1502.05666v6*, 2016. https://arxiv.org/abs/1502.05666.

Guillaume Van Dessel and François Glineur. Optimal inexactness schedules for tunable oracle based methods. *arXiv:2309.07787*, 2023. https://arxiv.org/abs/2309.07787.

Iyad Walwil and Olivier Fercoq. The smoothed duality gap as a stopping criterion. *Mathematical Programming Computation*, 17:653–697, 2025. doi: 10.1007/s12532-025-00284-0.

Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science*, pp. 222–227. IEEE, 1977. doi: 10.1109/SFCS.1977.24.

APPENDIX GUIDE

The main proofs are in Appendix H.1 (Theorem 4.1) and Appendix H.2 (Proposition 5.1). The local-message extension follows in Appendix H.3. The comparison bounds and their extensions are proved in Appendices B.1–C. Appendices D–E provide the recovery-based baselines. Experimental protocols are specified in Appendices F and I. Detailed results and replay checks are in Appendices G and J.

A PROTOCOL CONVENTIONS

Public information and local models. The public instance specifies the finite-dimensional spaces, linear maps, coupling constraint, query domain, protocol, and any explicitly available bounds. Each party accesses only its own objective and complete feasible set. Feasibility, attainment, and curvature are assumptions on the instance class. The coordinator is not presumed able to verify them by inspecting private problems. A response selection is fixed before interaction and measurable in its arguments. The threshold family specifies its tie rule explicitly.

Randomized transcripts. Let U be independent of the private instance and let $T_q = ((\lambda^j, y^j))_{j=1}^q$. A full-transcript query is $\lambda^q = \pi_q(I_{\mathrm{pub}}, U, T_{q-1})$. In the scalar results, the complete feasible output is a function of $(I_{\mathrm{pub}}, U, T_Q)$, and its estimate is $\widehat{\beta} = \widehat{x}_2$. Party 1 cannot directly transmit its coefficient or choose a different private-data-dependent terminal allocation outside this interface. A single response call returns the specified vector, not its objective value or the local solver trace.

Strict retained information.

$$\lambda^q = \pi_q(I_{\mathrm{pub}}, U, z^q), \qquad z^{q+1} = U_q(I_{\mathrm{pub}}, U, z^q, y^q), \qquad \widehat{x} = D(I_{\mathrm{pub}}, U, z^{Q+1}). \tag{26}$$

Eq. (26) replaces T_q by a state with at most 2^S values at every round boundary. After fixing U, a terminal decoder has at most 2^S possible outputs. A query, last response, local cache, output stream, or warm start that later conveys response-dependent information must be encoded in that state. Stateless response oracles do not carry information between calls. Public round-dependent code may reconstruct interval endpoints from a stored bit prefix. Temporary arithmetic is not charged as persistent state. These conventions are necessary for an output-resolution bound and differ from working-space complexity.

Stopping and local state. Protocols use at most their stated query budget. A stopped run may be padded with fixed ignored queries, while holding its terminal state unchanged. The stopping time is not an extra uncharged message. In the real-valued convex protocol, parties instead explicitly retain local running sums, and the coordinator retains a price and any required aggregate quantities. This protocol is outside the strict finite-bit class. If A_m is not injective, a running sum of $A_m x_m^q$ alone cannot recover the local average. The declared party-local sum supplies it.

Terminal recovery. A last response is the local tuple generated at the final queried price, not a new solve at the updated price. A running average uses only stored local decisions and is locally feasible for convex $\mathcal{X}_m$. Joint feasibility may require a separate recovery P. When P operates on public constraints and the transmitted decisions, it can be executed without private objectives. When it requires hidden feasible sets or objective-dependent completion, that work must be specified as a further distributed procedure. An offline evaluator's access is not online information. The regret bounds assume an available recovery with the stated objective-change bound. Numerical evaluations separately report raw violation and recovery displacement.

B SCALAR INFORMATION BOUNDS

B.1 EXACT RESPONSES AND FINITE CANDIDATES

Lemma B.1 (Coordination as estimation)**.** *The unique optimum of Eq.* (9) *is* $(1-\beta, \beta)$*. Every feasible output has the form* $\widehat{x} = (1-\widehat{\beta}, \widehat{\beta})$*, with* $\widehat{\beta} \in [0,1]$*, and* $\mathrm{Reg}_\beta(\widehat{x}) = \frac{1}{2}(\widehat{\beta} - \beta)^2$*.*

Let $\mathfrak{R}_{Q,S}$ and $\mathfrak{D}_{Q,S}$ be randomized and deterministic minimax regret for the exact comparison family. Let $\mathfrak{B}_{Q,S}$ be its optimal uniform-prior expected regret.

Theorem B.2 (Exact-response calibration)**.** *For integers $Q, S \geq 0$ and $n = \min\{Q, S\}$,*

$$\mathfrak{B}_{Q,S} = \frac{1}{24\,4^n}, \qquad \mathfrak{D}_{Q,S} = \frac{1}{8\,4^n}, \qquad \frac{1}{24\,4^n} \leq \mathfrak{R}_{Q,S} \leq \frac{1}{8\,4^n}. \tag{27}$$

For a full-transcript protocol, replace n by Q. A public set of B distinct candidate coefficients can be recovered with zero regret at every candidate exactly when $n \geq \lceil \log_2 B \rceil$.

At $\lambda = -p$, the local problems are

$$\min_{x_1 \in [0,1]} (\beta - p)x_1, \qquad \min_{x_2 \in [0,1]} \tfrac{1}{2}x_2^2 - px_2. \tag{28}$$

The first gives (10) under the zero tie rule. The second has minimizer p. On the feasible line,

$$F_\beta(1 - x_2, x_2) = \beta - \tfrac{1}{2}\beta^2 + \tfrac{1}{2}(x_2 - \beta)^2. \tag{29}$$

Thus the unique optimum is $(1 - \beta, \beta)$ and every admissible estimate in $[0, 1]$ produces a feasible decision with regret (11). This proves Lemma B.1 at interior and endpoint coefficients.

Lemma B.3 (Finite-output scalar quantization)**.** *If a measurable $g : [a, b] \to \mathbb{R}$ takes at most $N \geq 1$ values and $\ell = b - a > 0$, then*

$$\frac{1}{\ell}\int_a^b (g(t) - t)^2\, dt \geq \frac{\ell^2}{12N^2}, \qquad \sup_{t \in [a,b]} |g(t) - t| \geq \frac{\ell}{2N}. \tag{30}$$

Proof. Project the reproduction values onto $[a, b]$ and assign each point to its nearest value. Neither operation increases error. The nonempty Voronoi intervals have lengths ℓ_i, $i = 1, \ldots, K$, with $K \leq N$ and $\sum_i \ell_i = \ell$. On such an interval a constant reconstruction has integrated squared error at least $\ell_i^3/12$. Therefore the normalized error is at least

$$\frac{1}{12\ell}\sum_{i=1}^{K} \ell_i^3 \geq \frac{\ell^2}{12K^2} \geq \frac{\ell^2}{12N^2}. \tag{31}$$

For worst-case error r, intervals of radius r around the reproduction values must cover $[a, b]$. Their total length is at most $2Nr$, giving $r \geq \ell/(2N)$. Taking a limiting radius handles a supremum that is not attained. □

Proof of Theorem B.2. Fix the random tape. Only $\mathbf{1}\{p^q > \beta\}$ branches the transcript, so a depth-Q comparison tree has at most 2^Q nonempty leaves. Leaf sets are intervals because each response intersects the current consistent set with a half-interval. The terminal estimate is constant on each leaf and is also decoded from at most 2^S states. Its range therefore has cardinality at most $N = 2^{\min\{Q,S\}} = 2^n$.

Applying Lemma B.3 on $[0, 1]$ and multiplying by $1/2$ gives uniform-prior regret at least $1/(24N^2)$. This bound holds for each random tape and hence after averaging. Worst-case expected regret is at least this prior risk. For deterministic protocols, the second inequality in (30) gives worst-case regret at least $1/(8N^2)$. Midpoint bisection for n rounds partitions $[0, 1]$ into 2^n equal intervals. Retain the first n comparison bits and output the identified midpoint. Its worst-case distance is 2^{-n-1} and its uniform-prior squared error is $1/(12\,4^n)$. The corresponding regrets attain the two claimed exact values and give the randomized upper bound. At round q, the public round index determines the length of the stored prefix, so at most $2^{\min\{q,n\}} \leq 2^S$ states are used. Later responses may be ignored. This also covers $n = 0$.

For B public distinct candidates, zero regret requires identifying each coefficient. Under any deterministic tape successful on all candidates, the tree and terminal state need at least B distinguishable outputs, implying $Q, S \geq \lceil \log_2 B \rceil$. If a randomized protocol has zero expected nonnegative regret on each candidate, it succeeds almost surely on each one. Since the candidate set is finite, the intersection of those success events has probability one, giving a tape successful on all candidates. Necessity therefore holds for randomized protocols as well. A balanced comparison tree querying between adjacent candidates, with its leaf index encoded in the state, proves sufficiency. □

The Bayes value and deterministic minimax value differ by a factor of three. They do not determine the exact randomized minimax constant when $n > 0$. When $n = 0$, any output Z is independent of β. Averaging its regret at $\beta = 0$ and 1 gives

$$\tfrac{1}{4}\mathbb{E}[Z^2 + (1-Z)^2] = \tfrac{1}{8} + \tfrac{1}{2}\mathbb{E}[(Z - \tfrac{1}{2})^2] \geq \tfrac{1}{8}. \tag{32}$$

The constant estimate $1/2$ attains worst-case regret $1/8$.

B.2 Ambiguous responses

Robust bisection. Begin with $[L_0, U_0] = [0, 1]$. At round q, query $p^q = (L_{q-1} + U_{q-1})/2$. For returned bit a_q, update

$$[L_q, U_q] = \begin{cases} [L_{q-1}, \min\{U_{q-1}, p^q + \tau\}], & a_q = 1, \\ [\max\{L_{q-1}, p^q - \tau\}, U_{q-1}], & a_q = 0. \end{cases} \tag{33}$$

Admissibility implies $\beta \leq p^q + \tau$ after a one and $\beta \geq p^q - \tau$ after a zero. Thus the interval continues to contain β. Its width obeys $w_q \leq w_{q-1}/2 + \tau$, so

$$w_q \leq 2\tau + (1 - 2\tau)2^{-q}. \tag{34}$$

After n informative queries, return the midpoint. Equations (11) and (34) give regret at most $U_n(\tau)$. The first n bits, public τ, and round index reconstruct every endpoint. No uncharged real-valued state is needed.

Indistinguishable pairs. For each $c \in [\tau, 1 - \tau]$, use the selector $a_c(p) = \mathbf{1}\{p > c\}$. It is admissible both for $\beta_-(c) = c - \tau$ and for $\beta_+(c) = c + \tau$. At the ambiguity boundaries either bit is permitted. Coupling both runs with the same random tape makes every query, response, and state identical. Let $g_u(c)$ be their common estimate. For a fixed tape, it takes at most $N = 2^n$ values.

Proof of the deterministic identity in Theorem 4.2. For a deterministic protocol,

$$\max_{\sigma \in \{-1,1\}} |g(c) - (c + \sigma\tau)| = |g(c) - c| + \tau. \tag{35}$$

If $\tau < 1/2$, Lemma B.3 on $[\tau, 1-\tau]$ gives $\sup_c |g(c) - c| \geq (1-2\tau)/(2N)$. Both endpoint instances are allowed in the worst case, so regret is at least

$$\tfrac{1}{2}\left[\tau + \frac{1-2\tau}{2N}\right]^2 = U_n(\tau). \tag{36}$$

For $\tau = 1/2$, the common selector at $c = 1/2$ is admissible for $\beta = 0, 1$ and yields the lower bound $1/8 = U_n(1/2)$. Robust bisection attains these bounds. □

Proof of the randomized bounds in Theorem 4.2. For $\tau < 1/2$, draw C uniformly from $[\tau, 1 - \tau]$ and an independent fair sign Σ. Set $\beta = C + \Sigma\tau$ and use the selector a_C. This defines a distribution over admissible instance–selector pairs, independent of the protocol tape. At fixed tape u, averaging over the sign gives

$$\mathbb{E}_{C,\Sigma}\tfrac{1}{2}(g_u(C) - C - \Sigma\tau)^2 = \tfrac{1}{2}\mathbb{E}_C(g_u(C) - C)^2 + \tfrac{\tau^2}{2} \geq \frac{(1-2\tau)^2}{24N^2} + \frac{\tau^2}{2}. \tag{37}$$

Average over u. Worst-case expected regret is no smaller than this distributional risk. For $\tau = 1/2$, use $C = 1/2$ deterministically to obtain $1/8$. Independently, the exact selector $\mathbf{1}\{p > \beta\}$ is admissible for every τ, so the exact-response lower bound $1/(24N^2)$ also applies. Their maximum and robust bisection prove (14). In particular,

$$\frac{4^{-n} + \tau^2}{48} \leq \mathfrak{R}^{\mathrm{dz}}_{Q,S,\tau} \leq \frac{1}{4}\,4^{-n} + \tau^2. \tag{38}$$

For fixed τ, both the sharper lower bound and $U_n(\tau)$ approach $\tau^2/2$. At $n = 0$ the midpoint argument above gives exact randomized risk $1/8$. □

B.3 Rounded approximate solves

Suppose $\widetilde{x} \in [0,1]$ satisfies

$$(\beta - p)\widetilde{x} - \min_{x\in[0,1]} (\beta - p)x \le \delta, \qquad 0 \le \delta \le 1/4, \tag{39}$$

and the transmitted bit is $\widetilde{a} = \mathbf{1}\{\widetilde{x} > 1/2\}$. If $p > \beta$ but $\widetilde{a} = 0$, the conditional gap is $(p-\beta)(1-\widetilde{x}) \ge (p-\beta)/2$. If $p < \beta$ but $\widetilde{a} = 1$, it is $(\beta - p)\widetilde{x} > (\beta - p)/2$. A wrong bit is therefore impossible outside $|p - \beta| \le 2\delta$, proving $\widetilde{a} \in \mathcal{Y}_{2\delta}(p;\beta)$.

This inclusion transfers an upper bound to the specified rounded solver. It does not transfer the full-class minimax lower bound to every such solver: an exact solver also satisfies every positive-width inclusion. At a tie, every local point is conditionally optimal, so conditional suboptimality alone need not control distance to the selected exact response.

C Response Width and Conditional Work Bounds

An enriched proximal query minimizing $\beta x_1 + (x_1 - 1)^2$ on $[0,1]$ returns $1 - \beta/2$, revealing the coefficient in one call. The following extensions keep the returned information explicit.

C.1 Separable vector responses

For $\beta \in [0,1]^d$, consider

$$F_\beta(x_1, x_2) = d^{-1}(\beta^\top x_1 + \tfrac{1}{2}\|x_2\|^2), \quad x_1, x_2 \in [0,1]^d, \quad x_1 + x_2 = \mathbf{1}. \tag{40}$$

At $\lambda = -p/d$, one call returns $y_{1,j} = \mathbf{1}\{p_j > \beta_j\}$ and $y_{2,j} = p_j$. The optimum is $(\mathbf{1} - \beta, \beta)$ and feasible regret is $\left\|\widehat{\beta} - \beta\right\|^2/(2d)$. Let $\mathfrak{R}^{(d)}_{Q,S}$ be randomized minimax regret with Q vector calls per party and S total retained bits.

