OPTIMIZING THE BIAS–UNCERTAINTY TRADE-OFF WITH PARTIALLY ORDERED BIAS BOUNDS

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the problem of aggregating multiple estimators to optimize the bias–uncertainty trade-off, especially when latent bias bounds are known only through their ordering and estimator uncertainties are represented by intervals. This problem naturally arises in hyperparameter tuning when candidate estimators cannot be directly validated, as in off-policy evaluation and pointwise nonparametric function estimation. However, existing methods cannot straightforwardly handle multidimensional hyperparameters for general estimator families because the associated bias bounds are only partially ordered. To address this issue, we formalize the aggregation of partially ordered intervals by minimizing the worst-case oracle-approximation ratio, leading to a novel aggregation rule called *Optimal Aggregation of Confidence Intervals* (OACIS). Our main contribution is a comprehensive theoretical analysis of this minimax problem, including the boundedness of the oracle ratio, uniqueness of the solution when this ratio is finite, stability, statistical error guarantees, and a linear-time algorithm. We also evaluate OACIS in synthetic and real-data experiments.

1 INTRODUCTION

Many statistical estimators are indexed by hyperparameters that control the trade-off between bias and uncertainty. However, evaluating this bias–uncertainty trade-off is sometimes difficult because there is no direct criterion for validating candidate estimators. For example, validating off-policy evaluation (OPE) estimates is challenging because of the mismatch between the evaluation and logging policies (Su et al., 2020; Tucker & Lee, 2021). Similarly, validating pointwise nonparametric function estimates is nontrivial over a continuous domain because held-out observations do not directly reveal the exact function value at a fixed point (Lepskii, 1991; Lepski et al., 1997). In such settings, validation-free methods are desirable for hyperparameter tuning.

We study an abstraction of this problem in which the candidate estimators are represented by *intervals ordered according to their bias tendencies*. In other words, we assume that the uncertainty of each candidate is quantitatively known, while only an ordering of oracle upper bounds on the candidates' biases is known. This setting arises commonly in machine learning: the uncertainty of an estimator can be estimated using asymptotic analysis, bootstrap methods, or concentration inequalities, whereas numerical bias bounds are difficult to obtain even when their relative magnitudes are naturally (partially) ordered by the complexity of the underlying model, such as the order of a Markov chain, the bandwidth used in kernel smoothing, and regularization weights. Given only these ordered intervals, we seek to produce a single point aggregate that competes with an oracle that knows the bias bounds.[1]

Classical Lepski-type methods have addressed this type of aggregation problem when the order is *total* (Lepskii, 1991; Lepski et al., 1997; Su et al., 2020; Tucker & Lee, 2021), whereas existing partial-order extensions critically rely on algebraic structure of kernel-smoothing operators (Goldenshluger & Lepski, 2008; 2011). Therefore, they do not directly apply to general interval estimators with a partial bias order. However, partial bias orders arise naturally in a wide range of applications since estimators often have multiple hyperparameters in practice and each hyperparameter may control a

[1] Since the bias bounds are unknown to us, we compete simultaneously with all oracles whose bias bounds are compatible with the order.

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.

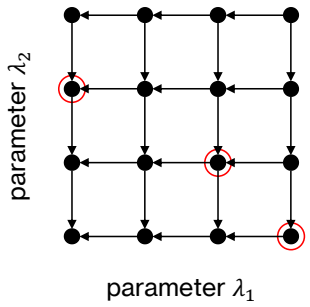


Figure 1.1: The Hasse diagram of a two-dimensional grid. Two candidates (nodes) are comparable if and only if they are connected with a directed path. The nodes marked with circles are mutually incomparable, showing that the order is not total.

distinct aspect of the bias–uncertainty trade-off, thereby inducing a grid partial order, as depicted in Fig. 1.1.

To close this gap, we propose an estimator-agnostic framework for aggregating partially ordered interval estimates. The key idea is to formulate the aggregation problem as a minimax optimization of the oracle-approximation ratio, treating any estimator-specific information not captured by the intervals and their bias order as unknown. We call the resulting minimax-optimal aggregation rule *Optimal Aggregation of Confidence Intervals* (OACIS) and the minimax oracle-approximation ratio the *minimax inefficiency*. Our main contributions are:

- closed-form expressions for OACIS and the minimax inefficiency;
- simple and sharp characterizations of both the boundedness of the minimax inefficiency and the stability of OACIS under input perturbations;
- oracle inequalities for the error quantiles and mean absolute error (MAE) of OACIS in a statistical regime; and
- a linear-time algorithm for computing both OACIS and the minimax inefficiency.

We also demonstrate the empirical performance of OACIS in synthetic and real-world experiments.

Paper organization. After reviewing related approaches in Section 2, we formalize the aggregation problem and its oracle-comparison criterion in Section 3. We first analyze the boundedness of the minimax inefficiency and the stability of OACIS under deterministic coverage in Section 4. We then extend these guarantees to the statistical setting with probabilistic coverage in Section 5. Next, Section 6 develops efficient algorithms for OACIS and its discretized variant, and Section 7 evaluates them on synthetic data and a real-world OPE application. Finally, Section 8 summarizes the findings and discusses future directions.

2 Related Work

Hyperparameter tuning and model selection. Standard hyperparameter-tuning and model-selection procedures rank candidate configurations using a computable objective, such as held-out loss or an information criterion (Stone, 1974; Akaike, 1974). Grid search, random search, and Bayesian optimization primarily address how to explore the configuration space efficiently once such an objective is available (Bergstra & Bengio, 2012; Snoek et al., 2012). In our setting, however, the estimation error of each candidate is not directly observable. The central difficulty is therefore candidate comparison rather than configuration search.

Estimator selection and aggregation for off-policy evaluation. Several OPE-specific methods construct surrogate comparison criteria for policy-value estimators without any trajectories of the target policy for validation. BVFT-PE and LSTD-Tournament select candidate value functions using Bellman-consistency tests (Zhang & Jiang, 2021; Liu et al., 2025), whereas duality-based residual estimation evaluates Bellman residuals through a dual formulation (Miyaguchi, 2026). PAS-IF

constructs pseudo-policy evaluation tasks to estimate the relative accuracy of candidate OPE estimators (Udagawa et al., 2023), while OPERA forms an affine combination using a bootstrap estimate of the cross-estimator MSE matrix (Nie et al., 2024). These methods address the absence of a direct validation criterion, but rely essentially on OPE-specific structure, such as Bellman equations and policy information, or strong regularity assumptions, such as uniform conditioning of feature matrices and bootstrap consistency. Their guarantees therefore do not directly transfer beyond their original settings. Lepski-type methods take a more estimator-agnostic approach based on confidence intervals and bias orderings.

Lepski-type and black-box estimator selection. Lepski's method was developed to adapt kernel-based nonparametric estimators to unknown smoothness and structural parameters (Lepskii, 1991; Lepski et al., 1997; Goldenshluger & Lepski, 2008; 2011). Abstracting from the kernel-based construction, SLOPE and its refinement SLOPE++ apply Lepski-type selection to totally ordered families of black-box estimates equipped with confidence radii (Su et al., 2020; Tucker & Lee, 2021). They establish finite-sample oracle inequalities under monotonicity assumptions on latent bias bounds and uncertainty radii, with SLOPE++ attaining the universal oracle-ratio bound. We extend this line of work from candidate selection to aggregation, allow arbitrary uncertainty radii, and replace the total order on latent bias bounds with a partial order.

3 Problem setup

Let $\theta^\star \in \mathbb{R}$ be an unknown target parameter and $I_v = [\ell_v, r_v]$ be an interval estimate of $\theta^\star$ for each hyperparameter configuration $v \in \mathcal{V}$, where the configuration set $\mathcal{V}$ satisfies $2 \le |\mathcal{V}| < \infty$. Also let $m_v := (\ell_v + r_v)/2$ and $u_v := (r_v - \ell_v)/2$ denote the midpoint and radius of the interval I_v, respectively. Our goal is to design an aggregation rule $\mathsf{Agg} : (I_v)_{v\in\mathcal{V}} \mapsto \theta \in \mathbb{R}$ so that $|\theta - \theta^\star|$ is guaranteed to be "small" under two mild assumptions we describe next.

Assumption I: Surrogate coverage of intervals. There exists a surrogate parameter $\eta_v \in \mathbb{R}$ for each $v \in \mathcal{V}$ such that every interval I_v is an approximate confidence interval for η_v, or, equivalently, $\eta_v \in [\ell_v - \epsilon, r_v + \epsilon]$ with high probability and a small error margin $\epsilon > 0$. We defer the precise probabilistic coverage and approximation conditions to Section 5. Until then, we assume the *exact* surrogate coverage

$$\eta_v \in I_v, \qquad v \in \mathcal{V}, \tag{1}$$

for simplicity of exposition.

Assumption II: Partially ordered bias bounds. There exists a (preferably tight) upper bound $b_v \ge 0$ on the bias of each surrogate parameter η_v relative to $\theta^\star$,

$$|\eta_v - \theta^\star| \le b_v, \qquad v \in \mathcal{V}, \tag{2}$$

where the numerical values of b_v are unknown, but a known partial order $\preceq$ on $\mathcal{V}$ satisfies

$$\forall v, w \in \mathcal{V}, \qquad w \preceq v \implies b_w \le b_v. \tag{3}$$

Consequently, the visible part of the problem parameters is identified with a poset of center-radius pairs, $\mathcal{I} := (m, u, (\preceq))$, while the remaining part, $\mathfrak{s} := (\theta^\star, \eta, b)$, is hidden. Here, $m = (m_v)_{v\in\mathcal{V}}$, $u = (u_v)_{v\in\mathcal{V}}$, $\eta = (\eta_v)_{v\in\mathcal{V}}$ and $b = (b_v)_{v\in\mathcal{V}}$ are the vectors of midpoints, radii, surrogate parameters and bias bounds, respectively.

The following example illustrates a typical setting in which this partial order arises.

Example 3.1 (Hyperparameter tuning over a grid). Let $\mathcal{V} = \mathcal{V}_1 \times \cdots \times \mathcal{V}_d$ be a finite d-dimensional grid in $\mathbb{R}^d$, $d \ge 1$, endowed with the usual product order as in Fig. 1.1: $(v_1, \ldots, v_d) \preceq (w_1, \ldots, w_d)$ if and only if $v_j \le w_j$ for every $j = 1, \ldots, d$. Then, the partial order $\preceq$ reflects the assumption that there exists a reasonable oracle bias bound b_v which is coordinatewise monotone. In other words, each coordinate $v_j \in \mathcal{V}_j$ controls a distinct aspect of η_v as an approximation of $\theta^\star$.

Objective: Minimax inefficiency. To measure how "small" the aggregation error $|\theta - \theta^\star|$ is, we need a baseline for comparison. To this end, consider an oracle that knows the bias bounds b. With

b known, the error of each candidate $v \in \mathcal{V}$ can be naturally bounded by the sum of the bias and uncertainty bounds,

$$|m_v - \theta^\star| \le |\eta_v - \theta^\star| + |m_v - \eta_v| \le b_v + u_v.$$

The oracle can exploit this inequality and achieve the *oracle error bound*

$$R^\star := \min_{v \in \mathcal{V}} (b_v + u_v), \tag{4}$$

by returning m_v for a candidate v minimizing $b_v + u_v$. Note that $R^\star$ is a typical oracle benchmark for hyperparameter tuning: the smallest bias-plus-uncertainty bound among the candidates. This motivates us to study *how well an aggregation rule can perform relative to Eq.* (4), which is quantified by the notion of *inefficiencies* we introduce next.

Definition 3.1 (Inefficiency)**.** *Write the oracle error bound as $R^\star(\mathfrak{s}, \mathcal{I}) := R^\star$ to make its dependence on $\mathfrak{s}$ and $\mathcal{I}$ explicit. For an aggregate $\theta \in \mathbb{R}$, define its inefficiency on $(\mathfrak{s}, \mathcal{I})$ as*

$$\alpha(\theta; \mathfrak{s}, \mathcal{I}) = \frac{|\theta - \theta^\star|}{R^\star(\mathfrak{s}, \mathcal{I})},$$

with the conventions $1/0 = \infty$ and $0/0 = 1$.

