FORECASTS THAT MOVE DECISIONS: TRUTHFULNESS VS. INFLUENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

A forecaster whose report changes a decision can benefit from misreporting even under a strictly proper score. We study protocols that separate a scored forecast from a bounded request to adjust the executed plan. With an independent identifying audit and plan-only operational routing, we characterize truthfulness for affine protocols and convex operational costs: the best attainable operational value must be minimized at the truthful report. A first-order certificate makes this condition independent of the weight on any differentiable strictly convex score. For quadratic costs, metric projection yields an exact vector solution and bounds on report distortion; in one dimension, truthfulness has a closed-form authority threshold. The principal's preferred authority can nevertheless be smaller, because honest reports may accompany systematically biased plans. We also derive the most informative feasible plan-preserving perturbation for detecting direct message effects. Binding requests prevent symmetric tests but permit one-sided tests, and prescribed symmetric-test power requires an explicit authority margin beyond truthfulness. Finally, we bound the propagation of estimation error through the equilibrium map, separating persistent report distortion from statistical and operational sources of decision loss. Simulations examine these results under bounded authority, censored feedback, and multi-item capacity constraints.

1 INTRODUCTION

A forecast used to choose an action has two roles: it describes uncertainty and it influences a decision. These roles conflict when the forecaster cares about the decision. A demand forecaster evaluated partly on stockouts, for example, may prefer an order above its own demand estimate. A proper score penalizes inaccurate reporting, but the operational benefit of a changed order can outweigh that penalty. This conflict is established in work on self-interested experts (Boutilier, 2012); our question is how to control it through the allocation of decision authority.

We study a protocol with two messages: a scored forecast and a bounded request to adjust the plan. The protocol publishes a single plan, and operational agents and future state transitions depend on that plan. The original messages are retained for audit. This structure is motivated by statistical baselines and documented overrides in demand planning (Fildes et al., 2009; 2025), and separates two design choices: how strongly the report moves the plan, and which adjustments the request can implement. An identifying audit, whose observation law is independent of the current messages, gives the forecast a stable statistical target.

Running example. Suppose mean demand is 100, but the forecaster prefers an order of 120. A request that can implement 120 while reporting 100 removes the incentive to distort the forecast. If the request is limited to ten units above or below the report, however, raising the report also raises the attainable orders. With score cost $2(q-100)^2$ and plan cost $(y-120)^2$, the optimal report is $q = 103\frac{1}{3}$ and the order is $y = 113\frac{1}{3}$. Anchoring the same adjustment band at a fixed benchmark of 100 gives a report-independent interval $[90, 110]$: the optimal report is 100 and the order is 110. Thus authority, report accuracy, and decision quality are distinct properties of a protocol.

Contributions. We develop a geometric account of these properties. First, for affine protocols with convex plan costs, we characterize truthfulness by the absence of any operational improvement

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.

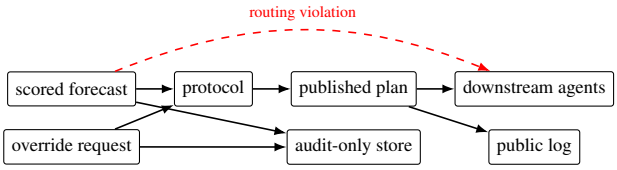


Figure 1: The two-channel protocol. Downstream agents read only the published plan; the dashed red edge is the routing violation ruled out by (P1), tested along feasible constant-plan directions in Section 6.

available through the report channel (Theorem 5.2). This extends beyond quadratic preferences and makes the sufficient plan-attainment condition necessary in the convex setting. For quadratic costs, a metric projection gives an exact vector solution and computable bounds on report distortion; the scalar case yields an explicit authority threshold (Theorems 4.1 and 5.3). Second, the principal's design problem trades information use against preference misalignment: authority can eliminate report bias while increasing operational loss. Third, we characterize the most informative feasible constant-plan perturbation, including binding requests, and derive the additional authority needed for a symmetric test with prescribed power (Theorem 6.2). Finally, we propagate statistical estimation error through the equilibrium map with explicit sensitivity constants and distinguish persistent report distortion from decision risk.

The results concern a fixed proper score without decision-contingent compensation. Such compensation is an alternative when the forecaster's utility is known and suitable transfers are feasible (Boutilier, 2012; Chen et al., 2014). Our main analysis is a committed protocol with a strategic forecaster; extensions allow plan-dependent downstream play, finite-horizon feedback, censored audits, and shared capacity. Simulations check the mathematical characterizations and illustrate their implications; they do not establish the information restrictions in a deployed institution.

2 Related Work

Self-interested experts and delegation. Strategic communication and delegation study how preference disagreement shapes information transmission and decision rights (Crawford & Sobel, 1982; Holmstrom, 1984). Boutilier (2012) characterizes compensation rules for experts with intrinsic decision utility and bounds the effects of utility misspecification. Decision scoring rules address action-dependent feedback and payments (Othman & Sandholm, 2010; Chen et al., 2014; Oesterheld & Conitzer, 2020). We retain a fixed audit score and design the reachable plans associated with each report. The new object is the joint geometry of report influence, bounded requests, and feasible routing tests; the underlying first-order and projection arguments are standard convex optimization tools.

Prediction that changes decisions or outcomes. Decision-focused learning optimizes forecasts for downstream loss (Donti et al., 2017; Elmachtoub & Grigas, 2022; Wilder et al., 2019; Mandi et al., 2024). Strategic classification and performative prediction study responses to a deployed predictor and the resulting distribution changes (Hardt et al., 2016; Perdomo et al., 2020; Mendler-Duenner et al., 2020; Hardt et al., 2022; Brown et al., 2022). Oesterheld et al. (2023) analyzes proper-score incentives when predictions change the predicted outcome, and Boeken et al. (2025) studies conditional forecasts and scoring constructions for performative predictions. Our main audit target is fixed conditional on the current information; strategic incentives arise from an additional operational preference. The plan may still change future information. Causal incentive analysis (Everitt et al., 2021) and information design (Kamenica & Gentzkow, 2011) provide related perspectives on information flow, but do not supply the bounded-request calculations here.

Forecast overrides and censored observations. Documented judgmental adjustments motivate separate forecast and request channels (Fildes et al., 2009; 2025); learning to defer instead selects which decision maker acts (Madras et al., 2018; Mozannar & Sontag, 2020). Inventory learning must account for censored observations (Ban & Rudin, 2019; Besbes & Muharremoglu, 2013; Huh & Rusmevichientong, 2009). The observed-data scoring framework of Jonkers et al. (2026) provides

proper scores for right-censored laws. We use that fixed-experiment construction and distinguish it from an experiment whose threshold the forecaster can change through the plan.

3 SETTING: REPORTS, REQUESTS, AND PLANS

3.1 INFORMATION AND OBJECTIVES

Fix public state x and public history $\mathcal{H}$. The forecaster F privately knows a statistical type $p \in \mathcal{P}$ and sends a report $q \in \mathcal{Q}$, where $\mathcal{P} \subseteq \mathcal{Q}$, and a request w in a specified feasible set. The type is a probability law in the general model and a conditional mean in the scalar model. The protocol publishes $y = \Pi(q, w; x)$ before downstream agents choose actions a. A separate audit produces observation o from an announced law $\mathsf{Q}_p(\cdot \mid x, \mathcal{H})$.

For an audit loss S, define the score regret

$$D_S(p, q) := \mathbb{E}_{o \sim \mathsf{Q}_p}[S(q, o) - S(p, o)]. \tag{1}$$

We assume these expectations are finite, $D_S(p, q) \geq 0$, and equality holds only at $q = p$ on the identifiable type class. This is strict propriety for probability reports, or strict consistency for the elicited statistic (Gneiting & Raftery, 2007; Savage, 1971). For scalar means, $S(q, o) = (q - o)^2/2$ gives $D_S(p, q) = (q - p)^2/2$ when $\mathbb{E}[o] = p$ and the second moment is finite.

Given a fixed downstream continuation rule, let $V_{\mathrm{F}}(y; p)$ be the forecaster's expected operational cost, including continuation costs. We suppress its type dependence when unambiguous. The forecaster minimizes

$$J_{\mathrm{F}}(q, w; p) = \alpha D_S(p, q) + V_{\mathrm{F}}(\Pi(q, w)), \qquad \alpha > 0. \tag{2}$$

The scalar specialization has preferred plan $p + \Delta_{\mathrm{F}}$ and curvature $c_{\mathrm{F}} > 0$:

$$V_{\mathrm{F}}(y) = \frac{c_{\mathrm{F}}}{2}(y - p - \Delta_{\mathrm{F}})^2. \tag{3}$$

The principal P commits to the protocol before observing p. When the plan is executed directly, its loss is

$$\mathcal{L}(q, y; p) = \tfrac{1}{2}(y - p)^2 + \tfrac{\lambda_q}{2}(q - p)^2, \qquad \lambda_q \geq 0. \tag{4}$$

The second term values the accuracy of the retained report. It is a principal loss, not a transfer to the forecaster and not permission for additional operational uses of the report. We call a minimizing report and request the equilibrium response to the committed protocol.

3.2 THE TWO-CHANNEL PROTOCOL

The scalar protocol is

$$y = \Pi_{\iota,\beta}(q, w) := \iota q + (1 - \iota)\bar{v} + \beta w, \quad |w| \leq U, \quad \iota \in [0, 1], \quad \beta > 0. \tag{5}$$

Here $\bar{v}$ is a public, report-independent benchmark, ι is report influence, β is request leverage, and $U \geq 0$ is the request budget. The authority radius in plan units is $B := \beta U$. The attainable plans form an interval of radius B centered at $\iota q + (1 - \iota)\bar{v}$. Report-anchored and benchmark-anchored bands correspond to $\iota = 1$ and $\iota = 0$, respectively. Even at $\iota = 0$, the request can convey private information through its choice within the fixed interval.

3.3 ROUTING AND VERIFICATION

(P1) Plan-only operational routing. Conditional on $(x, y, \mathcal{H})$, the downstream action law and the law of the next payoff-relevant state are independent of (q, w). Raw messages remain in an audit-only store. Downstream strategies and future data may depend on the plan. Future benchmarks and decision rules must satisfy the same restriction; a future use of the raw report is also an influence channel.

(P2) Plan-only operational costs. Apart from the announced score, operational charges depend on y, not separately on q or w. Message feasibility is specified by the protocol; physical constraints are evaluated on the plan or its downstream action. A sign-off cost proportional to $|w|$ violates this condition even when (P1) holds.

Assumption 3.1 (Independent identifying audit)**.** Conditional on the type and public information, the entire audit observation law is independent of (q, w, y), and its expected score has regret (1) with the stated identification property.

Independence of audit assignment alone is insufficient if the executed plan also changes the audited outcome. For censored demand, the assumption holds when a separate audit threshold is drawn independently of the messages and latent demand is unaffected by the operational order. The observed-data log score then has regret equal to an average of Kullback–Leibler divergences (Theorem B.2).

Lemma 3.1 (Score–plan separation)**.** *Under (P1)–(P2) and Assumption 3.1, additive score and operational costs give (2), up to a report-independent constant. This remains true at a finite-horizon continuation history when the continuation rule obeys (P1).*

The proof conditions expected operational cost on the plan and subtracts the expected truthful audit score. An explicitly score-dependent continuation is a separate extension: its expected cost must preserve a nonnegative score regret (Appendix D). A plan-centered audit threshold $e = y + \xi$, with independent noise ξ, also lies outside the assumption. Its truthful-score entropy must be included in the effective plan cost, and identification checked for every feasible plan (Appendix F).

4 Influence, Authority, and Principal Loss

4.1 An exact scalar threshold

Consider (3) and (5), with reports in $\mathbb{R}$. Define the displacement of the preferred plan from the truthful center, its uncovered part, and the required request budget by

$$G(\iota) = (1-\iota)(p-\bar{v}) + \Delta_{\mathrm{F}}, \quad D_B = \operatorname{sgn}(G)(|G| - B)_+, \quad U^\dagger = |G|/\beta, \tag{6}$$

where $(t)_+ = \max\{t, 0\}$ and $\operatorname{sgn}(0) = 0$. Write $\Lambda_q = c_{\mathrm{F}}\iota/(\alpha + c_{\mathrm{F}}\iota^2)$ and $\Lambda_y = \alpha/(\alpha + c_{\mathrm{F}}\iota^2)$.

Theorem 4.1 (Scalar equilibrium and truthfulness threshold)**.** *The unique equilibrium response is*

$$q^\star - p = \Lambda_q D_B, \qquad \beta w^\star = \operatorname{sgn}(G)\min\{|G|, B\}, \qquad y^\star - p = \Delta_{\mathrm{F}} - \Lambda_y D_B. \tag{7}$$

For $\iota > 0$, the report is truthful if and only if $B \geq |G|$, equivalently $U \geq U^\dagger$. In that case $y^\star = p + \Delta_{\mathrm{F}}$. For $\iota = 0$, the report is truthful for every budget. The same formulas hold on a constrained report domain whenever the unrestricted optimum is feasible.

Proof. Set $z = q - p$. The forecaster minimizes $\alpha z^2/2 + c_{\mathrm{F}}(\iota z + \beta w - G)^2/2$. For fixed w, the unique minimizing report has $z = \Lambda_q(G - \beta w)$. Substitution leaves a positive constant times $(G - \beta w)^2$, whose minimizer is the projection of G/β onto $[-U, U]$. Substituting this request proves the result. □

Corollary 4.2 (Anchoring determines report incentives)**.** *A radius-U band around the report has $(\iota, \beta) = (1, 1)$ and is truthful exactly when $U \geq |\Delta_{\mathrm{F}}|$. The same band around a report-independent benchmark has $(\iota, \beta) = (0, 1)$ and is truthful for every U.*

The operational shortfall and report influence jointly determine the bias. At a constrained report boundary, a profitable direction may be infeasible, so uncovered displacement alone need not imply misreporting. For fixed β, $U^\dagger \to |\Delta_{\mathrm{F}}|/\beta$ as $\iota \to 1$. A divergent request budget along $\beta = 1 - \iota$ reflects vanishing request leverage; the required authority in plan units remains $|G|$.

4.2 Choosing authority before observing the type

Truthful reporting need not minimize principal loss. On the covered branch, the plan implements the forecaster's offset and the principal pays $\Delta_{\mathrm{F}}^2/2$. For example, if $p \sim \mathrm{Unif}[-.1, .1]$, $\bar{v} = 0$, $\Delta_{\mathrm{F}} = 1$,

and $\iota = 0$, zero authority yields loss $1/600$, whereas full coverage yields $1/2$. Both protocols are truthful.

For a given type distribution and nonnegative authority cost $C_{\text{authority}}$, the full design objective is

$$\inf_{\iota\in[0,1],\,\bar{v}\in\mathbb{R},\,B\geq 0} \frac{1}{2}\mathbb{E}\left[(\Delta_{\mathrm{F}} - \Lambda_y D_B)^2 + \lambda_q(\Lambda_q D_B)^2\right] + C_{\text{authority}}(B) + C_{\text{audit}}, \tag{8}$$

where C_{audit} is the cost of the chosen audit design. Constraints on benchmarks or authority can be added to the feasible design set. This objective allows both covered and uncovered types. An optimum exists on a compact design set with continuous costs; unbounded benchmarks can yield only an infimum (Proposition D.18).

A tractable special case explains the role of private information. Let $m = \mathbb{E}[p]$, let $p_{\min}$ be the minimum of the compact type support, and assume $\Delta_{\mathrm{F}} > 0$ is constant.

Assumption 4.1 (Common binding branch)**.** Every type has $D_B = \Delta_{\mathrm{F}} + (1-\iota)(p-\bar{v}) - B > 0$.

Define $\omega_{\text{bias}} = \lambda_q\Lambda_q^2/(\Lambda_y^2 + \lambda_q\Lambda_q^2)$ and $\omega_{\text{track}} = (1-\iota)^2(\Lambda_y^2 + \lambda_q\Lambda_q^2)$.

Proposition 4.3 (Optimal centering on a binding branch)**.** *Fix $\iota < 1$, and put $\mu_D^\star = \Lambda_y\Delta_{\mathrm{F}}/(\Lambda_y^2 + \lambda_q\Lambda_q^2)$. If $\mu_D^\star > (1-\iota)(m - p_{\min})$, then*

$$\min_{\bar{v},B:\,\text{common binding}} \mathbb{E}\mathcal{L} = \tfrac{1}{2}\left\{\Delta_{\mathrm{F}}^2\omega_{\text{bias}} + \mathrm{Var}(p)\omega_{\text{track}}\right\}. \tag{9}$$

If the feasibility inequality holds near $\iota = 0$, positive type variance implies that some positive influence strictly improves this centered loss relative to zero influence.

Completing the square in $\mathbb{E}[D_B]$ gives the formula (Appendix D). At zero influence, $\omega'_{\text{bias}} = 0$ and $\omega'_{\text{track}} = -2$: the initial information benefit is first order and the report-bias cost is second order. This branch calculation does not solve the full design problem (8).

5 A Geometric Characterization of Truthfulness

For a report q, let $\mathcal{P}(q)$ be its reachable-plan set and define $\underline{V}_{\mathrm{F}}(q) = \min_{y\in\mathcal{P}(q)} V_{\mathrm{F}}(y)$, assuming the minimum is attained. Consider the condition

$$p \in \arg\min_{q\in\mathcal{Q}} \underline{V}_{\mathrm{F}}(q). \tag{T}$$

It says that truthful reporting can attain the lowest operational cost available through any report. It need not attain every reachable plan.

Theorem 5.1 (Truthfulness under plan-only routing)**.** *Under the score–plan separation of Lemma 3.1, Condition (T) implies that every optimal report is truthful. The conclusion holds at every history of a finite-horizon problem if (T) holds there for the Bellman operational value, including off-path histories.*

Proof. The minimized objective is $\alpha D_S(p,q) + \underline{V}_{\mathrm{F}}(q)$. At $q = p$ both terms attain their minimum; strict identification makes every other report more costly. Backward induction applies this comparison to the continuation values. □

A common minimizer of the operational value reachable under every report is sufficient. A guarantee uniform over all quadratic plan preferences with arbitrary targets and unbounded curvature requires identical reachable-plan sets, provided all reports are possible types and pairwise score regrets are finite (Theorem D.8). This robust requirement is stronger than type-specific truthfulness.

5.1 Convex costs: a necessary and sufficient certificate

Let reports and types lie in $\mathbb{R}^d$, plans in $\mathbb{R}^m$, and requests in a nonempty compact convex set $\mathcal{W} \subseteq \mathbb{R}^r$. Fix matrices $C \in \mathbb{R}^{m\times d}$ and $E \in \mathbb{R}^{m\times r}$, and a vector $\mathsf{c} \in \mathbb{R}^m$, with protocol

$$y = Cq + Ew + \mathsf{c}, \qquad w \in \mathcal{W}. \tag{10}$$

For $w \in \mathcal{W}$, define the outward normal cone $N_{\mathcal{W}}(w) = \{u : u^\top(v-w) \leq 0 \text{ for all } v \in \mathcal{W}\}$.