Proposition C.1 (Product response family)**.** *For integers $d \ge 1$ and $Q, S \ge 0$, put $r = \min\{Q, S/d\}$. Then*

$$\frac{4^{-r}}{12\pi e} \le \mathfrak{R}^{(d)}_{Q,S} \le \frac{4^{-r}}{2}. \tag{41}$$

With the coordinatewise ambiguity class (12)*, the risk is $\Theta(4^{-r} + \tau^2)$ with dimension-independent constants, and approaches $\tau^2/2$ as $r \to \infty$.*

Proof. Let v_d be the volume of the unit Euclidean ball. For a measurable cell E of volume v decoded to a, the layer-cake formula gives

$$\int_E \|b - a\|^2 \, db = \int_0^\infty 2t \, |\{b \in E : \|b - a\| > t\}| \, dt \ge \frac{d}{d+2} v_d^{-2/d} v^{1+2/d}. \tag{42}$$

For a decoder with at most N cells in the unit cube, sum the cell bounds and use convexity to obtain uniform-prior regret at least

$$c_d N^{-2/d}, \qquad c_d = \frac{v_d^{-2/d}}{2(d+2)} \ge \frac{1}{12\pi e}. \tag{43}$$

The last inequality follows from $(2\pi/d)^{d/2} = \int_{\mathbb{R}^d} e^{-d\|b\|^2/2}\, db \ge v_d e^{-d/2}$. After fixing the random tape, a Q-call tree has at most 2^{dQ} leaves and the terminal state has at most 2^S values. Hence $N \le 2^{\min\{dQ, S\}}$. The lower bound follows after averaging the tape. For the upper bound, distribute $T = \min\{dQ, S\}$ informative bits as evenly as possible: write $T = da + t$ with $0 \le t < d$ and set $n_j = a + 1$ for $j \le t$, $n_j = a$ otherwise. Query coordinate j for $n_j \le Q$ bisection rounds and store its prefix. The total prefix length is at most S. Coordinatewise midpoints give

$$\mathfrak{R}^{(d)}_{Q,S} \le \frac{1}{8d} \sum_{j=1}^{d} 4^{-n_j} \le \frac{1}{8}\, 4^{-\lfloor r \rfloor} \le \frac{1}{2}\, 4^{-r}. \tag{44}$$

For ambiguity, take c uniformly from $[\tau, 1-\tau]^d$ and independent signs $\epsilon \in \{-1, 1\}^d$, set $\beta = c + \tau\epsilon$, and use the selector $y_{1,j} = \mathbf{1}\{p_j > c_j\}$. Conditional on c and the tape the estimate is independent of the signs, so

$$\mathbb{E}_\epsilon \frac{\left\|\widehat{\beta} - c - \tau\epsilon\right\|^2}{2d} = \frac{\left\|\widehat{\beta} - c\right\|^2}{2d} + \frac{\tau^2}{2}. \tag{45}$$

Scaled-cube quantization and the admissibility of exact comparisons yield the lower bound $\max\{c_d 4^{-r},\ \tau^2/2 + c_d(1-2\tau)^2 4^{-r}\}$. Robust bisection with the same n_j gives the upper bound $\frac{1}{2d}\sum_{j=1}^d [\tau + (1/2 - \tau)2^{-n_j}]^2$. Both bounds give the stated order and limiting floor. At $\tau = 1/2$, use the random corners directly. At $Q = 0$ or $S = 0$ the exact minimax risk is $1/8$. □

The normalization by d makes the loss comparable across dimensions. This extension counts the comparison bits supplied by a vector call. It does not establish the same rate for coupled preferences or arbitrary real-valued response interfaces.

C.2 A CONDITIONAL WORK ALLOCATION

Suppose one call with $J \geq 1$ work units permits the full ambiguity class with $\tau_J = 2^{-J-1}$. Fix positive integers Q, J before protocol randomization, use the same J at every call, and impose $QJ \leq W$. Define the deterministic and randomized work-constrained risks by minimizing $\mathfrak{D}^{\mathrm{dz}}_{Q,S,2^{-J-1}}$ and $\mathfrak{R}^{\mathrm{dz}}_{Q,S,2^{-J-1}}$ over these allocations.

Corollary C.2 (Fixed work and retained state)**.** *For integers $W, S \geq 1$ and $k = \min\{S, \lfloor\sqrt{W}\rfloor\}$,*

$$\frac{4^{-k}}{8} \leq \mathfrak{D}^{\mathrm{work}}_{W,S} \leq \frac{4^{-k}}{2}, \qquad \frac{4^{-k}}{24} \leq \mathfrak{R}^{\mathrm{work}}_{W,S} \leq \frac{4^{-k}}{2}. \tag{46}$$

Both risks equal $1/8$ when $S = 0$.

Proof. For any allocation let $n = \min\{Q, S\}$ and $m = \min\{J, n\}$. Then $m \leq S$ and $m^2 \leq QJ \leq W$, so $m \leq k$. The deterministic ambiguity value is

$$\tfrac{1}{8}(2^{-J} + 2^{-n} - 2^{-(J+n)})^2 \geq \tfrac{1}{8}\, 4^{-m} \geq \tfrac{1}{8}\, 4^{-k}. \tag{47}$$

The randomized lower bounds give $\max\{4^{-n}/24, 4^{-J}/8\} \geq 4^{-k}/24$. For the upper bound choose $Q = k$ and $J = \lfloor W/k \rfloor \geq k$. Robust bisection has regret at most $(2^{-J} + 2^{-k})^2/8 \leq 4^{-k}/2$. The constant midpoint and the endpoint argument in Appendix B.1 handle $S = 0$. □

The lower bounds optimize against all admissible selectors, not a particular numerical solver. The model uses a fixed uniform fidelity and counts only the declared local work. Adaptive per-call work and total runtime are different resource models.

D CONVEX RESPONSES AND FEASIBLE RECOVERY

D.1 EXACT-RESPONSE GUARANTEES

Write $\mathcal{L}(x, \lambda) = F(x) + \langle \lambda, Ax - b\rangle$ and $\Phi = F + \iota_{\mathcal{X}}$. If $x(\lambda)$ is an attained minimizer, then for any λ',

$$d(\lambda') \leq \mathcal{L}(x(\lambda), \lambda') = d(\lambda) + \langle Ax(\lambda) - b, \lambda' - \lambda\rangle. \tag{48}$$

Thus the aggregate exact response supplies a supergradient of the concave dual.

General convex case. Let $\mathcal{D}$ be a closed convex price domain containing a maximizer $\lambda^\star$, assume $\|g^q\| \leq G$, and set $\lambda^{q+1} = \Pi_{\mathcal{D}}(\lambda^q + \eta_q g^q)$. Projection nonexpansiveness and the supergradient inequality give

$$\left\|\lambda^{q+1} - \lambda^\star\right\|^2 \leq \|\lambda^q - \lambda^\star\|^2 + \eta_q^2 G^2 - 2\eta_q[d(\lambda^\star) - d(\lambda^q)]. \tag{49}$$

Summing and applying concavity at $\bar{\lambda}_Q = (\sum_q \eta_q \lambda^q)/(\sum_q \eta_q)$ yields

$$d(\lambda^\star) - d(\bar{\lambda}_Q) \le \frac{\left\|\lambda^1 - \lambda^\star\right\|^2 + G^2 \sum_{q=1}^Q \eta_q^2}{2\sum_{q=1}^Q \eta_q}. \tag{50}$$

A public bound $R \ge \left\|\lambda^1 - \lambda^\star\right\|$ gives $RG/\sqrt{Q}$ with $\eta_q = R/(G\sqrt{Q})$ when $R, G > 0$. This is a dual guarantee. Strong duality and primal recovery require their own assumptions.

Smooth dual from strong primal curvature. For proper closed μ-strongly convex Φ, its conjugate Φ^* has a $1/\mu$-Lipschitz gradient. Since $h(\lambda) = \Phi^*(-A^\top \lambda) + \langle \lambda, b\rangle$, the gradient of h is $b - Ax(\lambda)$, with Lipschitz bound $L_d \ge \|A\|_{\mathrm{op}}^2 / \mu$. Assume an attained multiplier $\lambda^\star$ with $d(\lambda^\star) = F^\star$, and write $D_0 = \left\|\lambda^1 - \lambda^\star\right\|$. At exact step $1/L_d$, the standard descent and convexity inequalities give

$$h(\lambda^{q+1}) - h(\lambda^\star) \le \frac{L_d}{2}\left(\|\lambda^q - \lambda^\star\|^2 - \left\|\lambda^{q+1} - \lambda^\star\right\|^2\right). \tag{51}$$

The dual gaps are nonincreasing and distance to the optimum is nonincreasing. Hence, with $\bar{x}_Q = Q^{-1}\sum_q x(\lambda^q)$,

$$d(\lambda^\star) - d(\lambda^{Q+1}) \le \frac{L_d D_0^2}{2Q}, \tag{52}$$

$$\|A\bar{x}_Q - b\| = \frac{L_d}{Q}\left\|\lambda^{Q+1} - \lambda^1\right\| \le \frac{2L_d D_0}{Q}. \tag{53}$$

For the objective bound let $\Delta_q = F^\star - d(\lambda^q)$ and $D_q = \|\lambda^q - \lambda^\star\|$. Exact conditional minimization gives

$$\mathcal{L}(x(\lambda^q), \lambda^\star) - F^\star = -\Delta_q - \langle \lambda^q - \lambda^\star, g^q\rangle. \tag{54}$$

The squared-distance update and $\|g^q\|^2 \le 2L_d\Delta_q$ imply

$$0 \le \mathcal{L}(x(\lambda^q), \lambda^\star) - F^\star \le \tfrac{L_d}{2}(D_q^2 - D_{q+1}^2). \tag{55}$$

Convexity and summation yield

$$0 \le F(\bar{x}_Q) - F^\star + \langle \lambda^\star, A\bar{x}_Q - b\rangle \le \frac{L_d D_0^2}{2Q}. \tag{56}$$

Consequently, if an available recovery obeys $F(P\bar{x}_Q) - F(\bar{x}_Q) \le K\|A\bar{x}_Q - b\|$ for $K \ge 0$,

$$0 \le \mathrm{Reg}(P\bar{x}_Q) \le \frac{L_d D_0^2}{2Q} + \frac{2L_d D_0}{Q}(\|\lambda^\star\| + K). \tag{57}$$

These bounds concern the average of the Q queried responses. The updated price λ^{Q+1} does not supply an additional response for free.

D.2 Approximate responses and the implemented average

Write $\Phi = F + \iota_{\mathcal{X}}$ and let $L_d \ge \|A\|_{\mathrm{op}}^2 / \mu$ be positive. Under strong primal convexity, $\nabla h(\lambda) = b - Ax(\lambda)$ is L_d-Lipschitz. At query λ^q, let the parties return locally feasible $\widetilde{x}^q$ with aggregate conditional suboptimality bounded by δ_q:

$$F(\widetilde{x}^q) + \langle \lambda^q, A\widetilde{x}^q - b\rangle - d(\lambda^q) \le \delta_q. \tag{58}$$

The coordinator updates the price, and each party retains its own running average:

$$\widetilde{g}^q = A\widetilde{x}^q - b, \qquad \lambda^{q+1} = \lambda^q + \eta\widetilde{g}^q, \qquad \bar{x}_Q = \frac{1}{Q}\sum_{q=1}^Q \widetilde{x}^q, \quad r_Q = A\bar{x}_Q - b. \tag{59}$$

Only $A_m\widetilde{x}_m^q$ is needed for the price update. Each party can average its complete local decisions without transmitting them. After Q calls, λ^{Q+1} is available but unqueried. The last response is $\widetilde{x}^Q$, whereas $\bar{x}_Q$ uses the declared local averaging state.

Proposition D.1 (Averaged recovery with inexact responses)**.** *Suppose Φ is proper, closed, and μ-strongly convex for $\mu > 0$, the primal optimum is attained, and an optimal multiplier $\lambda^\star$ satisfies $d(\lambda^\star) = F^\star$. Let $L_d > 0$ satisfy the bound above, use $0 < \eta \le 1/(2L_d)$, and suppose every issued price lies in Λ. Put $D_0 = \left\|\lambda^1 - \lambda^\star\right\|$, $\bar{\delta}_Q = Q^{-1}\sum_{q=1}^{Q}\delta_q$, and*

$$C_Q = \frac{D_0^2}{2\eta Q} + 2\bar{\delta}_Q, \qquad B_Q = \frac{D_0 + \sqrt{D_0^2 + 4\eta Q\bar{\delta}_Q}}{\eta Q}. \tag{60}$$

Then

$$0 \le F(\bar{x}_Q) - F^\star + \langle \lambda^\star, r_Q\rangle \le C_Q, \qquad \|r_Q\| \le B_Q. \tag{61}$$

If an available recovery P makes $\bar{x}_Q$ jointly feasible and satisfies $F(P\bar{x}_Q) - F(\bar{x}_Q) \le K\|r_Q\|$ for $K \ge 0$, then

$$0 \le \mathrm{Reg}(P\bar{x}_Q) \le C_Q + (\|\lambda^\star\| + K)B_Q. \tag{62}$$

For fixed problem constants the bound is $O(Q^{-1} + \sqrt{\bar{\delta}_Q/Q} + \bar{\delta}_Q)$, equivalently $O(Q^{-1} + \bar{\delta}_Q)$ in order. Its proof specializes inexact-dual primal recovery (Nedić & Ozdaglar, 2009; Necoara & Nedelcu, 2014). Strong convexity gives $\|e^q\|^2 \le 2L_d\delta_q$, and a one-step estimate gives

$$0 \le \mathcal{L}(\widetilde{x}^q, \lambda^\star) - F^\star \le \frac{D_q^2 - D_{q+1}^2}{2\eta} + 2\delta_q, \quad D_q = \|\lambda^q - \lambda^\star\|, \tag{63}$$

where $\mathcal{L}(x,\lambda) = F(x) + \langle\lambda, Ax - b\rangle$. Summation and $r_Q = (\lambda^{Q+1} - \lambda^1)/(\eta Q)$ control the average's objective and balance. Recovery adds its stated cost. The proof follows below. No strong convexity of h is needed, so response saturation at box boundaries is allowed. Recovery must be executable. A distance bound $\|Px - x\| \le \kappa\|Ax - b\|$ and an objective Lipschitz bound L_F give $K = L_F\kappa$. For public boxes with coordinatewise balance, Euclidean projection satisfies $\|Px - x\|_\mathcal{X} \le \|Ax - b\|_2$. A public gradient bound therefore supplies K. Recovery involving hidden constraints requires a separate distributed procedure.

For the separable box solver below, J derivative bisections give conditional gap at most $c_{mt}u_{mt}^2/(8\,4^J)$ for an interior coordinate of curvature at most c_{mt}. Boundary returns are exact. Thus

$$\mathrm{Reg}(P\bar{x}_Q) = O\Big(Q^{-1} + 2^{-J}Q^{-1/2} + 4^{-J}\Big). \tag{64}$$

This is a concrete solver's upper bound, not the adversarial work lower bound of Appendix C. Appendix D.5 gives the constants, and exact responses recover $O(Q^{-1})$.

We prove Proposition D.1 directly. The argument follows inexact-dual primal-recovery analysis (Necoara & Nedelcu, 2014), with affine equalities, a strongly convex extended objective, varying pathwise conditional errors, and explicit feasible recovery. No strong dual curvature or centering of the errors is assumed.

At query λ^q, write $x^q = x(\lambda^q)$, $g^q = Ax^q - b$, $\widetilde{g}^q = A\widetilde{x}^q - b$, and $e^q = \widetilde{g}^q - g^q$. Let $\Delta_q = F^\star - d(\lambda^q) \ge 0$ and $D_q = \|\lambda^q - \lambda^\star\|$, so $D_1 = D_0$ in the proposition statement. Strong convexity at the conditional minimizer and (58) give

$$\|\widetilde{x}^q - x^q\|^2 \le \frac{2\delta_q}{\mu}, \qquad \|e^q\|^2 \le 2L_d\delta_q. \tag{65}$$

The local solutions lie in $\mathcal{X}$, so their average does too.