Definition 3.2 (Worst-case inefficiency)**.** *For an aggregate $\theta \in \mathbb{R}$, its worst-case inefficiency on $\mathcal{I}$ is given by*

$$\bar{\alpha}(\theta; \mathcal{I}) = \sup_{\mathfrak{s} \text{ compatible with } \mathcal{I}} \alpha(\theta; \mathfrak{s}, \mathcal{I}), \tag{5}$$

where $\mathfrak{s}$ is compatible with $\mathcal{I}$ if and only if Eqs. (1) *to* (3) *hold.*

Definition 3.3 (Minimax inefficiencies and optimal aggregates)**.** *The minimax inefficiency of the observed input $\mathcal{I}$ and its discretized variant are given by*

$$\bar{\alpha}^\star(\mathcal{I}) = \inf_{\theta \in \mathbb{R}} \bar{\alpha}(\theta; \mathcal{I}), \qquad \bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}) = \inf_{\theta \in \Theta_{\mathrm{d}}} \bar{\alpha}(\theta; \mathcal{I}),$$

where $\Theta_{\mathrm{d}} := \{m_v : v \in \mathcal{V}\}$ is the set of candidate midpoints. The corresponding minimax-optimal aggregates, if they exist, are denoted by

$$\bar{\theta}^\star(\mathcal{I}) \in \arg\min_{\theta \in \mathbb{R}} \bar{\alpha}(\theta; \mathcal{I}), \qquad \bar{\theta}^\star_{\mathrm{d}}(\mathcal{I}) \in \arg\min_{\theta \in \Theta_{\mathrm{d}}} \bar{\alpha}(\theta; \mathcal{I}),$$

where ties are resolved by an arbitrary fixed rule. We call these aggregates Optimal Aggregation of Confidence Intervals *(OACIS) and* discretized OACIS *(d-OACIS), respectively.*

This minimax formulation is suitable for modeling black-box candidate estimators, but relies on exact surrogate coverage (1). Therefore, we first analyze the minimax problem in Section 4, and then extend it to take into account the probabilistic and approximate nature of the coverage in Section 5.

4 Minimax analysis

Henceforth we fix the underlying poset $(\mathcal{V}, \preceq)$ and identify the input with its pair of center-radius vectors $\mathcal{I} = (m, u)$, suppressing the partial order $\preceq$ from the input notation. We also denote the class of all inputs by $\mathbb{I} := \mathbb{R}^{\mathcal{V}} \times \mathbb{R}^{\mathcal{V}}_{\ge 0}$. The following proposition characterizes the epigraph of $\bar{\alpha}(\cdot; \mathcal{I})$, which is the cornerstone of our entire analysis.

Proposition 4.1 (Epigraph characterization)**.** *For $(x, y) \in \mathbb{R} \times \mathbb{R}$,*

$$\bar{\alpha}(x; \mathcal{I}) \le y$$

if and only if

$$y \ge 1, \qquad |x - \mu_v| + \delta_v^+ \le h_v y, \qquad v \in \mathcal{V}, \tag{6}$$

where $L_v := \max_{w \preceq v} \ell_w$, $R_v := \min_{w \preceq v} r_w$, $\mu_v := (L_v + R_v)/2$, $\delta_v := (R_v - L_v)/2$, $\delta_v^+ := \max\{0, \delta_v\}$, $\delta_v^- := \max\{0, -\delta_v\}$, and $h_v := u_v + \delta_v^-$ for all $v \in \mathcal{V}$.

The proof is provided in Section C.2. Proposition 4.1 plays a key role in deriving the following closed-form expressions for the minimax inefficiency and the optimal aggregate.

Theorem 4.1 (Closed form). *Write $Z := \{v \in \mathcal{V} : h_v = 0\}$. The value $\bar{\alpha}^\star(\mathcal{I})$ is finite if and only if*

$$\mu_v = \mu_w \qquad \text{for every } v, w \in Z. \tag{7}$$

When this condition holds,

$$\bar{\alpha}^\star(\mathcal{I}) = \max\left\{1, \max_{\substack{\{v,w\}\in\binom{\mathcal{V}}{2} \\ h_v+h_w>0}} \frac{|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+}{h_v + h_w}\right\}, \tag{8}$$

where the inner maximum is omitted when there is no qualifying pair. Moreover, whenever $\bar{\alpha}^\star(\mathcal{I}) < \infty$, the minimax-optimal aggregate exists, is unique, and satisfies

$$\bar{\theta}^\star(\mathcal{I}) = \max_{v\in\mathcal{V}} \left\{\mu_v + \delta_v^+ - h_v\bar{\alpha}^\star(\mathcal{I})\right\} = \min_{v\in\mathcal{V}} \left\{\mu_v - \delta_v^+ + h_v\bar{\alpha}^\star(\mathcal{I})\right\}. \tag{9}$$

Proof sketch. Fix a candidate inefficiency level $y \geq 1$. By Proposition 4.1, an output x achieves $\bar{\alpha}(x;\mathcal{I}) \leq y$ if and only if it belongs to

$$C_v(y) := [\mu_v + \delta_v^+ - h_v y,\ \mu_v - \delta_v^+ + h_v y]$$

for every $v \in \mathcal{V}$. Thus the minimax problem asks for the smallest y at which the expanding intervals $\{C_v(y)\}_{v\in\mathcal{V}}$ have a common point. When such a finite level exists, the common point at the smallest level is unique and equals $\bar{\theta}^\star(\mathcal{I})$. The full proof, including detailed derivations of Eqs. (7) to (9), is given in Section C.3. □

Theorem 4.1 shows that $\bar{\alpha}(\cdot;\mathcal{I})$ is convex and has a unique minimizer if $\bar{\alpha}^\star(\mathcal{I}) < \infty$. Therefore, d-OACIS must be one of the candidate midpoints nearest to OACIS from below or above. Furthermore, it implies that the discretization only incurs a controlled increase in worst-case inefficiency. The proof is given in Section C.4.

Proposition 4.2 (Cost of discretization). *$\bar{\alpha}^\star(\mathcal{I}) \leq \bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}) \leq 1 + 2\bar{\alpha}^\star(\mathcal{I})$ for all $\mathcal{I} \in \mathbb{I}$.*

In the remainder of this section, we focus on non-discretized OACIS and its minimax inefficiency. In particular, we first study a certain class of structural conditions for bounding $\bar{\alpha}^\star(\mathcal{I})$ in Section 4.1 and then establish Lipschitz stability of $\bar{\theta}^\star(\mathcal{I})$ in Section 4.2.

4.1 Boundedness

Let $(\mathcal{V}, \preceq)$ be a meet-semilattice; every pair of elements $v, w \in \mathcal{V}$ has a meet $v \wedge w \in \mathcal{V}$, the greatest of their common predecessors. This includes, as special cases, grid structures with the product order (Example 3.1).

Definition 4.1 (κ-subadditivity). *For $\kappa \geq 1/2$, a radius vector $u \in \mathbb{R}^{\mathcal{V}}_{\geq 0}$ is κ-subadditive if*

$$u_{v\wedge w} \leq \kappa(u_v + u_w), \qquad v, w \in \mathcal{V}.$$

The class of all κ-subadditive radius vectors on $\mathcal{V}$ is denoted by $\mathbb{S}^{\mathcal{V}}_\kappa$, and the corresponding input class is denoted by $\mathbb{I}_\kappa := \mathbb{R}^{\mathcal{V}} \times \mathbb{S}^{\mathcal{V}}_\kappa$.

Intuitively, κ-subadditivity controls the growth of the uncertainty radii as one descends the poset with the meet operation. In fact, it is also useful for bounding $\bar{\alpha}^\star(\mathcal{I})$.

Theorem 4.2 (Sharp bound under subadditivity). *For every $\kappa \geq 1/2$,*

$$\sup_{\mathcal{I}\in\mathbb{I}_\kappa} \bar{\alpha}^\star(\mathcal{I}) = \begin{cases} 2\min\{\kappa, 1\}, & \text{if } \mathcal{V} \text{ is totally ordered}, \\ 2\kappa, & \text{otherwise}. \end{cases}$$

Proof sketch. Fix an arbitrary $\mathcal{I} \in \mathbb{I}_\kappa$ with $\kappa \geq 1/2$. The key inequality for the upper bound is

$$|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+ \leq 2u_{v\wedge w} + \delta_v^- + \delta_w^- \leq 2\kappa(h_v + h_w) \qquad v, w \in \mathcal{V}.$$

This inequality sharply bounds the pairwise quotients in (8) for partial orders; for total orders, the bound can be further tightened using the fact that $\mathbb{I} = \mathbb{I}_1$. Matching lower bounds follow from explicit constructions of intervals. The full proof is given in Section D. □

By letting $\kappa \to \infty$ in Theorem 4.2 and using the fact that $\mathbb{I} = \mathbb{I}_1$ when $\mathcal{V}$ is totally ordered, we obtain a clearer dichotomy between total and partial orders.

Corollary 4.1 (Sharp uniform bound)**.**

$$\sup_{\mathcal{I}\in\mathbb{I}} \bar{\alpha}^\star(\mathcal{I}) = \begin{cases} 2, & \textit{if } \mathcal{V} \textit{ is totally ordered}, \\ \infty, & \textit{otherwise}. \end{cases}$$

For grid-structured input classes, this implies that $\bar{\alpha}^\star(\cdot)$ is uniformly bounded by a constant if and only if the grid is one-dimensional. In this case, OACIS attains the optimal constant 2, and Proposition 4.2 gives the upper bound 5 for d-OACIS. The best previously known guarantee for Lepski-type selection rules is $\sqrt{6}+3 \approx 5.45$ (Tucker & Lee, 2021). See Remark D.1 for the precise relation between the two decision classes.

On the other hand, the unboundedness of the partially ordered case highlights the need for additional structural assumptions such as κ-subadditivity. The following remark provides an interpretation of κ-subadditivity in the context of multidimensional grids.

Remark 4.1 (A sufficient condition for κ-subadditivity on grids). Let $\mathcal{V}$ be as in Example 3.1 and fix $\kappa \geq 1/2$. Suppose there exist nonnegative functions $a_j : \mathcal{V}_j \to \mathbb{R}_{\geq 0}$ and a coordinatewise nondecreasing function $F : \mathbb{R}^d_{\geq 0} \to \mathbb{R}_{\geq 0}$ satisfying

$$F(x+y) \leq \kappa\{F(x)+F(y)\}, \qquad x, y \in \mathbb{R}^d_{\geq 0},$$

such that

$$u_v = F(a(v)), \qquad a(v) := \big(a_1(v_1), \ldots, a_d(v_d)\big), \qquad v \in \mathcal{V}.$$

Then u is κ-subadditive. The detailed proof is given in Section E.

In other words, the κ-subadditivity on a grid holds when the uncertainty radii admit the decomposition $v \mapsto a(v) \mapsto u_v$. The intermediate representation $a(v) \in \mathbb{R}^d_{\geq 0}$ can be understood as the vector of coordinatewise latent uncertainties, while the map F propagates these uncertainties to the overall uncertainty. A practically relevant instance is composite estimation with Lipschitz error propagation; see Example E.2.

4.2 Stability

To measure perturbations to the input $\mathcal{I} = (m, u)$, we introduce a block norm that treats changes in the centers and increases and decreases in the radii separately. For $\Delta\mathcal{I} = (\Delta m, \Delta u) \in \mathbb{R}^{\mathcal{V}} \times \mathbb{R}^{\mathcal{V}}$, write

$$\|\Delta\mathcal{I}\|_{\mathrm{blk}} := \|\Delta m\|_\infty + \|\Delta u_+\|_\infty + \|\Delta u_-\|_\infty,$$

where Δu_+ and Δu_- denote the coordinatewise positive and negative parts, respectively, so that $\Delta u = \Delta u_+ - \Delta u_-$. Although this block norm is equivalent to the standard ℓ_∞ norm, it yields a tighter stability bound for $\bar{\theta}^\star$.