Theorem 5.2 (Convex truthfulness certificate)**.** *Suppose $D_S(p,\cdot)$ is differentiable and strictly convex, with $\nabla_q D_S(p,p) = 0$, and V_{F} is finite, differentiable, and convex. For any finite $\alpha > 0$, the following are equivalent: (i) a truthful report is globally optimal; (ii) Condition (T) holds; (iii) there is $w_0 \in \mathcal{W}$, with $y_0 = Cp + Ew_0 + \mathsf{c}$, such that*

$$C^\top \nabla V_{\mathrm{F}}(y_0) = 0, \qquad -E^\top \nabla V_{\mathrm{F}}(y_0) \in N_{\mathcal{W}}(w_0). \tag{11}$$

Whenever these conditions hold, every optimal report equals p.

Proof. The joint objective is differentiable and convex on $\mathbb{R}^d \times \mathcal{W}$. Its necessary and sufficient first-order conditions at (p, w_0) are exactly (11). These are also the first-order conditions for minimizing the operational cost alone over (q, w), proving equivalence to (T). Strict convexity of the score excludes two distinct optimal reports. □

Thus changing the finite weight of a smooth score can change the amount of bias without changing whether truth is optimal. For a closed convex report domain at a boundary type, the first equality is replaced by $-C^\top \nabla V_{\mathrm{F}}(y_0) \in N_{\mathcal{Q}}(p)$. The result applies to convex nonquadratic costs and shared convex request constraints; nonconvex downstream responses retain the sufficient condition (T) but can invalidate its necessity.

5.2 Quadratic costs: projection, distortion, and stability

Take score cost $\frac{1}{2}(q-p)^\top H_q(q-p)$ and plan cost $\frac{1}{2}(y-y_{\mathrm{F}})^\top H_y(y-y_{\mathrm{F}})$, with symmetric positive definite matrices $H_q \in \mathbb{R}^{d\times d}$ and $H_y \in \mathbb{R}^{m\times m}$. Here y_{F} is the preferred plan. Define $\mathsf{g} = y_{\mathrm{F}} - Cp - \mathsf{c}$ and $K = H_q + C^\top H_y C$.

Theorem 5.3 (Vector equilibrium and quantitative certificate)**.** *The optimal report and plan are unique, although requests need not be. Every optimal request and its report satisfy*

$$w^\star \in \arg\min_{w\in\mathcal{W}} \tfrac{1}{2}(\mathsf{g} - Ew)^\top M(\mathsf{g} - Ew), \qquad q^\star - p = K^{-1}C^\top H_y(\mathsf{g} - Ew^\star), \tag{12}$$

where $M = (H_y^{-1} + CH_q^{-1}C^\top)^{-1} \succ 0$. Equivalently, project g onto $E\mathcal{W}$ in the H_y metric, let t_0 be that projection, and define $h = C^\top H_y(\mathsf{g} - t_0)$ and $\kappa = \|H_y^{1/2} C H_q^{-1/2}\|_2^2$, where $\|\cdot\|_2$ is the spectral matrix norm. Writing $\|u\|_H = (u^\top H u)^{1/2}$, we have

$$\frac{\|h\|_{H_q^{-1}}}{1+\kappa} \le \|q^\star - p\|_{H_q} \le \|h\|_{H_q^{-1}}. \tag{13}$$

In particular, truthfulness is equivalent to either certificate

$$h = 0 \quad \Longleftrightarrow \quad C^\top H_y(\mathsf{g} - Ew^\star) = 0. \tag{14}$$

If C has full row rank, this is equivalent to exact coverage $\mathsf{g} \in E\mathcal{W}$.

Eliminating the report gives the positive definite metric M. The bounds follow from strong convexity of the reduced objective and nonexpansiveness of projection (Appendix E). The truthful projection uses H_y, whereas the equilibrium request uses M; conflating these two projections generally gives a wrong solution. A rank-deficient report channel can be truthful with an unattained preferred plan, provided the remaining operational gradient is orthogonal to that channel. Appendix I gives minimal request-set completion, and Appendix D treats soft request costs and strategic overlay owners.

6 Testing Plan-Only Routing

Plan-only routing restricts the full conditional law of downstream actions. A feasible randomized intervention can falsify this restriction. We first consider a scalar downstream action and the affine protocol (5); vector actions can be tested componentwise with multiplicity control.

Definition 6.1 (Linear message effects)**.** Within a predeclared experimental stratum, suppose the conditional action mean at displayed messages (q', w') is

$$\mathbb{E}[a \mid q', w', x, y, \mathcal{H}] = \mathcal{A}(x, y, \mathcal{H}) + b_q(q' - \bar{q}) + b_w(w' - \bar{w}), \tag{15}$$

where $\bar{q}, \bar{w}$ are fixed reference messages and b_q, b_w are constant coefficients over the perturbation range. In the report-only special case, write $b_q = \lambda_{\mathrm{leak}}\Xi$ and $b_w = 0$, with $\lambda_{\mathrm{leak}} \ge 0$ and a fixed nonzero action scale Ξ.

A nonzero report or request coefficient need not violate routing if the two coefficients combine to depend only on the plan. The relevant restriction is their response along a constant-plan direction.

6.1 Identification by a compensated perturbation

After receiving (q, w), draw a randomized perturbation η, independent of potential action noise conditional on the unperturbed record, and display

$$(q', w') = (q + \eta,\ w - \iota\eta/\beta). \tag{16}$$

The original q remains the scored report. The protocol must apply the intervention to every operational copy of the messages and prevent a subsequent forecaster response. The published plan is unchanged exactly.

Theorem 6.1 (Constant-plan identification)**.** *Under (15), a feasible intervention with positive conditional variance identifies*

$$b_\eta = b_q - \frac{\iota}{\beta} b_w \tag{17}$$

by regression on η with a stratum intercept. For n balanced observations with known independent Gaussian residual variance σ_a^2 and centered design variance s_η^2, the slope variance is $\sigma_a^2/(ns_\eta^2)$. Let z_u denote the standard normal u-quantile. A two-sided level-a_{sig} test has power at least $1 - b_{\mathrm{pow}}$ against $|b_\eta| \geq \lambda_0 > 0$ if

$$n \geq \frac{(z_{1-a_{\mathrm{sig}}/2} + z_{1-b_{\mathrm{pow}}})^2 \sigma_a^2}{\lambda_0^2 s_\eta^2}. \tag{18}$$

Substitution proves plan invariance and the conditional linear response. The Gaussian slope distribution gives the power bound. Unknown residual variance requires Student inference or a large-sample approximation; heterogeneous slopes require conditional or stratum-specific analysis. A zero slope does not establish routing: nonlinear mean responses, variance changes, and delayed effects can remain undetected.

6.2 Optimal feasible probes and the authority margin

For a symmetric amplitude $\bar{\eta} > 0$, eligibility is evaluated before randomization:

$$|w| \leq U - \iota\bar{\eta}/\beta. \tag{19}$$

Binding requests fail this criterion when $\iota > 0$. They need not prevent an informative asymmetric intervention.

Theorem 6.2 (Information available within a plan fiber)**.** *Fix $\iota > 0$, an unperturbed feasible request w, unrestricted displayed reports, and an amplitude limit $h > 0$. The feasible perturbation interval is*

$$\ell = \max\{-h, \beta(w - U)/\iota\}, \qquad u = \min\{h, \beta(w + U)/\iota\}. \tag{20}$$

Every randomization on this interval has variance at most $(u - \ell)^2/4$. An equal allocation to its endpoints attains this bound and maximizes Gaussian slope information with an unknown intercept. With n even, the minimum variance of an unbiased slope estimator in this model is $4\sigma_a^2/[n(u - \ell)^2]$ when $u > \ell$.

At a covered scalar equilibrium, the largest symmetric amplitude is $(B - |G|)/\iota$. Thus an uncapped symmetric endpoint design satisfies (18) whenever

$$B \geq |G| + \frac{\iota\sigma_a(z_{1-a_{\mathrm{sig}}/2} + z_{1-b_{\mathrm{pow}}})}{\lambda_0\sqrt{n}}. \tag{21}$$

At $w = \pm U$ and $0 < h \leq 2B/\iota$, the maximum variance is $h^2/4$, whereas at an interior request permitting $[-h, h]$ it is h^2.

The interval follows directly from request feasibility. The variance bound and optimality follow from bounded support and the Gaussian information calculation; Appendix F gives the proof. At a binding request, an equally allocated zero/inward test therefore needs four times as many observations as a fully feasible symmetric test of the same maximum amplitude. This is a test-design cost, not a universal failure of auditability. At $B = 0$ and $\iota > 0$ the plan fiber is a singleton and no such intervention is possible. A shadow report shown only to the downstream interface tests a separate display channel. When $\iota = 0$, report perturbations require no request compensation.

7 Learning and Persistent Incentive Distortion

Let X have covariate law P_X, and let $f^\star(x)$ be the correctly specified statistical target under the independent audit. Write $\|f\| = (\mathbb{E}[f(X)^2])^{1/2}$. An estimator $\widehat{f}_n$ is trained on n audit observations. For a fixed protocol, define $\Gamma(f)(x)$ and $Y(f)(x)$ as the equilibrium report and plan when the forecaster's belief is $f(x)$. The preference offset and protocol parameters are held fixed when f changes. The incentive distortion is $\Delta_{\mathrm{inc}} := \|\Gamma(f^\star) - f^\star\|$.

Theorem 7.1 (Propagation of estimation error)**.** *Suppose $\|\widehat{f}_n - f^\star\| \le \varepsilon_n$ with probability at least $1-\delta$, and Γ is L_Γ-Lipschitz on a set containing $f^\star$ and every estimator realization in this event. Then, on the event,*

$$(\Delta_{\mathrm{inc}} - L_\Gamma \varepsilon_n)_+^2 \le \|\Gamma(\widehat{f}_n) - f^\star\|^2 \le (\Delta_{\mathrm{inc}} + L_\Gamma \varepsilon_n)^2. \tag{22}$$

If $\widehat{f}_n \to f^\star$ in probability and Γ is continuous there, report error converges in probability to Δ_{inc}^2. In the scalar model with fixed parameters,

$$L_\Gamma = 1 + \Lambda_q(1-\iota), \qquad L_Y = 1 \tag{23}$$

are global Lipschitz constants for Γ and Y, including threshold crossings. If $\iota = 1$, then $L_\Gamma = 1$ and $\Delta_{\mathrm{inc}} = \Lambda_q(|\Delta_{\mathrm{F}}| - B)_+$.

Proof. The triangle and reverse triangle inequalities give (22). In the scalar model, $\Gamma(f) = f + \Lambda_q D_B(f)$ and $Y(f) = f + \Delta_{\mathrm{F}} - \Lambda_y D_B(f)$, where $D_B(f) = \mathrm{sgn}(g_f)(|g_f| - B)_+$ and $g_f = (1-\iota)(f - \bar{v}) + \Delta_{\mathrm{F}}$. The scalar soft-threshold map is nondecreasing and 1-Lipschitz. The report slope is either 1 or $1 + \Lambda_q(1-\iota)$, and the plan slope is either 1 or $1 - \Lambda_y(1-\iota) \in [0,1]$. These pointwise bounds imply the stated $L^2(P_X)$ bounds. □

Define the principal's equilibrium risk, evaluated at the true type, by $\mathcal{R}_{\mathrm{str}}(f) = \mathbb{E}\mathcal{L}(\Gamma(f), Y(f); f^\star)$. Let $\mathcal{R}_{\mathrm{truth}}(f^\star)$ be its risk when $q = f^\star$ and the forecaster chooses its best request conditional on that report. Let $\mathcal{R}_{\mathrm{FB}}$ be the risk when the principal observes $f^\star$ and chooses the plan directly. Then

$$\begin{aligned}\mathcal{R}_{\mathrm{str}}(\widehat{f}_n) - \mathcal{R}_{\mathrm{FB}} = \underbrace{\mathcal{R}_{\mathrm{str}}(\widehat{f}_n) - \mathcal{R}_{\mathrm{str}}(f^\star)}_{\text{estimation effect}} + \underbrace{\mathcal{R}_{\mathrm{str}}(f^\star) - \mathcal{R}_{\mathrm{truth}}(f^\star)}_{\text{strategic effect}} \\ + \underbrace{\mathcal{R}_{\mathrm{truth}}(f^\star) - \mathcal{R}_{\mathrm{FB}}}_{\text{truthful operational gap}}.\end{aligned} \tag{24}$$

This is an accounting identity; neither of the first two differences need be nonnegative. Its statistical content comes from controlling the first term. For the scalar loss (4), put $A_Y = \|Y(f^\star) - f^\star\|$. On the event in Theorem 7.1, Cauchy–Schwarz gives

$$|\mathcal{R}_{\mathrm{str}}(\widehat{f}_n) - \mathcal{R}_{\mathrm{str}}(f^\star)| \le (A_Y L_Y + \lambda_q \Delta_{\mathrm{inc}} L_\Gamma)\varepsilon_n + \tfrac{1}{2}(L_Y^2 + \lambda_q L_\Gamma^2)\varepsilon_n^2. \tag{25}$$

A quadratic rate requires cancellation of the first-order term; smoothness alone does not imply it. Independent-audit empirical-risk bounds, censoring correction, and optimization error are treated in Appendix G. Increasing the sample size can remove estimation error while leaving both incentive distortion and truthful operational loss. Changing the protocol can affect all three terms.

8 Numerical Studies

Six suites use master seed 20260922 and are reproduced by `experiments/run_all.py`. E0 independently solves the objectives used to derive the closed forms; E1–E4 use those verified solutions to study design, diagnostics, learning, and downstream capacity. E5 checks the additional convex and information results. Full specifications, comparator definitions, and figures are in Appendix J.

Equilibrium geometry (E0, E5). Independent optimization on 2,000 scalar and 200 vector instances matches the formulas to maximum report discrepancies 7.7×10^{-6} and 4.4×10^{-5}. E5 tests the vector bounds and stability on 160 further instances, including redundant requests, and the convex certificate on 180 nonquadratic instances with a shared request constraint. All classifications agree within numerical tolerance. The largest vector bound violation is 1.4×10^{-8} (Figure 2).

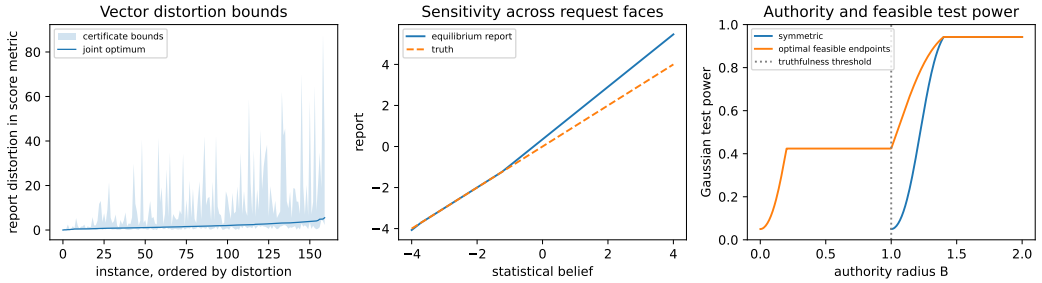


Figure 2: Geometry and test design. Left: independently optimized report errors lie within the projection-certificate bounds. Center: the scalar equilibrium map remains Lipschitz across request faces. Right: binding requests still permit informative asymmetric probes; symmetric probes need an authority margin beyond truthfulness. Parameters are in Appendix J.

Principal design and routing tests (E1, E2). E1 searches a finite (ι, B) grid with fixed benchmark over 52 parameter cells. The minimizing grid point has $B = 0$ in 42 of 48 homogeneous cells. In a representative cell, widening a report-anchored band removes report bias while increasing action loss from .5 to 2. E2 uses 1,000 randomizations per cell. Paired perturbations move the plan by at most 1.4×10^{-14} and recover the predicted negative slope of a direct request effect. Symmetric eligibility is zero at binding budgets, whereas inward interventions remain feasible. Nonlinear, variance-only, and delayed alternatives demonstrate limitations of a contemporaneous linear test. E5 additionally checks the predicted fourfold boundary variance cost in 20,000 Gaussian repetitions.

Learning and capacity (E3, E4). E3 learns demand mean $p(x) = 100 + 10 \sin x_1 + 3x_2$ from independent audit labels. Across 30 ridge-estimation seeds, estimation error decreases from 2.08 to .026 as sample size grows from 256 to 16,384, while report error under a binding report-anchored protocol approaches the predicted floor $16/9$ (observed 1.74). A benchmark-anchored protocol removes the report floor but retains operational loss 24.2; full coverage would give 32. A separate network experiment learns report and request heads, and a censoring experiment shows how observed-sales scoring can favor a downward-biased forecast. E4 applies quadratic plan preferences to twenty items, followed by physical capacity projection. Truthful reports still yield downstream regret (1.6% with slack capacity; 8.9% with binding capacity). A centered filter is compared with an uncentered filter to separate report direction from benchmark displacement. These are equilibria of preferences over the published plan, not of a forecaster optimizing the projected newsvendor action.

Limitations. The guarantees depend on enforceable routing and an identifying audit whose observation law is independent of the current messages. The exact general certificate requires convex operational value; downstream strategic responses can violate that condition. The scalar design calculation is a restricted mechanism class, and the numerical design search is finite. The experiments are synthetic and do not estimate institutional preferences or establish routing in deployment. The leakage test identifies specified message effects along feasible constant-plan directions; a null result cannot certify full conditional independence.

9 Conclusion

We have studied truthful forecasting as a problem of allocating decision authority. Under independent identifying audits and plan-only routing, separating a scored report from a bounded operational request yields exact truthfulness conditions for convex plan costs and explicit distortion formulas in quadratic models. The resulting design tradeoff is substantive: authority can remove incentives to misreport while increasing the principal's operational loss. Feasible plan-preserving perturbations characterize the information available for detecting direct message effects, and the learning bounds show why more audit data need not remove incentive distortion. These results connect truthful elicitation, operational performance, and statistical verification within a common protocol. A natural next step is to design information access, authority, and audits jointly when preferences and downstream responses must also be learned.

Reproducibility Statement

writes every reported figure to `figures-EH/` and every reported Appendix J records the complete data-generating processes, seeds, parameter grids, and the mapping from files to tables and figures.

Use of Large Language Models

Large language model assistance was used in mathematical analysis, code inspection, numerical verification, and exposition. The authors are responsible for the correctness and originality of the final submission.

References

Gah-Yi Ban and Cynthia Rudin. The big data newsvendor: Practical insights from machine learning. *Operations Research*, 2019.

Omar Besbes and Alp Muharremoglu. On implications of demand censoring in the newsvendor problem. *Management Science*, 2013.

Philip Boeken, Onno Zoeter, and Joris M. Mooij. Conditional forecasts and proper scoring rules for reliable and accurate performative predictions. In *Advances in Neural Information Processing Systems*, 2025.

Craig Boutilier. Eliciting forecasts from self-interested experts: Scoring rules for decision makers. In *International Conference on Autonomous Agents and Multiagent Systems*, 2012.

Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. In *International Conference on Artificial Intelligence and Statistics*, 2022.

Yiling Chen, Ian A. Kash, Michael Ruberry, and Victor Shnayder. Eliciting predictions and recommendations for decision making. *ACM Transactions on Economics and Computation*, 2014.

Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6): 1431–1451, 1982.

Priya Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, 2017.

Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 2022.

Tom Everitt et al. Agent incentives: A causal perspective. In *AAAI Conference on Artificial Intelligence*, 2021.