Proof of Proposition D.1. Conditional suboptimality and the unprojected update imply

$$\mathcal{L}(\widetilde{x}^q, \lambda^\star) - F^\star \le \delta_q - \Delta_q - \langle\lambda^q - \lambda^\star, \widetilde{g}^q\rangle = \frac{D_q^2 - D_{q+1}^2}{2\eta} + \delta_q - \Delta_q + \frac{\eta}{2}\|g^q + e^q\|^2. \tag{66}$$

Smooth convex duality gives $\|g^q\|^2 \le 2L_d\Delta_q$. Using $\|g^q + e^q\|^2 \le 2\|g^q\|^2 + 2\|e^q\|^2$ and $2\eta L_d \le 1$ yields

$$0 \le \mathcal{L}(\widetilde{x}^q, \lambda^\star) - F^\star \le \frac{D_q^2 - D_{q+1}^2}{2\eta} + \delta_q + \eta\|e^q\|^2 \le \frac{D_q^2 - D_{q+1}^2}{2\eta} + 2\delta_q. \tag{67}$$

The left inequality follows because $\inf_{x\in\mathcal{X}} \mathcal{L}(x,\lambda^\star) = F^\star$ under strong duality. Sum over q, discard the nonnegative terminal squared distance, and use convexity of $\mathcal{L}(\cdot,\lambda^\star)$ to obtain the saddle bound (61). The same sum and nonnegative left side give

$$D_{Q+1}^2 \le D_0^2 + 4\eta Q\bar{\delta}_Q. \tag{68}$$

The received responses telescope exactly:

$$r_Q = \frac{1}{Q}\sum_{q=1}^{Q} \widetilde{g}^q = \frac{\lambda^{Q+1}-\lambda^1}{\eta Q}. \tag{69}$$

The triangle inequality and (68) give $\|r_Q\| \le B_Q$. Finally,

$$F(\bar{x}_Q) - F^\star \le C_Q + \|\lambda^\star\|\,\|r_Q\|. \tag{70}$$

Add the recovery bound to prove (62). Feasibility gives nonnegative regret. Since

$$B_Q \le \frac{2D_0}{\eta Q} + 2\sqrt{\frac{\bar{\delta}_Q}{\eta Q}}, \tag{71}$$

the stated rate follows with fixed $\eta, D_0, \|\lambda^\star\|, K$. Every step is pathwise and permits errors depending on the complete preceding history. □

Uniform rates over an instance class additionally require uniform bounds on those constants. The step restriction is sufficient, not necessary. The proposition does not certify the approximate experimental runs at $\eta = 1/L_d$. For an offline check, (67) also permits the tighter term $Q^{-1}\sum_q(\delta_q + \eta\|e^q\|^2)$ instead of $2\bar{\delta}_Q$, provided those errors are evaluated. They need not be transmitted to execute the algorithm.

D.3 A FIXED-UPDATE ILLUSTRATION

Proposition D.2 (Transcript-identical fidelity obstruction for a fixed update)**.** *Fix $Q \ge 1$, $\eta > 0$, and $a \in (0, 1/2)$. For* (59) *initialized at $\lambda^1 = 0$, there are two strongly convex two-party worlds with identical public sets, the same declared aggregate conditional-error bound a^2, and the same complete transcript*

$$\lambda^q = 0, \qquad \widetilde{x}^q = (1/2 + a, 1/2 - a), \qquad q = 1, \dots, Q. \tag{72}$$

The common feasible average has regret 0 in world A and a^2 in world B. A transcript-measurable, possibly randomized rule that certifies this average with zero false-certification probability in every admissible world must reject at tolerance $\varepsilon < a^2$.

Sharpness of the objective-error term for this output rule. For the world-B objectives in Appendix D.3, the response at zero is $(1/2 + a, 1/2 - a)$, with conditional suboptimality $a^2/2$ per party and aggregate error $\delta = a^2$. Elsewhere the response is exact. This measurable oracle satisfies the aggregate conditional-error bound at every price. Its two errors cancel at zero, so by induction every queried price remains zero. All averages equal that same feasible point and have regret δ. Euclidean projection leaves it unchanged. The example shows that neither aggregate residual nor price error can remove the $O(\delta)$ term for this update, averaging, and recovery rule. It is not a lower bound over all response protocols.

Proof of Proposition D.2. In world A, let

$$f_1^A(x) = \tfrac{1}{2}(x - \tfrac{1}{2} - a)^2, \qquad f_2^A(x) = \tfrac{1}{2}(x - \tfrac{1}{2} + a)^2, \qquad x \in [0,1]. \tag{73}$$

Use exact responses at every price. The zero-price responses are $(1/2 + a, 1/2 - a)$. They already satisfy the coupling equality, so the price remains zero on every round. Their average is the unique feasible optimum and has regret zero.

In world B, let $f_1^B(x) = f_2^B(x) = \frac{1}{2}(x - \frac{1}{2})^2$. At zero price, let the admissible approximate oracle return the same pair $(1/2 + a, 1/2 - a)$. Define it to return exact minimizers at every nonzero price. Each observed local response has conditional error $a^2/2$, hence aggregate error a^2. The responses

again balance exactly, so induction gives the same all-zero price sequence and the same complete response transcript. The common average is feasible, but its regret in world B is $\frac{1}{2}a^2+\frac{1}{2}a^2 = a^2$. The public description, including the valid bound $\widehat{\delta}_q = a^2$, is the same in both worlds. Any acceptance event determined by this transcript and instance-independent randomness has the same probability in both. If that event certifies the prescribed common average and $\varepsilon < a^2$, every acceptance in world B is a false certification. A rule with zero false-certification probability in every admissible world must therefore reject with probability one in world B, and hence also in world A. This proves the fixed-output statement. □

Certification of the fixed output versus choosing another output. The threshold a^2 in Proposition D.2 concerns the common average, not every feasible decoding of the transcript. To see the distinction explicitly, restrict attention to the two worlds just constructed and write a feasible output as $x(s) = (1/2 + s, 1/2 - s)$, $s \in [-1/2, 1/2]$. Its regrets are

$$\mathrm{Reg}_A(x(s)) = (s-a)^2, \qquad \mathrm{Reg}_B(x(s)) = s^2. \tag{74}$$

For any randomized decoded scalar Z with the same law in both worlds,

$$\max\{\mathbb{E}(Z-a)^2, \mathbb{E}Z^2\} \geq \tfrac{1}{2}\mathbb{E}\big[(Z-a)^2 + Z^2\big] = \mathbb{E}(Z-a/2)^2 + a^2/4 \geq a^2/4. \tag{75}$$

The deterministic choice $Z = a/2$ attains equality. Thus the two-world subproblem has optimal worst-case regret $a^2/4$ when a different output is allowed. This does not guarantee that value on the full admissible oracle class. Similarly, additional nonzero-price queries can distinguish these particular worlds. The proposition fixes both the update and the output being certified.

D.4 Public stopping information

Keeping the terminal price term in the same proof removes the unknown optimal multiplier from a computable upper bound.

Corollary D.3 (Response-based certificate)**.** *Under Proposition D.1, suppose valid bounds $\widehat{\delta}_q \geq \delta_q$ and the recovery constant K are available to the coordinator. Then*

$$0 \leq \mathrm{Reg}(P\bar{x}_Q) \leq H_Q := \frac{\left\|\lambda^1\right\|^2 - \left\|\lambda^{Q+1}\right\|^2}{2\eta Q} + \frac{2}{Q}\sum_{q=1}^{Q}\widehat{\delta}_q + K\left\|r_Q\right\|. \tag{76}$$

Hence $H_Q \leq \varepsilon$ certifies a feasible ε-optimal output.

The price trajectory supplies r_Q. Public curvature and a specified solver can supply the error bounds. No private objective values, optimal multiplier, or new response are required. The inequality holds at every prefix, so it also holds at its first crossing. Failure to cross is not evidence that the output is inaccurate: the public bounds may be conservative. The proof below establishes the statement. Appendix D.5 specifies the public constants. For the separable box solver below, complete local decisions are returned, so the same transcript also identifies interior coordinates and the recovery displacement. With public curvature bounds c_i and absolute gradient bounds G_i, indexed by $i = (m, t)$, set

$$\widehat{\delta}_q^{\mathrm{act}} = \sum_{i \in \mathcal{I}_q} \frac{c_i u_i^2}{8\,4^J}, \qquad R_Q^{\mathrm{coord}} = \sum_i G_i \left|(P\bar{x}_Q - \bar{x}_Q)_i\right|, \tag{77}$$

where $\mathcal{I}_q$ contains the strictly interior responses. Boundary returns are exact. Substitution in the same proof yields

$$0 \leq \mathrm{Reg}(P\bar{x}_Q) \leq H_Q^{\mathrm{coord}} := \frac{\left\|\lambda^1\right\|^2 - \left\|\lambda^{Q+1}\right\|^2}{2\eta Q} + \frac{2}{Q}\sum_{q=1}^{Q}\widehat{\delta}_q^{\mathrm{act}} + R_Q^{\mathrm{coord}}. \tag{78}$$

For our public constants, $\widehat{\delta}_q^{\mathrm{act}} \leq \widehat{\delta}_q$ and $R_Q^{\mathrm{coord}} \leq \|G\|_2 \left\|P\bar{x}_Q - \bar{x}_Q\right\|_2 \leq K\left\|r_Q\right\|$, with $K = \|G\|_2$. Hence $H_Q^{\mathrm{coord}} \leq H_Q$ at every prefix, without extra information. Appendix D.4 proves validity and this comparison separately. The recovery-based bound has a different limitation. For zero

initialization, fixed J, and bounded prices, with $V = \|u\|_F^2$ and the uniform public curvature bound c, it satisfies

$$\lim_{Q\to\infty} \frac{H_Q}{V} = \frac{c}{4\,4^J}. \tag{79}$$

At $J = 4$, this limit is $1/512$ for quadratics ($c = 2$) and $1/256$ for log-cosh ($c = 4$), above the 10^{-3} target. This is a resolution limit of the recovery-based sufficient bound, not a regret lower bound or a limit on every certificate. It does not exclude a transient crossing. Appendix D.5 gives the full expression.

Proof of Corollary D.3. Conditional suboptimality gives

$$F(\widetilde{x}^q) - F^\star \le \delta_q - \Delta_q - \langle \lambda^q, \widetilde{g}^q \rangle. \tag{80}$$

Expanding the squared norm of the update, and using the same estimates as in (67), yields

$$F(\widetilde{x}^q) - F^\star \le \frac{\|\lambda^q\|^2 - \left\|\lambda^{q+1}\right\|^2}{2\eta} + \delta_q - \Delta_q + \frac{\eta}{2}\left\|g^q + e^q\right\|^2 \le \frac{\|\lambda^q\|^2 - \left\|\lambda^{q+1}\right\|^2}{2\eta} + 2\delta_q. \tag{81}$$

Sum, divide by Q, apply convexity, replace δ_q by the available $\widehat{\delta}_q$, and add $K\,\|r_Q\|$. This gives Eq. (76). Recovered feasibility gives its lower bound. The inequality holds on every prefix, including the first one at which $H_Q \le \varepsilon$. □

The coordinator obtains r_Q from Eq. (69) and retains a running sum of the accessible error bounds. A public solver specification can provide these bounds without additional private messages. A first crossing authorizes the already specified recovery, not another private-objective solve. In exact arithmetic $H_Q \ge 0$ under the assumptions. A negative numerical value indicates numerical error or a violated assumption. Certified floating-point execution would also need bounds on solve, update, and recovery roundoff. For (78), the solver's boundary checks make every returned boundary coordinate an exact conditional minimizer. Summing the public interior-coordinate bounds therefore gives $\widehat{\delta}_q^{\mathrm{act}} \ge \delta_q$. Along the segment from $\bar{x}_Q$ to $P\bar{x}_Q$, the coordinatewise gradient bounds give

$$F(P\bar{x}_Q) - F(\bar{x}_Q) \le \sum_i G_i\,|(P\bar{x}_Q - \bar{x}_Q)_i| = R_Q^{\mathrm{coord}}. \tag{82}$$

Substituting these two bounds in the preceding proof establishes validity of Eq. (78). In this box experiment, A_m is the identity, so the full local decision is the transmitted response. Consequently both $\mathcal{I}_q$ and $P\bar{x}_Q - \bar{x}_Q$ are computable from the same history and public constraints. For a general noninjective A_m, these quantities need not be observable. No same-information claim is made in that setting. To compare the two valid certificates, choose $0 \le c_i \le c$, set $\widehat{\delta}_q = c\sum_i u_i^2/(8\,4^J)$, and let $K = \|G\|_2$. Then

$$\widehat{\delta}_q^{\mathrm{act}} \le \widehat{\delta}_q, \qquad R_Q^{\mathrm{coord}} \le \|G\|_2\,\|P\bar{x}_Q - \bar{x}_Q\|_2 \le K\,\|r_Q\|_2\,. \tag{83}$$

The first inequality sums over a subset of coordinates. The second uses Cauchy–Schwarz and the box projection bound (84). The quadratic constants are $c_i = c = 2$, $G_i = 2u_i$, and $K = 2\,\|u\|_2$. For log-cosh they are $c_i = c = 4$, $G_i = 2u_i + 2\tanh u_i$, and $K = \|2u + 2\tanh u\|_2$. Thus the exact public constants used in the experiment give $H_Q^{\mathrm{coord}} \le H_Q$ at every prefix. This pointwise comparison requires these compatible choices. Validity alone does not order certificates with unrelated constants.

D.5 Recovery and solver constants for public boxes

Suppose $0 \le x_{mt} \le u_{mt}$ and $0 \le d_t \le \sum_m u_{mt}$. For one coordinate, correct a surplus by decreasing entries only, or a deficit by increasing entries only. The available capacity suffices, and the total absolute change equals $|\sum_m x_{mt} - d_t|$. Its Euclidean norm is no larger. Combining coordinates constructs a feasible point within $\|Ax - d\|_2$ of x, so Euclidean projection P satisfies

$$\|Px - x\|_F \le \|Ax - d\|_2\,. \tag{84}$$

A uniform objective gradient bound $\|\nabla F(x)\|_F \le K$ on the box then implies $F(Px) - F(x) \le K\,\|Ax - d\|_2$. For the quadratic family, $a_{mt} \le 2$ and $0 \le p_{mt} \le u_{mt}$ give $K = 2\,\|u\|_F$. For the log-cosh family with $\gamma_{mt} \le 2$, use $K = \|2u + 2\tanh u\|_F$, with entrywise $\tanh$. After J

bisections, an interior local minimizer lies within $u_{mt}2^{-(J+1)}$ of the returned midpoint. With coordinate curvature bounded above by c_{mt}, its conditional gap is at most $c_{mt}u_{mt}^2/(8\,4^J)$. Exact boundary responses contribute zero. Summing over interior coordinates provides a valid δ_q. For quadratic costs $c_{mt} = a_{mt}$. For log-cosh costs use $a_{mt} + \gamma_{mt}$. Substitution into Proposition D.1 gives $O(Q^{-1} + 2^{-J}Q^{-1/2} + 4^{-J})$ feasible regret, with fixed problem constants and the specified step. For an online certificate, let $V = \|u\|_F^2$ and use the public curvature upper bound $c = 2$ for quadratics or $c = 4$ for log-cosh. Then $\widehat{\delta}_q = cV/(8\,4^J)$ is valid without revealing realized coefficients or interior-coordinate indicators. At zero initialization,

$$\frac{H_Q}{V} = -\frac{\left\|\lambda^{Q+1}\right\|^2}{2\eta QV} + \frac{c}{4\,4^J} + \frac{K\left\|\lambda^{Q+1}\right\|}{\eta QV}. \tag{85}$$

The public normalization is essential for response-only stopping. The private scale $F(0)$ used in offline plots does not enter this rule. If $\sup_Q \left\|\lambda^{Q+1}\right\| < \infty$, the first and third terms in (85) vanish as $Q \to \infty$, proving (79). This is an asymptotic resolution statement for the sufficient bound: it does not lower-bound the actual regret, and it does not exclude a transient crossing on a particular trajectory.