Theorem 4.3 (Lipschitz stability)**.** *Let $\mathcal{C}$ be a convex subset of $\mathbb{I}$ and $M := \sup_{\mathcal{I}\in\mathcal{C}} \bar{\alpha}^\star(\mathcal{I})$. If $M < \infty$, then for every $\mathcal{I}, \mathcal{I}' \in \mathcal{C}$,*

$$|\bar{\theta}^\star(\mathcal{I}) - \bar{\theta}^\star(\mathcal{I}')| \leq M\|\mathcal{I} - \mathcal{I}'\|_{\mathrm{blk}}. \tag{10}$$

Proof sketch. The key result is Lemma F.1, which gives blockwise bounds on the gradient of $\bar{\theta}^\star$ almost everywhere on the interior of $\mathbb{I}$,

$$\|\nabla_m \bar{\theta}^\star(\mathcal{I})\|_1 \leq \bar{\alpha}^\star(\mathcal{I}), \qquad \|(\nabla_u \bar{\theta}^\star(\mathcal{I}))_+\|_1 \leq \bar{\alpha}^\star(\mathcal{I}), \qquad \|(\nabla_u \bar{\theta}^\star(\mathcal{I}))_-\|_1 \leq \bar{\alpha}^\star(\mathcal{I}).$$

Together with the local Lipschitz condition implied by the closed-form characterization (9), these gradient bounds yield the desired stability claim. The full proof is given in Section F.2. □

Remark 4.2 (Stability under subadditivity). When $\mathcal{V}$ is a meet-semilattice, Theorem 4.3 applies to $\mathcal{C} = \mathbb{I}_\kappa$ and $M = 2\kappa$ with $\kappa \geq 1/2$, thanks to Theorem 4.2 and the convexity of $\mathbb{I}_\kappa$. In particular, OACIS is globally 2-Lipschitz with respect to $\|\cdot\|_{\mathrm{blk}}$ when $\mathcal{V}$ is totally ordered.

Remark 4.3 (Instability of selection rules). Selection rules, including d-OACIS and Lepski-type methods, choose a single candidate midpoint and are generally unstable. More precisely, any selection rule that has bounded inefficiency on $\mathbb{I}$ must be discontinuous at some input $\mathcal{I} \in \mathbb{I}$, even if $\mathcal{V}$ is totally ordered. To see this, let $\mathcal{V} = \{1, 2\}$, $1 \preceq 2$, and consider the input comprising $I_1 = [-1, 1]$ and $I_2 = [t, t]$, where t increases from 0 to ∞. For $0 \le t \le 1$, any selection rule with bounded inefficiency must choose $m_2 = t$, whereas for sufficiently large t, it must choose $m_1 = 0$. Hence the switch from $m_2 = t$ to $m_1 = 0$ must occur at some $t > 1$, creating a discontinuity.

5 Statistical analysis

We now extend the minimax theory in Section 4 by relaxing the exact surrogate coverage condition (1). Fix a finite meet-semilattice $(\mathcal{V}, \preceq)$ again and suppose that the candidate-estimand vector η has an unknown bias-bound vector $b \in \mathbb{R}^{\mathcal{V}}_{\ge 0}$ satisfying Eqs. (2) and (3). Let $\widehat{m} \in \mathbb{R}^{\mathcal{V}}$ be a random estimator of η. For each $t \in (0, 1)$, suppose that there is an unknown nonnegative radius vector $u(t) \in \mathbb{R}^{\mathcal{V}}_{\ge 0}$ satisfying the probabilistic surrogate coverage condition

$$\Pr\{\eta_v \in [\widehat{m}_v - u_v(t), \widehat{m}_v + u_v(t)],\ \forall v \in \mathcal{V}\} \ge 1 - t. \tag{11}$$

In practice, however, the exact radii $u(t)$ satisfying Eq. (11) may be unknown so that we only observe a random approximation $\widehat{u}(t) \in \mathbb{R}^{\mathcal{V}}_{\ge 0}$. Thus, the goal of this section is to provide statistical guarantees for the natural aggregate $\widehat{\theta}_t := \bar{\theta}^\star(\widehat{m}, \widehat{u}(t))$ by leveraging the results established in Section 4.

To simplify the exposition, we assume that the true and estimated radii share two known profiles $q, z : (0, 1) \to [0, \infty)$ that are common to all candidates, nonincreasing, integrable, and differentiable, such that

$$u_v(t) = \sigma_v q(t) + \gamma_v z(t), \qquad \widehat{u}_v(t) = \widehat{\sigma}_v q(t) + \gamma_v z(t), \qquad v \in \mathcal{V},$$

where $\sigma, \widehat{\sigma} \in \mathbb{R}^{\mathcal{V}}_{>0}$ denote an unknown scaling vector and its random estimator, respectively, while $\gamma \in \mathbb{R}^{\mathcal{V}}_{>0}$ is a known coefficient vector. This is a natural assumption in many applications, including Hoeffding-, Wald-, and Bernstein-type confidence intervals, as summarized in Table G.1. More specifically, we require below that there exist $\kappa \ge 1/2$, $c > 0$, $C < \infty$, $\epsilon \ge 0$, $0 < \delta < 1$, and $L > 0$ such that

i) σ and γ are κ-subadditive,
ii) $\Pr\{\max_{v \in \mathcal{V}} |\ln \widehat{\sigma}_v - \ln \sigma_v| \le \epsilon\} \ge 1 - \delta$,
iii) $\sigma_v, \widehat{\sigma}_v \in [c, C]$ almost surely for all $v \in \mathcal{V}$, and
iv) $|q'(t)| \le (L/t) q(t)$ and $|z'(t)| \le (L/t) z(t)$ for all $t \in (0, 1/2]$.

Note that conditions i)–iv) are satisfied by all examples in Table G.1 under suitable conditions on $\widehat{\sigma}$ for the Wald and Bernstein profiles; see Section G.4 for a detailed Bernstein example. We further simplify the subsequent analysis by considering the asymptotic regime in which $\epsilon \to 0$, $\delta (C/c)^{1+L} \to 0$, and $\kappa, \rho, L = O(1)$, where $\rho := \max\{\sigma_{\max}/\sigma_{\min}, \gamma_{\max}/\gamma_{\min}\}$.[2] All $o(1)$ terms below are understood along any sequence of problem instances satisfying this regime, with the finite meet-semilattice $(\mathcal{V}, \preceq)$ fixed. Explicit nonasymptotic versions of the bounds, including their dependence on ϵ and δ, are given in Eqs. (24) and (26) in Section G.

The following proposition gives an oracle inequality for the error quantiles of $\widehat{\theta}_t$.

Proposition 5.1 (Quantile bound)**.** *Suppose conditions i) and ii) above hold. Then, for $t \in (0, 1)$,*

$$\Pr\left\{|\widehat{\theta}_t - \theta^\star| \le \{2\kappa + o(1)\} \min_{v \in \mathcal{V}} \{b_v + u_v(t)\}\right\} \ge 1 - t - o(1).$$

A similar bound for the mean absolute error (MAE) can be obtained by further exploiting stability with respect to perturbations of the confidence parameter t. To this end, we introduce the tail-calibrated radius

$$\bar{u}_v(t) = u_v(t) + (1 + \rho) \int_0^t \{u_v(x) - u_v(t)\}\, dx, \qquad v \in \mathcal{V}.$$

[2] For a vector $s \in \mathbb{R}^{\mathcal{V}}$, let $s_{\min}$ and $s_{\max}$ denote its minimum and maximum coordinates, respectively.

Proposition 5.2 (MAE bound)**.** *Suppose conditions i)–iii) above hold. Then, for* $t \in (0, 1)$*,*

$$\mathbb{E}|\widehat{\theta}_t - \theta^\star| \leq \{2\kappa + o(1)\} \min_{v \in \mathcal{V}} \{b_v + \bar{u}_v(t)\}. \tag{12}$$

In particular, the right-hand side is minimized at $t^\star := 1/(1+\rho)$*. Moreover, if condition iv) also holds, the minimized asymptotic bound is attained using the plug-in confidence parameter* $\hat{t} := 1/(1+\widehat{\rho})$*, where* $\widehat{\rho} := \max\{\widehat{\sigma}_{\max}/\widehat{\sigma}_{\min}, \gamma_{\max}/\gamma_{\min}\}$*.*

The proofs are given in Section G. Both propositions show that the estimation error of OACIS is bounded by a constant multiple of the corresponding oracle error bound, while Proposition 5.2 also gives a guide for choosing the confidence parameter t to minimize the MAE upper bound.

6 Algorithms

We now present linear-time algorithms for OACIS and d-OACIS, where the latter is a relatively thin wrapper around the former. Both algorithms assume that the partial order is represented by a directed acyclic graph $G = (\mathcal{V}, E)$: every comparable pair is connected by a directed path, and vice versa. They then decide whether $\bar{\alpha}^\star(\mathcal{I}) < \infty$ and, if so, return $\bar{\theta}^\star(\mathcal{I})$ or $\bar{\theta}^\star_{\mathrm{d}}(\mathcal{I})$. The key idea is to formulate the minimax problem as a bivariate linear program

$$\min_{(x,y) \in \mathbb{R}^2} y \quad \text{s.t.} \quad y \geq \bar{\alpha}(x; \mathcal{I}),$$

where the constraint can be evaluated efficiently using Proposition 4.1. The resulting pseudocode is given in Algorithms H.1 and H.2 and runs in $O(|\mathcal{V}| + |E|)$ time. See Section H for the detailed derivation and Section H.1 for the formal statements and the proofs of their correctness and computational complexity.

Remark 6.1 (Comparison with direct closed-form evaluation)*.* Directly evaluating (8) and then (9) examines $\binom{|\mathcal{V}|}{2}$ pairs and takes $O(|\mathcal{V}|^2)$ time, which is significantly slower than the algorithm in Algorithm H.1 when $|\mathcal{V}|^2 \gg |E|$. This is the case for d-dimensional grids, whose Hasse graphs have only $|E| = O(|\mathcal{V}| \log |\mathcal{V}|)$ edges.

7 Experiments

We evaluate OACIS in two settings. The synthetic-data experiment illustrates the performance of OACIS under controlled bias–uncertainty trade-offs, while the real-data experiment reports the effectiveness of OACIS for OPE on a web recommendation system. The full experimental details, including additional results, are given in Sections A and B, respectively.

7.1 Synthetic-data experiment

Experimental overview. We test OACIS in three synthetic conditions: a one-dimensional chain with same-sign and alternating-sign deviations $\eta_v - \theta^\star$ (conditions 1 and 2, respectively), and a two-dimensional grid combining the two structures (condition 3). The methods receive Wald confidence intervals and the associated partial order. In every condition, the true scaling vector σ is subadditive with $\kappa = 1$.

Methods and protocol. We compare Random interval, SLOPE (Su et al., 2020), SLOPE++ (Tucker & Lee, 2021), OACIS, d-OACIS, and a hindsight oracle. For condition 3, the SLOPE variants operate on a topological sort of the grid, whereas the OACIS variants operate directly on its partial order. We use 100 independent repetitions and report the mean absolute error (MAE) and its standard error for each method and condition.

Results. Table 7.1 reports the results. Among the deployable methods, OACIS achieves the lowest or nearly the lowest MAE in all three conditions. Both OACIS variants substantially outperform the hindsight oracle in condition 2, indicating the advantage of data-dependent selection and aggregation over a fixed candidate. Complementary visual comparisons appear in Fig. A.1. The sensitivity results in Fig. A.2 also show OACIS's robustness to the choice of the confidence parameter t.

Table 7.1: Synthetic estimator results. Entries report MAE (standard error).