Robert Fildes, Paul Goodwin, Michael Lawrence, and Konstantinos Nikolopoulos. Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, 25(1):3–23, 2009. doi: 10.1016/j.ijforecast.2008.11.010.

Robert Fildes, Paul Goodwin, and Shari De Baets. Forecast value added in demand planning. *International Journal of Forecasting*, 41(2):649–669, 2025. doi: 10.1016/j.ijforecast.2024.07.006.

Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *ACM Conference on Innovations in Theoretical Computer Science*, 2016.

Moritz Hardt, Meena Jagadeesan, and Celestine Mendler-Duenner. Performative prediction in a dynamic environment. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, 2022.

Bengt Holmstrom. On the theory of delegation. *Bayesian Models in Economic Theory*, pp. 115–141, 1984.

Woonghee Tim Huh and Paat Rusmevichientong. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research*, 2009.

Jef Jonkers, Glenn Van Wallendael, Luc Duchateau, and Sofie Van Hoecke. Proper scoring rules for right-censored survival data. *arXiv preprint*, 2026. arXiv:2606.06393.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101 (6):2590–2615, 2011.

David Madras, Toniann Pitassi, and Richard Zemel. Predict responsibly: Improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems*, 2018.

Jayanta Mandi et al. Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*, 2024.

Celestine Mendler-Duenner, Juan C. Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. In *Advances in Neural Information Processing Systems*, 2020.

Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. In *International Conference on Machine Learning*, 2020.

Caspar Oesterheld and Vincent Conitzer. Decision scoring rules. In *International Workshop on Internet and Network Economics*, 2020.

Caspar Oesterheld, Johannes Treutlein, Emery Cooper, and Rubi Hudson. Incentivizing honest performative predictions with proper scoring rules. In *Conference on Uncertainty in Artificial Intelligence*, 2023.

Abraham Othman and Tuomas Sandholm. Decision rules and decision markets. In *International Conference on Autonomous Agents and Multiagent Systems*, 2010.

Juan C. Perdomo, Tijana Zrnic, Celestine Mendler-Duenner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, 2020.

Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.

Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *AAAI Conference on Artificial Intelligence*, 2019.

A NOTATION AND GENERAL-GAME FOUNDATIONS

Table 1 indexes selected notation used throughout the paper. Symbols local to a proof are defined at first use. Table 2 gives the mapping to deployed workflows; in every row, raw messages remain audit-only while the published plan is the sole operational input.

Table 1: Main notation. Symbols used only in extensions are defined locally.

symbol	meaning	defined	symbol	meaning	defined
p	private belief	§3.1	q	scored report	§3.1
w, U	request and budget	§3.1	y	published plan	§3.1
α, D_S	score weight and regret	§3.1	$J_{\mathrm{F}}, V_{\mathrm{F}}$	forecaster and plan costs	§3.1
$c_{\mathrm{F}}, \Delta_{\mathrm{F}}$	forecaster curvature and offset	§3.1	$\lambda_q, \mathcal{L}$	forecast weight and principal loss	§3.1
$\Pi_{\iota,\beta}$	two-channel protocol	§3.2	ι, β	influence and leverage	§3.2
$\bar{v}$	public benchmark	§3.2	$G(\iota), U^{\dagger}$	displacement and threshold	§4.1
D_B	uncovered displacement	§4.1	$\Xi, \lambda_{\mathrm{leak}}$	action scale and leakage	§6
n, η	count and perturbation	§6	Δ_{inc}	incentive bias	§7

Table 2: The same information-flow abstraction in deployed workflows.

workflow	scored report q	request w	published plan y
demand planning	statistical baseline	reason-coded overlay	approved consensus forecast
model governance	model output	documented override	approved decision
tool-using agent	model prediction	structured action proposal	validated tool call

Table 3: Positioning of plan-routed feedback.

framework	objectives	downstream response	future data	report channel
predict-then-optimize	usually integrated	passive optimizer	usually exogenous	direct forecast
decision-focused learning	integrated downstream loss	differentiable layer	usually fixed in training	direct forecast
performative prediction	usually one learner loss	implicit response	prediction changes law	direct prediction
self-interested expert	expert and principal differ	decision-dependent score	selected outcomes	direct message
this work	forecaster, operator, principal	strategic, plan-measurable	plan may change law	published plan only

A.1 FINITE-HORIZON CONTINUATION VALUES

Let periods be $k = 0, \ldots, T$. Use an augmented information state s containing the public history and the forecaster's current statistical type. Fix a downstream response rule. It induces a controlled transition kernel for the next information state depending on the messages only through the current plan. A Markov policy in this augmented state can therefore be history-dependent in the original variables.

For a public state-history pair $(x, \mathcal{H})$, let

$$\mathcal{P}_k(x, \mathcal{H}; q) = \{\Pi_k(q, w; x) : w \in \mathcal{W}_k(x, \mathcal{H}, q)\}$$

be the reachable-plan set. Let $\Upsilon_{k,p}(x, \mathcal{H}, y)$ be the forecaster's current operational cost plus optimal expected continuation cost. A sufficient common-minimizer condition is

$$\arg \min_{y \in \bigcup_q \mathcal{P}_k(x,\mathcal{H};q)} \Upsilon_{k,p}(x, \mathcal{H}, y) \cap \bigcap_q \mathcal{P}_k(x, \mathcal{H}; q) \neq \varnothing. \tag{26}$$

Only one common minimizer is required. Identical reachable sets make this condition automatic for every attained continuation value; overlap alone, without an appropriate minimizer in the intersection, does not.

Theorem A.1 (Existence and truthful optimal policies)**.** *Consider the induced finite-horizon control problem after a downstream response rule has been fixed. Suppose its augmented state and feasible message spaces are compact metric spaces, feasible-message correspondences are nonempty, compact valued, and continuous, expected stage costs are continuous, and controlled transition kernels are weakly continuous. Then a measurable deterministic optimal policy exists. If score–plan separation, strict identification, and Condition (T) for the Bellman operational value hold at every state and period, every optimal policy reports truthfully there. Condition (26) is sufficient for (T).*

Proof. The terminal value is continuous. Inductively, integrating a continuous bounded next-period value against the weakly continuous controlled kernel produces a continuous objective. Berge's maximum theorem gives a continuous minimum value and a nonempty compact argmin, and the measurable maximum theorem gives a measurable selector. This constructs a deterministic optimal policy by backward induction. At each state the Bellman objective is a report score regret plus the operational value. Under (T), a truthful report attains its lower bound; any false report adds a strictly positive regret. Apply this comparison backward at every state. □

A continuous selected downstream equilibrium is one way to obtain these reduced continuity assumptions. Existence of downstream best responses alone does not supply such a selection or establish a simultaneous-game equilibrium. This theorem is conditional on the specified downstream rule.

B Verification Foundations

Proper scoring is a statement about a fixed data-generating experiment. In a decentralized predict-then-optimize system, the report can select the experiment through the operator, so verification must be stated separately.

B.1 Fixed-action observed-data scoring

Let ϑ be a latent state and a an audit action. An action a induces observation kernel $K_a(do \mid \vartheta)$. For belief p, write $(K_a p)(do) = \int K_a(do \mid \vartheta)p(d\vartheta)$ and assume all induced laws have densities with respect to a common measure μ_a. Define the observed-data logarithmic loss

$$S_a(q, o) = -\log \frac{d(K_a q)}{d\mu_a}(o). \tag{27}$$

Proposition B.1 (Proper scoring under a fixed experiment)**.** *For a fixed action a,*

$$\mathbb{E}_{o\sim K_a p} S_a(q, o) - \mathbb{E}_{o\sim K_a p} S_a(p, o) = \mathrm{KL}(K_a p \| K_a q) \geq 0. \tag{28}$$

The score is strictly proper for beliefs in $\mathcal{P}$ if and only if $p \mapsto K_a p$ is injective on $\mathcal{P}$.

Interpretation: at a fixed experiment, expected log-score regret is exactly a Kullback–Leibler divergence, and strictness requires observational identification.

Proof. The expected score is the cross entropy of $K_a p$ relative to $K_a q$. Subtracting its value at $q = p$ gives the displayed Kullback–Leibler divergence. Equality holds exactly when the induced laws agree. □

B.2 Why operator-selected scores can fail

If the report induces action $A(q)$, the expected observed-data log loss is

$$\mathcal{E}_p(q) = H(K_{A(q)}p) + \mathrm{KL}(K_{A(q)}p \| K_{A(q)}q), \tag{29}$$

where H is entropy relative to the dominating measure. At truth this is $H(K_{A(p)}p)$, but the entropy term changes with q. Therefore a report can induce an easier-to-predict or less informative verification experiment and beat the truthful report even though the conditional score is proper. Plan-only routing removes the report's experiment selection, but the request retains it: under plan-centered operational scoring the experiment-selection incentive relocates to the published plan, and Condition (T) must then be checked on the entropy-augmented plan value (Appendix F).

B.3 Report-independent audit mechanism

Let an audit action $\mathsf{a}^{\mathrm{aud}}$ be sampled from $\mathsf{Q}(\cdot \mid x, \mathcal{H})$ independently of the current report, request, and published plan, and suppose the audit observation conditional on a has law $K_a p$, unaffected by the operational plan. Score the report only on this observation.

Theorem B.2 (Audit-based strict propriety)**.** *The expected audit-score regret is*

$$\begin{aligned}&\mathbb{E}_{\mathsf{a}^{\mathrm{aud}}\sim\mathsf{Q}}\mathbb{E}_{o\sim K_{\mathsf{a}^{\mathrm{aud}}}p}[S_{\mathsf{a}^{\mathrm{aud}}}(q,o)-S_{\mathsf{a}^{\mathrm{aud}}}(p,o)]\\&\quad=\mathbb{E}_{\mathsf{a}^{\mathrm{aud}}\sim\mathsf{Q}}\,\mathrm{KL}(K_{\mathsf{a}^{\mathrm{aud}}}p\|K_{\mathsf{a}^{\mathrm{aud}}}q)\geq 0.\end{aligned}\tag{30}$$

Truth is unique if and only if

$$p\neq q\Longrightarrow \mathsf{Q}(\{a:K_ap\neq K_aq\}\mid x,\mathcal{H})>0.\tag{31}$$

Interpretation: a report-independent mixture of audits is strictly proper exactly when it distinguishes every pair of beliefs with positive probability.

Proof. Condition on $\mathsf{a}^{\mathrm{aud}}=a$ and apply Proposition B.1; integrate the nonnegative conditional regret over Q. Its expectation is zero exactly when the induced laws agree for Q-almost every audit action. □

The architecture should distinguish a forecast channel, evaluated on report-independent audit data; an intervention channel, through which an agent requests exploration; and the operator's physical action. The forecast should not choose its own verification experiment.

Assumption 3.1 is itself falsifiable: under the plan-holding perturbation of Section 6, regress the realized audit action on the assigned report perturbation while controlling for the plan. A nonzero coefficient rejects report-independent audit assignment. The converse does not hold: a plan-centered rule $e=y+\xi$ passes every plan-holding intervention although its law depends on the plan, so a zero coefficient cannot certify independence (Appendix F).

The published plan can still select future information. Appendix C states the corresponding Blackwell ordering and proves that, absent direct action preferences, a forecaster weakly favors plans that induce more informative experiments. With direct action preferences the information and action gradients can reinforce or oppose one another, so Blackwell dominance alone does not determine the sign of report distortion.

C Verification and Blackwell Proofs

C.1 Fixed-experiment scoring

Using the observation kernel K_a and dominating measure μ_a defined in Appendix B, for fixed a, the expected observed-data score is

$$-\int\log\frac{d(K_aq)}{d\mu_a}\,d(K_ap).$$

Subtracting its value at $q=p$ gives $\mathrm{KL}(K_ap\|K_aq)$, proving Proposition B.1. Equality is exactly observational equivalence, so injectivity of $p\mapsto K_ap$ is necessary and sufficient for strict propriety.

C.2 Audit theorem

Conditioning on an audit action $\mathsf{a}^{\mathrm{aud}}=a$ and applying the fixed-experiment identity gives the conditional regret $\mathrm{KL}(K_ap\|K_aq)$. Integrating over a report-independent audit distribution proves Theorem B.2; the expected nonnegative quantity is zero precisely when induced laws agree for almost every audit action.

C.3 Blackwell ordering and strategic exploration

Let ϑ be the latent state, π its current posterior, and π_{a,O_a} the posterior after observation O_a. An action a selects experiment $O_a\sim K_a(\cdot\mid\vartheta)$ and posterior π_{a,O_a}. For an integrable terminal loss $\ell(y,\vartheta)$, assume finite Bayes risks and measurable ε-optimal decision rules for every $\varepsilon>0$. Define

$$\mathfrak{r}(\pi)=\inf_y\int\ell(y,\vartheta)\pi(d\vartheta),\qquad \mathsf{R}(\pi,a)=\mathbb{E}[\mathfrak{r}(\pi_{a,O_a})\mid\pi,a].$$

Definition C.1 (Blackwell dominance)**.** $K_{a_2} \succeq_B K_{a_1}$ if a state-independent Markov kernel G satisfies

$$K_{a_1}(do_1 \mid \vartheta) = \int G(do_1 \mid o_2) K_{a_2}(do_2 \mid \vartheta)$$

for every ϑ.

Theorem C.1 (Blackwell dominance lowers posterior Bayes risk)**.** *If $K_{a_2} \succeq_B K_{a_1}$, then $\mathsf{R}(\pi, a_2) \leq \mathsf{R}(\pi, a_1)$ for every prior and every loss in the displayed Bayes-risk definition.*

Interpretation: a more informative experiment cannot increase the smallest expected loss attainable after observing its outcome.

Write the preposterior Bayes risk as

$$\mathsf{R}(\pi, a) = \inf_{\delta} \mathbb{E}[\ell(\delta(O_a), \vartheta)],$$

where δ ranges over measurable decision rules. If $K_{a_2} \succeq_B K_{a_1}$, take any rule for O_{a_1}, garble O_{a_2} through the state-independent kernel in the displayed Blackwell equation, and apply that rule. The resulting joint law of $(\vartheta, \widetilde{O}_{a_1})$ equals that of (ϑ, O_{a_1}), so the more informative experiment can attain the less-informative risk. Taking the infimum proves Theorem C.1.

Corollary C.2 (Strategic exploration without direct action preferences)**.** *Let y denote a scalar report, $L(y)$ its expected score cost, and $\varrho(y)$ its induced experimental action. Suppose their domains are compact intervals, L has unique minimizer y^0, the plan map $\varrho(y)$ is nondecreasing, and $a_2 \geq a_1$ implies $K_{a_2} \succeq_B K_{a_1}$, and there is no direct action cost. Let $\theta_F \geq 0$ be the continuation weight. If the forecaster minimizes*

$$H(y) = L(y) + \theta_F \mathsf{R}(\pi, \varrho(y)),$$

then every minimizer satisfies $y^\star \geq y^0$. If L and the continuation term are differentiable at an interior y^0, $\theta_F > 0$, $\varrho'(y^0) > 0$, and $\partial_a \mathsf{R}(\pi, \varrho(y^0)) < 0$, then $y^\star > y^0$.

Interpretation: even without a direct action preference, a forecaster may raise its report to induce a more informative future experiment.

Proof. For $y < y^0$, the score is strictly larger and Blackwell monotonicity makes the continuation risk no smaller. Under the strict assumptions, $H'(y^0) < 0$. □

C.4 Strictness and common information

Blackwell dominance is weak. Strict reduction requires that the additional signal change an optimal decision on a set of positive probability and that the loss be strict there; a strict derivative additionally requires smooth dependence on the experiment parameter. It does not follow from weak Blackwell dominance. The strategic-exploration corollary holds with a fixed common posterior; the report is an influence instrument, not evidence that changes beliefs.

D Influence and Truthfulness Details

This appendix retains the dynamic decomposition and the local results moved out of the main reading path. Period and role subscripts are restored here.

D.1 Dynamic score–plan separation

Let $\bar{S}_k(p, q)$ be expected audit loss and let $D_{S,k}(p, q) = \bar{S}_k(p, q) - \bar{S}_k(p, p)$. Under strict (P1), the continuation law depends only on the plan, so its expected value is already part of V_{F}. The following extension instead allows an explicitly announced dependence on the realized score.

Lemma D.1 (Affine expected score continuation)**.** *Suppose the discounted expected continuation charge conditional on the audit observation o is $a_k(y) + b_k(y) S_k(q, o)$, with the coefficients fixed before o is drawn, and the audit law is independent of (q, w, y). If $A_k(y) := \alpha_k + b_k(y) > 0$, the total expected objective is*

$$\mathcal{J}_k(q, w) = A_k(y) D_{S,k}(p, q) + \Upsilon_{k,p}(y), \qquad y = \Pi_k(q, w), \tag{32}$$

where $\Upsilon_{k,p}(y)$ includes operational cost, $a_k(y)$, and $A_k(y)\bar{S}_k(p,p)$. If truth attains a global minimum of this augmented plan value, every optimum is truthful.

Proof. Take expectation over the independent audit and write $\bar{S}_k(p,q) = D_{S,k}(p,q) + \bar{S}_k(p,p)$. The first term in (32) is nonnegative, and is positive at any false report under strict identification. The stated truthful plan attains the lower bound on the second term. □

An affine rule for a future authority variable is not enough: the induced *expected continuation value* must have the stated affine form. If A_k depends on y, this is not the constant-weight objective used in the scalar and vector closed forms.

Remark D.2 (Nonlinear score continuation). A nonlinear transformation can destroy propriety. For a Bernoulli audit outcome with mean $p \in (0,1)$ and squared loss, squaring the realized loss gives expected cost $(1-p)q^4 + p(1-q)^4$. Its minimizer is $p^{1/3}/(p^{1/3} + (1-p)^{1/3})$, which differs from p unless $p = 1/2$. This counterexample shows that arbitrary nonlinear continuation charges are not covered by the separation argument.

D.2 Plan geometry

For a bounded report interval $[\underline{q}, \overline{q}]$, let $\mathcal{P}(q)$ be the plans reachable at report q under the affine protocol.

Lemma D.3 (Exact reachable-plan geometry)**.** *With request interval $[-U, U]$ and benchmark $\bar{v}$,*

$$\bigcup_q \mathcal{P}(q) = [\iota\underline{q} + (1-\iota)\bar{v} - \beta U, \iota\overline{q} + (1-\iota)\bar{v} + \beta U],$$
$$\bigcap_q \mathcal{P}(q) = [\iota\overline{q} + (1-\iota)\bar{v} - \beta U, \iota\underline{q} + (1-\iota)\bar{v} + \beta U].$$

The intersection is nonempty exactly when $\iota(\overline{q} - \underline{q}) \le 2\beta U$.

Interpretation: influence shifts the reachable interval with the report, while the override budget widens it; overlap requires widening to dominate the shift.

Proof. Each reachable set is an interval of radius βU centered at $\iota q + (1-\iota)\bar{v}$. Taking extreme endpoints gives the union and intersection. Comparing the intersection endpoints gives the condition. □

Definition D.1 (Quadratic plan-preference class)**.** For an open plan set $\mathcal{N}$ and a curvature bound $\bar{c}_F$, the class contains plan values $c_F\|y - a_F\|^2/2 + b(y)$ with $a_F \in \mathcal{N}$, $0 < c_F \le \bar{c}_F$, and plan-only cost b.