E Response Errors and Stable Price Feedback

E.1 Conditional error, aggregate error, and direct perturbations

For a proper closed μ_m-strongly convex conditional objective with minimizer $x_m^\star$, a locally feasible δ_m-suboptimal solution satisfies

$$\|\widetilde{x}_m - x_m^\star\| \le \sqrt{2\delta_m/\mu_m}, \qquad \|A_m\widetilde{x}_m - A_m x_m^\star\| \le \|A_m\|_{\mathrm{op}}\sqrt{2\delta_m/\mu_m}. \tag{86}$$

Aggregate interface error is bounded by the sum of these terms. Without curvature, especially at a tie, no distance bound to a fixed selected minimizer follows from conditional suboptimality alone. For unprojected updates with variable step η_q,

$$\lambda^{Q+1} = \lambda^1 + \sum_{q=1}^{Q}\eta_q[Ax(\lambda^q) - b] + \sum_{q=1}^{Q}\eta_q e^q. \tag{87}$$

The exact-response term is evaluated along the actual, error-dependent price trajectory. Thus the last sum is not generally the difference between an approximate run and a separate exact run. A bias may contribute in a fixed direction, but this identity by itself does not prove growth of total error. If η_q is measurable before call q, $\mathbb{E}[e^q \mid \mathcal{H}_{q-1}] = 0$, and $\mathbb{E}[\|e^q\|^2 \mid \mathcal{H}_{q-1}] \le \sigma_q^2$, orthogonality of martingale differences gives

$$\mathbb{E}\left\|\sum_{q=1}^{Q}\eta_q e^q\right\|^2 \le \sum_{q=1}^{Q}\mathbb{E}[\eta_q^2\sigma_q^2]. \tag{88}$$

This statement concerns the direct perturbation sum. Numerical bisection errors are deterministic functions of the query and are not assumed centered.

E.2 Projected dual ascent with bounded interface error

Let a public closed convex set $\mathcal{D}$ contain a dual maximizer $\lambda^\star$, with $R_{\mathcal{D}} = \sup_{\lambda\in\mathcal{D}}\|\lambda - \lambda^\star\| < \infty$. Suppose $\widetilde{g}^q = g^q + e^q$, $g^q \in \partial^+ d(\lambda^q)$, $\|\widetilde{g}^q\| \le \widetilde{G}$, and $\|e^q\| \le \epsilon_q$. For $\lambda^{q+1} = \Pi_{\mathcal{D}}(\lambda^q + \eta_q\widetilde{g}^q)$ and the weighted price average $\bar{\lambda}_Q$,

$$d(\lambda^\star) - d(\bar{\lambda}_Q) \le \frac{\left\|\lambda^1 - \lambda^\star\right\|^2 + \widetilde{G}^2\sum_q\eta_q^2}{2\sum_q\eta_q} + R_{\mathcal{D}}\frac{\sum_q\eta_q\epsilon_q}{\sum_q\eta_q}. \tag{89}$$

To prove this, expand the squared distance after projection as in (50). The additional term is $2\eta_q\langle e^q, \lambda^q - \lambda^\star\rangle$, bounded by $2\eta_q R_{\mathcal{D}}\epsilon_q$. Sum and apply concavity. This is a dual-gap statement on a bounded price domain, separate from the unprojected feasible-recovery guarantee.

E.3 Contraction under strong dual curvature

Proposition E.1 (Stable inexact price feedback). *Suppose $h = -d$ is μ_d-strongly convex with L_d-Lipschitz gradient on the entire price space, $0 < \mu_d \le L_d$. Let $\lambda^\star$ minimize h, $D_0 = \left\|\lambda^1 - \lambda^\star\right\|$, and $\rho = 1 - \mu_d/L_d$. Use $\lambda^{q+1} = \lambda^q - \nabla h(\lambda^q)/L_d + e^q/L_d$, with issued prices in Λ. If $\|e^q\| \le \epsilon$, then*

$$\left\|\lambda^Q - \lambda^\star\right\| \le \rho^{Q-1} D_0 + \frac{\epsilon}{\mu_d}(1 - \rho^{Q-1}), \tag{90}$$

$$\left\|\widetilde{g}^Q\right\| \le L_d \rho^{Q-1} D_0 + \epsilon \left[1 + \frac{L_d}{\mu_d}(1 - \rho^{Q-1})\right]. \tag{91}$$

Alternatively, if λ^1 is deterministic and the errors satisfy $\mathbb{E}[e^q \mid \mathcal{H}_{q-1}] = 0$ and $\mathbb{E}[\|e^q\|^2 \mid \mathcal{H}_{q-1}] \le \sigma^2$, then

$$\mathbb{E}\left\|\lambda^Q - \lambda^\star\right\|^2 \le \rho^{2(Q-1)} D_0^2 + \frac{\sigma^2[1 - \rho^{2(Q-1)}]}{\mu_d(2L_d - \mu_d)}. \tag{92}$$

Use $\rho^0 = 1$, including when $\rho = 0$.

Proof. Write $L = L_d$, $\mu = \mu_d$, and $T = I - \nabla h/L$. For $a = u - v$ and $b = \nabla h(u) - \nabla h(v)$, strengthened cocoercivity gives

$$\|b\|^2 \le (L + \mu)\langle a, b\rangle - L\mu \|a\|^2. \tag{93}$$

For $\mu < L$, this follows by applying cocoercivity to $h - \mu \|\cdot\|^2/2$, whose gradient is $(L-\mu)$-Lipschitz. For $\mu = L$, $b = La$. Therefore, using $\langle a, b\rangle \ge \mu \|a\|^2$,

$$\|T(u) - T(v)\|^2 \le (1 - \mu/L)\|a\|^2 - \frac{L - \mu}{L^2}\langle a, b\rangle \le (1 - \mu/L)^2 \|a\|^2. \tag{94}$$

Since $T(\lambda^\star) = \lambda^\star$, the update gives $D_{q+1} \le \rho D_q + \epsilon/L_d$. Iteration proves (90). Lipschitz continuity gives $\left\|\widetilde{g}^Q\right\| \le L_d D_Q + \epsilon$, proving (91). Under conditional centering, the cross term vanishes and $\mathbb{E}[D_{q+1}^2 \mid \mathcal{H}_{q-1}] \le \rho^2 D_q^2 + \sigma^2/L_d^2$. Sum the geometric series and use $L_d^2(1 - \rho^2) = \mu_d(2L_d - \mu_d)$ to obtain (92). □

The deterministic floor is realized by $h(\lambda) = \mu_d \lambda^2/2$ with persistent error $e^q = \epsilon$: the price converges to ϵ/μ_d, although its approximate residual can approach zero. With centered scalar errors of variance σ^2, the same quadratic has stationary mean squared price error $\sigma^2/[\mu_d(2L_d - \mu_d)]$.

Where dual curvature holds. For unconstrained $F \in C^2$ with $\mu_p I \preceq \nabla^2 F \preceq L_p I$ and full-row-rank A,

$$\nabla^2 h(\lambda) = A[\nabla^2 F(x(\lambda))]^{-1} A^\top, \qquad \frac{\lambda_{\min}(AA^\top)}{L_p} I \preceq \nabla^2 h(\lambda) \preceq \frac{\|A\|_{\mathrm{op}}^2}{\mu_p} I. \tag{95}$$

Box constraints can destroy this property. With $\Phi(x) = x^2/2 + \iota_{[0,1]}(x)$, the response $x(\lambda) = [-\lambda]_{[0,1]}$ is constant in saturated price regions, making ∇h constant there. The experimental box constraints therefore do not imply global strong dual convexity. Their feasible-recovery bounds use Proposition D.1, which does not require it.

F Controlled Allocation Protocols

F.1 Instance generation and study blocks

There are $M = 4$ parties and $T = 24$ coordinates. The coordinates can represent time slots but have no temporal coupling. For each seed, `numpy.random.default_rng` draws arrays of shape $(20, 4, 24)$ in this order:

$$u_{mt} \sim \mathrm{Unif}[0.8, 1.2], \qquad a_{mt} \sim \mathrm{Unif}[0.5, 2], \qquad p_{mt} = u_{mt} v_{mt}, \quad v_{mt} \sim \mathrm{Unif}[0.1, 0.9]. \tag{96}$$

A final array of shape $(20, 24)$ gives $d_t = w_t \sum_m u_{mt}$ with $w_t \sim \mathrm{Unif}[0.4, 0.65]$. All underlying uniform draws are independent. The boxes u and demands d are public. a, p remain local. Strictly interior aggregate demand and positive curvature ensure a unique primal optimum. The depth, averaging, and work-grid study uses seeds 202700–202709. Fresh-grid comparisons and offline inexact-recovery checks use seeds 202800–202809. Stopping uses seeds 202900–202909. Each block contains 200 instances and each comparison is paired within its block. The work-grid selections are fixed before the fresh-grid evaluation. These studies have distinct roles and are not pooled into one sample. For the nonquadratic family, after drawing a quadratic instance, draw independent $\gamma_{mt} \sim \mathrm{Unif}[0.5, 2]$ and add $\gamma_{mt} \log\cosh(x_{mt} - p_{mt})$ to each coordinate cost. Its derivative and curvature are

$$a_{mt}(x_{mt} - p_{mt}) + \gamma_{mt} \tanh(x_{mt} - p_{mt}) + \lambda_t, \qquad a_{mt} + \gamma_{mt} \operatorname{sech}^2(x_{mt} - p_{mt}). \tag{97}$$

Thus $\mu \geq 0.5$ and $L_d = 8$ are valid for both families. The paired stopping families share the boxes, quadratic coefficients, offsets, and demand. The second additionally uses γ.

F.2 Local responses and exact call counts

For a quadratic coordinate, the exact response is $[p_{mt} - \lambda_t/a_{mt}]_{[0,u_{mt}]}$. A finite-work response first evaluates the derivatives at 0 and u_{mt}. A nonnegative derivative at zero returns zero. A nonpositive derivative at the upper bound returns u_{mt}. Otherwise start with $[\ell, h] = [0, u_{mt}]$. Repeat J times: set $c = (\ell + h)/2$, replace h by c if the derivative at c is nonnegative, and replace ℓ otherwise. Return $(\ell + h)/2$ from the final bracket. The nonquadratic routine uses its corresponding derivative with the same bracketing rule.

Boundary responses are exact. Each interior response differs from its exact minimizer by at most $u_{mt} 2^{-(J+1)}$. The vectorized routine executes all J midpoint passes and both boundary comparisons, including when a boundary is selected. Consequently one call performs $J + 2$ local derivative comparisons per coordinate. For Q calls, $C = Q(J + 2)$ per coordinate, TC per party, and MTC in total. Communication, price arithmetic, local averaging, recovery, and offline evaluation are separate costs. A scalar implementation that skips midpoint passes at boundaries would use a different work count. The price starts at zero and is updated using the received response at that price. The depth and work studies use $\eta = 1/8$. The inexact-bound and stopping studies use $\eta = 1/16$. Each run returns the Qth response or the declared average of its first Q responses, not a further solve at λ^{Q+1}. The reset ablation queries zero on every call. Since the response rule is deterministic, its invariance in Q follows by construction. It is not a lower bound for every memoryless protocol.

F.3 Public recovery and offline reference

For a received decision x, the public Euclidean box–balance projection has the form

$$z_{mt} = [x_{mt} - \nu_t]_{[0,u_{mt}]}, \qquad \sum_m [x_{mt} - \nu_t]_{[0,u_{mt}]} = d_t. \tag{98}$$

The specified numerical projection uses 80 bisections with initial endpoints $\min_m(x_{mt} - u_{mt}) - 1$ and $\max_m x_{mt} + 1$. It accesses only x, u, d. The objective is evaluated on these repaired vectors directly, with no local objective-minimizing completion. Eq. (84) and Appendix D.5 establish the required recovery constant. The quadratic offline optimum uses 80 price bisections on $\sum_m [p_{mt} - \lambda_t/a_{mt}]_{[0,u_{mt}]} = d_t$, with endpoints $\min_m a_{mt}(p_{mt} - u_{mt}) - 1$ and $\max_m a_{mt} p_{mt} + 1$. The nonquadratic reference uses 52 local bisections inside 72 price bisections. These high-accuracy references and evaluations of exact responses at queried prices are offline diagnostics. They do not influence the online policy or count as its responses.

F.4 Metrics and numerical interpretation

Let $z = Px$ and use the family's own objective F. Per-instance normalized quantities are

$$\text{mismatch} = \frac{\|Ax - d\|_2}{\|d\|_2}, \qquad \text{regret} = \frac{F(z) - F^\star}{\max\{F(0), 10^{-12}\}}, \tag{99}$$

$$\text{recovery distance} = \frac{\|z - x\|_F}{\|u\|_F}, \qquad \text{last-response error} = \frac{\|\widetilde{x}^Q - x(\lambda^Q)\|_F}{\sqrt{MT}}. \tag{100}$$

The last metric measures numerical response fidelity at an issued price, not objective performance. For averaging, the identity (69) is checked using the received responses. Each reported mean first averages the 20 instances in a seed. The final mean and sample standard deviation are computed across the ten seed means. Instance-level medians and quantiles, when reported, summarize the 200 instances rather than the seed means. These standard deviations are not confidence intervals. The displayed tables retain the precision of the reported summaries. High-accuracy reference and recovery checks are numerical consistency checks, not formal interval certificates. Tiny negative objective gaps at roundoff scale are clipped only after checking a 10^{-9} lower tolerance. No such clipping changes a decision rule. Stopping uses $V = \|u\|_F^2$, not the private objective normalization in (99).

G Empirical Results and Additional Comparisons

G.1 Depth and Output Choice

Figure 3 reports last-response performance for $Q \in \{1, 2, 4, 8, 16, 32\}$ with $J \in \{2, 4, 8, 16\}$ and exact responses. Every call uses its stated fidelity, including the first call at zero. The reset control has the same first response as retained-price feedback and repeats that response on every subsequent query. With coarse responses, extra updates may preserve or worsen objective error. These trajectories do not define the optimal achievable risk at that budget. The averaging extension retains the same instances

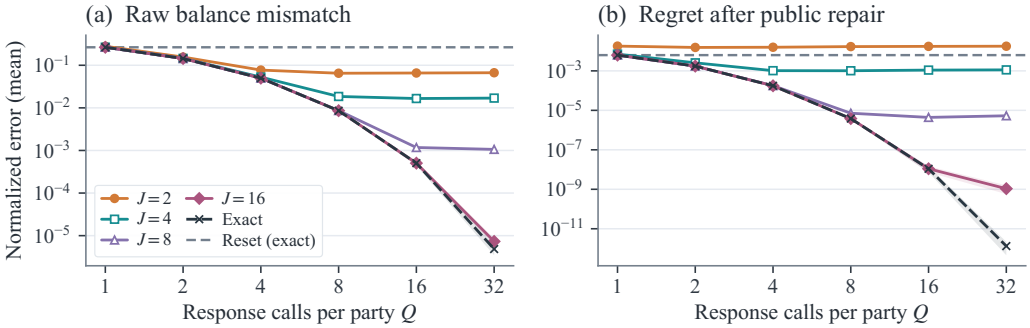


Figure 3: Last-response depth study on the 200 depth-study instances. Left: raw mismatch. Right: feasible regret after public recovery. Lines are means and shading is one sample standard deviation across ten seed means. The reset control repeats the same zero-price query.

and follows trajectories to $Q = 512$. Table 4 separates the changed decision rule from the local fidelity. At $J = 2$, the average's mismatch falls substantially while its regret remains positive. At $Q = 512$, its instance-level regret has median 0.01014 and 10th–90th percentiles $[0.00796, 0.01247]$. With exact responses, the last response can be much more accurate than an average that still includes initial transients. Recovery displacement decreases for the coarse average, so its objective floor is not merely a consequence of a large terminal projection.