Condition	Random	SLOPE	SLOPE++	OACIS	d-OACIS	Hindsight
1	3.904 (0.510)	0.720 (0.033)	0.440 (0.028)	**0.305** (0.022)	**0.297** (0.022)	0.351 (0.028)
2	1.344 (0.072)	0.963 (0.055)	0.923 (0.067)	**0.435** (0.034)	0.537 (0.036)	0.727 (0.053)
3	4.818 (0.584)	2.728 (0.155)	1.541 (0.100)	**0.558** (0.044)	**0.530** (0.044)	0.205 (0.016)

7.2 Real-data experiment

Experimental overview. The task is to estimate the effect of a new recommendation policy on daily click rates in a web recommendation system. We have conducted an A/B test with two groups of users across 25 deployment dates. Each method receives a grid of partially ordered interval OPE estimators for each date. The output of aggregation is then compared with the observed difference in the A/B test. The candidate estimators are modified partial-history IPW estimators (Hu & Wager, 2023) indexed by history length and impression-effect duration. These choices form a two-dimensional partial order comprising 56 candidates and an abstention candidate fixed at zero. The resulting interval-radius vector was subadditive with $\kappa = 1$ for all dates.[3]

Methods and evaluation. We compare SLOPE, SLOPE++, OACIS, and d-OACIS, all applied to the same candidate grid, with two fixed baselines: No OPE and Naive Bandit. The SLOPE variants linearize the partial order as in the synthetic-data experiment. Performance is evaluated over the 25 dates using MAE normalized by that of Naive Bandit, along with Pearson's r between the predictions and the observed A/B differences.

Results. Table 7.2 reports the results. OACIS has the smallest NMAE, 0.754, with a standard error of 0.074, and the strongest Pearson correlation, 0.746 ($p = 1.90 \times 10^{-5}$).

Table 7.2: Full-grid click results for the real-data experiment. Each NMAE is followed by its standard error, and each Pearson correlation is followed by its p-value. A dash denotes an undefined correlation.

Method	NMAE (SE)	Pearson r (p)
No OPE	3.635 (0.111)	—
Naive Bandit	1.000 (0.087)	0.620 (9.53×10^{-4})
SLOPE	0.977 (0.089)	0.551 (4.30×10^{-3})
SLOPE++	0.843 (0.078)	0.587 (2.05×10^{-3})
OACIS	**0.754** (0.074)	**0.746** (1.90×10^{-5})
d-OACIS	0.903 (0.095)	0.567 (3.14×10^{-3})

8 Conclusion

We introduced OACIS, a method of estimator aggregation based on partially ordered bias bounds, designed for validation-free multidimensional hyperparameter tuning. Its minimax formulation yields closed-form characterizations, sharp conditions for boundedness of the oracle ratio, and stability guarantees under input perturbations. We further translated these deterministic results into quantile and mean-absolute-error oracle inequalities and developed linear-time algorithms for OACIS and its discretized selection variant.

Across the synthetic conditions and a real-world OPE application, OACIS compared favorably with Lepski-type baselines while using the partial order directly. OACIS in its current form is however not readily applicable beyond one-dimensional uncertainty sets $I_v \subset \mathbb{R}$, so extending the framework to higher-dimensional (or more general) uncertainty sets is an important direction for future work.

[3] $\kappa = 1$ here is tight due to the presence of the abstention.

AI USE STATEMENT

In this work, we used generative AI tools for literature survey assistance, generating code for experiments including method implementations, and polishing the manuscript structurally and stylistically. We have not used generative AI tools for translation, proposing or refining research hypotheses, generating or manipulating datasets, or developing core theoretical results, their proofs, algorithmic and experimental designs, or analyses. We have reviewed all AI-assisted work, and made necessary edits to ensure accuracy and clarity. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

The detailed experimental procedures, including the data-generating processes, hyperparameter configurations, and evaluation protocols, are provided in Sections A and B. The algorithms for OACIS and d-OACIS are described in Section H.

REFERENCES

Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974. doi: 10.1109/TAC.1974.1100705.

James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012. URL https://www.jmlr.org/papers/v13/bergstra12a.html.

Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer, Berlin, Heidelberg, 3rd edition, 2008. doi: 10.1007/978-3-540-77974-2.

Alexander Goldenshluger and Oleg Lepski. Universal pointwise selection rule in multivariate function estimation. *Bernoulli*, 14(4):1150–1190, 2008. doi: 10.3150/08-BEJ144.

Alexander Goldenshluger and Oleg Lepski. Bandwidth selection in kernel density estimation: Oracle inequalities and adaptive minimax optimality. *The Annals of Statistics*, 39(3):1608–1632, 2011. doi: 10.1214/11-AOS883.

Yuchen Hu and Stefan Wager. Off-policy evaluation in partially observed markov decision processes under sequential ignorability. *The Annals of Statistics*, 51(4):1561–1585, 2023.

Oleg V. Lepski, Enno Mammen, and Vladimir G. Spokoiny. Optimal spatial adaptation to inhomogeneous smoothness: An approach based on kernel estimates with variable bandwidth selectors. *The Annals of Statistics*, 25(3):929–947, 1997.

Oleg V. Lepskii. On a problem of adaptive estimation in gaussian white noise. *Theory of Probability and Its Applications*, 35(3):454–466, 1991. doi: 10.1137/1135065.

Pai Liu, Lingfeng Zhao, Shivangi Agarwal, Jinghan Liu, Audrey Huang, Philip Amortila, and Nan Jiang. Model selection for off-policy evaluation: New algorithms and experimental protocol. In *Advances in Neural Information Processing Systems*, volume 38, pp. 47912–47946, 2025. doi: 10.52202/085713-1430.

Nimrod Megiddo. Linear-time algorithms for linear programming in $\mathbb{R}^3$ and related problems. *SIAM Journal on Computing*, 12(4):759–776, 1983. doi: 10.1137/0212052.

Kohei Miyaguchi. Duality-based residual estimation for fully offline value-based reinforcement learning. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, 2026. URL https://openreview.net/forum?id=5ymJ1ylWNd.

Allen Nie, Yash Chandak, Christina J. Yuan, Anirudhan Badrinath, Yannis Flet-Berliac, and Emma Brunskill. OPERA: Automatic offline policy evaluation with re-weighted aggregates of multiple estimators. In *Advances in Neural Information Processing Systems*, volume 37, pp. 103652–103680, 2024. doi: 10.52202/079017-3292.

Raimund Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete & Computational Geometry*, 6(3):423–434, 1991. doi: 10.1007/BF02574699.

Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, volume 25, 2012. URL https://proceedings.neurips.cc/paper/2012/hash/05311655a15b75fab86956663e1819cd-Abstract.html.

Mervyn Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2):111–133, 1974. doi: 10.1111/j.2517-6161.1974.tb00994.x.

Yi Su, Pavithra Srinath, and Akshay Krishnamurthy. Adaptive estimator selection for off-policy evaluation. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9196–9205, 2020. URL https://proceedings.mlr.press/v119/su20d.html.

George Tucker and Jonathan Lee. Improved estimator selection for off-policy evaluation. In *Workshop on Reinforcement Learning Theory at the 38th International Conference on Machine Learning*, 2021. URL https://icml.cc/virtual/2021/12874.

Takuma Udagawa, Haruka Kiyohara, Yusuke Narita, Yuta Saito, and Kei Tateno. Policy-adaptive estimator selection for off-policy evaluation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(8):10025–10033, 2023. doi: 10.1609/aaai.v37i8.26195.

Siyuan Zhang and Nan Jiang. Towards hyperparameter-free policy selection for offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 34, pp. 12864–12875, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/6add07cf50424b14fdf649da87843d01-Abstract.html.

A Additional details for the synthetic experiment

A.1 Data-generating process

We consider a Gaussian generative model on a two-dimensional grid $\mathcal{V} \subset \mathbb{Z}^2$. For a grid point $v = (i, j)$, define the surrogate estimand and the standard error of its estimator by

$$\eta_{i,j} = a_1 R_1^i + a_2 R_2^j, \qquad \sigma_{i,j} = \frac{a_1 Q_1^i + a_2 Q_2^j}{4}, \tag{13}$$

where

$$R_1 = 2, \qquad R_2 = -5/4, \qquad Q_1 = 5/6, \qquad Q_2 = 2/3.$$

The target value is the limit $\theta^\star = \eta_{-\infty,-\infty} = 0$. Then each $\eta_{i,j}$ satisfies (2) with an isotonic bias bound $b_{i,j} = a_1|R_1|^i + a_2|R_2|^j$. With the exception of the grid partial order, all these quantities are unknown to the methods under evaluation.

For each grid point, we independently draw

$$Z_{i,j} \sim N(0,1), \qquad W_{i,j} \sim \chi^2_\nu, \qquad \nu = 100,$$

and define a Gaussian estimate of $\eta_{i,j}$ and an estimated standard error by

$$\widehat{m}_{i,j} = \eta_{i,j} + \sigma_{i,j} Z_{i,j}, \qquad \widehat{\sigma}_{i,j} = \sigma_{i,j}\sqrt{W_{i,j}/\nu},$$

thereby simulating statistics from a dataset of size $\nu + 1$ with Gaussian noise. The input intervals are then formed as

$$I_{i,j} = [\widehat{m}_{i,j} - z\widehat{\sigma}_{i,j}, \widehat{m}_{i,j} + z\widehat{\sigma}_{i,j}]$$

with a critical value $z > 0$ specified below.

A.2 Experimental conditions

The three experimental conditions are listed in Table A.1 from easiest to hardest. The first two are chains and the last is a two-dimensional partial order. In each condition, the scaling vector σ is subadditive with $\kappa = 1$.

Table A.1: **Synthetic experimental conditions**. $S = \{-4, -3, \ldots, 4\}$.

Condition	a_1	a_2	Grid $\mathcal{V}$	$\lvert\mathcal{V}\rvert$
Condition 1	1	0	$\{(i,0) : i \in S\}$	9
Condition 2	0	1	$\{(0,j) : j \in S\}$	9
Condition 3	1	1	$S \times S$	81

A.3 Methods and configurations

We compare the following six estimators.

- *Random interval* samples one grid point uniformly and returns its center $\widehat{m}_v$.
- *SLOPE* (Su et al., 2020) is a selection rule designed for one-dimensional inputs. For multidimensional inputs, we first linearize the grid DAG using a deterministic topological sort that places wider intervals before narrower ones, and then apply SLOPE to the resulting chain. The default value for z is 2 as specified in the original paper.
- *SLOPE++* (Tucker & Lee, 2021) is a refinement of SLOPE, with the same linearization technique. The default value for z is 1 as specified in the original implementation.
- *OACIS* is applied directly to the grid partial order, without linearizing incomparable candidates. For each repetition, set

$$\widehat{\rho} = \frac{\max_v \widehat{\sigma}_v}{\min_v \widehat{\sigma}_v}, \qquad \widehat{z} = \Phi^{-1}\left(1 - \frac{1}{2|\mathcal{V}|(1+\widehat{\rho})}\right), \tag{14}$$

  and supply OACIS with the intervals $[\widehat{m}_v - \widehat{z}\widehat{\sigma}_v, \widehat{m}_v + \widehat{z}\widehat{\sigma}_v]$ following the Wald profile and plug-in confidence-parameter rule in Table G.1 and Proposition 5.2.
- *d-OACIS* uses the same intervals and partial order as OACIS, but restricts the output to the candidate midpoints, as in Definition 3.3.
- *Hindsight oracle* chooses one fixed grid point per condition that minimizes mean absolute error over all repetitions, and returns that point's center in every repetition. It is an in-sample reference rather than a deployable selector.

A.4 Evaluation protocol

We use 100 independent repetitions. For each method and condition, we report the mean absolute error (MAE) and its standard error. For OACIS and d-OACIS, we additionally summarize their minimax inefficiencies $\bar{\alpha}^\star(\mathcal{I})$ and $\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I})$, respectively. We also evaluate sensitivity to the critical value z on the logarithmic grid

$$z \in \{2^{k/10} : k = -10, -9, \ldots, 20\},$$

ranging from 0.5 to 4 and containing the reference values $z = 1$ and $z = 2$ exactly.

A.5 MINIMAX INEFFICIENCIES

Table A.2 reports the minimax inefficiencies of OACIS and d-OACIS across the synthetic conditions.

Table A.2: Synthetic minimax inefficiencies. Entries report means (standard deviations) over 100 independent repetitions.