D.3 Local influence gradients and necessity

Suppose the plans reachable at report q are described by smooth inequalities $g(q,y) \le 0$. Let $\mu^\star$ be the KKT multiplier at a locally unique optimum, and assume linear independence of active constraint gradients, together with differentiability of the local optimized value and the corresponding envelope formula. Let T_p be the tangent space of the report domain at p, and let proj_{T_p} denote orthogonal projection.

Theorem D.4 (Influence-gradient necessity)**.** *Assume the score is differentiable with zero report gradient at truth. At an interior truthful report, local optimality requires*

$$\mathrm{proj}_{T_p}\left[(\nabla_q g(p, y^\star))^\top \mu^\star\right] = 0. \tag{33}$$

A nonzero expression rules out local truthfulness. If the reduced objective is locally convex, the condition is also sufficient for local optimality.

Interpretation: a binding plan constraint biases the report only through how that constraint moves when the report changes.

Proof. Differentiate the optimized plan-value Lagrangian using the assumed envelope formula. Properness makes the score gradient vanish at truth, leaving (33). Local convexity upgrades stationarity to local optimality. □

Theorem D.5 (Two-sided robustness)**.** *Let $\mathcal{Q}$ be convex under norm $\|\cdot\|$. Suppose p is interior, $D_S(p,\cdot)$ is differentiable with $\nabla_q D_S(p,p) = 0$, and $D_S(p,\cdot)$ is μ_S-strongly convex. If $\underline{V}_{\mathrm{F}}$ is L_V-Lipschitz, and a truthful request is within $\varepsilon_{\mathrm{plan}}$ of the globally best plan cost, every interior optimal report satisfies*

$$\|q^\star - p\| \le \min\left\{\frac{L_V}{\alpha\mu_S}, \sqrt{\frac{2\varepsilon_{\mathrm{plan}}}{\alpha\mu_S}}\right\}. \tag{34}$$

For the scalar affine protocol with a report-independent request set, if V_{F} is K_V-Lipschitz in the plan, then $\underline{V}_{\mathrm{F}}$ is L_V-Lipschitz with $L_V \le \iota K_V$; in particular $\iota = 0$ gives $L_V = 0$ and exact truthfulness.

Proof. Here subgradients of the locally Lipschitz operational value are Clarke subgradients, with the dual norm used for covectors. At an interior optimum, $0 \in \alpha\partial_q D_S(p,q^\star) + \partial\underline{V}_{\mathrm{F}}(q^\star)$. Strong convexity and the fact that p minimizes $D_S(p,\cdot)$ give every score subgradient norm at least $\mu_S\|q^\star - p\|$, while Lipschitzness bounds every plan-value subgradient by L_V; this proves the first bound. Optimality compared with $q = p$ and $\varepsilon_{\mathrm{plan}}$-attainability gives $\alpha D_S(p,q^\star) \le \varepsilon_{\mathrm{plan}}$. Strong convexity then gives the second bound. □

Corollary D.6 (Influence-scaled robustness)**.** *If plan value is K-Lipschitz and the protocol is affine, the smooth constant in Theorem D.5 satisfies $L_V \le \iota K$. For the scalar score $D_S(p,q) = (q-p)^2/2$, $|q^\star - p| \le \iota K/\alpha$.*

Interpretation: smooth misspecification creates no report bias at zero influence and at most linear bias near that corner.

Corollary D.7 (Loss relative to a fixed truthful response)**.** *Suppose $q = p + z$ and a downstream scalar action differs from its truthful value a_0 by at most $L_a|z|$. For principal loss $(a-p)^2/2 + \lambda_q(q-p)^2/2$, the absolute loss change is at most $|a_0 - p|L_a|z| + (L_a^2 + \lambda_q)z^2/2$. One may substitute the bound of Theorem D.5 for $|z|$. The linear term vanishes if $a_0 = p$.*

Proof. Expand $(a-p)^2 - (a_0-p)^2 = 2(a_0-p)(a-a_0) + (a-a_0)^2$ and apply the assumed response bound. □

Theorem D.8 (Necessity over a quadratic preference class)**.** *Fix an open plan set $\mathcal{N}$ and nonempty compact reachable-plan sets. Assume every report is a possible type and every pairwise score regret $D_S(q_1,q_0)$ is finite. If truthfulness holds for every quadratic plan preference with target in $\mathcal{N}$ and unbounded curvature, every reachable plan in $\mathcal{N}$ must be reachable under every report. With curvature bounded by $\bar{c}_F$, truthfulness fails whenever*

$$\bar{c}_F > \frac{2\alpha D_S(q_1,q_0)}{\mathrm{dist}(y_0, \mathcal{P}(q_1))^2}$$

for a plan $y_0 \in \mathcal{N}$ reachable at q_0 but not at q_1. If $\mathcal{N}$ contains the union of reachable sets, uniform truthfulness is therefore equivalent to identical reachable-plan sets (under attainment).

Interpretation: a guarantee uniform over arbitrarily strong plan preferences requires plan-set overlap; Condition (T) is weaker because it is specific to the actual preference.

Proof. Choose type $p = q_1$ and make y_0 its preferred plan. Truth incurs at least half the curvature times squared distance to $\mathcal{P}(q_1)$, whereas report q_0 reaches y_0 at score cost $\alpha D_S(q_1,q_0)$. The displayed inequality makes the deviation strictly profitable. □

D.4 Plan constraints and hidden intertemporal influence

Proposition D.9 (Report-anchored constraints)**.** *For a scalar constraint $g(q,y) = y - [(1-\omega)\bar{v} + \omega q] - \delta \le 0$, where $\omega \in [0,1]$ and $\delta \ge 0$, an optimal multiplier $\mu \ge 0$ contributes $-\omega\mu$ to the report derivative of the optimized plan value. A slack constraint has zero contribution.*

Proof. The envelope contribution is $\mu\partial_q g = -\omega\mu$. Complementary slackness makes $\mu = 0$ at a slack constraint. □

Corollary D.10 (A report-anchored interval is an affine protocol)**.** *A report-anchored interval* $[(1-\omega)\bar{v} + \omega q - \delta, (1-\omega)\bar{v} + \omega q + \delta]$ *is the two-parameter protocol with influence* ω*, leverage one, and override budget* δ*. It contains the forecaster's preferred plan at a truthful report exactly when* $\delta \geq |(1-\omega)(p-\bar{v}) + \Delta_{\mathrm{F}}|$.

Interpretation: report anchoring and direct pass-through have the same threshold geometry.

Proof. Write every feasible plan as $\omega q + (1-\omega)\bar{v} + w$ with $|w| \leq \delta$ and evaluate it at $q = p$. □

Proposition D.11 (Hidden intertemporal influence)**.** *Let a future plan be* $a(q) = (1-\theta)\bar{v}' + \theta q$*, where* $\theta \in [0,1]$ *and* $\bar{v}'$ *is fixed. A forecaster minimizing* $\alpha(q-p)^2/2 + \kappa[a(q) - a_F]^2/2$*, with future target* a_F *and* $\kappa > 0$*, reports*

$$q^\star - p = \frac{\kappa\theta[a_F - (1-\theta)\bar{v}' - \theta p]}{\alpha + \kappa\theta^2}.$$

Thus a future benchmark that reads the current report can cause distortion even when the current action is report-independent.

Proof. Differentiate the strictly convex quadratic and solve its first-order condition. The numerator identifies precisely when truth is optimal. □

D.5 Scalar extensions

Proof of Theorem 4.1. Write $z = q - p$. The plan residual relative to $p + \Delta_{\mathrm{F}}$ is $\iota z + \beta w - G(\iota)$. If $G(\iota)/\beta$ is feasible, the request alone makes this residual zero and both terms in J_{F} attain their strict minima at $z = 0$. Otherwise the request projects onto the nearer endpoint. The remaining strictly convex quadratic in z has minimizer (7); the branch is self-consistent because $c_{\mathrm{F}}\iota^2/(\alpha + c_{\mathrm{F}}\iota^2) < 1$. At $\iota = 0$, the report enters only through the score. □

Example D.1 (Cheap authority can be worthless)**.** Let $p \sim \mathrm{Unif}[-.1, .1]$, $\bar{v} = 0$, $\Delta_{\mathrm{F}} = 1$, $\iota = 0$, $\lambda_q = 0$. With no authority the equilibrium is truthful with the plan pinned at the benchmark, $\mathbb{E}\mathcal{L} \approx .00167$; with authority covering every preferred plan ($\beta U \geq 1.1$), the equilibrium remains truthful but $y^\star = p + 1$, so $\mathbb{E}\mathcal{L} = 1/2$ — three hundred times worse, even when authority is free.

Proposition D.12 (Strategic overlay owner)**.** *Suppose the forecaster has no request channel. An owner knows the type* p *independently of the report, observes* q*, publishes* $y = q + w_{\mathrm{O}}$ *with* $|w_{\mathrm{O}}| \leq U_{\mathrm{O}}$*, and has a quadratic cost minimized at* $p + \Delta_{\mathrm{O}}$*, where* Δ_{O} *is its preferred-plan offset. Truthful reporting is optimal when* $|\Delta_{\mathrm{O}}| \leq U_{\mathrm{O}}$ *and* $|\Delta_{\mathrm{O}} - \Delta_{\mathrm{F}}| \leq \sqrt{\alpha/(\alpha + c_{\mathrm{F}})}(U_{\mathrm{O}} - |\Delta_{\mathrm{F}}|)_+$*. Proposition D.15 gives the exact piecewise comparison and Remark D.17 the sequential-owner qualification.*

Proof. At truth the owner attains its target because $|\Delta_{\mathrm{O}}| \leq U_{\mathrm{O}}$, so the forecaster's cost is $c_{\mathrm{F}}(\Delta_{\mathrm{O}} - \Delta_{\mathrm{F}})^2/2$. Reports in the owner's nonbinding interval preserve that plan and only increase score cost. On a binding branch, minimizing the corresponding quadratic even without its branch constraint costs at least $\alpha c_{\mathrm{F}}(U_{\mathrm{O}} - |\Delta_{\mathrm{F}}|)^2/[2(\alpha + c_{\mathrm{F}})]$ when $U_{\mathrm{O}} \geq |\Delta_{\mathrm{F}}|$. The displayed condition compares these costs. If $U_{\mathrm{O}} < |\Delta_{\mathrm{F}}|$, its premises cannot hold together. □

Corollary D.13 (Opposing displacements)**.** *If* $(p - \bar{v})\Delta_{\mathrm{F}} < 0$ *and* $|p - \bar{v}| > |\Delta_{\mathrm{F}}|$*, then*

$$\iota_0 = 1 + \frac{\Delta_{\mathrm{F}}}{p - \bar{v}} \in (0,1), \qquad U^\dagger(\iota_0, \beta; p) = 0.$$

For fixed β*,* $U^\dagger(\cdot, \beta; p)$ *is V-shaped and has its unique minimum at* ι_0*.*

Proof. The affine function $G(\iota) = (1-\iota)(p-\bar{v}) + \Delta_{\mathrm{F}}$ changes sign once on $(0,1)$ under the stated conditions. Since $U^\dagger = |G|/\beta$, the claim follows. □

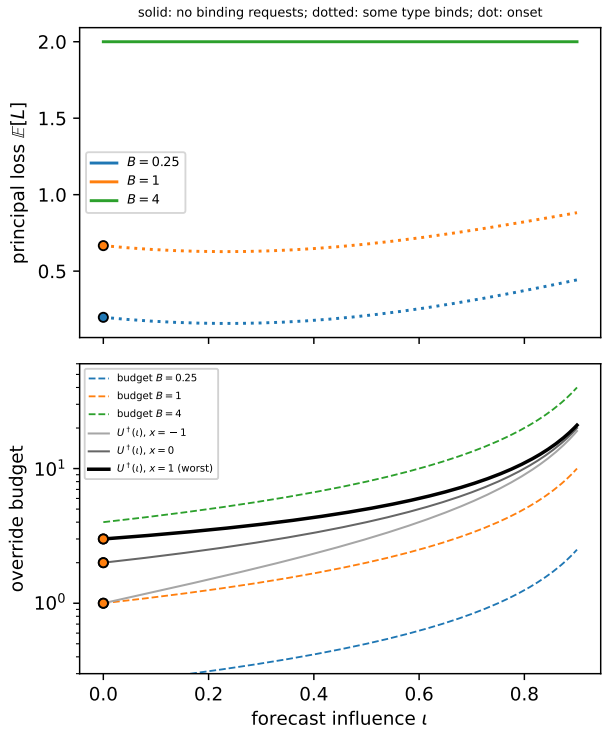


Figure 3: Along the diagonal $\beta = 1 - \iota$ (so $\iota = 1$ is excluded), with $p - \bar{v} \sim \mathrm{Unif}[-1, 1]$, $\Delta_{\mathrm{F}} = 2$, $\alpha = 4$, $c_{\mathrm{F}} = 2$, $\lambda_q = 1$: expected principal loss (top) at authority $B \in \{.25, 1, 4\}$, solid while no type's request binds and dotted once some type binds (dots mark the onset); and the truthful-report threshold $U^\dagger(\iota, 1 - \iota; p) = |G(\iota)|/(1 - \iota)$ for representative types $p - \bar{v} \in \{-1, 0, 1\}$ against the available budgets $B/(1 - \iota)$ (bottom, log scale). The threshold diverges as $\iota \to 1$: the required budget in request units grows, because request leverage $\beta = 1 - \iota$ vanishes. The authority threshold $B^\dagger = |G|$ remains finite, and the available budget $B/(1 - \iota)$ grows on the same scale.

Proposition D.14 (Quadratically penalized requests)**.** *Replace the hard request constraint by the penalty $\rho_w w^2/2$, where $\rho_w > 0$, and allow $w \in \mathbb{R}$. In the scalar model the unique optimum is*

$$q^\star - p = \frac{c_{\mathrm{F}} \iota \rho_w \, G(\iota)}{\alpha \rho_w + c_{\mathrm{F}} \rho_w \iota^2 + \alpha c_{\mathrm{F}} \beta^2}, \tag{35}$$

$$w^\star = \frac{c_{\mathrm{F}} \beta \alpha \, G(\iota)}{\alpha \rho_w + c_{\mathrm{F}} \rho_w \iota^2 + \alpha c_{\mathrm{F}} \beta^2}. \tag{36}$$

Thus a finite quadratic penalty has no nontrivial exact-truthful branch: when $\iota > 0$, the report is truthful if and only if $G(\iota) = 0$. Truthfulness still holds for all ρ_w at $\iota = 0$.

Proof. Set $z = q - p$ and $e = \iota z + \beta w - G(\iota)$. The first-order conditions are $\alpha z + c_{\mathrm{F}} \iota e = 0$ and $\rho_w w + c_{\mathrm{F}} \beta e = 0$. Substitution gives

$$e = -\frac{\alpha \rho_w \, G(\iota)}{\alpha \rho_w + c_{\mathrm{F}} \rho_w \iota^2 + \alpha c_{\mathrm{F}} \beta^2},$$

which yields (35)–(36). The Hessian is positive definite because $\alpha, \rho_w > 0$, so the stationary point is the unique global minimizer. □

Proposition D.15 (Exact comparison for a strategic overlay owner)**.** *Let $z = q - p$ and let the overlay owner best respond according to*

$$T(z) := y - p = z + \mathrm{clip}(\Delta_{\mathrm{O}} - z, -U_{\mathrm{O}}, U_{\mathrm{O}}).$$

Put $\ell = \Delta_{\mathrm{O}} - U_{\mathrm{O}}$, $u = \Delta_{\mathrm{O}} + U_{\mathrm{O}}$ and

$$z_- = \min\left\{\frac{c_{\mathrm{F}}(\Delta_{\mathrm{F}} - U_{\mathrm{O}})}{\alpha + c_{\mathrm{F}}}, \ell\right\}, \quad z_0 = \mathrm{proj}_{[\ell, u]}(0), \quad z_+ = \max\left\{\frac{c_{\mathrm{F}}(\Delta_{\mathrm{F}} + U_{\mathrm{O}})}{\alpha + c_{\mathrm{F}}}, u\right\}.$$

Define $V_j = \alpha z_j^2/2 + c_{\mathrm{F}}[T(z_j) - \Delta_{\mathrm{F}}]^2/2$ *for* $j \in \{-, 0, +\}$. *The forecaster's equilibrium distortions are precisely*

$$\arg \min_{z \in \{z_-, z_0, z_+\}} V(z).$$

Consequently truthful reporting is optimal if and only if $V(0) = \min\{V_-, V_0, V_+\}$, *where this equality is possible only when* $0 \in \{z_-, z_0, z_+\}$.

Proof. The overlay response equals $T(z) = z + U_{\mathrm{O}}$ on $(-\infty, \ell]$, $T(z) = \Delta_{\mathrm{O}}$ on $[\ell, u]$, and $T(z) = z - U_{\mathrm{O}}$ on $[u, \infty)$. On each closed interval the forecaster's objective

$$V(z) = \frac{\alpha}{2} z^2 + \frac{c_{\mathrm{F}}}{2}[T(z) - \Delta_{\mathrm{F}}]^2$$

is a strictly convex quadratic. Projecting its unconstrained minimizer onto the corresponding interval gives z_-, z_0, and z_+. Comparing their values is therefore necessary and sufficient globally. □

Corollary D.16 (Binding delta overlays retain pass-through)**.** *On either binding branch of* $y = q + w_{\mathrm{O}}$, $dy/dq = 1$. *A separate protocol with an interval fixed around a report-independent benchmark has zero report pass-through; its report-conditional attainable sets generally differ from those of the delta protocol.*

Remark D.17 (Sequential overlay owners). With sequential owners, the composed plan map is piecewise affine. Its local slope is zero when a downstream owner with slack absorbs a report change and one when every downstream request is constrained. A global characterization with multiple owners requires tracking all active-set sequences.

D.6 Completion of squares for ex-ante influence

On the common upper-endpoint branch, let

$$D_\beta := \Delta_{\mathrm{F}} + (1 - \iota)(p - \bar{v}) - \beta U$$

and define

$$A(\iota) = \frac{\alpha}{\alpha + c_{\mathrm{F}}\iota^2}, \qquad K(\iota) = \frac{c_{\mathrm{F}}\iota}{\alpha + c_{\mathrm{F}}\iota^2}.$$

Then

$$\begin{aligned} q^\star - p &= K D_\beta, \\ y^\star - p &= \Delta_{\mathrm{F}} - A D_\beta. \end{aligned}$$

Writing $\mu_D = \mathbb{E} D_\beta$ and using $\mathrm{Var}(D_\beta) = (1 - \iota)^2 \mathrm{Var}(p)$, the expected principal loss is

$$\begin{aligned} \mathbb{E}\mathcal{L} = \frac{1}{2}\left[(\Delta_{\mathrm{F}} - A\mu_D)^2 + A^2 \mathrm{Var}(D_\beta)\right] \\ + \frac{\lambda_q}{2} K^2 \left[\mu_D^2 + \mathrm{Var}(D_\beta)\right]. \end{aligned}$$

Differentiation gives

$$\mu_D^\star = \frac{A\Delta_{\mathrm{F}}}{A^2 + \lambda_q K^2}.$$

At this value, the squared-mean terms equal

$$\frac{\lambda_q K^2 \Delta_{\mathrm{F}}^2}{2(A^2 + \lambda_q K^2)},$$

and adding the variance terms yields (9). This also shows why feasibility of the unconstrained $\mu_D^\star$ for the common-binding support condition must be stated rather than assumed.