G.2 Fixed Local Work and Fresh-Instance Checks

At $C = 64$, the grid is $(Q, J) \in \{(2, 30), (4, 14), (8, 6), (16, 2)\}$. At $C = 128$, it is $\{(4, 30), (8, 14), (16, 6), (32, 2)\}$. Both use the last-response output and public recovery. Table 5 contains every selection-grid cell. The selected pairs are $(8, 6)$ and $(8, 14)$. At $C = 64$, the per-seed mean difference in regret between $(4, 14)$ and $(8, 6)$ is 1.146×10^{-4}, with sample standard deviation 1.021×10^{-5}. The latter is better on $199/200$ paired instances. At $C = 128$, the difference between $(16, 6)$ and $(8, 14)$ is 6.585×10^{-5}, with seed-mean standard deviation 1.799×10^{-6}, and $(8, 14)$ is better on every instance. Nevertheless, $(16, 6)$ has lower mismatch. Selection and these comparisons use the same selection-stage instances, so they are descriptive.

The fresh-instance study fixes the selected pairs and evaluates the full feedback grid using seeds 202800–202809 (Table 6). The selected $C = 64$ pair beats $(4, 14)$ on $199/200$ instances and the

Table 4: Last and averaged outputs on the same 200 instances at $\eta = 1/8$. Entries are normalized means. Recovery denotes displacement, not the objective change caused by projection.

J	Q	Output	Mismatch	Regret	Recovery
2	32	Last	$6.677\,10^{-2}$	$1.776\,10^{-2}$	$3.528\,10^{-2}$
2	32	Average	$1.970\,10^{-2}$	$9.040\,10^{-3}$	$1.044\,10^{-2}$
2	512	Last	$6.674\,10^{-2}$	$1.759\,10^{-2}$	$3.529\,10^{-2}$
2	512	Average	$1.252\,10^{-3}$	$1.023\,10^{-2}$	$6.640\,10^{-4}$
Exact	32	Last	$4.851\,10^{-6}$	$1.323\,10^{-12}$	$2.922\,10^{-6}$
Exact	32	Average	$1.923\,10^{-2}$	$3.042\,10^{-5}$	$1.021\,10^{-2}$
Exact	512	Last	$5.284\,10^{-17}$	$3.068\,10^{-18}$	$7.898\,10^{-17}$
Exact	512	Average	$1.202\,10^{-3}$	$2.024\,10^{-7}$	$6.382\,10^{-4}$

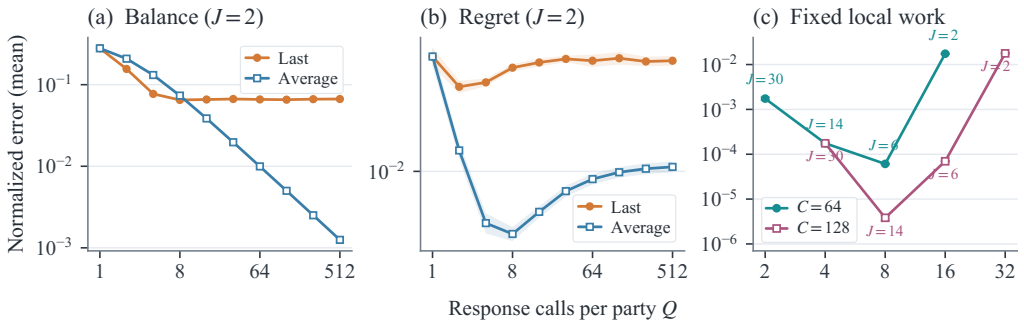


Figure 4: Output choice and the feedback work grid. Left and middle: last and averaged responses at $J = 2$. Right: tested feedback allocations at $C = Q(J + 2)$. Regret is normalized by $F(0)$ after public recovery. Bands and bars show sample standard deviations across ten seed means. This grid does not include model identification. Figure 5 provides that comparison on the fresh instances.

other two cells on every instance. The selected $C = 128$ pair beats all other cells on every instance. All ten seed-mean differences favor the selected pair in each comparison. This tests transfer within the specified generator, not distribution-free optimality.

Same-interface price bisection. The public generator bounds imply that $[-2.4, 2.4]$ brackets an equilibrium price in each coordinate for the quadratic family. The midpoint is queried with the same J-step local solver. Positive mismatch replaces the lower endpoint, and nonpositive mismatch replaces the upper endpoint. With approximate responses, the selected half need not contain the exact equilibrium. At the fixed pairs $(8, 6)$ and $(8, 14)$, this reference uses the same call count, local work, last-response output, and public recovery as gradient feedback. It stores two price endpoints rather

Table 5: All feedback work-grid cells, with $C = Q(J + 2)$ and $\eta = 1/8$. Regret is mean $\pm$ sample standard deviation across ten seed means. Quantiles are over individual instances. All metrics are normalized.

C	Q	J	Mismatch	Regret	Median [10th,90th]
64	2	30	0.14401	$(1.740 \pm 0.116)10^{-3}$	$1.650[1.018, 2.604]10^{-3}$
64	4	14	0.04984	$(1.755 \pm 0.102)10^{-4}$	$1.648[0.996, 2.608]10^{-4}$
64	8	6	0.00939	$(6.089 \pm 0.160)10^{-5}$	$5.972[4.965, 7.335]10^{-5}$
64	16	2	0.06581	$(1.744 \pm 0.036)10^{-2}$	$1.721[1.448, 2.103]10^{-2}$
128	4	30	0.04984	$(1.755 \pm 0.102)10^{-4}$	$1.648[0.995, 2.610]10^{-4}$
128	8	14	0.00859	$(3.823 \pm 0.498)10^{-6}$	$3.252[1.493, 6.466]10^{-6}$
128	16	6	0.00406	$(6.967 \pm 0.169)10^{-5}$	$6.823[5.746, 8.180]10^{-5}$
128	32	2	0.06677	$(1.776 \pm 0.029)10^{-2}$	$1.745[1.459, 2.105]10^{-2}$

Table 6: All fresh-instance work-grid cells at $\eta = 1/8$. The two selected pairs were fixed using the feedback grid. Regret is mean $\pm$ sample standard deviation across ten seed means.

C	Q	J	Mismatch	Feasible regret
64	2	30	0.14267	$(1.720 \pm 0.120)10^{-3}$
64	4	14	0.04951	$(1.765 \pm 0.196)10^{-4}$
64	8	6	0.00962	$(6.217 \pm 0.201)10^{-5}$
64	16	2	0.06660	$(1.749 \pm 0.062)10^{-2}$
128	4	30	0.04951	$(1.765 \pm 0.196)10^{-4}$
128	8	14	0.00860	$(4.030 \pm 0.674)10^{-6}$
128	16	6	0.00413	$(6.920 \pm 0.183)10^{-5}$
128	32	2	0.06534	$(1.763 \pm 0.052)10^{-2}$

Table 7: Fresh-instance comparison at the two previously selected allocations, both with $Q = 8$. Means are normalized. Local response work is matched, but tuning and coordinator state are not.

C	J	Coordinator	Mismatch	Feasible regret
64	6	Gradient	0.00962	$(6.217 \pm 0.201)10^{-5}$
64	6	Price bisection	0.01958	$(1.022 \pm 0.088)10^{-4}$
128	14	Gradient	0.00860	$(4.030 \pm 0.674)10^{-6}$
128	14	Price bisection	0.01873	$(4.926 \pm 0.678)10^{-5}$

than one price and is not separately tuned over the full work grid. Gradient feedback has lower regret on $190/200$ and $200/200$ paired instances at those two operating points. Since the allocations were selected for that method and only two reference settings are evaluated, these observations do not give a general ranking of optimizers.

G.3 Parametric identification through the same interface

The response vector can identify the declared separable models more directly than either price update above. For the quadratic family, the public ranges imply $a_{mt}p_{mt} \geq 0.5(0.1)(0.8) = 0.04$. Hence the responses at $\lambda = 0$ and $\lambda = 0.02\mathbf{1}$ are interior and satisfy (23). The baseline clips finite-precision coefficient estimates to $[0.5, 2]$, solves the inferred public-coupling problem by high-accuracy price bisection, and returns that feasible allocation. The solve uses only the two received vectors, public boxes and demand, and the declared parametric family. It makes no further private call.

The log-cosh family is also finitely identifiable. Its zero-price response is p. Use scalar probe levels $\ell_1 = -0.04$ and $\ell_2 = 0.075$, broadcast $\lambda_k = \ell_k\mathbf{1}$, and set $d_k = y(\lambda_k) - p$. The public ranges keep both probes interior. Coordinatewise,

$$\begin{bmatrix} d_1 & \tanh d_1 \\ d_2 & \tanh d_2 \end{bmatrix} \begin{bmatrix} a \\ \gamma \end{bmatrix} = \begin{bmatrix} -\ell_1 \\ -\ell_2 \end{bmatrix}. \tag{101}$$

Both sides of p have at least 0.08 distance to the box boundary. At that distance the derivative magnitude is at least $\frac{1}{2}(0.08 + \tanh 0.08) > 0.075$, proving that both nonzero probes remain interior. The odd, strictly increasing map $v \mapsto av + \gamma \tanh v$ implies $d_1 > 0$, $d_2 < 0$, and $|d_2| > |d_1|$. Since $\tanh v/v$ is strictly decreasing for $v > 0$, the determinant in (101) is nonzero. Together with the zero query, this identifies (p, a, γ) from three exact response vectors. With finite bisection, the baseline solves (101), replaces numerically singular estimates by the public-range midpoint, and clips both curvatures to $[0.5, 2]$. This inversion is ill-conditioned at coarse fidelity because d and $\tanh d$ are nearly collinear near zero. The finite-J results reflect that sensitivity.

For the quadratic rows, parameter error is mean relative error in a. For log-cosh it is the average of the relative errors in a and γ. The near-zero regrets are numerical-reference values, not symbolic zeros. This baseline is intentionally family-aware. It shows that the response interface, not merely the number of calls, determines the information available to a coordinator. It also prevents the iterative work grid from being interpreted as a lower bound on query complexity.

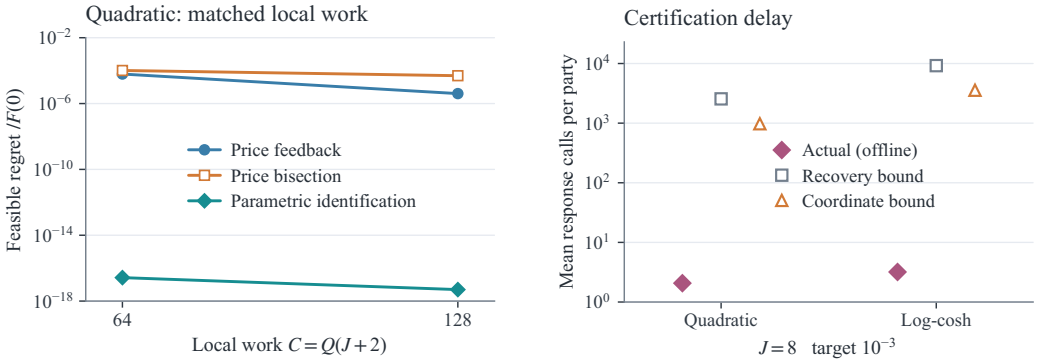


Figure 5: Information use and certification are different limitations. Left: fresh quadratic instances at matched local work. Identification uses two calls and the declared model family. Values near 10^{-17} are numerical-reference values. Right: mean first-target and certificate-crossing depths at $J = 8$, with target $\mathrm{Reg}/V \le 10^{-3}$. The actual target crossing is offline. Points show reported means. Available variability and paired diagnostics appear in Appendices G.2, G.3, and G.6.

Table 8: Identify-then-coordinate on the 200 fresh instances. Work is $C = Q(J + 2)$ derivative comparisons per coordinate. Regret is mean $\pm$ sample standard deviation across ten seed means. "Exact" uses analytic/high-accuracy responses and has no finite-work count.

Family	Q	J	C	Mean parameter error	Feasible regret
Quadratic	2	8	20	$7.34\,10^{-2}$	$(3.786 \pm 0.360)10^{-4}$
Quadratic	2	14	32	$1.27\,10^{-3}$	$(1.143 \pm 0.126)10^{-7}$
Quadratic	2	30	64	$1.95\,10^{-8}$	$(2.660 \pm 0.349)10^{-17}$
Quadratic	2	Exact	—	$1.19\,10^{-15}$	$4.756\,10^{-18}$
Log-cosh	3	19	63	$3.25\,10^{-1}$	$(6.714 \pm 0.565)10^{-4}$
Log-cosh	3	30	96	$3.07\,10^{-4}$	$(5.525 \pm 6.105)10^{-12}$
Log-cosh	3	Exact	—	$5.84\,10^{-11}$	$1.663\,10^{-17}$

G.4 Coupled-local-variable stress test

We draw a new block of 200 instances with seeds 203000–203009 and the same public box and balance generator. Party m has the private objective

$$f_m(x_m) = \frac{1}{2}\sum_{t=1}^{T} a_{mt}(x_{mt} - p_{mt})^2 + \frac{\kappa_m}{2}\sum_{t=1}^{T-1}\left[(x_{m,t+1} - p_{m,t+1}) - (x_{mt} - p_{mt})\right]^2, \qquad (102)$$

where $\kappa_m \sim \mathrm{Unif}[0.5, 2]$ is private. The banded Hessian is positive definite, and every conditional response couples all 24 coordinates. A response is computed from p_m by J projected-gradient steps of size 0.1 on the private box. The coordinator uses the same price update with $\eta = 1/8$ and the same public recovery. A 2000-step projected-gradient solve of the full public-coupling problem supplies the offline reference. Its largest projected-gradient fixed-point residual is 2.78×10^{-16}. Increasing depth improves the high-fidelity responses, whereas low local work leaves an objective floor even after the raw mismatch is small. The experiment is a mechanism stress test rather than a claim of optimal query complexity: the response family remains parametric, and other probing strategies may exploit that structure.

G.5 Inexact-recovery bounds on executed outputs

The offline bound study uses $\eta = 1/16$, $Q \in \{8, 32, 128, 512\}$, and $J \in \{2, 4, 8\}$ on 200 paired instances of each family. These are 4800 repeated outputs, not 4800 independent instances. For the

Table 9: Private banded-quadratic stress test on 200 fresh instances. Entries use the last response followed by public recovery. Regret and mismatch are normalized as in (99).