Condition	Continuous $\bar{\alpha}^\star(\mathcal{I})$	Discretized $\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I})$
1	$1.234\ (0.097)$	$1.361\ (0.182)$
2	$1.292\ (0.110)$	$1.446\ (0.127)$
3	$1.260\ (0.084)$	$1.275\ (0.084)$

A.6 ADDITIONAL RELATIVE COMPARISON

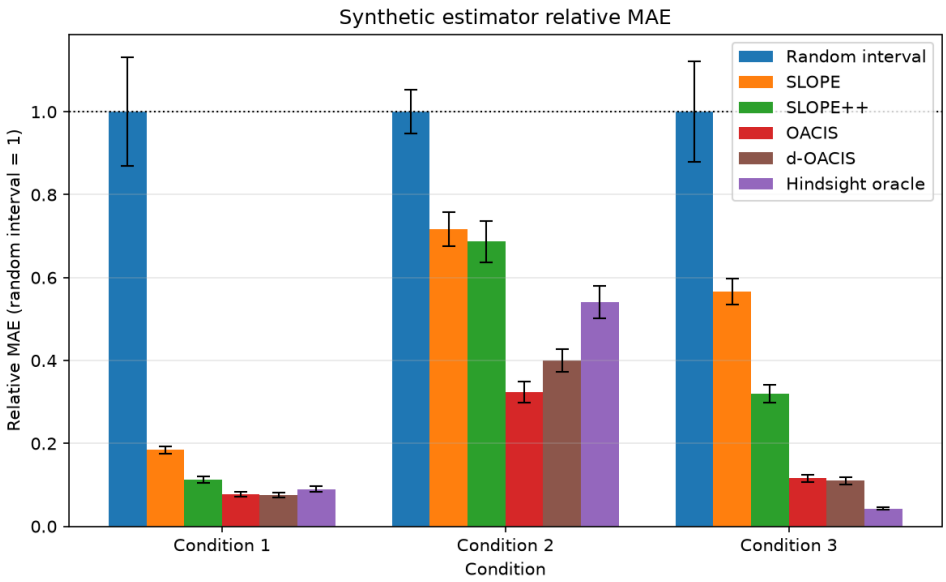


Figure A.1: Synthetic estimator MAE relative to the random-interval baseline. Error bars are the standard errors of the unnormalized MAEs divided by the random-interval MAE.

A.7 CRITICAL-VALUE SENSITIVITY

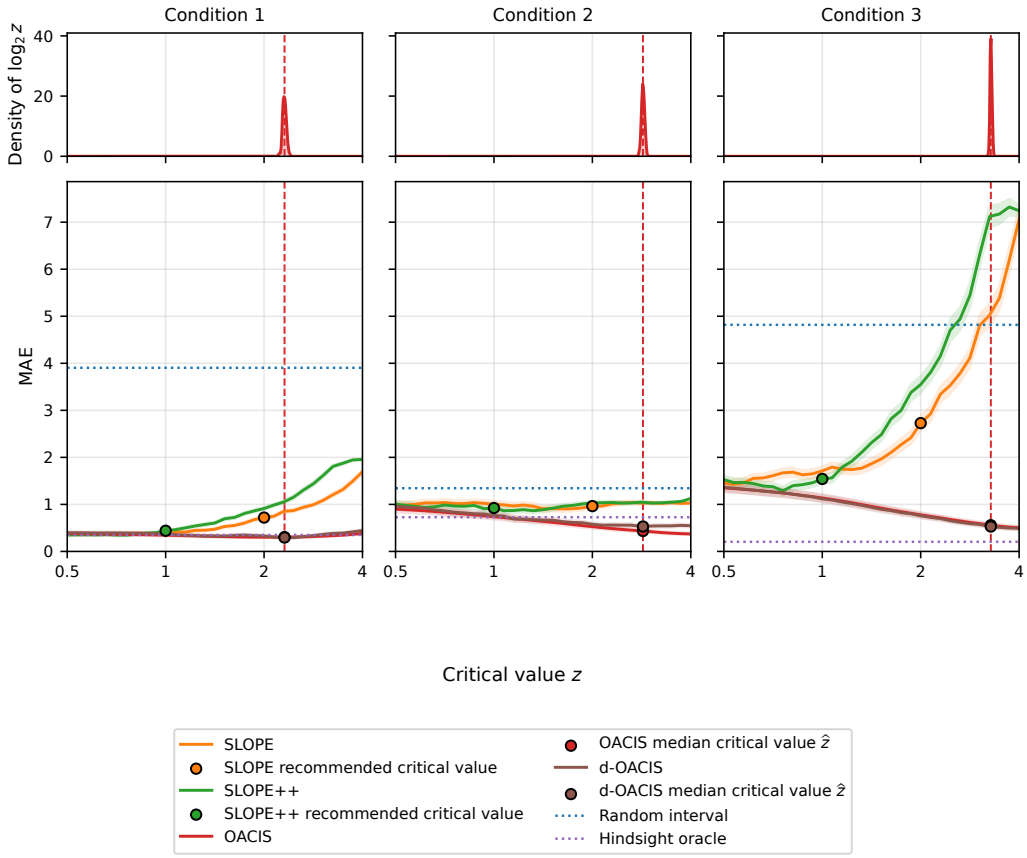


Figure A.2: Sensitivity to the critical value z. The columns correspond to the three conditions and use common vertical scales within each row. The upper row shows kernel density estimates of the plug-in OACIS critical value on the log scale; the lower row reports MAE at fixed critical values z for SLOPE, SLOPE++, OACIS, and d-OACIS with pointwise standard-error bands. The dashed vertical line in each panel marks the median plug-in critical value, and the red and brown points in each lower panel mark the OACIS and d-OACIS MAE values, obtained by linear interpolation with respect to $\log z$.

B ADDITIONAL DETAILS FOR THE REAL-DATA EXPERIMENT

B.1 DATA AND ESTIMAND

We consider the problem of off-policy evaluation (OPE) for user-routing policies. On an online platform, these policies select, for each user, a service from a fixed finite set $\mathcal{A}$ of candidate services to maximize user engagement. We use data from two buckets of an online experiment conducted from January 2 to February 8, 2026. Group A supplies the behavior-policy data, while Group B supplies the target-policy data. The groups contain approximately equal numbers of users.

We treat each date as a separate evaluation unit and use the per-user click rate as the outcome. Dates without policy deployment are excluded from evaluation since clicks are not recorded, leaving 25 dates for evaluation. For each date t, the prediction target $\theta_t^\star$ is the daily difference in expected per-user outcomes, approximated by the difference y_t between the Group B and Group A sample means:

$$\theta_t^\star := \mathbb{E}^{\pi_e}[r_t] - \mathbb{E}^{\pi_b}[r_t] \approx y_t := \frac{1}{|N^B|}\sum_{i\in N^B} r_{i,t} - \frac{1}{|N^A|}\sum_{i\in N^A} r_{i,t},$$

where π_e and π_b are the target and behavior policies, N^A and N^B are the sets of users in Groups A and B, respectively, and $r_{i,t}$ is the outcome for user i on date t.

B.2 CANDIDATE ESTIMATORS

The candidate point estimators $\widehat{m}_{t,v}$ are constructed using a modification of the partial-history inverse-propensity weighting (IPW) method (Hu & Wager, 2023), with hyperparameters $v = (h, \ell)$ consisting of history length $h \geq 0$ and impression-effect duration $\ell \geq 0$. Suppressing the candidate index v for notational simplicity, the estimator is

$$\widehat{m}_t := \frac{1}{|N^A|}\sum_{i\in N^A} (\widetilde{r}_{i,t} - \overline{r}_t)(\widetilde{\rho}_{i,t} - 1),$$

where

$$\widetilde{r}_{i,t} := E_{i,t}^{\text{imp}} r_{i,t}, \qquad \overline{r}_t := \frac{1}{|N^A|}\sum_{i\in N^A} \widetilde{r}_{i,t},$$

are the modified reward and its baseline,

$$E_{i,t}^{\text{imp}} := \begin{cases} 1, & \text{if } \ell = \infty \text{ or user } i \text{ receives at least one impression during } [t-\ell, t], \\ 0, & \text{otherwise,} \end{cases}$$

is the exposure indicator, and

$$\widetilde{\rho}_{i,t} := \prod_{s=t-h}^{t} \frac{\pi_e(a_{i,s} \mid x_{i,s})}{\pi_b(a_{i,s} \mid x_{i,s})}$$

is the partial-history importance weight. Here, $a_{i,s} \in \mathcal{A}$ is the action taken for user i at time s, and $x_{i,s}$ is the corresponding context, including the user and item features.

The corresponding candidate estimand is the population functional

$$\eta_{t,v} := \mathbb{E}^{\pi_b}\left[(\widetilde{r}_{t,v} - \mathbb{E}^{\pi_b}[\widetilde{r}_{t,v}])\,(\widetilde{\rho}_{t,v} - 1)\right],$$

where the expectation is over a user drawn under the behavior policy. The estimator $\widehat{m}_{t,v}$ is the corresponding sample analogue, with $\overline{r}_{t,v}$ replacing the population mean.

B.3 CONFIDENCE INTERVALS

To construct confidence intervals, we define the influence-function values

$$\widehat{\text{IF}}_{i,t,v} := (\widetilde{r}_{i,t,v} - \overline{r}_{t,v})(\widetilde{\rho}_{i,t,v} - \overline{\rho}_{t,v}) - \widehat{m}_{t,v},$$

where $\overline{\rho}_{t,v} := \frac{1}{|N^A|}\sum_{i\in N^A} \widetilde{\rho}_{i,t,v}$ is the mean partial-history weight. The standard error of $\widehat{m}_{t,v}$ is estimated by

$$\widehat{\sigma}_{t,v} := \left\{ \frac{1}{|N^A|(|N^A|-1)} \sum_{i\in N^A} \widehat{\text{IF}}_{i,t,v}^{\,2} \right\}^{1/2}.$$

The resulting Wald-type interval estimates of $\eta_{t,v}$ are

$$I_{t,v} := [\widehat{m}_{t,v} - z\widehat{\sigma}_{t,v}, \widehat{m}_{t,v} + z\widehat{\sigma}_{t,v}],$$

where $z > 0$ is set according to each method's configuration, as described in Section A.

B.4 Hyperparameter grid and partial order

The hyperparameters are summarized in Table B.1. The history length h reflects our belief about the history dependence of user behavior. A finite ℓ is justified by our *a priori* expectation that the lift in clicks is mostly mediated by impressions, whose effect should decay over time.

Table B.1: Hyperparameter grids for the real-data experiment. Values in each grid are listed from smaller to larger assumed bias bounds.

Hyperparameter	Notation	Grid
IPW history length	h	$\{6, 5, \ldots, 0\}$
Impression effect duration	ℓ	$\{\infty, 6, 5, \ldots, 0\}$
Non-abstention grid dimension	—	2
Candidates per date (plus one for abstention)	—	$56 + 1$

The partial order $\preceq$ is specified so that the hyperparameter values on the left precede those on the right in each dimension. Abstention is included by adding the candidate estimand $\eta_{t,\top} := 0$ and its constant zero-width interval $I_\top := [0, 0]$ as the greatest element $\top$.

B.5 Methods and evaluation protocol

We compare two fixed baseline estimators (*No OPE*, *Naive Bandit*) and four adaptive estimators (SLOPE, SLOPE++, OACIS, d-OACIS). No OPE predicts zero for all dates, corresponding to the abstention $\top$. Naive Bandit uses the fixed hyperparameter combination $h = 0$ and $\ell = \infty$. The adaptive methods use the same configurations as in the synthetic experiment (Section A.3) and are applied to the full (h, ℓ) grid of 56 candidates plus abstention. Here, OACIS and d-OACIS choose their critical values $\widehat{z}$ by estimating $\widehat{\rho}$ as described in (14) with the abstention candidate $v = \top$ excluded from the maximum and minimum.

For the observed A/B differences y_t and predictions $\widehat{\theta}_t$, we report MAE normalized by the Naive Bandit MAE (NMAE) and its standard error across dates, together with the Pearson correlation coefficient and its p-value. Both metrics are computed over the 25 evaluation dates.

C PROOFS OF THE EPIGRAPH CHARACTERIZATION, CLOSED FORMS, AND DISCRETIZATION BOUNDS

This appendix first establishes elementary facts about the least compatible bias bound. It then proves the finite epigraph characterization in Proposition 4.1, derives the closed-form solution in Theorem 4.1, and compares continuous and discretized minimax inefficiency in Proposition 4.2.