For $\iota < 1$, any desired mean displacement can be attained with $B = 0$ by taking $\bar{v} = m + (\Delta_{\mathrm{F}} - \mu_D^\star)/(1 - \iota)$. All types remain strictly on the upper branch exactly when $\mu_D^\star > (1 - \iota)(m - p_{\min})$. This proves the feasibility condition in Proposition 4.3. The derivative at zero of its minimized value is $-\mathrm{Var}(p)$ whenever the branch is feasible in a neighborhood of zero. With a degenerate type, zero influence, $B = 0$, and $\bar{v} = p$ already give the full-information loss zero.

Proposition D.18 (Unbounded benchmarks can prevent an optimal design)**.** *Let the scalar type have positive variance and compact interval support, let $\Delta_{\mathrm{F}} > 0$, and take $\lambda_q = 0$ and zero design costs. With fixed $\alpha, c_{\mathrm{F}} > 0$, unrestricted finite benchmarks, and finite $B \geq 0$, the infimum principal loss is zero but is not attained.*

Proof. Write $m = \mathbb{E}[p]$. For $\iota < 1$, choose $B = 0$ and

$$\bar{v} = m - \frac{c_{\mathrm{F}} \iota^2 \Delta_{\mathrm{F}}}{\alpha(1-\iota)}.$$

For ι sufficiently close to one all types lie on the positive branch, and substitution gives $y^\star - p = -\Lambda_y(1-\iota)(p-m)$. Expected action loss tends to zero. For any finite design with $\iota < 1$, zero loss would require the soft-thresholded displacement $D_B(p)$ to equal the positive constant $\Delta_{\mathrm{F}}/\Lambda_y$ throughout a nondegenerate interval. Its only flat portion is zero and its nonzero branches have slope $1 - \iota > 0$, so this is impossible. At $\iota = 1$, $0 \leq D_B \leq \Delta_{\mathrm{F}}$ gives $y^\star - p \geq \Delta_{\mathrm{F}}(1 - \Lambda_y) > 0$. Thus no finite design attains zero. The improving sequence uses $\bar{v} \to -\infty$. □

A compact admissible benchmark range rules out this construction and gives an attained minimum when the remaining design domain is compact and costs are continuous. The finite design search in E1 fixes the benchmark.

E Vector Projection and Stability Proofs

E.1 Eliminating the report

Write $d = q - p$ and $t = Ew$. The request image $\mathcal{T} = E\mathcal{W} \subseteq \mathbb{R}^m$ is nonempty, compact, and convex. The forecaster minimizes

$$J(d,t) = \tfrac{1}{2} d^\top H_q d + \tfrac{1}{2}(Cd + t - \mathrm{g})^\top H_y (Cd + t - \mathrm{g}).$$

For fixed t, stationarity gives $Kd = C^\top H_y(\mathrm{g} - t)$, with $K = H_q + C^\top H_y C \succ 0$. Substitution gives exactly $\tfrac{1}{2}(\mathrm{g} - t)^\top M(\mathrm{g} - t)$, where the matrix inversion lemma yields

$$M = H_y - H_y C K^{-1} C^\top H_y = (H_y^{-1} + C H_q^{-1} C^\top)^{-1} \succ 0.$$

No rank assumption on C or E is required. Strict convexity in t gives a unique optimal request image $t^\star$, hence unique $d^\star$ and $y^\star$. Requests implementing $t^\star$ can differ by elements of $\ker E$. If E is injective on $\mathcal{W}$, the request is also unique. This proves (12).

E.2 A certificate computed at truth

For a closed convex set $\mathcal{T}$ and a positive definite matrix H, define its metric projection and distance by

$$P_{\mathcal{T}}^H(g) = \arg\min_{t \in \mathcal{T}} \|g - t\|_H, \qquad \mathrm{dist}_H(g, \mathcal{T}) = \|g - P_{\mathcal{T}}^H(g)\|_H.$$

The minimized operational value at displacement d is

$$\psi(d) = \tfrac{1}{2} \mathrm{dist}^2_{H_y}(\mathrm{g} - Cd, \mathcal{T}).$$

Squared distance to a closed convex set is continuously differentiable, including at active-set changes. Transforming to Euclidean coordinates gives

$$\nabla\psi(d) = -C^\top H_y[\mathrm{g} - Cd - P_{\mathcal{T}}^{H_y}(\mathrm{g} - Cd)].$$

In particular, $\nabla\psi(0) = -h$. Since the reduced objective $F(d) = d^\top H_q d/2 + \psi(d)$ is strictly convex, $d^\star = 0$ if and only if $h = 0$. Stationarity of the joint problem gives the equivalent equilibrium certificate in (14). If C has full row rank, $C^\top$ is injective and H_y invertible, so $h = 0$ if and only if $\mathrm{g} = t_0 \in \mathcal{T}$.

E.3 Quantitative report distortion

Let $z = H_q^{1/2} d$ and $\widetilde{F}(z) = \|z\|_2^2/2 + \psi(H_q^{-1/2} z)$. The residual map $I - P$ of a Euclidean convex projection is nonexpansive. Consequently the gradient of its second term is Lipschitz with constant $\kappa = \|H_y^{1/2} C H_q^{-1/2}\|_2^2$. The function $\widetilde{F}$ is 1-strongly convex and its gradient is $(1+\kappa)$-Lipschitz. Its minimizer $z^\star$ satisfies $\nabla \widetilde{F}(z^\star) = 0$ and $\nabla \widetilde{F}(0) = -H_q^{-1/2} h$. Strong monotonicity gives

$$\|z^\star\|_2^2 \le \langle \nabla \widetilde{F}(0), -z^\star \rangle \le \|H_q^{-1/2} h\|_2 \|z^\star\|_2,$$

and gradient Lipschitzness gives $\|H_q^{-1/2} h\|_2 \le (1+\kappa)\|z^\star\|_2$. These prove both bounds in (13). For a scalar protocol with no requests, the lower bound is an equality. The bounds remain valid when request faces change, unlike a derivative computed within a fixed active set.

Proposition E.1 (Global stability under target perturbations)**.** *Hold $C, E, \mathcal{W}, H_q, H_y$ fixed. For displacement targets $\mathsf{g}_1, \mathsf{g}_2$, let (d_j, t_j) minimize J above. Then*

$$\|d_1 - d_2\|_{H_q}^2 + \|C(d_1 - d_2) + t_1 - t_2\|_{H_y}^2 \le \|\mathsf{g}_1 - \mathsf{g}_2\|_{H_y}^2.$$

Proof. The pair $(H_q^{1/2} d, H_y^{1/2}(Cd + t))$ ranges over a closed convex set as $d \in \mathbb{R}^d$ and $t \in \mathcal{T}$. Minimizing J projects $(0, H_y^{1/2}\mathsf{g})$ onto this set. Nonexpansiveness of Euclidean projection proves the inequality. Closedness follows because convergence of the first coordinate bounds d and $\mathcal{T}$ is compact. □

Incomplete coverage without distortion. Let $C = (1,0)^\top$, $E = (1,0)^\top$, $H_q = 1$, $H_y = I_2$, and $\mathcal{W} = [-1, 1]$. With $\mathsf{g} = (1/2, 2)^\top$, the projection at truth is $t_0 = (1/2, 0)^\top$, so the residual is nonzero but $h = 0$. The second plan coordinate cannot be moved by reporting, and truth is optimal for every positive score weight. Adding an independent request coordinate that does not change Ew creates nonunique requests without changing this unique report or plan.

Numerical checks. Experiment E0 compares the reduced and joint quadratic programs. The additional geometry checks in `experiments/e5_geometry.py` include rank-deficient report channels, redundant requests, active-set crossings, the distortion bounds, and the stability inequality.

F Leak-Test and Randomized-Threshold Details

F.1 Finite-sample inference

After partialling out predeclared plan-measurable controls, write

$$a_k = \beta_\eta \eta_k + \varepsilon_{a,k}, \qquad \beta_\eta = b_q - \frac{\iota}{\beta} b_w,$$

which equals $\lambda_{\text{leak}} \Xi$ under the leak model (15). Eligibility is determined from the unperturbed request before η_k is drawn. Conditional on a balanced gated design and independent $\varepsilon_{a,k} \sim \mathcal{N}(0, \sigma_a^2)$,

$$\widehat{\beta}_\eta = \frac{\sum_k \eta_k a_k}{\sum_k \eta_k^2} \sim \mathcal{N}\left(\beta_\eta, \frac{\sigma_a^2}{\sum_k \eta_k^2}\right).$$

With unknown variance and a fitted intercept, the studentized statistic has a noncentral $t_{n_{\text{elig}}-2}$ distribution under the alternative. The planning formula is

$$n_{\text{elig}} \ge \frac{(z_{1-a_{\text{sig}}/2} + z_{1-b_{\text{pow}}})^2 \sigma_a^2}{\lambda_0^2 s_\eta^2}, \tag{37}$$

which is a sufficient bound for known Gaussian variance and a normal planning approximation otherwise. The exact two-sided Gaussian power at noncentrality $d = |b_\eta| \sqrt{n_{\text{elig}}} s_\eta / \sigma_a$ is $1 - \Phi(z_{1-a_{\text{sig}}/2} - d) + \Phi(-z_{1-a_{\text{sig}}/2} - d)$, where Φ is the standard normal distribution function. For $\sigma_a = .8$, $s_\eta = \Xi = 1$, level .05, and power .8 against $\lambda_0 = .1$, $n_{\text{elig}} \ge 503$ eligible periods suffice.

Proposition F.1 (Observational leakage screen). *Let $\widetilde{q} = q - \mathbb{E}[q \mid x, y, \mathcal{H}]$ and assume $\mathrm{Var}(\widetilde{q}) > 0$ and $\bar{q}$ is measurable with respect to $(x, y, \mathcal{H})$. Under the report-only case of (15), write ζ_a for its conditionally mean-zero action residual. Then*

$$\frac{\mathrm{Cov}(a, \widetilde{q})}{\mathrm{Var}(\widetilde{q})} = \lambda_{\mathrm{leak}} \Xi$$

if $\mathbb{E}[\mathcal{A}(x, y, \mathcal{H})\widetilde{q}] = 0$ and $\mathbb{E}[\zeta_a \widetilde{q}] = 0$.

Proof. Multiply (15) by $\widetilde{q}$, take expectations, and use the two orthogonality conditions. □

Theorem F.2 (Leakage degradation). *Suppose every plan attainable by a deviation is also attainable under truth, the only direct message effect is $\lambda_{\mathrm{leak}}\Xi(q - \bar{q})$ with plan-measurable $\bar{q}$, and reports are scalar in a convex interval. Suppose operational cost is L_G-Lipschitz over the compared downstream actions, and $D_S(p, \cdot)$ is μ_S-strongly convex. Every interior equilibrium report satisfies*

$$\|q^\star - p\| \le \frac{2 L_G \lambda_{\mathrm{leak}} |\Xi|}{\alpha \mu_S}.$$

If local principal loss is at most $M_{\mathrm{P}}\|q - p\|^2/2$, the associated loss is at most $M_{\mathrm{P}}/2$ times the square of this bound.

Proof. Compare a deviation with a truthful request implementing the same plan. For $t = \|q - p\|$, leakage can improve continuation cost relative to truth by at most $L_G \lambda_{\mathrm{leak}} |\Xi| t$, while strong properness charges at least $\alpha \mu_S t^2/2$. Optimality relative to the truthful deviation implies $\alpha \mu_S t^2/2 \le L_G \lambda_{\mathrm{leak}} |\Xi| t$. The claim is immediate for $t = 0$ and follows by division otherwise. □

F.2 Dependence and multiple downstream agents

With serial dependence, retain random assignment and use the HAC covariance

$$(X^\top X)^{-1} \left(\sum_{|\ell| \le L} k(\ell/L) \sum_t x_t x_{t-\ell}^\top \widehat{\varepsilon}_{a,t} \widehat{\varepsilon}_{a,t-\ell} \right) (X^\top X)^{-1}.$$

Here X is the regression design matrix, x_t its column-form row, $\widehat{\varepsilon}_{a,t}$ the fitted residual, k a covariance kernel, and L a lag bandwidth. Consistency requires standard mixing and moment conditions together with $L \to \infty$ and $L/n_{\mathrm{elig}} \to 0$. For m downstream agents, a joint Wald test uses their cross-agent covariance. Alternatively, Holm's step-down rule compares the ordered values $p_{(j)}$ with $a_{\mathrm{sig}}/(m - j + 1)$ and controls family-wise error under arbitrary dependence.

F.3 Estimating influence and leverage

Use an independent report perturbation η_k^q with the request fixed and an independent request perturbation η_k^w with the report fixed:

$$\widehat{\iota} = \frac{\mathrm{Cov}(y_k, \eta_k^q)}{\mathrm{Var}(\eta_k^q)}, \qquad \widehat{\beta} = \frac{\mathrm{Cov}(y_k, \eta_k^w)}{\mathrm{Var}(\eta_k^w)}.$$

For nonlinear protocols these are intervention-specific regression slopes; interpreting them as derivative averages requires additional conditions on the perturbation law. Sample splitting or joint delta-method inference addresses estimation uncertainty, but does not make the plan exactly invariant. The residual plan change equals $[\iota - \beta\widehat{\iota}/\widehat{\beta}]\eta$ and can contaminate the leakage estimand. Exact compensation requires the known protocol or a separately validated plan-holding intervention.

F.4 Boundary-compatible interventions

The symmetric design (16) requires slack $U - |w| \ge \iota\bar{\eta}/\beta$, which fails exactly on the uncovered branch of Theorem 4.1, where requests saturate. The boundary-compatible variant randomizes $\eta \in \{0, +\bar{\eta}\}$ at $w = U$ and $\eta \in \{0, -\bar{\eta}\}$ at $w = -U$, compensating the request as before and

stratifying by boundary sign. The plan is again unchanged, and the contrast between arms identifies b_η of (17) under the linear response model, or a finite difference in the constant-plan direction otherwise. Inference is the same paired regression run within strata. At $U = 0$ with $\iota > 0$ no feasible plan-preserving direction exists, and we report only the shadow-interface design, in which the downstream-visible copy of q is perturbed while the scored report and plan are held fixed; that channel requires interface access and is analyzed as a separate mechanism, never pooled with paired-perturbation cells.

F.5 Independent audits versus plan-centered operational scoring

For a predictive density g_q and CDF G_q, the observed-data log loss at threshold e is

$$S_e(q, O_e) = -\mathbf{1}\{Z < e\} \log g_q(Z) - \mathbf{1}\{Z = e\} \log\{1 - G_q(e)\},$$

where R is continuous latent demand, $Z = \min\{R, e\}$, and the threshold is conditionally independent of R given the audited information. It is proper for the censored law because its excess risk is the KL divergence between the true and reported censored distributions.

Under Route A the threshold is drawn from a report-independent law, $e \sim \nu(\cdot \mid x, \mathcal{H})$, and the censored score inherits the guarantees of Theorem B.2: expected regret is the ν-average of nonnegative KL terms, strictly positive whenever the threshold family identifies the belief over the sampled interval. The next result records the identification content and the cost of the randomization.

Proposition F.3 (Independent randomized threshold)**.** *Let $e \sim \nu(\cdot \mid x, \mathcal{H})$ be drawn independently of (q, w, y) with a known density giving positive mass to every open interval of a threshold range covering the allowed plans. Suppose the latent model is identified by its censored laws over every open threshold interval. Then the randomized audit is strictly identifying. If the principal's threshold cost $C(e)$ is twice continuously differentiable on the audit support around a stationary reference $y^\star$ fixed independently of the current messages and type, then*

$$\mathbb{E}[C(e) - C(y^\star)] = \frac{1}{2}\mathbb{E}\big[C''(\widetilde{y})(e - y^\star)^2\big], \tag{38}$$

which is $O(h^2)$ when ν is supported on an interval of half-width h around $y^\star$. Operational observations can be transported to this audit target only under conditional ignorability and overlap of the threshold laws; known propensities do not recover an unobserved tail by themselves.

Proof. Agreement of the censored laws on an interval implies agreement of their sub-threshold densities and tail probabilities there; the stated identification assumption then identifies the latent model. The expansion (38) is Taylor's theorem with the first-order term vanishing at the stationary plan. Independence from the messages is what places the design inside Theorem B.2. □

Under Route B the threshold is plan-centered, $e = y + \xi$ with ξ independent noise. The noise is independent of (q, w); the threshold is not, because $y = \Pi(q, w)$. The forecaster's expected score at belief p decomposes as

$$\mathbb{E}_\xi\big[H(K_{y+\xi}p) + \mathrm{KL}(K_{y+\xi}p \,\|\, K_{y+\xi}q)\big], \tag{39}$$

where H is the entropy of the truthful censored law. The KL term is minimized by truth at a fixed plan, but the entropy term varies with y and therefore enters the forecaster's effective plan value: the analysis proceeds with the augmented value $\widetilde{V}_{\mathrm{F}}(y) := V_{\mathrm{F}}(y) + \alpha\, \mathbb{E}_\xi H(K_{y+\xi}p)$, and Condition (T) must be imposed on $\widetilde{V}_{\mathrm{F}}$, not on V_{F} alone. Identification must likewise hold uniformly over the allowed plans, not merely at the realized one. With those two changes the plan-routed truthfulness argument goes through; without them, plan-centered scoring is a different mechanism from the independent audit, and claims proven for one do not transfer to the other. A plan-holding perturbation check cannot certify audit independence: Route B passes every intervention that holds y fixed, although its threshold law depends on y.

F.6 Proof of Theorem 6.2

The inequality $-U \le w - \iota\eta/\beta \le U$ is equivalent to $\beta(w - U)/\iota \le \eta \le \beta(w + U)/\iota$. Intersecting with $[-h, h]$ gives (20). For any random variable $Z \in [\ell, u]$, the pointwise inequality $(Z - \ell)(u - Z) \ge 0$ implies

$$\mathrm{Var}(Z) \le (\mathbb{E}Z - \ell)(u - \mathbb{E}Z) \le (u - \ell)^2/4.$$

Equal masses at ℓ and u attain equality. The same argument applies to the empirical distribution of a fixed design. For Gaussian observations $a_i = \theta + b_\eta \eta_i + \varepsilon_i$ with unknown intercept θ and independent $\varepsilon_i \sim \mathcal{N}(0, \sigma_a^2)$, eliminating the intercept from the information matrix leaves slope information $\sum_i (\eta_i - \bar{\eta}_n)^2/\sigma_a^2$, where $\bar{\eta}_n = n^{-1}\sum_i \eta_i$. Its maximum is $n(u-\ell)^2/(4\sigma_a^2)$ for even n. Ordinary least squares attains the inverse-information variance, proving optimality among unbiased estimators in this Gaussian model.

At a covered equilibrium, $\beta w^\star = G$ and the largest symmetric support has radius $(B - |G|)/\iota$. Substitute its squared radius into (18) to obtain (21). If a predeclared amplitude cap is present, the required amplitude must also be at most that cap. At $w = U$, the interval is $[0, h]$ when $h \le 2B/\iota$; at $w = -U$ it is $[-h, 0]$. Both have maximal variance $h^2/4$.