Local steps J	Q	Local gradient steps	Mismatch	Feasible regret
4	8	32	0.10886	$4.138\,10^{-3}$
4	32	128	0.00995	$1.885\,10^{-3}$
4	64	256	$6.70\,10^{-4}$	$1.740\,10^{-3}$
16	8	128	0.08665	$1.670\,10^{-3}$
16	32	512	0.00750	$9.457\,10^{-5}$
16	64	1024	$4.59\,10^{-4}$	$7.081\,10^{-5}$
64	8	512	0.08592	$1.490\,10^{-3}$
64	32	2048	0.00746	$1.682\,10^{-5}$
64	64	4096	$4.53\,10^{-4}$	$2.389\,10^{-7}$

interior coordinates $\mathcal{I}_q$ of each query, the conditional-error bound is

$$\delta_q = \sum_{(m,t)\in\mathcal{I}_q} \frac{(a_{mt}+\gamma_{mt})u_{mt}^2}{8\,4^J}, \tag{103}$$

with $\gamma = 0$ for the quadratic family. The offline analysis can use realized curvature and an optimal multiplier. The online stopping study below instead uses only public constants. Table 10 reports the displayed $Q = 512$ summaries. Saddle values and bounds, as well as regret and its recovery bound, use the family's own $F(0)$ normalization. No reported conditional-error, saddle, residual, price-distance, or recovery inequality is violated. The largest averaging-identity error is 2.93×10^{-14}, and the largest recovery imbalance is 8.88×10^{-16}. Nonquadratic reference balance and projected KKT residuals are at most 1.11×10^{-15} and 8.88×10^{-16}, respectively. These numerical checks support consistency of the implementation described in the study, not exact-real certification or tightness of the theoretical bounds.

Table 10: Mean normalized values at $Q = 512$, $\eta = 1/16$. Saddle error is the nonnegative expression in (61). The recovery bound is the right side of (62).

Family	J	Saddle error	Saddle bound	Regret	Recovery bound
Quadratic	2	0.01154	0.10301	0.01153	0.51602
Quadratic	4	$7.15\,10^{-4}$	0.00698	$7.04\,10^{-4}$	0.13509
Quadratic	8	$1.34\,10^{-5}$	$5.76\,10^{-4}$	$3.22\,10^{-6}$	0.05698
Log-cosh	2	0.01217	0.10761	0.01213	0.66139
Log-cosh	4	$7.59\,10^{-4}$	0.00789	$7.28\,10^{-4}$	0.19617
Log-cosh	8	$3.39\,10^{-5}$	0.00123	$3.66\,10^{-6}$	0.10709

G.6 Stopping with public information

Table 11: Stopping outcomes on 200 instances per family, target $\mathrm{Reg}\,/V \le 10^{-3}$, cap 16384. Entries are crossing counts/mean first depth conditional on crossing. A dash denotes no crossing. Actual uses private objectives offline. Residual failures count accepted outputs above target.

Family	J	Actual	Recovery	Coordinatewise	Residual failures
Quadratic	2	5/7.6	0/—	0/—	200
Quadratic	4	200/2.4	0/—	0/—	0
Quadratic	8	200/2.1	200/2553.4	200/976.8	0
Log-cosh	2	0/—	0/—	0/—	200
Log-cosh	4	200/4.1	0/—	0/—	0
Log-cosh	8	200/3.2	200/9215.6	200/3566.3	0

On seeds 202900–202909, use $\eta = 1/16$, $J \in \{2, 4, 8\}$, and target $\mathrm{Reg}(P\bar{x}_Q)/V \le 10^{-3}$ with public $V = \|u\|_F^2$. The recovery-based certificate uses (85), the coordinatewise certificate uses (78),

and the residual control stops when $K \|r_Q\| / V \leq 10^{-3}$. All three operate on the same trajectories and the same average-and-recovery output. The residual control deliberately omits local conditional error and is not asserted to be an optimality certificate. None of the three stopping rules receives $F(0)$, $F^\star$, objective values, realized private curvature, or an optimal multiplier. For each rule, replay records its first crossing and the corresponding local average. A stopped run would execute exactly that prefix. Further simulated calls permit the other rules and the cap to be evaluated. They are not part of the recorded stopping decision. A Q-call prefix uses $Q(J+2)$ derivative comparisons per coordinate.

Table 12: Extended-horizon stopping outcomes. Each row represents 200 paired instances. N is the number of crossings and E the number of crossed outputs above target. Mean Q conditions on crossing. A dash denotes no crossing, not zero cost.

Family	Rule	J	N	E	Mean Q
Quadratic	Certificate	2	0	0	—
Quadratic	Residual	2	200	200	2661.5
Quadratic	Certificate	4	0	0	—
Quadratic	Residual	4	200	0	2589.7
Quadratic	Certificate	8	200	0	2553.4
Quadratic	Residual	8	200	0	2587.6
Log-cosh	Certificate	2	0	0	—
Log-cosh	Residual	2	200	200	9662.3
Log-cosh	Certificate	4	0	0	—
Log-cosh	Residual	4	200	0	9303.0
Log-cosh	Certificate	8	200	0	9215.6
Log-cosh	Residual	8	200	0	9299.5

Table 13: Certificate delay at $J = 8$. "Actual" uses private objectives only as an offline first-target diagnostic. Delay ratio is the median of certificate depth divided by actual first-target depth over paired instances.

Family	Actual mean Q	Recovery mean Q	Coord. mean Q	Recovery ratio	Coord. ratio
Quadratic	2.07	2553.37	976.78	1347.0	516.5
Log-cosh	3.19	9215.58	3566.29	3603.3	1399.8

Horizon change. The initial cap $Q = 2048$ certified 20 quadratic $J = 8$ instances and no log-cosh instances. The residual rule stopped on ten quadratic $J = 2$ instances, all above target. After observing this censoring, the cap was extended to 16384 without changing seeds, target, rules, or accuracy levels. This extension is exploratory.

Table 12 includes every reported extended-horizon cell and non-crossing. The recovery-based certificate's long-run limitation at $J = 4$ has an analytic explanation. Eq. (79) gives asymptotic normalized floors $1/512 \approx 1.95 \times 10^{-3}$ for the quadratic family and $1/256 \approx 3.91 \times 10^{-3}$ for log-cosh. Both exceed the target. Under the bounded-price condition, the recovery-based certificate therefore eventually remains above the target. This asymptotic result does not exclude an earlier transient crossing. Neither certificate crosses in the reported runs, even though all capped outputs satisfy the target. Their maximum normalized regrets are 1.88×10^{-4} and 3.78×10^{-4}. At $J = 8$, recovery-based certified depths range from 1328 to 3868 for the quadratic family and 5200 to 13639 for log-cosh. Mean regret at those crossings, normalized by V, is 5.36×10^{-7} and 1.10×10^{-6}, far below the 10^{-3} target. The coordinatewise certificate reduces mean depth by more than 60% without an observed false acceptance, but Table 13 shows that it remains highly conservative. Non-crossing therefore does not establish insufficient accuracy or impossibility, and crossing can occur long after the evaluated output first meets the target.

Table 14 decomposes the two sufficient bounds at their own first crossings for $J = 8$. The fidelity contribution is below 1.6×10^{-5} in each case. The dominant positive contribution is the public recovery bound, at about 10^{-3}, partially offset by the negative price term. The coordinatewise rule changes both the error and recovery estimates, but the error contribution is nearly unchanged at these

crossings. This diagnoses recovery-bound conservatism in these two certificates on these trajectories. It does not lower-bound the delay of every certificate using the same information. At $J = 4$, by contrast, the recovery-bound fidelity term alone determines a limiting value above the target.

Table 14: Mean normalized certificate components at each rule's first crossing for $J = 8$. Components sum to the corresponding certificate value. The price telescope can be negative.

Family	Certificate	Price telescope	Fidelity	Recovery
Quadratic	Recovery	$-2.05\,10^{-5}$	$7.63\,10^{-6}$	$1.013\,10^{-3}$
Quadratic	Coordinatewise	$-5.37\,10^{-5}$	$7.54\,10^{-6}$	$1.046\,10^{-3}$
Log-cosh	Recovery	$-2.38\,10^{-5}$	$1.53\,10^{-5}$	$1.009\,10^{-3}$
Log-cosh	Coordinatewise	$-6.18\,10^{-5}$	$1.51\,10^{-5}$	$1.046\,10^{-3}$

No reported stopped or capped output violates its computed H_Q bound. The maximum recovery imbalance is 8.88×10^{-16} and the largest price/average telescoping discrepancy is 9.36×10^{-13}. These descriptive counts are not population error-rate confidence bounds.

H All-Price Ambiguity and Response-Position Certification

H.1 Proof of Theorem 4.1

The baseline conditional objective is strongly convex and coercive on the closed nonempty set $\mathcal{X}$. It therefore has a unique minimizer, and the response is continuous in λ. Write

$$\phi_0(x;\lambda) = \tfrac{1}{2}\|x-c\|_B^2 + \langle \lambda, Ax-b\rangle. \tag{104}$$

For $y = R_0(\lambda)$, the variational inequality at y and the exact quadratic expansion imply

$$\phi_0(x;\lambda) \geq \phi_0(y;\lambda) + \tfrac{1}{2}\|x-y\|_B^2 \qquad (x \in \mathcal{X}). \tag{105}$$

Since $\phi_\sigma(x;\lambda) = \phi_0(x;\lambda) - \sigma\langle Bv, x-c\rangle + \delta$,

$$\phi_\sigma(y;\lambda) - \phi_\sigma(x;\lambda) \leq -\tfrac{1}{2}\|x-y\|_B^2 + \sigma\langle Bv, x-y\rangle = \tfrac{1}{2}\|v\|_B^2 - \tfrac{1}{2}\|x-y-\sigma v\|_B^2 \leq \delta. \tag{106}$$

The same argument on party m gives δ_m, so this is a valid local oracle, not only an aggregate error construction. It is fixed before interaction and works on the entire price domain. For the two-world experiment, c, v, B, the sets, and the error caps may all be public. Only the sign is hidden. Fixing the protocol's random tape, induction on query depth shows that both worlds produce identical prices and responses. The terminal output therefore has the same distribution. The feasible optima $c \pm v$ have zero objective value, so

$$\frac{\mathbb{E}\,\mathrm{Reg}_+(\widehat{x}) + \mathbb{E}\,\mathrm{Reg}_-(\widehat{x})}{2} = \tfrac{1}{2}\mathbb{E}\,\|\widehat{x}-c\|_B^2 + \tfrac{1}{2}\|v\|_B^2 \geq \delta. \tag{107}$$

This also covers data-dependent stopping, since a common transcript produces the same stopping decision. The constant output c attains regret δ in each world. This proves the exact value for the specified pair and a lower bound for any larger oracle class containing it.

A coupled box example. Take two parties, $x_m \in [0,1]^2$, $A = [I\ I]$, $b = (1,1)^\top$, $c_m = (1/2, 1/2)^\top$, $v_1 = (a,a)^\top$, and $v_2 = -v_1$, with $0 < a < 1/2$. The local matrix $B_m = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$ has eigenvalues $1, 3$ and couples both coordinates. All assumptions hold and $\delta = 6a^2$. Thus complete two-dimensional responses, box boundaries, and arbitrary vector prices do not remove this obstruction. The hidden sign shifts the preferred decision within $\ker A$, and the common oracle can conceal that shift within its allowed error.

Sharp certificate reference. Suppose the transcript contains the zero-price response c and declares aggregate conditional error at most δ. In every compatible world, c is feasible and

$$F(c) - F^\star \leq F(c) - \min_{x \in \mathcal{X}} F(x) \leq \delta. \tag{108}$$

The two worlds above attain δ. Hence $C(T, c) = \delta$ for any compatible public class containing them. Eq. (17) has displacement zero and equals δ. This example identifies an information floor exactly. The numerical certificate comparisons elsewhere establish an improvement over particular sufficient bounds.

H.2 Proof and computation of Proposition 5.1

Write $g = \nabla\psi_\lambda(y)$. The public gradient caps imply the coordinate bounds in (16). For every w with $y + w \in \mathcal{X}$, approximate optimality and smoothness give

$$-\widehat{\delta} \le \psi_\lambda(y + w) - \psi_\lambda(y) \le \langle g, w\rangle + \tfrac{L}{2}\|w\|^2 . \tag{109}$$

Maximizing the negative quadratic expression proves $h_y(g) \le \widehat{\delta}$. For feasible z, weak duality gives

$$F(z) - F^\star \le \psi_\lambda(z) - \min_{x\in\mathcal{X}} \psi_\lambda(x) \le \widehat{\delta} + \langle g, z - y\rangle + \tfrac{L}{2}\|z - y\|^2 , \tag{110}$$

which proves (17). On $\mathbb{R}^n$, taking $w = -g/L$ shows $\|g\|^2 \le 2L\widehat{\delta}$. Cauchy–Schwarz yields (18). On a box that w may be infeasible, explaining why its boundary cannot be dropped from the calculation.

Separate local budgets. For $\mathcal{X} = \prod_m [0, u_m]$, known L_m, and local errors $\widehat{\delta}_m$, use $g_m = \nabla f_m(y_m) + A_m^\top\lambda$ and form $\mathcal{G}_{y_m}$ with each local budget. Then

$$\mathrm{Reg}(z) \le \sum_m \left[\widehat{\delta}_m + \sup_{g_m\in\mathcal{G}_{y_m}} \langle g_m, z_m - y_m\rangle + \frac{L_m}{2}\|z_m - y_m\|^2\right]. \tag{111}$$

Here y_m denotes the complete local response used in the position calculation, not just its interface image. Candidates can come from another query, averaging, or model fitting. They need only be locally and jointly feasible. Neither a specified projection nor the older recovery constant K enters this bound. No independence or coordinate separability of f_m is used.

Certified support upper bounds. For one party suppress its subscript, put $d_z = z - y$, and set $\ell_i = (A_m^\top\lambda)_i - G_i$, $r_i = (A_m^\top\lambda)_i + G_i$. The scalar contribution to h_y is

$$h_i(g_i) = -g_i w_i - \frac{L}{2}w_i^2, \qquad w_i = \mathrm{clip}(-g_i/L, -y_i, u_i - y_i). \tag{112}$$

For every $t \ge 0$, weak Lagrange duality gives

$$\sup_{g\in\mathcal{G}_y} \langle g, d_z\rangle \le t\widehat{\delta} + \sum_i \max_{\ell_i\le g_i\le r_i} \{(d_z)_i g_i - t h_i(g_i)\}. \tag{113}$$

Each h_i is convex with derivative $\mathrm{clip}(g_i/L, -(u_i - y_i), y_i)$. For $t > 0$, the scalar maximum is attained at one of ℓ_i, r_i, $\mathrm{clip}(L(d_z)_i/t, \ell_i, r_i)$. These candidates also cover flat tails through their endpoints. At $t = 0$, the maximum is $\max\{\ell_i(d_z)_i, r_i(d_z)_i\}$. The reference implementation takes the minimum over $t = 0$ and seven multiples $2^k t_0$, $k = -3, \dots, 3$, where $t_0 = \|d_z\|\sqrt{L/(2\widehat{\delta})}$. A positive numerical floor is used when $d_z = 0$. Every tested multiplier yields a valid upper bound. No convergence or strong-duality assumption is needed for this calculation.

Expanded multiplier search. Write the right-hand side of (113) as $U(t)$. It is convex in t. For maximizing scalar choices $g_i(t)$, a subgradient is $\widehat{\delta} - \sum_i h_i(g_i(t))$. The expanded search retains all reference candidates. Starting at $[t_0/8, 8t_0]$ with numerical positive floors, it makes eight bracketing passes: halve the left endpoint when its subgradient is positive and double the right endpoint when its subgradient is negative. It then makes 24 geometric-midpoint passes, updating the bracket by the subgradient sign. The returned value is the smallest evaluated upper bound, not a primal maximization value or a claim of dual convergence. Consequently the expanded value is no larger than the grid value in exact arithmetic. The added computation uses no response query.