C.1 AUXILIARY LEMMAS FOR PROPOSITION 4.1

Throughout this appendix, fix $\mathcal{I}$ and use the notation introduced before Proposition 4.1. For a target location $t \in \mathbb{R}$, define

$$\eta_v^\star(t) := \Pi_{I_v}(t), \qquad b_v^\star(t) := \max_{w \preceq v} \operatorname{dist}(t, I_w),$$

where $\Pi_{I_v}(t)$ is the Euclidean projection of t onto I_v, and the corresponding conditional worst-case instance

$$\mathfrak{s}^\star(t) = \big(t,\ (\eta_v^\star(t))_{v\in\mathcal{V}},\ (b_v^\star(t))_{v\in\mathcal{V}}\big).$$

Lemma C.1 (Compatibility of the conditional worst-case instance)**.** *For every $t \in \mathbb{R}$, the instance $\mathfrak{s}^\star(t)$ is compatible with $\mathcal{I}$.*

Proof. By the projection property, $\eta_v^\star(t) \in I_v$ and

$$|\eta_v^\star(t) - t| = \operatorname{dist}(t, I_v) \le b_v^\star(t).$$

If $w \preceq v$, then $\{a : a \preceq w\} \subseteq \{a : a \preceq v\}$, so $b_w^\star(t) \le b_v^\star(t)$. □

Lemma C.2 (Least compatible bias bound)**.** *For any compatible instance $\mathfrak{s} = (t, \eta, b)$ and any $v \in \mathcal{V}$,*

$$b_v \ge b_v^\star(t).$$

Proof. For every $w \preceq v$, compatibility gives

$$\operatorname{dist}(t, I_w) \le |\eta_w - t| \le b_w \le b_v,$$

where the first inequality uses $\eta_w \in I_w$. Taking the maximum over $w \preceq v$ gives the claim. □

Lemma C.3 (Unit lower bound)**.** *Fix $x \in \mathbb{R}$ and $y < \infty$. If $\bar\alpha(x; \mathcal{I}) \le y$, then $y \ge 1$.*

Proof. Choose a minimal element $v \in \mathcal{V}$. If $u_v = 0$, write $I_v = \{t\}$. By Lemma C.1, $\mathfrak{s}^\star(t)$ is compatible. For this instance, $R^\star(\mathfrak{s}^\star(t), \mathcal{I}) = 0$. Since $\bar\alpha(x; \mathcal{I}) \le y < \infty$, this forces $x = t$, and the convention $0/0 = 1$ then gives $\bar\alpha(x; \mathcal{I}) \ge 1$.

If $u_v > 0$, restrict the supremum in $\bar\alpha(x; \mathcal{I})$ to the two instances $\mathfrak{s}^\star(\ell_v)$ and $\mathfrak{s}^\star(r_v)$. Minimality gives $b_v^\star(\ell_v) = b_v^\star(r_v) = 0$, hence

$$\bar\alpha(x; \mathcal{I}) \ge \max_{t \in \{\ell_v, r_v\}} \frac{|x - t|}{u_v} \ge 1.$$

□

C.2 PROOF OF PROPOSITION 4.1

Proof of Proposition 4.1. For every $v \in \mathcal{V}$ and $t \in \mathbb{R}$, the least compatible bias bound satisfies

$$\begin{aligned} b_v^\star(t) &= \max_{w \preceq v} \operatorname{dist}(t, I_w) \\ &= \max\{L_v - t, 0, t - R_v\} \\ &= \delta_v^- + \big(|t - \mu_v| - \delta_v^+\big)_+ . \end{aligned}$$

Suppose first that $\bar\alpha(x; \mathcal{I}) \le y$. By Lemma C.3, $y \ge 1$. Fix $v \in \mathcal{V}$ and $t \in \mathbb{R}$. Applying the assumed bound to the compatible instance $\mathfrak{s}^\star(t)$ gives

$$|x - t| \le y R^\star(\mathfrak{s}^\star(t), \mathcal{I}) \le y\{u_v + b_v^\star(t)\}.$$

The first inequality remains valid when the oracle error is zero: finiteness of the ratio then forces $x = t$. At $t = \mu_v \pm \delta_v^+$, the displayed identity gives $b_v^\star(t) = \delta_v^-$, and hence

$$|x - \mu_v \mp \delta_v^+| \le y h_v.$$

Taking the maximum of these two inequalities and using

$$\max_{\tau \in \{-1,+1\}} |x - \mu_v - \tau \delta_v^+| = |x - \mu_v| + \delta_v^+$$

proves the necessity of (6).

Conversely, suppose that (6) holds and fix a compatible instance $\mathfrak{s} = (t, \eta, b)$. For each $v \in \mathcal{V}$, Lemma C.2, the triangle inequality, and $h_v = u_v + \delta_v^-$ give

$$\begin{aligned} |x - t| &\le |x - \mu_v| + |t - \mu_v| \\ &\le y h_v - \delta_v^+ + |t - \mu_v| \\ &\le y(u_v + \delta_v^-) + y\left(|t - \mu_v| - \delta_v^+\right)_+ \\ &= y\{u_v + b_v^\star(t)\} \\ &\le y(u_v + b_v). \end{aligned}$$

Here, the third line uses $a \le y\,(a)_+$ for every $a \in \mathbb{R}$ when $y \ge 1$. Taking the minimum over v yields

$$|x - t| \le y R^\star(\mathfrak{s}, \mathcal{I}).$$

If the oracle error is positive, dividing both sides by it yields $\alpha(x; \mathfrak{s}, \mathcal{I}) \le y$. If it is zero, the displayed inequality forces $x = t$, so $\alpha(x; \mathfrak{s}, \mathcal{I}) = 1 \le y$ by the convention $0/0 = 1$. Taking the supremum over compatible instances completes the proof. □

C.3 Proof of Theorem 4.1

Proof of Theorem 4.1. Because $L_v \ge \ell_v$ and $R_v \le r_v$,

$$\delta_v = \frac{R_v - L_v}{2} \le \frac{r_v - \ell_v}{2} = u_v.$$

Since $u_v \ge 0$, this implies

$$0 \le \delta_v^+ \le u_v \le h_v. \tag{15}$$

For $y \ge 1$, the constraints in Proposition 4.1 are

$$|x - \mu_v| + \delta_v^+ \le h_v y, \qquad v \in \mathcal{V}.$$

By (15), each constraint is equivalent to

$$\mu_v + \delta_v^+ - h_v y \le x \le \mu_v - \delta_v^+ + h_v y. \tag{16}$$

The intervals in (16) have a common point if and only if they intersect pairwise. For distinct v, w with $h_v + h_w > 0$, the pairwise-intersection condition is

$$|\mu_v - \mu_w| \le (h_v + h_w) y - \delta_v^+ - \delta_w^+,$$

or equivalently

$$y \ge \frac{|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+}{h_v + h_w}. \tag{17}$$

If $v, w \in Z$, then (15) gives $\delta_v^+ = \delta_w^+ = 0$, and their feasible intervals are the singletons $\{\mu_v\}$ and $\{\mu_w\}$. Hence, feasibility for some finite y requires (7). Conversely, under that condition, all intervals corresponding to indices with $h_v = 0$ coincide, while each pair with a positive denominator intersects whenever y satisfies its finite lower bound above. Because the family is finite, a finite y satisfying all these bounds exists. This proves the finiteness characterization.

The smallest $y \ge 1$ satisfying all pairwise lower bounds is exactly (8). At this value, the common intersection of the intervals in (16) is

$$\left[\max_{v \in \mathcal{V}} \{\mu_v + \delta_v^+ - h_v \bar{\alpha}^\star(\mathcal{I})\}, \qquad \min_{v \in \mathcal{V}} \{\mu_v - \delta_v^+ + h_v \bar{\alpha}^\star(\mathcal{I})\}\right].$$

By Proposition 4.1, this is precisely the set of minimax-optimal outputs. If $\bar{\alpha}^\star(\mathcal{I}) > 1$, at least one qualifying pair attains the inner maximum in (8). A direct calculation using (16) shows that the corresponding two feasible intervals meet at a single point, so the full intersection is a singleton.

If $\bar{\alpha}^\star(\mathcal{I}) = 1$, choose a minimal element $v_0 \in \mathcal{V}$. Then $L_{v_0} = \ell_{v_0}$ and $R_{v_0} = r_{v_0}$, so $\delta^+_{v_0} = h_{v_0} = u_{v_0}$. Its feasible interval at $y = 1$ is the singleton $\{\mu_{v_0}\}$. The full intersection is nonempty and contained in this interval, so it is the same singleton. Thus the minimax-optimal output is unique. Equating the largest left endpoint with the smallest right endpoint proves (9). □

C.4 Proof of the bounds on discretized minimax inefficiency

Proof of Proposition 4.2. Because $\{m_v : v \in \mathcal{V}\} \subset \mathbb{R}$, minimizing over all real outputs gives the lower bound $\bar{\alpha}^\star(\mathcal{I}) \le \bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I})$.

For the upper bound, we consider the selection rule that outputs m_q with $q \in \arg\min_{v\in\mathcal{V}} h_v$. Put $y^\star = \bar{\alpha}^\star(\mathcal{I})$ for brevity. If $y^\star = \infty$, the upper bound is immediate; hence suppose $y^\star < \infty$. Now, every candidate midpoint satisfies

$$|m_v - \mu_v| + \delta^+_v \le h_v. \tag{18}$$

Indeed, using $\mu_v - \delta_v = L_v \ge \ell_v = m_v - u_v$ and $\mu_v + \delta_v = R_v \le r_v = m_v + u_v$, we have

$$|m_v - \mu_v| \le u_v - \delta_v.$$

Adding δ^+_v to both sides and using $-\delta_v + \delta^+_v = \delta^-_v$ and $h_v = u_v + \delta^-_v$ yields (18). We next split cases depending on whether h_q is positive or zero.

Case 1 ($h_q > 0$). For every $w \ne q$, the closed-form characterization (8) gives

$$|\mu_q - \mu_w| + \delta^+_q + \delta^+_w \le y^\star(h_q + h_w).$$

Adding (18) with $v = q$ to both sides and using the triangle inequality gives

$$\begin{aligned} |m_q - \mu_w| + \delta^+_w &\le h_q + y^\star(h_q + h_w) - 2\delta^+_q \\ &\le h_q + y^\star(h_q + h_w) \\ &\le (1 + 2y^\star)h_w, \end{aligned}$$

where the last step uses $h_q \le h_w$ by the choice of q. For $w = q$, the same final bound follows directly from (18). Hence Proposition 4.1 implies

$$\bar{\alpha}(m_q; \mathcal{I}) \le 1 + 2y^\star.$$

Case 2 ($h_q = 0$). According to Eq. (6) with $v = q$, every output with finite worst-case inefficiency, specifically $x = \bar{\theta}^\star(\mathcal{I})$, must be equal to μ_q. Similarly, Eq. (18) forces $m_q = \mu_q$. Thus, we have $m_q = \mu_q = \bar{\theta}^\star(\mathcal{I})$ and $\bar{\alpha}(m_q; \mathcal{I}) = \bar{\alpha}^\star(\mathcal{I})$.

In both cases, the optimality of d-OACIS among the input midpoints gives

$$\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}) \le \bar{\alpha}(m_q; \mathcal{I}) \le 1 + 2\bar{\alpha}^\star(\mathcal{I}),$$

as claimed. □

D Proof of Theorem 4.2

Proof of Theorem 4.2. Fix a finite $\kappa \in [1/2, \infty)$, $\mathcal{I} = (m, u) \in \mathbb{R}^{\mathcal{V}} \times \mathbb{S}^{\mathcal{V}}_\kappa$, and $v, w \in \mathcal{V}$. The identities

$$\mu_v + \delta^+_v = R_v + \delta^-_v, \qquad \mu_v - \delta^+_v = L_v - \delta^-_v$$

give

$$|\mu_v - \mu_w| + \delta^+_v + \delta^+_w = \delta^-_v + \delta^-_w + \max\{R_v - L_w, R_w - L_v\}. \tag{19}$$

Because $v \wedge w \preceq v$ and $v \wedge w \preceq w$, the definitions of L_v and R_v imply

$$R_v \le m_{v\wedge w} + u_{v\wedge w}, \qquad L_w \ge m_{v\wedge w} - u_{v\wedge w},$$

and hence $R_v - L_w \le 2u_{v\wedge w}$. Interchanging v and w gives $R_w - L_v \le 2u_{v\wedge w}$. Therefore

$$\begin{aligned}|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+ &\le \delta_v^- + \delta_w^- + 2u_{v\wedge w}\\ &\le \delta_v^- + \delta_w^- + 2\kappa(u_v + u_w)\\ &\le 2\kappa(h_v + h_w),\end{aligned}$$

where the last step uses $h_v + h_w = u_v + u_w + \delta_v^- + \delta_w^-$ and $2\kappa \ge 1$.