These statements concern the specified affine mean model within a stratum. Without linearity, an endpoint design identifies a finite contrast rather than a derivative. With varying unperturbed records, fixed effects or appropriate residualization must account for their intercepts; random assignment identifies an average response but does not eliminate residual heterogeneity from the variance. In E2, subtracting the published plan leaves the linear-alternative residual variance $\sigma_a^2 + \lambda_{\text{leak}}^2 \operatorname{Var}(q)$, because the unperturbed report is not additionally controlled for. Its power therefore need not equal a calculation using σ_a^2 alone.

G Learning-Bound Details

Let Y denote continuous latent demand and e an audit threshold. Let $O_e = (Z_e, I_e)$ denote the censored observation at threshold e: the observed quantity is $Z_e = \min\{Y, e\}$ and the censoring indicator is $I_e = \mathbf{1}\{Y \ge e\}$. If $g_f(\cdot \mid x)$ and $G_f(\cdot \mid x)$ are the candidate conditional density and distribution function, respectively, the observed-data log loss is

$$\widetilde{\ell}_f(x, e, O_e) := -(1 - I_e)\log g_f(Z_e \mid x) - I_e \log\{1 - G_f(e \mid x)\}, \qquad \ell_f := \min\{\widetilde{\ell}_f, \mathsf{B}_{\text{clip}}\}. \tag{40}$$

For raw and clipped target risks R^{raw} and R_{B}, assume the declared uniform clipping control

$$\sup_{f \in \mathcal{F}} |\mathsf{R}^{\text{raw}}(f) - \mathsf{R}_{\mathsf{B}}(f)| \le b_{\mathsf{B}}, \qquad b_{\mathsf{B}} \downarrow 0 \quad \text{as } \mathsf{B}_{\text{clip}} \uparrow \infty. \tag{41}$$

Write $f^\star$ for the target-risk minimizer. Strong identification means that, on the model class under consideration, some $\mu_{\text{id}} > 0$ satisfies

$$\mathsf{R}^{\text{raw}}(f) - \mathsf{R}^{\text{raw}}(f^\star) \ge \frac{\mu_{\text{id}}}{2} \|f - f^\star\|^2_{L^2(P_X)}. \tag{42}$$

The Gaussian location family below verifies this inequality locally by a uniform Fisher-information bound and globally by compact separation.

G.1 Uniform deviation and propagation

For clarity, take independent and identically distributed audit records $Z_i = (X_i, E_i, O_i)$, $i = 1, \dots, n$. Conditional on $X = x$, the logging threshold law has density $\pi(e \mid x)$ relative to a declared target probability measure $\nu(de \mid x)$, and $\pi \ge \pi_0 > 0$ on its support. The conditional law of O_i at (X_i, E_i) is the target audit law. If the logging law is itself the target law, then $\pi = 1$. Assume the raw loss has a uniform lower bound and shift it by a constant so that $\widetilde{\ell}_f \ge 0$; this does not change raw-risk minimizers. For clipping level $B > 0$, put

$$h_{f,B}(Z) = \frac{\min\{\widetilde{\ell}_f(Z), B\}}{\pi(E \mid X)}, \qquad \mathcal{G}_B = \{h_{f,B} : f \in \mathcal{F}\}.$$

Then $0 \le h_{f,B} \le B/\pi_0$ and $\mathbb{E} h_{f,B} = \mathsf{R}_B(f)$, the clipped risk under ν. Let $P_n h = n^{-1}\sum_i h(Z_i)$ and let

$$\mathfrak{R}_n(\mathcal{G}_B) = \mathbb{E}_{Z,\epsilon} \sup_{h \in \mathcal{G}_B} n^{-1} \sum_i \epsilon_i h(Z_i)$$

be the expected Rademacher complexity, where the ϵ_i are independent uniform signs. Assume the usual measurability of these suprema. Symmetrization and bounded differences, applied to the upper and lower deviations, imply that with probability at least $1 - \delta$,

$$\sup_f |P_n h_{f,B} - \mathsf{R}_B(f)| \le u_n := 2\mathfrak{R}_n(\mathcal{G}_B) + \frac{B}{\pi_0}\sqrt{\frac{\log(2/\delta)}{2n}}.$$

If $\widehat{f}_n$ minimizes $P_n h_{f,B}$ over a class containing $f^\star$, the empirical minimizer comparison and (41) give

$$\mathsf{R}^{\mathrm{raw}}(\widehat{f}_n) - \mathsf{R}^{\mathrm{raw}}(f^\star) \le r_n := 2u_n + 2b_B, \qquad \|\widehat{f}_n - f^\star\| \le \sqrt{2r_n/\mu_{\mathrm{id}}}.$$

Here μ_{id} is the strong-identification constant in (42). Substituting into (25) yields

$$|\mathcal{R}_{\mathrm{str}}(\widehat{f}_n) - \mathcal{R}_{\mathrm{str}}(f^\star)| \le K_1\sqrt{2r_n/\mu_{\mathrm{id}}} + K_2 r_n/\mu_{\mathrm{id}}, \tag{43}$$

where $K_1 = A_Y L_Y + \lambda_q \Delta_{\mathrm{inc}} L_\Gamma$ and $K_2 = L_Y^2 + \lambda_q L_\Gamma^2$. For increasing clipping levels B_n, consistency requires all three quantities $\mathfrak{R}_n(\mathcal{G}_{B_n})$, $B_n\sqrt{\log(2/\delta_n)/n}/\pi_0$, and b_{B_n} to vanish, with $\delta_n \to 0$. Clipping control alone is insufficient. Clipping is used to learn the audit target; a clipped score used as the agent's actual incentive need not remain strictly proper.

Proposition G.1 (Approximate equilibrium computation)**.** *In the scalar quadratic model, suppose a feasible learned pair $(\widetilde{q}, \widetilde{w})$ has expected forecaster objective at most $\varepsilon_{\mathrm{opt}}$ above the pointwise equilibrium objective at belief $\widehat{f}_n$. Then*

$$\|\widetilde{q} - \Gamma(\widehat{f}_n)\| \le \sqrt{2\varepsilon_{\mathrm{opt}}/\alpha}.$$

The report-error bounds in Theorem 7.1 therefore remain valid with $L_\Gamma \varepsilon_n$ replaced by $L_\Gamma \varepsilon_n + \sqrt{2\varepsilon_{\mathrm{opt}}/\alpha}$.

Proof. Minimization over requests preserves convexity of the operational value. Adding $\alpha(q - \widehat{f}_n)^2/2$ makes the reduced objective α-strongly convex. Its objective gap lower-bounds $\alpha|\widetilde{q} - \Gamma(\widehat{f}_n)|^2/2$ pointwise. Integrate and apply the triangle inequalities. □

This statement requires an objective-gap certificate. Training a nonconvex network for a fixed number of steps alone supplies no such certificate.

Proposition G.2 (Bounded score weight forces a report floor)**.** *In the scalar quadratic model of Section 4.1 with a single report channel and no request channel ($y = q$ up to a public constant), suppose the score weight is at most $\bar{\alpha}$ and the type's plan displacement satisfies $|G| \ge \underline{\mathrm{g}} > 0$. Then every interior equilibrium report satisfies $|q^\star - p| \ge c_{\mathrm{F}}\underline{\mathrm{g}}/(\bar{\alpha} + c_{\mathrm{F}})$.*

Proof. Without a request channel, the forecaster minimizes $\alpha(q-p)^2/2 + c_{\mathrm{F}}(q - p - G)^2/2$ over interior q, giving $q^\star - p = c_{\mathrm{F}}G/(\alpha + c_{\mathrm{F}})$. The bound follows from $\alpha \le \bar{\alpha}$ and $|G| \ge \underline{\mathrm{g}}$. □

This is a statement about the scalar quadratic class; it does not assert a numerical distortion invariant across model classes or report alphabets.

G.2 Strong identification for a censored Gaussian family

Let $Y \sim \mathcal{N}(\mu, v^2)$ with known $v > 0$, and observe $(\min\{Y, e\}, \mathbf{1}\{Y \ge e\})$. Write $z = (e - \mu)/v$, let ϕ and Φ be the standard normal density and distribution function, and set $\overline{\Phi} = 1 - \Phi$. Direct calculation gives Fisher information for μ equal to

$$I(\mu, e, v) = v^{-2}\left[\Phi(z) - z\phi(z) + \frac{\phi(z)^2}{\overline{\Phi}(z)}\right] > 0.$$

Indeed, the uncensored location score is $(Y - \mu)/v^2$ and the censored score is $\phi(z)/(v\overline{\Phi}(z))$. Their squared expectations give the two displayed contributions, using $\int_{-\infty}^{z} t^2\phi(t)\,dt = \Phi(z) - z\phi(z)$. For true and candidate means in a compact interval, thresholds in a compact interval, and known variance in $[\underline{v}^2, \overline{v}^2] \subset (0, \infty)$, this information is continuous and uniformly positive. The censored Gaussian family is identifiable at every finite threshold. A uniform Taylor expansion of its KL divergence near the true mean and compact separation away from it therefore give a uniform quadratic lower bound. Applying the pointwise bound conditionally on X and integrating proves (42) for this correctly specified location class. If thresholds tend to $-\infty$, the information tends to zero; identification at each fixed threshold does not give a uniform learning rate over arbitrarily severe censoring.

The raw Gaussian log loss is unbounded, which is why (40) is the score used in the learning bound. To verify (41), restrict true and candidate locations, thresholds, and variances to the compact sets above. The stockout loss $-\log \overline{\Phi}((e-\mu)/v)$ is then uniformly bounded. On a non-stockout, the shifted negative log density is bounded above and below by affine functions of $1+|Y|^2$. Uniform Gaussian tail bounds therefore give

$$\sup_f \mathbb{E}[(\widetilde{\ell}_f - \mathsf{B}_{\mathrm{clip}})_+] \le C(1+\mathsf{B}_{\mathrm{clip}})e^{-c\mathsf{B}_{\mathrm{clip}}} \le C' e^{-c'\mathsf{B}_{\mathrm{clip}}}$$

for constants $C, C', c > c' > 0$. This is exactly the uniform clipping-bias order used in the main theorem; no bounded-loss rate is attributed to the raw score.

G.3 Doubly robust target risk

Let $m_f(x,e) = \mathbb{E}[\ell_f(x,e,O_e) \mid X = x]$ and let $\widehat{m}_f, \widehat{\pi}$ be cross-fitted nuisance estimates, with $\widehat{\pi}$ a density relative to the same target measure ν. A doubly robust criterion is

$$\widehat{\mathsf{R}}_{\mathrm{DR}}(f) = \frac{1}{n}\sum_{i=1}^{n}\Bigg[\int \widehat{m}_f(X_i,e)\,\nu(de \mid X_i) + \frac{\ell_f(X_i,E_i,O_i) - \widehat{m}_f(X_i,E_i)}{\widehat{\pi}(E_i \mid X_i)}\Bigg].$$

Its population bias is zero if either the logging density or the conditional loss regression is correct, subject to overlap and cross-fitting regularity. In the principal-owned randomized-threshold design the propensity is known, but augmentation can still reduce variance.

G.4 Why complete cases are biased

In the following two paragraphs, write R for latent demand, with conditional mean $\mu(x)$ and standard deviation $\sigma(x)$. Dropping stockouts conditions on $R < e$, whose probability depends on the threshold and latent demand. Even in a Gaussian model,

$$\mathbb{E}[R \mid R < e, X = x] = \mu(x) - \sigma(x)\frac{\phi((e-\mu(x))/\sigma(x))}{\Phi((e-\mu(x))/\sigma(x))},$$

which is not $\mu(x)$ and varies with the deployed policy. The induced bias is a population selection effect and therefore does not shrink with sample size. Censored likelihood or off-policy observed-data risk is required.

G.5 Censoring-corrected forecast value added

Corollary G.3 (Censoring correction restores the target)**.** *Let $\chi \in \{0,1\}$ index the audit target. Squared-error scoring on observed sales $Z = \min\{R,e\}$ ($\chi = 0$) targets the law of Z. If $\chi = 1$, replace Z by*

$$Z + \mathbb{E}[(R-e)_+ \mid R \ge e, X, e]\,\mathbf{1}\{R \ge e\}.$$

Assume R is conditionally independent of the threshold e given X, and the imputation law is fixed independently of the candidate report. Whenever this conditional expectation is correctly specified, the corrected target has conditional mean $\mathbb{E}[R \mid X]$ and removes the action-dependent drift. The observed-data censored likelihood in (40) instead targets the identifiable distribution. Neither conditional-mean imputation nor clipping is an identity for the full raw squared-error risk.

Proof. Since $R - Z = (R-e)_+$ and this difference is zero off the censoring event, iterated expectations show that the displayed correction has conditional mean $\mathbb{E}[R \mid X]$. Properness of the censored likelihood follows from its KL-risk representation. □

H Capacity-Constrained Multi-Item Extension

This appendix gives the analytic fixed-face calculation behind the multi-item experiment of Section 8 (E4), which adds active-set changes, nonnegativity, and heterogeneous preferences numerically.

For d products, let the forecaster report $q \in \mathbb{R}^d$ and let a fixed binding capacity face be $c^\top a = \mathsf{B}$, with $c \in \mathbb{R}^d_{++}$. The local quadratic operator problem is

$$\min_a \left\{ \tfrac{1}{2}(a - a_{\mathrm{O}})^\top H(a - a_{\mathrm{O}}) + \tfrac{\lambda_{\mathrm{cmp}}}{2}\|a - q\|_2^2 \right\} \quad \text{subject to } c^\top a = \mathsf{B}, \tag{44}$$

where $a_{\mathrm{O}} \in \mathbb{R}^d$ is the operator's preferred action, $H \succ 0$ its curvature, and $\lambda_{\mathrm{cmp}} > 0$ its tracking weight. Write $A(q)$ for the unique minimizing action.

Proposition H.1 (Capacity creates cross-item pass-through)**.** *With $K = H + \lambda_{\mathrm{cmp}} I$, the derivative on the fixed binding face is*

$$D_q A = \lambda_{\mathrm{cmp}} \left[K^{-1} - K^{-1}c(c^\top K^{-1} c)^{-1} c^\top K^{-1} \right], \qquad c^\top D_q A = 0. \tag{45}$$

If $K = \mathrm{diag}(k_1, \ldots, k_d)$, then for $i \neq j$,

$$\frac{\partial A_i}{\partial q_j} = -\lambda_{\mathrm{cmp}} \frac{(c_i/k_i)(c_j/k_j)}{\sum_{\ell=1}^d c_\ell^2/k_\ell} < 0. \tag{46}$$

Interpretation: on a binding shared-capacity face, increasing one item's forecast raises its own allocation by reducing other items' allocations.

Proof. The Lagrangian is $\mathcal{L} = \frac{1}{2}(a - a_{\mathrm{O}})^\top H(a - a_{\mathrm{O}}) + \frac{\lambda_{\mathrm{cmp}}}{2}\|a - q\|_2^2 + \pi_{\mathrm{cap}}(c^\top a - \mathsf{B})$. Differentiating the KKT equations gives

$$K D_q A + c D_q \pi_{\mathrm{cap}} = \lambda_{\mathrm{cmp}} I, \qquad c^\top D_q A = 0.$$

Solving the second equation after premultiplying the first by $c^\top K^{-1}$ and substituting proves (45); diagonal K gives (46). □

Let q^0 be the audit-optimal report. With forecaster report loss $\ell_F(q) = \frac{1}{2}\|q - q^0\|_2^2$ and continuation $W_F(A(q)) = -\upsilon^\top A(q)$, where $\upsilon \in \mathbb{R}^d$ is an action-value vector and $\gamma_F \geq 0$ its continuation weight, the operator response is affine on a fixed face, so the forecaster's objective is an exact quadratic in q and the best response is

$$q^\star - q^0 = \gamma_F \, (D_q A)^\top \upsilon, \tag{47}$$

the projected strategic gradient. In the binding regime, $c^\top[A(q^\star) - A(q^0)] = 0$: strategic reports reallocate orders rather than increase total capacity.

H.1 Direction-filtered pass-through

The same fixed-face model can filter the strategic component of cross-item sensitivity. Let $B_0 = D_q A$ be (45), estimated with principal-owned report perturbations, and let $\widehat{g}_F$ be the measured miscalibration direction relative to an audit-identified report. After a fixed change of coordinates making the principal's local quadratic metric the identity, define $\widetilde{g} = B_0(B_0^\top B_0)^\dagger \widehat{g}_F$, the minimum-norm solution to $B_0^\top \widetilde{g} = \mathrm{proj}_{\mathrm{range}(B_0^\top)} \widehat{g}_F$, where $\dagger$ denotes the Moore–Penrose inverse, and set

$$B^\perp = \left[I - \frac{\widetilde{g}\widetilde{g}^\top}{\|\widetilde{g}\|^2} \right] B_0, \tag{48}$$

whenever $\widetilde{g} \neq 0$.

Proposition H.2 (Direction filter)**.** *$B^\perp$ satisfies $(B^\perp)^\top \widetilde{g} = 0$ and coincides with B_0 on every report direction υ satisfying $B_0 \upsilon \perp \widetilde{g}$. Only the projection of the strategic direction onto* $\mathrm{range}(B_0)$ *is identifiable from report perturbations; the zero-pass-through benchmark sets $B_0 = 0$.*

Proof. $(B^\perp)^\top \widetilde{g} = B_0^\top \widetilde{g} - B_0^\top \widetilde{g}\,(\widetilde{g}^\top \widetilde{g})/\|\widetilde{g}\|^2 = 0$, and when $\widetilde{g}^\top B_0 \upsilon = 0$, the correction to $B_0 \upsilon$ vanishes. Identifiability is immediate: report perturbations identify B_0 exactly, and $\widehat{g}_F$ enters only through $\widetilde{g}$. □

This is a local construction on a fixed face. It does not imply global truthfulness when the operational gradient or active face changes. E4 evaluates simpler matrix filters in its affine *plan* map and then projects plans onto physical capacity. The forecaster's objective in that experiment is quadratic in the pre-projection plan; the fixed-face calculation above instead concerns preferences over the downstream action. The two models should not be conflated.

I RETROFIT DETAILS

I.1 IMPLEMENTATION

The retrofit can be implemented as the following deterministic transformation of an institutional specification.

Algorithm 1 Retrofit procedure

Require: Protocol Π, report set $\mathcal{Q}$, request set $\mathcal{W}$, plan cost cost, physical rules, audit rule, and operational score weight.

1: Create an audit-only store for (q, w) and expose only public state, history, and published plan to operational agents.
2: Replace every physical read of (q, w) by a read of $(x, \mathcal{H}, y)$; reject compilation if the corresponding data-flow edge cannot be removed.
3: Either set $\Pi^{\mathrm{rec}}(q, w) = w$ and $\mathcal{W}^{\mathrm{rec}} = \bigcup_q \mathcal{P}(q)$, or solve the completion problem below within the desired affine protocol family.
4: Set $\mathrm{cost}^{\mathrm{rec}}(y) = \inf\{\mathrm{cost}(q, w) : \Pi(q, w) = y\}$.
5: Install a report-independent audit law, remove operator-selected scoring, and assign threshold randomization to the principal.
6: **return** The retrofitted specification and its enlarged physical request set.

The explicit rejection in line 2 matters: retrofitting is an institutional change, not a relabeling of a channel that remains visible downstream.