History at a fixed candidate. For each retained common-price response, Proposition 5.1 applies to the same feasible z_Q. Taking the minimum proves (22). No independence between queries is required. One cannot instead minimize each party's bound over different prices, since the price terms then need not cancel under the coupling equality. Nor can a stored bound for z_q certify z_Q. The implementation retains full response/price/error triples and recomputes the complete sum for each selected triple. Retained history belongs to the protocol's state and is not free in a strict state budget. When $\widehat{\delta} = 0$, $h_i(g_i) = 0$ at every coordinate: $g_i = 0$ in the interior, $g_i \geq 0$ at $y_i = 0$, and $g_i \leq 0$ at $y_i = u_i$. Intersecting these intervals with the caps computes the support exactly. The code uses exact boundary equality rather than rounding interior positions to a boundary. Float64 evaluation adds $10^{-11}(1+|H|)$ to the aggregate bound. This is a numerical cushion, not interval arithmetic. The reported unit checks cover 100 scalar support problems, 600 boundary/interior conditional-gap problems, expanded-search domination, and historical-response recomputation at a changed candidate. They also check the midpoint error bound in $25,600$ local quadratic cases. Experimental bounds are compared with offline references. Mathematical validity comes from the assumptions and proof, not the absence of numerical failures.

Relation to certificate frameworks. Gladin et al. (2023) construct accuracy certificates from inexact first-order information and use dual certificates for primal recovery, with a declared multiplier region in that construction. Our calculation fixes a feasible candidate and bounds its gap from complete conditional responses and public curvature and error bounds. Smooth-convex interpolation (Taylor et al., 2016) characterizes consistency from sampled values and gradients. Here the gradient is unknown: $\mathcal{G}_y$ imposes necessary constraints and need not characterize all functions consistent with a multi-query record. Additional cross-query restrictions could tighten the relaxation.

H.3 A RESPONSE-PLUS-LOCAL-CERTIFICATE INTERFACE

The sum in (111) can be computed locally. This changes the messages available to the coordinator. It is separate from the response-only comparisons in Section 6.1.

Corollary H.1 (Local certificate summaries)**.** *At a common price λ, party m has a locally feasible candidate z_m and supplies a valid scalar h_m satisfying*

$$f_m(z_m) + \langle \lambda, A_m z_m \rangle - \inf_{x_m \in \mathcal{X}_m} \{ f_m(x_m) + \langle \lambda, A_m x_m \rangle \} \leq h_m. \tag{114}$$

If $\sum_m A_m z_m = b$, then $\mathrm{Reg}(z) \leq \sum_m h_m$.

Proof. The sum of local conditional infima, minus $\langle \lambda, b \rangle$, is a dual lower bound on $F^\star$. At a feasible candidate the price terms cancel. Subtracting that lower bound from $F(z)$ gives the inequality. □

For box-constrained smooth objectives, a party can compute h_m using its term in (111), retaining the response, box, and gradient-compatible calculation locally. The coordinator receives an interface vector and a certificate scalar rather than the full local response. A more general private convex set needs its own valid local bound computation. The scalar is an additional message, with local computation and any validation calls charged explicitly. Validity presumes cooperative execution of the declared calculation. The message alone is not a cryptographically verifiable proof.

Feasible combinations without central local constraints. Suppose party m retains feasible responses $r_m^1, \ldots, r_m^Q \in \mathcal{X}_m$, where $\mathcal{X}_m$ is convex, and the coordinator knows only $a_m^q = A_m r_m^q$. A linear feasibility problem finds weights

$$\alpha_{mq} \geq 0, \qquad \sum_q \alpha_{mq} = 1, \qquad \sum_{m,q} \alpha_{mq} a_m^q = b. \tag{115}$$

If it succeeds, party m forms $z_m = \sum_q \alpha_{mq} r_m^q$. Convexity gives local feasibility and (115) gives joint feasibility. Such weights exist exactly when $b \in \sum_m \mathrm{conv}\{a_m^q : q \leq Q\}$. This is the feasible-column combination principle of decomposition (Dantzig & Wolfe, 1960), not a new candidate-generation algorithm. Failure of the condition is a recovery failure. Successful recovery followed by $\sum_m h_m > \varepsilon$ is an inconclusive quality certificate. With coupling dimension p, n_m local coordinates, and Q retained responses, the coordinator's response history contains QMp scalars and the parties

retain $Q\sum_m n_m$ scalars. Transmitting dense combination weights costs a further MQ scalars, and validation returns M certificate scalars plus any newly requested interface responses. Prices, local certificate work, and further private solves are additional recorded costs. These resources fall outside a strict fixed-S implementation unless explicitly budgeted. This corollary supplies an implementation route. The experiments retain the central full-response interface.

I Identification Stability and Structural Misspecification

I.1 From response errors to feasible regret

An interior log-cosh response satisfies $a(y-p)+\gamma\tanh(y-p)=-\ell$. At zero price its exact value is p. If each of the three observations has coordinate error at most ϵ, each observed difference $\widehat{d}_j$ differs from d_j by at most 2ϵ. Since $\tanh$ is 1-Lipschitz, all four entries of $E=\widehat{T}-T$ have magnitude at most 2ϵ. Hence $\|E\|_{\mathrm{op}}\le\|E\|_F\le 4\epsilon$. If $4\epsilon<\sigma_{\min}(T)$, the identity

$$\widehat{\theta}-\theta=-\widehat{T}^{-1}E\theta \tag{116}$$

and the smallest-singular-value bound prove (24). Euclidean projection onto the known parameter box does not increase parameter error. An observable alternative, when $\widehat{T}$ is invertible, is

$$\left\|\widehat{\theta}-\theta\right\|\le\min\left\{\frac{4\epsilon\sqrt{8}}{\sigma_{\min}(\widehat{T})},\ 1.5\sqrt{2}\right\}=e_\theta \tag{117}$$

for $a,\gamma\in[0.5,2]$ and a clipped fit. At a singular matrix only the box-diameter term is used. A pseudoinverse followed by clipping supplies the candidate then. For $x,p,\widehat{p}\in[0,u]$ and $|p-\widehat{p}|\le\epsilon$,

$$|\widehat{f}'(x)-f'(x)|\le e_\theta\sqrt{u^2+\tanh^2 u}+4\epsilon=:e_g. \tag{118}$$

The first term bounds changes in (a,γ) at a fixed center. The second bounds the center change using $a+\gamma\le 4$. Summing squared coordinate bounds gives $\epsilon_g^2=\sum_i e_{g,i}^2$. For bisection, $\epsilon_i=u_i/2^{J+1}$. Let the true objective F be μ-strongly convex, and let a convex fitted objective have gradient error at most ϵ_g throughout their common convex feasible set. Its feasible minimizer $\widehat{x}$ satisfies

$$\mathrm{Reg}(\widehat{x})\le\epsilon_g^2/(2\mu) \tag{119}$$

To prove this bound, let $x^\star$ minimize F over that set. Strong convexity at $\widehat{x}$ and optimality of the fitted convex problem give

$$F(\widehat{x})-F(x^\star)\le\langle\nabla F(\widehat{x}),\widehat{x}-x^\star\rangle-\tfrac{\mu}{2}\left\|\widehat{x}-x^\star\right\|^2\le\epsilon_g\left\|\widehat{x}-x^\star\right\|-\tfrac{\mu}{2}\left\|\widehat{x}-x^\star\right\|^2\le\epsilon_g^2/(2\mu). \tag{120}$$

If the fitted solution has variational residual at most ξ, meaning $\langle\nabla\widehat{F}(\widehat{x}),\widehat{x}-x^\star\rangle\le\xi$, add ξ. For a misspecified class, uniform gradient discrepancy from approximation must be added to numerical estimation error before squaring. This is a sensitivity argument. Inverse optimization already studies decision-relevant losses (Chen & Kılınç-Karzan, 2021).

I.2 Safe three-query probing

The paired study uses seeds 203100–203109, 20 instances each, with the log-cosh distributions in Appendix F. Both policies know the exact family and its public ranges. At $J\in\{8,14,19,30\}$, both use one zero-price call and two common-price calls. The fixed probes are $-0.04,0.075$. After observing $\widehat{p}$, the adaptive policy computes at each time coordinate

$$r_-=\min_m(\widehat{p}_m-\epsilon_m)_+,\qquad r_+=\min_m(u_m-\widehat{p}_m-\epsilon_m)_+, \tag{121}$$

and limits $b_+=0.9(r_-+\tanh r_-)/2$, $b_-=0.9(r_++\tanh r_+)/2$. Lower parameter bounds ensure the prices remain nonsaturating. Use the larger-magnitude limit first and the opposite sign second, capped at half the first magnitude. This avoids symmetric probes, whose rows would be dependent in the centered odd derivative model. Every party receives the same vector price. No individual private center is used. Cost remains $Q=3$, or $3(J+2)$ derivative comparisons per coordinate. The fitted allocation is solved to the same numerical accuracy for both policies.

I.3 A COUPLED NONLINEAR SERIES

Each instance has four parties with $T = 8$ decisions, boxes $u_{mt} \sim \mathrm{Unif}[0.8, 1.2]$, centers $p_{mt}/u_{mt} \sim \mathrm{Unif}[0.3, 0.7]$, and demand $d_t / \sum_m u_{mt} \sim \mathrm{Unif}[0.4, 0.65]$. With first-difference matrix D, set

$$f_m(x) = \tfrac{1}{2}(x - p_m)^\top B_m (x - p_m) + \zeta \sum_{r=1}^{3} w_{mr} \log\cosh(v_{mr}^\top (x - p_m)),$$
$$B_m = \mathrm{diag}(a_m) + \kappa_m D^\top D. \tag{122}$$

Here $a_{mt} \sim \mathrm{Unif}[0.5, 2]$, $w_{mr} \sim \mathrm{Unif}[0.3, 1]$, and v_{mr} are independently normalized Gaussian directions. Settings are diagonal quadratic ($\kappa = \zeta = 0$), coupled quadratic ($\kappa_m \sim \mathrm{Unif}[0.25, 0.5]$, $\zeta = 0$), and the same coupled construction with $\zeta \in \{0.5, 1\}$. Five seeds 203200–203204, 20 instances each, are paired across settings and protocols. These are controlled structural changes, not an external-domain benchmark. Every coordinator receives the same public box, balance, $\mu = 0.5$, $L = 8$, absolute coordinate-gradient cap $G_i = 8$, and solver error formula. The bounds hold because $\|B_m\|_{\mathrm{op}} \le 4$, nonlinear Hessian norm is at most 3, and $\|B_m\|_\infty u_{\max} + 3 \le 7.8$. Realized matrices, centers, directions, and weights are private. The dense fitting baseline can represent cross-coordinate quadratic responses even when the diagonal fit cannot.

Local response accuracy. Every call starts projected gradient descent at $u_m/2$, with step $1/L$ and $J \in \{64, 128, 256\}$ gradient evaluations. Its public error cap is (25). One proximal-gradient step bounds the gap by $L \left\|x^0 - x^\star\right\|^2 / 2$. Since $x^0 = u_m/2$ and $x^\star \in [0, u_m]$, the squared distance is at most $\|u_m\|^2 / 4$. For subsequent steps, minimize the quadratic upper model and compare with $x + \alpha(x^\star - x)$, $\alpha = \mu/L$. Strong convexity cancels the squared-distance term, giving gap contraction by $1 - \mu/L$. Adding a price changes neither curvature constant. This proves the midpoint cap uniformly over prices and at boundaries. The diameter cap $L \|u_m\|^2 (1 - \mu/L)^{J-1}/2$ remains only a paired reference. Their normalized aggregate values at $J = 128$ are 2.756×10^{-4} and 1.103×10^{-3}. Every method uses the improved cap. This changes neither the local solver nor the separate bisection error formula.

Candidate generation and charged validation. The diagonal fit queries 0 and $0.04\mathbf{1}$, estimates each inverse slope, and clips fitted curvature to $[\mu, L]$. The dense fit queries 0, $\pm 0.04\mathbf{1}$, and $\pm 0.04 e_t$, for $2T + 3 = 19$ calls. It regresses $y(0) - y(\lambda)$ on λ, symmetrizes the compliance matrix, clips its eigenvalues to $[1/L, 1/\mu]$, and inverts it. Both solve the fitted feasible quadratic by 600 public projected-gradient iterations. No objective value is fitted. Feedback uses 19 calls, step $1/16$, and its recovered average or last response. Every candidate receives an additional validation call, counted in Q and local work. The fit uses its implied multiplier from interior-coordinate stationarity of the fitted problem. Feedback repeats its last queried multiplier. Four nested refinements keep the candidate and responses fixed: diameter cap with grid search; midpoint cap with grid; midpoint with expanded search; and the minimum expanded bound over validation and retained proposal responses. The last rule uses (22). No candidate is regenerated. The diagonal fit uses $Q = 3$, and other methods $Q = 20$. Work is QJ local vector-gradient evaluations per party, not scalar derivative comparisons. Retaining 2 or 19 proposal triples costs 704 or 6688 bytes per instance in float64 ($MT + T + M = 44$ scalars per triple), excluding current-response workspace and shared public arrays. Fitting time includes the feasible central solve. Recovery timing covers the two feedback projections together. Certificate components are timed separately. Every rule uses actual responses, not a fitted objective's error claim.

Scoring and verification. Regret is divided by public $V = \|u\|_F^2$, with target 10^{-3}. Offline references use 1200 projected-gradient iterations on the true objective. Their projected-step residual is checked below 10^{-9}. Every candidate is checked for balance and bounds. Reported runs use NumPy float64 on a remote CPU host with one BLAS thread. Timings are batched throughput per instance, not network latency. Instance records include seeds, calls, work, objective error, certificate value, acceptance, feasibility, and computation times. Good candidates may be rejected when the error cap is loose or the validation price poorly supports them.

J Additional Response-Information Results

J.1 Same-output certification records

The position study uses certificate seeds 202900–202909, 20 instances each, for both separable families and $J \in \{4, 8\}$. Every first-crossing rule is evaluated on the same prefix and projection of the running average. The 512-call cap is shared by this comparison. The 16384-call results for recovery-based certificates are a separate longer run. First depths agree with the offline target attainment. At $J = 8$, the new delay beyond attainment averages 6.515 calls for quadratics and 17.130 for log-cosh. At $J = 4$, position-certificate crossings are censored at the cap, not assigned a depth of 512. For each instance, family, and fidelity, we evaluated five rules: offline target attainment, the recovery-based, coordinate, and position certificates, and the minimum of the last two. Non-crossings were stored with depth -1 and no crossing regret. Bounds and regret were evaluated at every prefix through 512 calls. All 409600 prefix/configuration evaluations in the four cells were checked numerically. The combined rule has the same first crossings as the position rule in these cells, not a pointwise domination theorem. Reference position calculation costs 52–56 microseconds per instance in batch, excluding local solves and recovery.

Table 15: Acceptance (%) at selected prefixes on the separable families, 200 instances per cell. All rules use the specified bisection error caps. History retains earlier power-of-two prefixes. These are fixed-prefix rates, not first-crossing depths.