If $h_v = h_w = 0$, the preceding inequality forces the nonnegative left-hand side to be zero, and in particular $\mu_v = \mu_w$. Thus (7) holds. If $h_v + h_w > 0$, division by $h_v + h_w$ shows that the corresponding pairwise quotient in (8) is at most 2κ. Since $2\kappa \ge 1$, Theorem 4.1 now gives

$$\bar{\alpha}^\star(\mathcal{I}) \le 2\kappa.$$

Suppose next that $(\mathcal{V}, \preceq)$ is totally ordered. Then $v \wedge w \in \{v, w\}$, so (19) also gives

$$\begin{aligned}|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+ &\le \delta_v^- + \delta_w^- + 2(u_v + u_w)\\ &\le 2(h_v + h_w).\end{aligned}$$

Applying the same reasoning to pairs with $h_v + h_w = 0$ and invoking Theorem 4.1 gives $\bar{\alpha}^\star(\mathcal{I}) \le 2$. Combining this inequality with the general upper bound yields

$$\bar{\alpha}^\star(\mathcal{I}) \le \min\{2\kappa, 2\}.$$

We next prove the matching lower bounds. First, suppose that $(\mathcal{V}, \preceq)$ is totally ordered, let v_0 be its least element, and put $\lambda := \min\{\kappa, 1\}$. Define

$$I_{v_0} = [0, \lambda], \qquad I_x = [0, 1 - \lambda], \quad \forall x \ne v_0.$$

The radii are $u_{v_0} = \lambda/2$ and $u_x = (1 - \lambda)/2$ for $x \ne v_0$. For the pair (v_0, v_0), κ-subadditivity follows from $1 \le 2\kappa$. For a pair (v_0, x) with $x \ne v_0$, the condition reduces to $\lambda \le \kappa$. For a pair not containing v_0, it reduces to $1 - \lambda \le 2\kappa(1 - \lambda)$. Thus the radius vector is κ-subadditive. Moreover, $[0, 1 - \lambda] \subseteq [0, \lambda]$ because $\lambda \ge 1/2$. Thus,

$$\bigcap_{w \preceq v_0} I_w = I_{v_0} = [0, \lambda], \qquad \bigcap_{w \preceq x} I_w = I_x = [0, 1 - \lambda], \quad x \ne v_0.$$

Consequently,

$$\mu_{v_0} = \delta_{v_0}^+ = h_{v_0} = \frac{\lambda}{2}, \qquad \mu_x = \delta_x^+ = h_x = \frac{1 - \lambda}{2}.$$

For any $x \ne v_0$, the quotient for the pair (v_0, x) in (8) is

$$\frac{\lambda}{\lambda/2 + (1 - \lambda)/2} = 2\lambda = \min\{2\kappa, 2\}.$$

This proves the total-order lower bound.

Finally, suppose that $(\mathcal{V}, \preceq)$ is not totally ordered. Choose incomparable $v, w \in \mathcal{V}$, put $q := v \wedge w$, and write $M := 2\kappa$. Define

$$I_x = \begin{cases} [-M, M], & x \preceq q,\\ [M - 2, M], & x \preceq v,\ x \not\preceq q,\\ [-M, -M + 2], & x \preceq w,\ x \not\preceq q,\\ [-1, 1], & \text{otherwise.}\end{cases}$$

The second and third cases are mutually exclusive, since any common predecessor of v and w is below q. The corresponding radius vector is

$$u_x = \begin{cases} M, & x \preceq q,\\ 1, & x \not\preceq q.\end{cases}$$

To verify κ-subadditivity, take any $a, b \in \mathcal{V}$. If $a \wedge b \preceq q$, then

$$u_{a\wedge b} = M = 2\kappa \le \kappa(u_a + u_b),$$

because every radius is at least one. Otherwise, neither a nor b is below q, and

$$u_{a\wedge b} = u_a = u_b = 1 \le 2\kappa.$$

For the fixed incomparable elements v, w, we have $v \not\preceq q$ and $w \not\preceq q$, so $I_v = [M-2, M]$ and $I_w = [-M, -M+2]$. Every predecessor of v is assigned either I_v or $[-M, M]$, and every predecessor of w is assigned either I_w or $[-M, M]$. Since $M \ge 1$, both I_v and I_w are contained in $[-M, M]$. Therefore,

$$\bigcap_{x \preceq v} I_x = I_v = [M-2, M], \qquad \bigcap_{x \preceq w} I_x = I_w = [-M, -M+2].$$

It follows that

$$(\mu_v, \delta_v^+, h_v) = (M-1, 1, 1), \qquad (\mu_w, \delta_w^+, h_w) = (-M+1, 1, 1).$$

The quotient for the pair (v, w) in (8) is

$$\frac{|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+}{h_v + h_w} = \frac{(2M-2)+1+1}{1+1} = M = 2\kappa.$$

Together with the upper bounds, these constructions prove both cases in the theorem for every finite $\kappa \in [1/2, \infty)$. □

Remark D.1 (Relation to previous one-dimensional results). Lepski-type methods (Lepskii, 1991; Lepski et al., 1997), including the off-policy evaluation methods SLOPE (Su et al., 2020) and SLOPE++ (Tucker & Lee, 2021), also address settings in which $\mathcal{V}$ is totally ordered. The essential difference is that we consider arbitrary aggregation rules $\mathcal{I} \mapsto \mathsf{Agg}(\mathcal{I})$, whereas Lepski-type methods primarily select the midpoint of a candidate interval.

To our knowledge, the smallest previously known upper bound on the worst-case inefficiency of a selection rule is $\sqrt{6} + 3 \approx 5.45$, attained by SLOPE++ (Tucker & Lee, 2021). The monotone-uncertainty condition used in the SLOPE++ guarantee can be enforced without loss of generality. Indeed, if $w \preceq v$ and $u_w \le u_v$, then v is oracle-dominated by w, because every isotonic bias bound satisfies $b_w + u_w \le b_v + u_v$. Iteratively deleting each candidate whose radius is no smaller than that of one of its predecessors therefore preserves $\min_{v\in\mathcal{V}}(b_v + u_v)$ for every isotonic b, while the retained chain has decreasing radii. The SLOPE++ guarantee applies to this reduced chain and therefore yields a guarantee for a selection rule on the original chain.

Since selection rules are special cases of aggregation rules, this guarantee also gives an upper bound for the general aggregation class. Corollary 4.1 sharpens this upper bound to the optimal constant 2. Combining this result with Proposition 4.2 also gives $\bar{\alpha}_{\mathrm{d}}^{\star}(\mathcal{I}) \le 5$ for every input on a total order, improving the previous upper bound for selection rules.

E Details for Remark 4.1

Verification of Remark 4.1. Write

$$a(v) := \big(a_1(v_1), \dots, a_d(v_d)\big).$$

For each coordinate j, $(v \wedge w)_j$ equals either v_j or w_j. Since the functions a_j are nonnegative, this gives the coordinatewise inequality

$$a(v \wedge w) \le a(v) + a(w).$$

The coordinatewise monotonicity and κ-subadditivity of F therefore give

$$\begin{aligned} u_{v\wedge w} &= F\big(a(v \wedge w)\big) \\ &\le F\big(a(v) + a(w)\big) \\ &\le \kappa\{F\big(a(v)\big) + F\big(a(w)\big)\} \\ &= \kappa(u_v + u_w). \end{aligned}$$

Thus u is κ-subadditive. □

Example E.1 (Subadditive aggregation functions). For $\kappa = 1$, the conditions on F in Remark 4.1 are that it is nonnegative, coordinatewise nondecreasing, and subadditive. We call maps satisfying these conditions *admissible*. Representative choices of F include weighted ℓ_p aggregations of the following forms:

$$F(x) = \left(\sum_{j=1}^{d} (\lambda_j x_j)^p \right)^{1/p}, \quad 1 \le p < \infty, \qquad F(x) = \max_{1 \le j \le d} \lambda_j x_j,$$

where $\lambda_j \ge 0$. The case $p = 1$ includes additive separability.

Moreover, the class of admissible choices for F is closed under composition in the following sense. If $\Phi : \mathbb{R}^m_{\ge 0} \to \mathbb{R}_{\ge 0}$ and $G_k : \mathbb{R}^d_{\ge 0} \to \mathbb{R}_{\ge 0}$, $k = 1, \dots, m$, are admissible, then the following function is also admissible:

$$H(x) = \Phi\big(G_1(x), \dots, G_m(x)\big).$$

This yields, for example, $H(x) = c + F(x), \quad c \ge 0$, $H(x) = \max_{1 \le k \le m} c_k^\top x$, $c_k \in \mathbb{R}^d_{\ge 0}$, and $H(x) = (\sum_{j=1}^d \lambda_j x_j)^q$, $\lambda_j \ge 0$, $0 < q \le 1$, as admissible choices.

Proof of Example E.1. For $1 \le p < \infty$, the weighted ℓ_p aggregation is coordinatewise nondecreasing on $\mathbb{R}^d_{\ge 0}$, and its subadditivity follows from the triangle inequality for the ℓ_p norm.

To prove closure under composition, let $G = (G_1, \dots, G_m)$. Coordinatewise monotonicity of all component maps and of Φ implies that $H = \Phi \circ G$ is coordinatewise nondecreasing. Moreover, componentwise subadditivity gives $G(x + y) \le G(x) + G(y)$, and hence

$$\begin{aligned} H(x + y) &= \Phi\big(G(x + y)\big) \\ &\le \Phi\big(G(x) + G(y)\big) \\ &\le \Phi\big(G(x)\big) + \Phi\big(G(y)\big) \\ &= H(x) + H(y). \end{aligned}$$

This proves closure under the stated compositions. Nonnegative constant maps, coordinatewise maximum, nonnegative linear maps, and $t \mapsto t^q$ for $0 < q \le 1$ are also coordinatewise nondecreasing and subadditive, so the displayed examples follow from the composition result. □

Example E.2 (Composite estimation). Let $\mathcal{V}$ be as in Example 3.1. Define the target and candidate estimands in composite form by

$$\theta^\star = g(\xi_1^\star, \dots, \xi_d^\star), \qquad \eta_v = g(\xi_{1,v_1}, \dots, \xi_{d,v_d}),$$

where $g : \mathbb{R}^d \to \mathbb{R}$ is a composition map, $\xi_j^\star$ is a component target, and $(\xi_{j,t})_{t \in \mathcal{V}_j}$ is a sequence of candidate estimands for $1 \le j \le d$. Suppose that each $\xi_{j,t}$ is individually estimated using an interval

$$J_{j,t} = [\widehat{\xi}_{j,t} - a_j(t), \widehat{\xi}_{j,t} + a_j(t)], \qquad \widehat{\xi}_{j,t} \in \mathbb{R},\ a_j(t) \ge 0,$$

and that the composition map g satisfies the error-propagation bound

$$|g(x) - g(y)| \le F(|x - y|), \qquad x, y \in \mathbb{R}^d,$$

where $F : \mathbb{R}^d_{\ge 0} \to \mathbb{R}_{\ge 0}$ is coordinatewise nondecreasing and subadditive, and where the absolute value is taken coordinatewise. Given an approximation $\widehat{g} : \mathbb{R}^d \to \mathbb{R}$ satisfying the uniform error bound $\|g - \widehat{g}\|_\infty \le c$, it is natural to construct the composite interval estimator $I_v = [m_v - u_v, m_v + u_v]$ by setting

$$m_v = \widehat{g}(\widehat{\xi}_{1,v_1}, \dots, \widehat{\xi}_{d,v_d}), \qquad u_v = c + F(a(v)).$$

The error-propagation bound and monotonicity of F then imply $|m_v - \eta_v| \le u_v$ whenever every $J_{j,t}$ covers $\xi_{j,t}$. The resulting radius vector u is subadditive because Remark 4.1 with $\kappa = 1$ applies with F replaced by $c + F$, which remains subadditive. A canonical choice for F is the scaled ℓ_p norm $F(x) = L(\sum_j |x_j|^p)^{1/p}$ for $1 \le p < \infty$, corresponding to the L-Lipschitz continuity of g with respect to the ℓ_p norm. See Example E.1 in Section E for other admissible choices.