I.2 THE RETROFIT PROCEDURE AND MINIMUM COMPLETION

An existing pipeline may expose raw forecasts, attach costs to messages, or score on operator-selected observations. The retrofit: (i) seal raw messages and expose only the plan (P1); (ii) complete the request set to preserve required plans; (iii) attach operational costs and constraints only to the plan (P2); (iv) precommit a principal-owned identifying audit. Suppose reports lie in $[\underline{q}, \overline{q}]$ and every plan in $[\underline{y}, \overline{y}]$ must remain reachable under every report.

Proposition I.1 (Minimum override-set completion)**.** *For the affine protocol (5), the smallest such budget is*

$$U^{\mathrm{comp}}(\iota, \beta) = \frac{1}{\beta} \max\left\{\iota\overline{q} + (1-\iota)\overline{v} - \underline{y}, \overline{y} - \iota\underline{q} - (1-\iota)\overline{v}, 0\right\}. \tag{49}$$

It is finite for every $\iota \in [0, 1]$ and $\beta > 0$.

Intersecting the radius-βU intervals over reports gives (49); the polyhedral extension is below, and the price of compilation (52) — the change in action regret minus the original miscalibration — can be unbounded when report bias accidentally cancels operator bias (the cancellation vanishes after operator correction).

I.3 GENERAL POLYHEDRAL COMPLETION

Consider

$$\Pi(q, w) = Cq + Ew + c, \qquad \mathcal{W}(t) = \{w : Hw \le r + t\mathbf{1}\},$$

Here C, E, c define the affine plan map, H has one row for each request inequality, and r collects its thresholds; $t \ge 0$ is a common relaxation of the request faces. For compact polyhedra $\mathcal{Q}$ and $\mathcal{N}$, define

$$t^{\mathrm{comp}} := \sup_{y \in \mathcal{N},\ q \in \mathcal{Q}} \inf_{w : Ew = y - Cq - c} \max_j (Hw - r)_j. \tag{50}$$

The convention is that the inner value is $+\infty$ when the equality is infeasible.

Proposition I.2 (Exact polyhedral completion)**.** *The smallest common face relaxation for which $\mathcal{N} \subseteq \bigcap_q \Pi(q, \mathcal{W}(t))$ is $\max\{0, t^{\mathrm{comp}}\}$. If E is square and nonsingular, this becomes the maximum of finitely many linear-program values (one for each face j)*

$$\max_{y \in \mathcal{N}, q \in \mathcal{Q}, j} \left[HE^{-1}(y - Cq - c) - r\right]_j.$$

For the scalar affine influence protocol it reduces to Proposition I.1.

Interpretation: the formula gives the smallest uniform relaxation of the request constraints that preserves every required plan under every report.

Proof. For fixed (y, q), plan y is reachable under report q and budget t if and only if there is a solution of $Ew = y - Cq - c$ satisfying $Hw - r \le t\mathbf{1}$. The smallest such t is the inner optimization in (50). Requiring this reachability for every $(y, q) \in \mathcal{N} \times \mathcal{Q}$ takes the supremum, and $t \ge 0$ takes the maximum with zero. Nonsingular E makes the feasible request unique. For $C = \iota$, $E = \beta$, and the two scalar box faces, maximizing the two required relaxations gives exactly (49). □

I.4 A SHARED CAPACITY FACE

Let the multi-item request set be $\mathcal{W}_{\mathsf{K}} = \{w : \underline{w} \le w \le \overline{w}, c^\top w \le \mathsf{K}\}$ and let $\Pi_{\iota,\beta}(q, w) = \iota q + (1 - \iota)\bar{v} + \beta w$ coordinatewise. Provided the box faces contain $(y - \iota q - (1 - \iota)\bar{v})/\beta$ for every $(y, q) \in \mathcal{N} \times \mathcal{Q}$, the exact minimal shared capacity is

$$\mathsf{K}^{\mathrm{comp}}(\iota, \beta) = \sup_{y \in \mathcal{N}, q \in \mathcal{Q}} \frac{c^\top (y - \iota q - (1 - \iota)\bar{v})}{\beta}. \tag{51}$$

If the box condition fails, no enlargement of the shared face alone completes the core. Expression (51) is the completion counterpart of the fixed-face pass-through calculation in Appendix H.

I.5 DYNAMIC PERFORMANCE IDENTITY

With $\mathrm{Reg}^j_{\mathrm{act}}$ the expected discounted action regret under pipeline j, evaluated under its own state distribution, and $\mathcal{M}^{\mathrm{orig}}$ the original miscalibration, the price of compilation is

$$\mathcal{L}^{\mathrm{retro}} - \mathcal{L}^{\mathrm{orig}} = \underbrace{(\mathrm{Reg}^{\mathrm{retro}}_{\mathrm{act}} - \mathrm{Reg}^{\mathrm{orig}}_{\mathrm{act}})}_{\Delta \mathrm{Reg}_{\mathrm{act}}} - \mathcal{M}^{\mathrm{orig}}. \tag{52}$$

Let periods run from $k = 0$ to K. Let s_k be the principal's payoff-relevant state, $c_{\mathrm{P},k}(s, a)$ its resource cost, and $\gamma_{\mathrm{P}} \in (0, 1]$ the discount factor. Set $V^\star_{\mathrm{P},K+1} = 0$ and define

$$V^\star_{\mathrm{P},k}(s) = \min_a \{c_{\mathrm{P},k}(s, a) + \gamma_{\mathrm{P}} \mathbb{E}[V^\star_{\mathrm{P},k+1}(S') \mid s, a]\}.$$

The Bellman advantage is

$$\mathrm{Adv}_{\mathrm{P},k}(s, a) := c_{\mathrm{P},k}(s, a) + \gamma_{\mathrm{P}} \mathbb{E}[V^\star_{\mathrm{P},k+1}(S') \mid s, a] - V^\star_{\mathrm{P},k}(s) \ge 0.$$

For a pipeline j, define

$$\mathrm{Reg}^j_{\mathrm{act}} := \mathbb{E}_j \sum_{k=0}^{K} \gamma^k_{\mathrm{P}} \mathrm{Adv}_{\mathrm{P},k}(s_k, a^j_k), \qquad \mathcal{M}^j := \mathbb{E}_j \sum_{k=0}^{K} \gamma^k_{\mathrm{P}} m_k(p_k, q^j_k),$$

where $m_k(p, q) \ge 0$ and $m_k(p, p) = 0$.

Theorem I.3 (Dynamic performance-difference identity)**.** *For pipelines with a common initial state law and the same controlled physical dynamics, define $\mathcal{L}^j$ as expected discounted resource cost plus $\mathcal{M}^j$ and $\mathcal{L}^{\mathrm{FB}} = \mathbb{E}[V^\star_{\mathrm{P},0}(s_0)]$. Then for every pipeline j,*

$$\mathcal{L}^j - \mathcal{L}^{\mathrm{FB}} = \mathrm{Reg}^j_{\mathrm{act}} + \mathcal{M}^j.$$

If the retrofit is truthful, so $\mathcal{M}^{\mathrm{retro}} = 0$, then

$$\mathcal{L}^{\mathrm{retro}} - \mathcal{L}^{\mathrm{orig}} = \mathrm{Reg}^{\mathrm{retro}}_{\mathrm{act}} - \mathrm{Reg}^{\mathrm{orig}}_{\mathrm{act}} - \mathcal{M}^{\mathrm{orig}}. \tag{53}$$

Each expectation is taken under that pipeline's own state distribution.

Proof. Multiply the Bellman-advantage identity by γ^k_{P}, sum over k, and take expectations under pipeline j. The continuation terms telescope, leaving expected resource cost minus $\mathbb{E}[V^\star_{\mathrm{P},0}(s_0)]$. Adding $\mathcal{M}^j$ proves the first identity; subtracting its two instances and using retrofit truthfulness proves (52). □

Proposition I.4 (Unbounded retrofit cost through accidental cancellation)**.** *There are one-period scalar quadratic instances with $\lambda_q = 0$ for which a truthful plan-only retrofit increases principal loss by an arbitrarily large amount.*

Proof. Fix p and an operator offset Δ_{O}. In the original protocol, $y = q$ and the operator executes $a = y + \Delta_{\mathrm{O}}$. Take $\alpha = c_{\mathrm{F}} > 0$ and let the forecaster prefer action $p + \Delta_{\mathrm{F}}$ with $\Delta_{\mathrm{F}} = -\Delta_{\mathrm{O}}$. Its cost is

$$\tfrac{\alpha}{2}(q - p)^2 + \tfrac{\alpha}{2}(q + \Delta_{\mathrm{O}} - p - \Delta_{\mathrm{F}})^2.$$

The optimum is $q = p - \Delta_{\mathrm{O}}$, so $a = p$ and principal action loss is zero. Retrofit to $y = w$ with a report-independent request interval large enough to contain $p - 2\Delta_{\mathrm{O}}$, retaining the same operator response. The forecaster reports $q = p$ and chooses $w = p - 2\Delta_{\mathrm{O}}$, giving $a = p - \Delta_{\mathrm{O}}$ and principal loss $\Delta_{\mathrm{O}}^2/2$. All operational costs depend on y, and the audit is unchanged. Taking $|\Delta_{\mathrm{O}}| \to \infty$ proves the claim. The retrofit preserves available actions, not the action selected in equilibrium; it removes an accidental cancellation of preferences. □

Corollary I.5 (Equal action paths)**.** *If an original and a truthful retrofitted protocol induce the same law of physical state-action paths, their action regrets coincide and the retrofit changes total loss by $-\mathcal{M}^{\mathrm{orig}} \leq 0$.*

Proof. Apply Theorem I.3 under the common state-action law. □

An operator that corrects for report distortion changes the game itself. Whether it produces these equal action laws requires recomputing the forecaster's response to that correction.

J EXPERIMENTAL DETAILS AND REPRODUCIBILITY

The command `.venv/bin/python experiments/run_all.py` runs E0–E5 from the repository root using `experiments/requirements.txt`. The master seed is 20260922, and each suite uses deterministically named random streams. Numerical results are reproducible up to floating-point and solver differences; elapsed times, absolute output paths, and PDF metadata need not be identical. E0 compares closed forms with direct numerical optimization. E1–E4 use those formulas to generate mechanism responses; their learning data and intervention noise are sampled separately.

J.1 E0: SCALAR AND VECTOR VALIDATION

The 2,000 scalar configurations draw $p \sim \mathcal{N}(100, 10^2)$, $\bar{v} \sim \mathcal{N}(100, 5^2)$, $\Delta_{\mathrm{F}} \sim \mathrm{Unif}[-8, 8]$, $\iota = 0$ with probability $1/4$ and otherwise $\iota \sim \mathrm{Unif}[.05, 1]$, $\beta \sim \mathrm{Unif}[.3, 3]$, and independent $\alpha, c_{\mathrm{F}} \sim 10^{\mathrm{Unif}[-1,1.2]}$. Set $B = \theta|G|$ with $\theta \sim 10^{\mathrm{Unif}[-1,1.3]}$. Half impose $q \in [p - r, p + r]$ with $r \sim \mathrm{Unif}[.5, 5]$. Independent multi-start L-BFGS-B solves the direct report–request objective. The vector suite uses 200 instances with dimensions 2–6, positive definite random curvatures, and tight, wide, or mixed box request sets. Most report matrices have full row rank; rank-deficient examples and redundant request directions are added in E5. Outputs are `e0_validation.json` and `e0_discrepancies.csv`. Tolerances and optimizer settings are specified in `e0_validation.py`.

Table 4: Analytic regression targets at $(\iota, \beta, \alpha, c_{\mathrm{F}}, \Delta_{\mathrm{F}}, \lambda_q) = (1, 1, 4, 2, -8, 1)$. Symmetric eligibility uses amplitude .25. Inward tests are feasible at $U = 4$; $U = 0$ requires a separate shadow interface.

U	$\lvert q^\star - p\rvert$	$w^\star$	$y^\star - p$	$\mathcal{L}$	symmetric	inward
0	$8/3$	0	$-8/3$	$64/9$	0	0
4	$4/3$	-4	$-16/3$	$136/9$	0	1
12	0	-8	-8	32	1	1

J.2 E1: A FINITE DESIGN SEARCH

Let $p - \bar{v} \sim \mathrm{Unif}[-s, s]$ with fixed $\bar{v} = 0$ and $s \in \{.25, 1, 4\}$. Cross $\Delta_{\mathrm{F}} \in \{-2, -.5, .5, 2\}$ and $\alpha/c_{\mathrm{F}} \in \{.25, 1, 4, 16\}$ with $c_{\mathrm{F}} = 1$, giving 48 cells. Four additional cells use the equally weighted mixture of the offsets at $s = 1$. The design grid is $\iota \in \{0, .1, .25, .5, .75, 1\}$ and $B \in \{0, .1, .25, .5, 1, 2, 4, 8, 16\}$, with $\lambda_q = 1$ and $C_{\mathrm{authority}}(B) = .05B$. Expectations use a 2001-point midpoint quantile grid. Thus "optimized" means best on the declared grid with a fixed benchmark, not a global optimum of (8).

The nine comparator keys in the code represent: (1) a single report with $B = 0$; (2) a report-anchored band; (3) a benchmark-anchored band; (4) no influence and no authority; (5) the grid optimum; (6) an idealized known-utility compensation benchmark with $q = y = p$; (7) a residual-preference surrogate with offset $-.5$ and $\iota = 0$; (8) a report-anchored protocol with an iterated mean-bias correction; and (9) full information. Comparator 7 retains the data key `misspecified_compensation`; no transfer mechanism is implemented for that row. Comparator 6 excludes transfer costs and is an oracle benchmark. Mean-bias correction is recomputed with the forecaster's response, and its convergence diagnostics are recorded. Outputs are `e1_frontier.json` and `e1_frontier_cells.csv`.

J.3 E2: FEASIBLE RANDOMIZED ROUTING TESTS

Draw $p \sim \mathcal{N}(100, 7^2)$ with $(\bar{v}, \Delta_{\mathrm{F}}, \alpha, c_{\mathrm{F}}, \iota, \beta) = (100, -8, 4, 2, 1, 1)$. Use $U \in \{0, 4, 12\}$, Gaussian action noise standard deviation .8, amplitude .25, eligible sample counts $250, 500, 1000, 2000$, and 1,000 repetitions per cell. Designs are symmetric paired perturbations, inward boundary perturbations, a multi-amplitude design with a zero arm, and a separate zero-budget shadow interface. Alternatives are no direct message effect, linear report, direct request, quadratic report, variance-only, one-period-delayed, and binding-only effects. The effect strength grid is $0, .05, .1, .25$.

The primary regression includes an intercept and subtracts the published plan from the action. Under linear report effects, this removes the plan component but leaves random unperturbed-report variation; its residual variance is larger than the action-noise variance alone. The test uses a normal critical value of 1.96, so finite-sample coverage is assessed by simulation. Eligibility is decided before randomization. Zero eligible counts are reported as infeasible tests, not zero rejection probabilities. A reported "power per total period" in the result file is the descriptive product of conditional rejection probability and eligible fraction; it is not the power of a test with a fixed total-period budget. Outputs are `e2_leakage.json` and `e2_cells.csv`.

J.4 E3: LEARNED FORECASTS AND CENSORED EVALUATION

Draw $x_1 \sim \mathrm{Unif}[0, 2\pi]$ and $x_2 \sim \mathcal{N}(0, 1)$ independently, with $p(x) = 100 + 10 \sin x_1 + 3x_2$ and demand noise $\mathcal{N}(0, 12^2)$. A second noise law is $.8\mathcal{N}(0, 12^2) + .2\mathcal{N}(6, 24^2)$ shifted by -1.2 to preserve the conditional mean. Training uses independent *uncensored* audit labels. Training sample sizes are $256, 1024, 4096, 16384$; evaluation uses $50{,}000$ held-out covariates.

Track 1 fits ridge regression on $[1, \sin x_1, x_2]$ with penalty 10^{-6} and 30 seeds. Track 2 uses two hidden tanh layers of width eight, report and clipped-request heads, and full-batch Adam for $2{,}000$ steps with learning rate .01, over ten seeds. Its empirical objective is score cost plus quadratic plan cost; no global optimization certificate is assumed. Parameters are $\alpha = 4, c_{\mathrm{F}} = 2, \Delta_{\mathrm{F}} = -8, \beta = 1, \lambda_q = 1$, and $\bar{v} = 100$. The principal arms are a report-anchored band with $U = 4$, a benchmark-anchored band with $U = 12$, an oracle $q = y = f$ comparator, and single-report influence with $U = 0$. A residual-preference surrogate changes the offset to -2 at $\iota = 0, U = 12$; it is truthful in reports and retains plan error. Its code key is `misspecified_compensation`; it does not exhibit a report floor.

The benchmark-anchored arm clips preferred plans on approximately 36% of covariates. Its operational loss is approximately 24.2, while the hypothetical fully covered plan would cost $(-8)^2/2 = 32$. For the binding report-anchored arm, the population decomposition is $136/9 = 0 + 64/9 + 8$.

The separate censoring ablation fixes a common threshold $e = p(x) - 8$ for all candidate forecasts and observes $Z = \min\{R, e\}$, where R is demand. This is a fixed-policy evaluation experiment, not training under a candidate-selected threshold. At a censored observation it substitutes the modeled conditional squared residual $\mathbb{E}[(R - f)^2 \mid R \geq e, X]$. Oracle, estimated-mean, split-fit-averaged mean, and misspecified-variance corrections are compared. The split-fit-average curve (data key `cross_fitted`) averages two disjoint-half fits and evaluates on an independent held-out sample; it is not the doubly robust cross-fitting estimator of Appendix G. Forecast value added is the score of truth minus the candidate's score, so positive values favor the candidate. Outputs are `e3_learning.json` and `e3_cells.csv`.

J.5 E4: PLANS FOLLOWED BY CAPACITY-CONSTRAINED EXECUTION

There are 20 items and 30 paired seeds. Draw benchmark means $\bar{v}_i \sim \mathrm{Unif}[80, 120]$ and private means $p_i = \bar{v}_i + \nu_i$ with $\nu_i \sim \mathcal{N}(0, 4^2)$; offsets belong to $\{-4, -1, 1, 4\}$. Demand standard deviations are uniform on $[8, 16]$. The forecaster has $H_q = 4I$ and diagonal H_y entries uniform on $[1, 3]$, and optimizes the quadratic cost of the published plan. The operator then executes its Euclidean projection onto $\{a \geq 0 : r^\top a \leq K\}$, with $r_i \sim \mathrm{Unif}[.5, 1.5]$ and $K = 1.5 r^\top p$ or $.7 r^\top p$. Thus the forecaster does not internalize the capacity projection in this experiment.

Principal cost is Gaussian newsvendor loss with holding costs uniform on $[1, 3]$, shortage costs uniform on $[4, 9]$, and regularization $.02\|a\|^2/2$. The benchmark solves the exact Gaussian expected-cost problem; reported costs use a shared 512-draw normal grid, so differences from this benchmark are numerical regret estimates. Direct numerical validation of the quadratic plan objective uses ten seeds. The deviation check is a coordinate grid with requests held fixed; joint global optimality of the affine quadratic rows follows from convexity and the vector theorem, not from this deviation grid alone.