Family	J	Q	Offline target	Grid	Expanded	History
Quadratic	4	8	100.0	0.0	0.0	0.0
Quadratic	4	32	100.0	0.0	0.0	0.0
Quadratic	4	512	100.0	0.0	0.0	0.0
Quadratic	8	8	100.0	52.5	52.5	52.5
Quadratic	8	32	100.0	99.0	99.0	100.0
Quadratic	8	512	100.0	100.0	100.0	100.0
Log-cosh	4	8	94.5	0.0	0.0	0.0
Log-cosh	4	32	100.0	0.0	0.0	0.0
Log-cosh	4	512	100.0	0.0	0.0	0.0
Log-cosh	8	8	96.0	0.0	0.0	0.0
Log-cosh	8	32	100.0	99.0	99.0	99.0
Log-cosh	8	512	100.0	100.0	100.0	100.0

Expanded search and retained prefixes. Grid, expanded-search, and history rules are also evaluated at $Q \in \{1, 2, 4, 8, 16, 32, 64, 128, 256, 512\}$. History contains earlier power-of-two responses, recomputed at the current recovered average. The current response is included. At most nine earlier triples are retained, costing 8928 bytes per instance in float64: each triple has $MT + T + M = 124$ scalars. The midpoint projected-gradient cap does not apply to this bisection experiment. Across 8000 paired checkpoint/configuration records, expanded search lowers 1432 bounds by more than 10^{-12}, but adds no acceptances. History lowers a further 2367 and adds three accepted checkpoint records. For quadratic $J = 8, Q = 32$, acceptance rises from 99% to 100% (Table 15). These counts concern repeated configurations, not distinct successful trials or earlier stopping times. At $J = 4$, no tested checkpoint passes. Expanded evaluation costs a median 339 microseconds per instance. History adds a median 1.57 milliseconds. The small acceptance gain therefore requires central work despite using no additional private response.

J.2 Stable-probe outcomes

Here $V = \sum_m \|u_m\|_2^2$ is the public, instance-specific normalization. *Fixed (%)* and *Adaptive (%)* report the fraction of 200 instances satisfying $\mathrm{Reg}(z)/V \le 10^{-3}$ under fixed and adaptive probing, respectively. These are offline target-attainment rates, not certificate acceptance rates. Each seed contributes 20 instances. Every observed probe is interior in these runs. For the known family, the observed-matrix bound (117) and decision bound (119) hold on each checked instance to numerical tolerance. They are sufficient sensitivity bounds and can be much looser than actual error. The checks track the smallest observed singular values, parameter error, decision error, and its bound to distinguish conditioning from decision quality. Neither policy uses true coefficients to choose prices.

Table 16: Stable probing with the same three calls. Normalized regret $\mathrm{Reg}\,/V$ is reported as mean $\pm$ sample standard deviation across ten seed means. Both policies know the same log-cosh family and parameter ranges.

J	Fixed-probe regret	Adaptive-probe regret	Fixed (%)	Adaptive (%)
8	$1.04 \times 10^{-3} \pm 1.30 \times 10^{-4}$	$8.05 \times 10^{-4} \pm 8.69 \times 10^{-5}$	52.5	78.0
14	$8.95 \times 10^{-4} \pm 1.01 \times 10^{-4}$	$1.44 \times 10^{-4} \pm 1.88 \times 10^{-5}$	69.5	100.0
19	$2.25 \times 10^{-4} \pm 3.18 \times 10^{-5}$	$2.91 \times 10^{-6} \pm 1.82 \times 10^{-6}$	100.0	100.0
30	$2.54 \times 10^{-12} \pm 6.55 \times 10^{-12}$	$5.70 \times 10^{-16} \pm 3.26 \times 10^{-16}$	100.0	100.0

J.3 Full structural grid and cost

Table 3 reports the $J = 256$ comparison, while Table 17 gives every tested fidelity. All four settings use the same five seed blocks and public constants. Fitting procedures receive no additional private information. Dense fitting is a stronger approximation class than diagonal fitting and uses a larger displayed query budget. Its 19 proposal calls match feedback. The two-query fit is a cheaper reference, not a budget-matched competitor. An additional validation call is counted for every method.

Table 17: All structural configurations. Acceptance uses the midpoint error bound and reference multiplier grid. Each row has 100 paired instances. $Q = 3$ for diagonal fitting and 20 otherwise. Regret is normalized by V.

Setting	Candidate	J	Mean regret	Seed SD	Target (%)	Accept (%)
Diagonal quadratic	Diagonal fit	64	3.80×10^{-8}	1.30×10^{-8}	100	0
Diagonal quadratic	Dense SPD fit	64	3.80×10^{-8}	1.30×10^{-8}	100	0
Diagonal quadratic	Feedback average	64	2.72×10^{-5}	5.29×10^{-6}	100	0
Diagonal quadratic	Feedback last	64	1.70×10^{-7}	3.32×10^{-8}	100	0
Diagonal quadratic	Diagonal fit	128	4.90×10^{-12}	2.43×10^{-12}	100	100
Diagonal quadratic	Dense SPD fit	128	4.90×10^{-12}	2.43×10^{-12}	100	100
Diagonal quadratic	Feedback average	128	2.64×10^{-5}	5.09×10^{-6}	100	72
Diagonal quadratic	Feedback last	128	1.07×10^{-7}	1.70×10^{-8}	100	100
Diagonal quadratic	Diagonal fit	256	8.44×10^{-19}	2.18×10^{-19}	100	100
Diagonal quadratic	Dense SPD fit	256	9.31×10^{-19}	2.15×10^{-19}	100	100
Diagonal quadratic	Feedback average	256	2.64×10^{-5}	5.09×10^{-6}	100	100
Diagonal quadratic	Feedback last	256	1.06×10^{-7}	1.69×10^{-8}	100	100
Coupled quadratic	Diagonal fit	64	9.55×10^{-5}	1.23×10^{-5}	100	0
Coupled quadratic	Dense SPD fit	64	7.88×10^{-10}	6.85×10^{-10}	100	0
Coupled quadratic	Feedback average	64	4.69×10^{-5}	7.88×10^{-6}	100	0
Coupled quadratic	Feedback last	64	3.38×10^{-6}	4.34×10^{-7}	100	0
Coupled quadratic	Diagonal fit	128	9.58×10^{-5}	1.24×10^{-5}	100	1
Coupled quadratic	Dense SPD fit	128	1.14×10^{-14}	1.50×10^{-14}	100	100
Coupled quadratic	Feedback average	128	4.69×10^{-5}	7.87×10^{-6}	100	21
Coupled quadratic	Feedback last	128	3.38×10^{-6}	4.31×10^{-7}	100	43
Coupled quadratic	Diagonal fit	256	9.58×10^{-5}	1.24×10^{-5}	100	1
Coupled quadratic	Dense SPD fit	256	9.57×10^{-19}	3.74×10^{-19}	100	100
Coupled quadratic	Feedback average	256	4.69×10^{-5}	7.87×10^{-6}	100	93
Coupled quadratic	Feedback last	256	3.38×10^{-6}	4.31×10^{-7}	100	96
Coupled, $\zeta = 0.5$	Diagonal fit	64	1.91×10^{-4}	3.09×10^{-5}	99	0
Coupled, $\zeta = 0.5$	Dense SPD fit	64	5.14×10^{-9}	3.30×10^{-9}	100	0
Coupled, $\zeta = 0.5$	Feedback average	64	5.69×10^{-5}	8.88×10^{-6}	100	0
Coupled, $\zeta = 0.5$	Feedback last	64	4.56×10^{-6}	5.76×10^{-7}	100	0
Coupled, $\zeta = 0.5$	Diagonal fit	128	1.92×10^{-4}	3.09×10^{-5}	99	1
Coupled, $\zeta = 0.5$	Dense SPD fit	128	4.85×10^{-9}	3.25×10^{-9}	100	100
Coupled, $\zeta = 0.5$	Feedback average	128	5.69×10^{-5}	8.87×10^{-6}	100	11
Coupled, $\zeta = 0.5$	Feedback last	128	4.55×10^{-6}	5.71×10^{-7}	100	30
Coupled, $\zeta = 0.5$	Diagonal fit	256	1.92×10^{-4}	3.09×10^{-5}	99	3
Coupled, $\zeta = 0.5$	Dense SPD fit	256	4.85×10^{-9}	3.25×10^{-9}	100	100
Coupled, $\zeta = 0.5$	Feedback average	256	5.69×10^{-5}	8.87×10^{-6}	100	90
Coupled, $\zeta = 0.5$	Feedback last	256	4.55×10^{-6}	5.71×10^{-7}	100	95
Coupled, $\zeta = 1$	Diagonal fit	64	4.17×10^{-4}	6.24×10^{-5}	95	0
Coupled, $\zeta = 1$	Dense SPD fit	64	8.34×10^{-9}	5.12×10^{-9}	100	0
Coupled, $\zeta = 1$	Feedback average	64	8.22×10^{-5}	1.15×10^{-5}	100	0
Coupled, $\zeta = 1$	Feedback last	64	7.30×10^{-6}	9.74×10^{-7}	100	0
Coupled, $\zeta = 1$	Diagonal fit	128	4.17×10^{-4}	6.24×10^{-5}	95	0
Coupled, $\zeta = 1$	Dense SPD fit	128	8.19×10^{-9}	5.19×10^{-9}	100	100
Coupled, $\zeta = 1$	Feedback average	128	8.22×10^{-5}	1.15×10^{-5}	100	7
Coupled, $\zeta = 1$	Feedback last	128	7.29×10^{-6}	9.68×10^{-7}	100	22
Coupled, $\zeta = 1$	Diagonal fit	256	4.17×10^{-4}	6.24×10^{-5}	95	3
Coupled, $\zeta = 1$	Dense SPD fit	256	8.19×10^{-9}	5.19×10^{-9}	100	100
Coupled, $\zeta = 1$	Feedback average	256	8.22×10^{-5}	1.15×10^{-5}	100	81
Coupled, $\zeta = 1$	Feedback last	256	7.29×10^{-6}	9.68×10^{-7}	100	93

J.4 WHICH REFINEMENT CHANGES ACCEPTANCE?

The four columns in Table 18 hold every candidate, response, price, and offline regret fixed. The diameter-grid column uses the unrefined public error bound. The midpoint initialization bound alone changes acceptance in 811 of the 4800 candidate/configuration records: 808 at $J = 128$ and three at $J = 256$. In particular, dense-fit acceptance at $J = 128$ increases from zero to 100% in each setting. This removes avoidable conservatism in the public error bound without changing the local solver.

Table 18: Acceptance (%) on identical structural candidates. D: diameter bound and reference grid. M: midpoint bound and reference grid. S: midpoint bound with expanded search. H: additionally reuse proposal responses at the same candidate. Target is the offline target fraction. Each row has 100 paired instances.

Setting	Candidate	J	Target	D	M	S	H
Diagonal quadratic	Diagonal fit	64	100	0	0	0	0
Diagonal quadratic	Dense SPD fit	64	100	0	0	0	0
Diagonal quadratic	Feedback average	64	100	0	0	0	0
Diagonal quadratic	Feedback last	64	100	0	0	0	0
Diagonal quadratic	Diagonal fit	128	100	0	100	100	100
Diagonal quadratic	Dense SPD fit	128	100	0	100	100	100
Diagonal quadratic	Feedback average	128	100	0	72	72	72
Diagonal quadratic	Feedback last	128	100	0	100	100	100
Diagonal quadratic	Diagonal fit	256	100	100	100	100	100
Diagonal quadratic	Dense SPD fit	256	100	100	100	100	100
Diagonal quadratic	Feedback average	256	100	100	100	100	100
Diagonal quadratic	Feedback last	256	100	100	100	100	100
Coupled quadratic	Diagonal fit	64	100	0	0	0	0
Coupled quadratic	Dense SPD fit	64	100	0	0	0	0
Coupled quadratic	Feedback average	64	100	0	0	0	0
Coupled quadratic	Feedback last	64	100	0	0	0	0
Coupled quadratic	Diagonal fit	128	100	0	1	1	1
Coupled quadratic	Dense SPD fit	128	100	0	100	100	100
Coupled quadratic	Feedback average	128	100	0	21	21	21
Coupled quadratic	Feedback last	128	100	0	43	43	43
Coupled quadratic	Diagonal fit	256	100	1	1	1	1
Coupled quadratic	Dense SPD fit	256	100	100	100	100	100
Coupled quadratic	Feedback average	256	100	93	93	93	93
Coupled quadratic	Feedback last	256	100	96	96	96	96
Coupled, $\zeta = 0.5$	Diagonal fit	64	99	0	0	0	0
Coupled, $\zeta = 0.5$	Dense SPD fit	64	100	0	0	0	0
Coupled, $\zeta = 0.5$	Feedback average	64	100	0	0	0	0
Coupled, $\zeta = 0.5$	Feedback last	64	100	0	0	0	0
Coupled, $\zeta = 0.5$	Diagonal fit	128	99	0	1	1	1
Coupled, $\zeta = 0.5$	Dense SPD fit	128	100	0	100	100	100
Coupled, $\zeta = 0.5$	Feedback average	128	100	0	11	11	11
Coupled, $\zeta = 0.5$	Feedback last	128	100	0	30	30	30
Coupled, $\zeta = 0.5$	Diagonal fit	256	99	3	3	3	3
Coupled, $\zeta = 0.5$	Dense SPD fit	256	100	100	100	100	100
Coupled, $\zeta = 0.5$	Feedback average	256	100	89	90	90	90
Coupled, $\zeta = 0.5$	Feedback last	256	100	95	95	95	95
Coupled, $\zeta = 1$	Diagonal fit	64	95	0	0	0	0
Coupled, $\zeta = 1$	Dense SPD fit	64	100	0	0	0	0
Coupled, $\zeta = 1$	Feedback average	64	100	0	0	0	0
Coupled, $\zeta = 1$	Feedback last	64	100	0	0	0	0
Coupled, $\zeta = 1$	Diagonal fit	128	95	0	0	0	0
Coupled, $\zeta = 1$	Dense SPD fit	128	100	0	100	100	100
Coupled, $\zeta = 1$	Feedback average	128	100	0	7	7	7
Coupled, $\zeta = 1$	Feedback last	128	100	0	22	22	22
Coupled, $\zeta = 1$	Diagonal fit	256	95	2	3	3	3
Coupled, $\zeta = 1$	Dense SPD fit	256	100	100	100	100	100
Coupled, $\zeta = 1$	Feedback average	256	100	80	81	81	81
Coupled, $\zeta = 1$	Feedback last	256	100	93	93	93	93

Expanded search strictly lowers 16 structural bounds and history lowers another 108 (threshold 10^{-12}), but neither adds an acceptance beyond the midpoint-grid rule. The cheap rule and full refinement thus have identical acceptance on this grid although their numerical bounds differ. Every rule remains inconclusive at $J = 64$. Some accurate diagonal-fit candidates remain uncertified at $J = 256$. A poor validation price can leave a good candidate far from its conditional response. This is different from candidate-generation failure, and the reference multiplier grid is sufficient for most comparisons here.

Computation and numerical checks. The fitted central solves cost approximately 1.8–2.9 milliseconds per instance. Across the 48 batch/configuration measurements, midpoint-grid certification has median cost 18 microseconds (range 14–41). Expanded search has median 109 microseconds (range 94–182). History adds a median 2.09 milliseconds, with the two-response fit requiring less work than methods retaining 19 proposals. Full expanded-plus-history evaluation ranges from 0.29 to 3.06 milliseconds per instance. These are batch CPU throughput measurements, not end-to-end latency. Local response, validation, fitting, recovery, and certification times are recorded separately. All balances have absolute violation below 9×10^{-16} and reference projected-step residuals pass the 10^{-9} check.

Internal replay check. A separate replay of the certificate calculations used recorded structural candidates, public bounds, validation responses, and retained proposal responses, without accessing private objective parameters. It recomputed all 19200 bounds (four rules on 4800 paired records). The recomputed values matched the recorded values exactly on that host. Same-output crossing records and identification results also matched the preceding run, while structural candidate regrets remained unchanged. These float64 checks assess internal numerical consistency, not interval-certified guarantees.