The native and report-relative rows coincide ($C = E = I$, request radius six). Two zero-influence rows use radii six and two. The uncentered filter sets $C = I - dd^\top$, with $d = \Delta_{\mathrm{F}}/\|\Delta_{\mathrm{F}}\|$, but retains $\mathsf{c} = \bar{v}$. The centered filter instead uses $\mathsf{c} = (I - C)\bar{v}$, making the plan center $\bar{v} + C(q - \bar{v})$. This distinguishes filtering from a large unintended benchmark shift. Neither filter is optimized

over all matrices. The known-utility comparator executes the principal's optimal action. The `compensation_miss` row is a frozen-response quadratic surrogate followed by a new principal best response; it is not a Stackelberg compensation equilibrium. Its formula is explicitly recorded in the code and result metadata. Sensitivity measures refer to the plan before projection for the affine rows. Outputs are `e4_multiitem.json` and `e4_cells.csv`.

J.6 E5: CONVEXITY, STABILITY, AND INFORMATION

The 160 quadratic instances have report dimension three, plan dimension four, and request dimension five. One request column duplicates another; half the report matrices have a zero last row. Requests have box radius .05 or .6. Gaussian random matrices generate positive definite score and plan metrics by adding the identity to a Gram matrix. A direct joint solver is compared with the projection certificate bounds and with the stability inequality after a target perturbation. Tolerance is 2×10^{-5}.

For 180 nonquadratic instances, let $C = [I_2; 0]$, $E = I_3$, and $\mathcal{W} = \{w : |w_i| \le .3,\ w_1 + w_2 \le .4\}$. Writing $d = q - p$ and $e = Cd + w - g$, the objective is $\alpha \sum_i (d_i^2/2 + .025 d_i^4) + \sum_j (e_j^2/2 + .05 e_j^4)$, with $\alpha \in \{.2, 1, 8\}$. The nonquadratic score regret is realizable by independent Gaussian audit coordinates $O_i \sim \mathcal{N}(p_i, 1)$ and loss $S(q, O) = \sum_i [.35(q_i - O_i)^2 + .025(q_i - O_i)^4]$: subtracting its truthful expectation gives $\sum_i (d_i^2/2 + .025 d_i^4)$. Half the instances place the first two target coordinates in $[-.12, .12]$ and the third in $[1.5, 2.5]$; the other half set the first target coordinate to 1.5. The former are truthful with an uncovered but report-inaccessible third coordinate, and the latter are biased. SLSQP solves the joint constrained objective.

A further 30,000 pairs of beliefs verify the global scalar Lipschitz bounds. The information experiment uses 20,000 Gaussian repetitions, $n = 200$, noise standard deviation .8, and amplitude .25. The empirical symmetric slope variance is .0501 versus the analytic .0512; rescaling the centered design from $\{-h, h\}$ to $\{0, h\}$ gives exactly the predicted factor four. Figure 2's sensitivity panel uses $\iota = .6, \alpha = 1, c_{\mathrm{F}} = 2, \Delta_{\mathrm{F}} = 1, B = .5, \bar{v} = 0$. Its power panel uses $\iota = 1, G = 1, h = .4, n = 200, \sigma_a = .8$ and effect .5 at level .05. Additional deterministic checks verify the nonattainment sequence in Proposition D.18 and the censored Gaussian information formula by finite differences of KL divergence. Outputs are `e5_geometry.json` and `e5_geometry.pdf`.

J.7 ADDITIONAL EXPERIMENTAL FIGURES

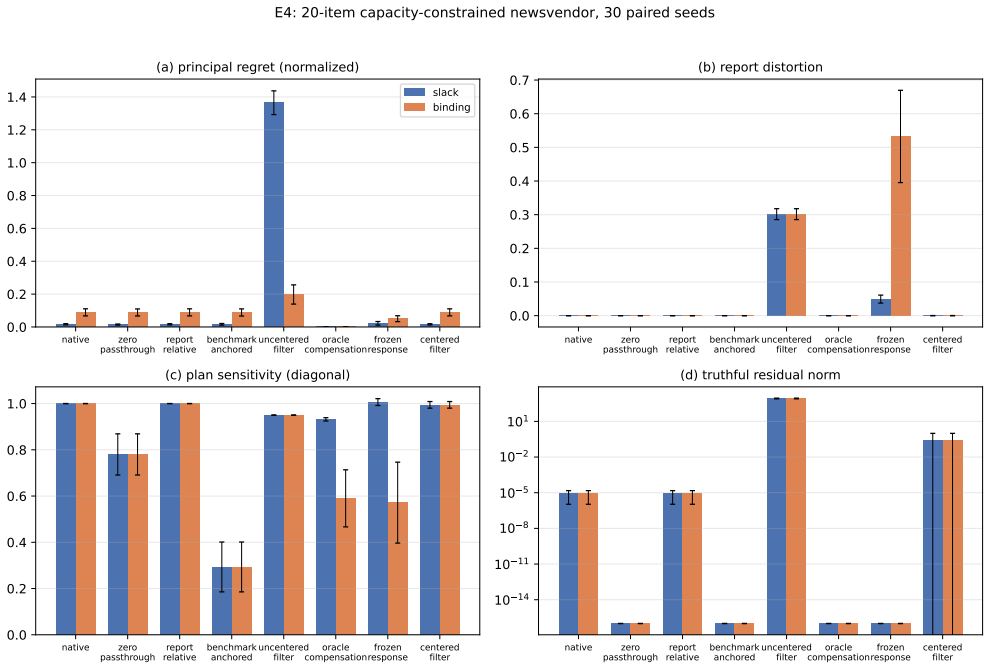


Figure 4: Twenty-item results. Error bars show standard deviations across 30 seeds. Plans are optimized before physical capacity projection. The two filter rows separate direction filtering from benchmark centering; the frozen-response row is a surrogate, not a compensation equilibrium.

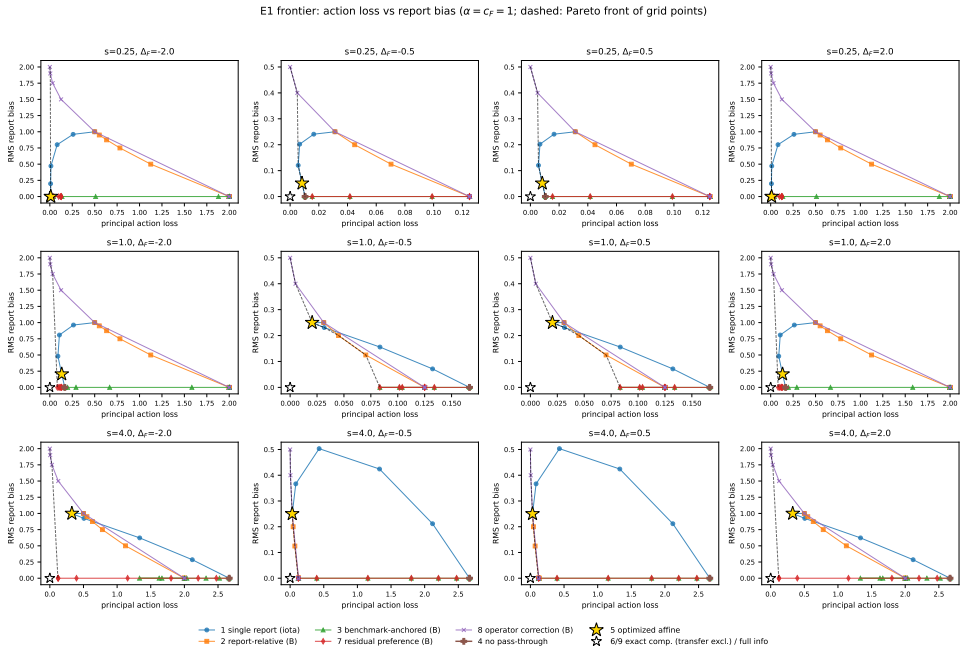


Figure 5: Report error and principal action loss across affine designs. Stars are optima on the declared finite grid. Increased authority can improve report accuracy while increasing action loss.

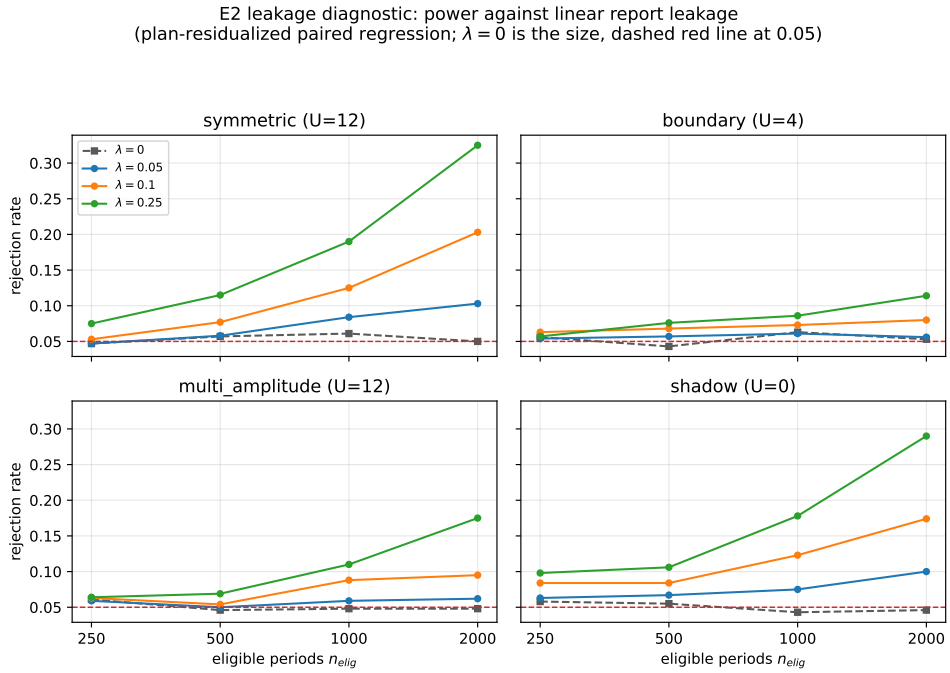


Figure 6: Routing-test power against linear message effects by design and eligible sample size. The dashed line is nominal level .05. Symmetric tests require slack; inward tests can operate at binding positive budgets.

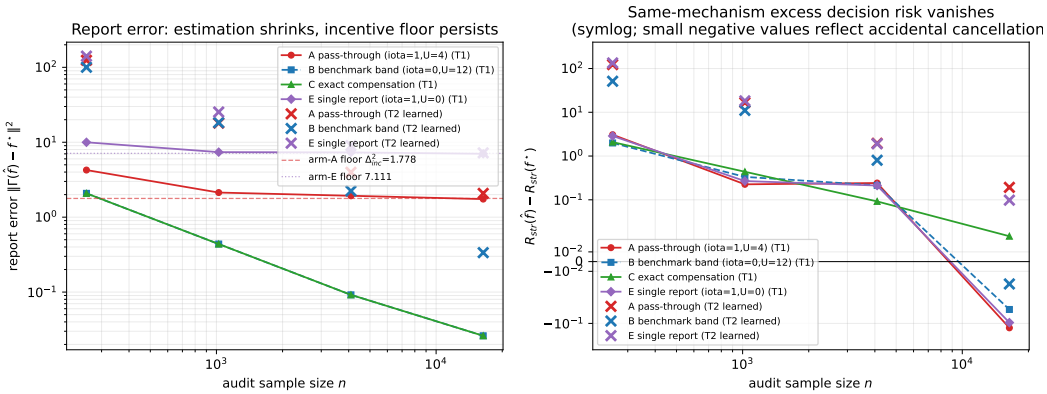


Figure 7: Report error and same-mechanism risk differences versus training size. The dashed report-error level $16/9$ is the binding protocol's incentive floor. Track 2 uses joint neural report and request heads and has no certified global optimization guarantee.

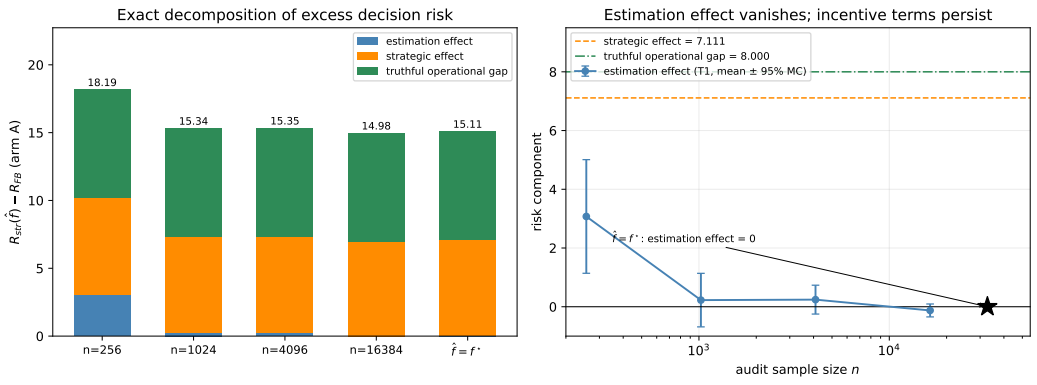


Figure 8: Risk decomposition for E3's binding report-anchored arm. The population strategic effect is $64/9$ and the truthful operational gap is 8. The estimation effect is zero when the estimated belief equals the population belief. Error bars show 95% Monte Carlo intervals.

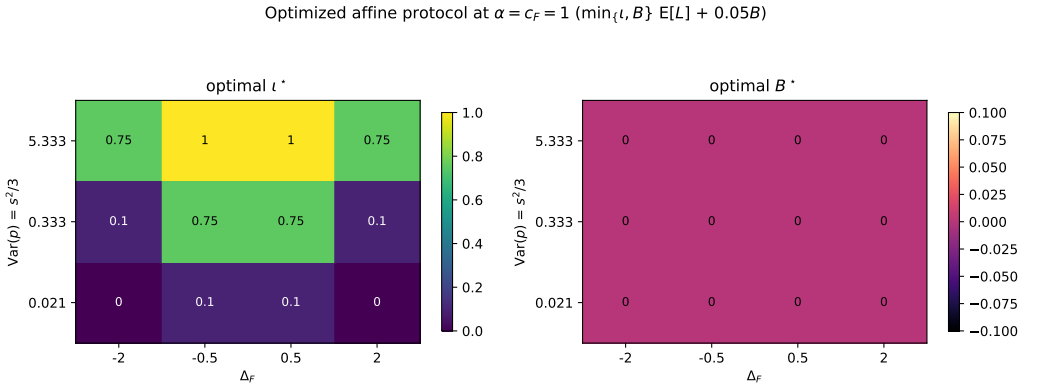


Figure 9: Selected influence and authority on the E1 grid, with fixed benchmark and positive linear authority cost. These results do not solve continuous mechanism optimization.

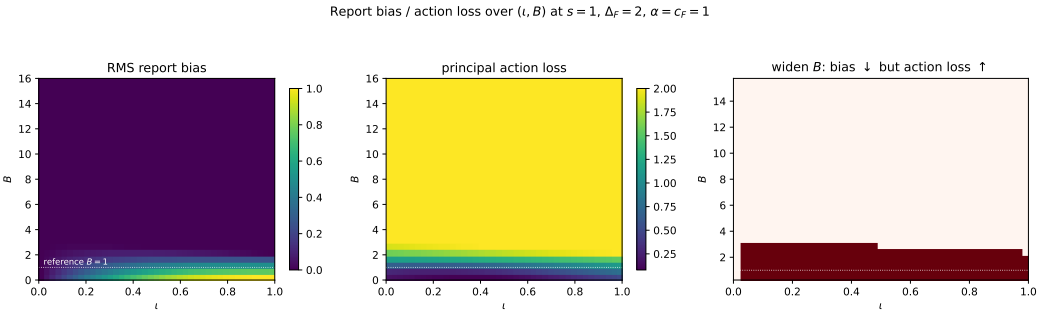


Figure 10: Local changes in report error and action loss when authority increases, in the cell $s = 1, \Delta_{\mathrm{F}} = 2, \alpha = 1$.

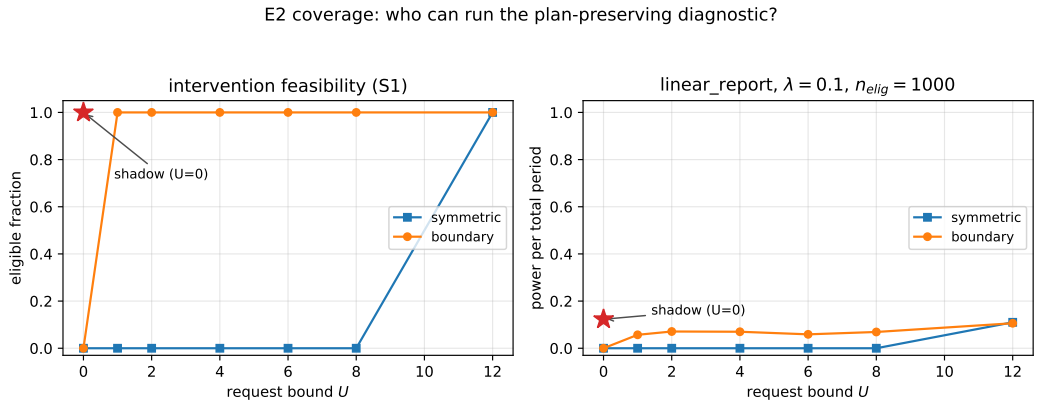


Figure 11: Eligibility and conditional power for symmetric and inward tests. Positive binding budgets prevent symmetric tests but permit inward ones. The shadow-interface point at zero budget is a separate intervention.

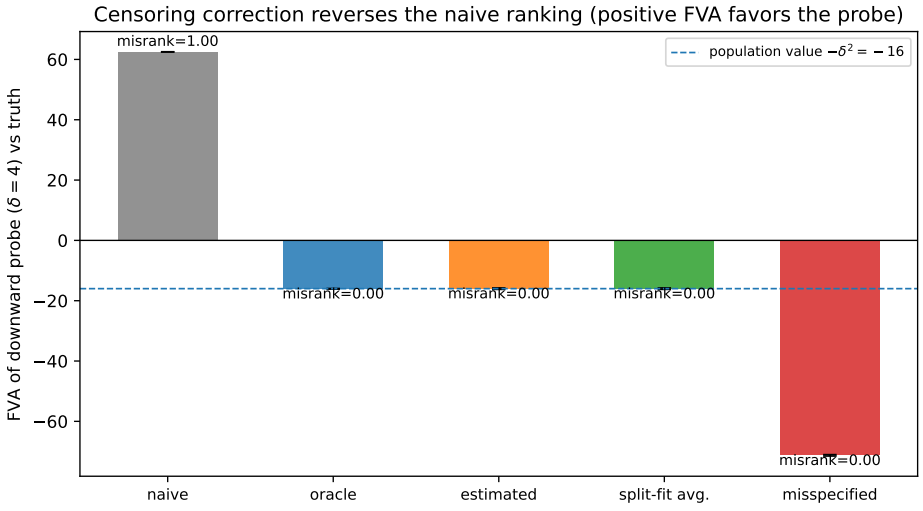


Figure 12: Evaluation under fixed demand censoring. Positive forecast value added favors the downward probe. The oracle correction yields the population value -16; the split-fit-average correction uses independent training and evaluation samples.