F PROOFS OF THE STABILITY BOUNDS

This appendix proves the gradient-control lemma and uses it to derive Lipschitz continuity from bounded inefficiency.

F.1 PROOF OF GRADIENT CONTROL

Lemma F.1 (Gradient control)**.** *For every $\mathcal{I}$ in the (Euclidean) interior of $\mathbb{I}$, the value $\bar{\alpha}^\star(\mathcal{I})$ is finite and the minimax output $\bar{\theta}^\star(\mathcal{I})$ is uniquely defined. Moreover, both maps are locally Lipschitz, and*

$$\|\nabla_m \bar{\theta}^\star(\mathcal{I})\|_1 \le \bar{\alpha}^\star(\mathcal{I}), \qquad \|(\nabla_u \bar{\theta}^\star(\mathcal{I}))_+\|_1 \le \bar{\alpha}^\star(\mathcal{I}), \qquad \|(\nabla_u \bar{\theta}^\star(\mathcal{I}))_-\|_1 \le \bar{\alpha}^\star(\mathcal{I}) \tag{20}$$

for almost every $\mathcal{I}$ in the interior of $\mathbb{I}$. Consequently, $\|\nabla_{(m,u)} \bar{\theta}^\star(\mathcal{I})\|_1 \le 3\bar{\alpha}^\star(\mathcal{I})$ almost everywhere there.

Proof. Each affine map $\mathcal{I} \mapsto m_w \pm u_w$ is 1-Lipschitz with respect to $\|\cdot\|_\infty$ in each of the m- and u-blocks. Finite maxima and minima preserve these bounds. Since

$$\delta_v^- = \frac{1}{2}(L_v - R_v)_+,$$

the map δ_v^- is also 1-Lipschitz in each block. Put

$$\widetilde{\alpha}(\mathcal{I}) := \bar{\alpha}^\star(\mathcal{I}) - 1, \qquad P_v := R_v - u_v, \qquad Q_v := L_v + u_v.$$

It follows that P_v, Q_v, and $h_v = u_v + \delta_v^-$ are 1-Lipschitz in the m-block and 2-Lipschitz in the u-block, and hence 3-Lipschitz in $\mathcal{I}$.

On the interior of $\mathbb{I}$, every h_v is positive. Hence the closed forms (8) and (9) show that $\bar{\alpha}^\star(\mathcal{I})$ is finite, $\bar{\theta}^\star(\mathcal{I})$ is uniquely defined, and both maps are locally Lipschitz. By Rademacher's theorem, there is a full-measure subset of the interior of $\mathbb{I}$ on which $\bar{\alpha}^\star$, $\bar{\theta}^\star$, and the finitely many maps $L_v, R_v, \delta_v^-, P_v, Q_v, h_v$ are simultaneously differentiable.

Fix a point in this subset. Since $\mu_v + \delta_v^+ = R_v + \delta_v^-$, $\mu_v - \delta_v^+ = L_v - \delta_v^-$, and $h_v = u_v + \delta_v^-$, the two representations of $\bar{\theta}^\star(\mathcal{I})$ become

$$\bar{\theta}^\star(\mathcal{I}) = \max_{v \in \mathcal{V}} \{P_v - \widetilde{\alpha}(\mathcal{I}) h_v\} = \min_{v \in \mathcal{V}} \{Q_v + \widetilde{\alpha}(\mathcal{I}) h_v\}.$$

Choose v_+ attaining the maximum and v_- attaining the minimum at this point. The nonnegative functions $\bar{\theta}^\star - (P_{v_+} - \widetilde{\alpha} h_{v_+})$ and $(Q_{v_-} + \widetilde{\alpha} h_{v_-}) - \bar{\theta}^\star$ both attain a local minimum of zero there. Differentiability therefore yields

$$\begin{aligned}
\nabla_{(m,u)} \bar{\theta}^\star(\mathcal{I}) &= \nabla_{(m,u)} P_{v_+} - \widetilde{\alpha}(\mathcal{I}) \nabla_{(m,u)} h_{v_+} - h_{v_+} \nabla_{(m,u)} \widetilde{\alpha}(\mathcal{I}), \\
\nabla_{(m,u)} \bar{\theta}^\star(\mathcal{I}) &= \nabla_{(m,u)} Q_{v_-} + \widetilde{\alpha}(\mathcal{I}) \nabla_{(m,u)} h_{v_-} + h_{v_-} \nabla_{(m,u)} \widetilde{\alpha}(\mathcal{I}).
\end{aligned}$$

Let

$$D := h_{v_+} + h_{v_-}, \qquad w_+ := \frac{h_{v_-}}{D}, \qquad w_- := \frac{h_{v_+}}{D}.$$

Then $w_+, w_- \ge 0$, $w_+ + w_- = 1$, and $w_+ h_{v_+} = w_- h_{v_-}$. Multiply the first gradient identity by w_+ and the second by w_-, and then add the resulting identities. The $\nabla_{(m,u)} \widetilde{\alpha}(\mathcal{I})$ terms cancel, yielding

$$\begin{aligned}
\nabla_{(m,u)} \bar{\theta}^\star(\mathcal{I}) = w_+ &\left(\nabla_{(m,u)} P_{v_+} - \widetilde{\alpha}(\mathcal{I}) \nabla_{(m,u)} h_{v_+}\right) \\
&+ w_- \left(\nabla_{(m,u)} Q_{v_-} + \widetilde{\alpha}(\mathcal{I}) \nabla_{(m,u)} h_{v_-}\right).
\end{aligned} \tag{21}$$

We next derive sharper blockwise bounds. The preceding Lipschitz estimates give

$$\|\nabla_m P_v\|_1 \le 1, \qquad \|\nabla_m Q_v\|_1 \le 1, \qquad \|\nabla_m h_v\|_1 \le 1.$$

For the u-block, $\nabla_u L_v$ is coordinatewise nonpositive and $\nabla_u R_v$ is coordinatewise nonnegative, both with ℓ_1 norm at most one. It follows that, for each $S_v \in \{P_v, Q_v\}$,

$$\|(\nabla_u S_v)_+\|_1 \le 1, \qquad \|(\nabla_u S_v)_-\|_1 \le 1.$$

If $\delta_v^- = 0$, differentiability implies that $\nabla_u h_v$ equals the v-th coordinate vector. If $\delta_v^- > 0$, then

$$\nabla_u h_v = \nabla_u u_v + \frac{1}{2}(\nabla_u L_v - \nabla_u R_v),$$

where the second term is coordinatewise nonpositive and has ℓ_1 norm at most one. Thus, in either case,

$$\|(\nabla_u h_v)_+\|_1 \le 1, \qquad \|(\nabla_u h_v)_-\|_1 \le 1.$$

Apply these bounds to (21) in each block. Since $\widetilde{\alpha}(\mathcal{I}) = \bar{\alpha}^\star(\mathcal{I}) - 1 \ge 0$, the triangle inequality and $w_+ + w_- = 1$ give

$$\|\nabla_m \bar{\theta}^\star(\mathcal{I})\|_1 \le 1 + \widetilde{\alpha}(\mathcal{I}) = \bar{\alpha}^\star(\mathcal{I}),$$

and, for each sign $\pm$,

$$\|(\nabla_u \bar{\theta}^\star(\mathcal{I}))_\pm\|_1 \le 1 + \widetilde{\alpha}(\mathcal{I}) = \bar{\alpha}^\star(\mathcal{I}).$$

Adding the three bounds gives $\|\nabla_{(m,u)} \bar{\theta}^\star(\mathcal{I})\|_1 \le 3\bar{\alpha}^\star(\mathcal{I})$. All three bounds hold on a full-measure subset of the interior of $\mathbb{I}$, which completes the proof. □

F.2 Proof of Lipschitz stability

Proof of Theorem 4.3. Since $\bar{\alpha}^\star(\mathcal{I}) < \infty$ for every $\mathcal{I} \in \mathcal{C}$, Theorem 4.1 shows that $\bar{\theta}^\star(\mathcal{I})$ is uniquely defined throughout $\mathcal{C}$.

We first examine the effect of uniformly inflating the radii. For $\mathcal{I} = (m, u) \in \mathbb{I}$ with $\bar{\alpha}^\star(\mathcal{I}) < \infty$ and $s \ge 0$, write

$$\mathcal{I}^{(s)} := (m, u + s\mathbf{1})$$

and use a superscript $^{(s)}$ for quantities associated with this input. Their definitions imply

$$L_v^{(s)} = L_v - s, \qquad R_v^{(s)} = R_v + s, \qquad \mu_v^{(s)} = \mu_v, \qquad \delta_v^{(s)} = \delta_v + s.$$

If

$$a_v(s) := \left(s - \delta_v^-\right)_+,$$

then

$$(\delta_v^+)^{(s)} = \delta_v^+ + a_v(s), \qquad h_v^{(s)} = h_v + a_v(s). \tag{22}$$

Hence, for a pair with $h_v + h_w > 0$, its quotient in (8) after inflation is

$$\frac{|\mu_v - \mu_w| + \delta_v^+ + \delta_w^+ + a_v(s) + a_w(s)}{h_v + h_w + a_v(s) + a_w(s)},$$

which lies between its pre-inflation value and 1. If $h_v + h_w = 0$, finiteness of $\bar{\alpha}^\star(\mathcal{I})$ gives $\mu_v = \mu_w$ and $\delta_v^\pm = \delta_w^\pm = 0$. For $s > 0$, the inflated quotient is then 1. The closed form therefore gives

$$\bar{\alpha}^\star(\mathcal{I}^{(s)}) \le \bar{\alpha}^\star(\mathcal{I}). \tag{23}$$

Fix $\mathcal{I}, \mathcal{I}' \in \mathcal{C}$ and put

$$\mathcal{I}_t := (1-t)\mathcal{I} + t\mathcal{I}', \qquad 0 \le t \le 1.$$

Convexity of $\mathcal{C}$ and the hypothesis give $\bar{\alpha}^\star(\mathcal{I}_t) \le M$. For $s > 0$, the inflated segment

$$K_s := \{\mathcal{I}_t^{(s)} : 0 \le t \le 1\}$$

is a compact subset of the interior of $\mathbb{I}$, and (23) gives $\bar{\alpha}^\star(x) \le M$ for every $x \in K_s$.

Fix $\varepsilon > 0$. The gradient-control lemma implies, in particular, that the map $\mathcal{I} \mapsto \bar{\alpha}^\star(\mathcal{I})$ is continuous on the interior of $\mathbb{I}$. Compactness of K_s therefore yields $r \in (0, s)$ such that the tubular neighborhood

$$\mathcal{D}_{s,r} := K_s + \{z \in \mathbb{R}^{\mathcal{V}} \times \mathbb{R}^{\mathcal{V}} : \|z\|_\infty < r\}$$

is an open convex subset of the interior of $\mathbb{I}$ and satisfies $\bar{\alpha}^\star(x) \le M + \varepsilon$ for every $x \in \mathcal{D}_{s,r}$. By Lemma F.1, the map $\mathcal{I} \mapsto \bar{\theta}^\star(\mathcal{I})$ is locally Lipschitz on $\mathcal{D}_{s,r}$. Moreover, the three gradient norms in (20) are at most $M + \varepsilon$ almost everywhere there. Put

$$\Delta m := m' - m, \qquad \Delta u := u' - u, \qquad \Delta u_+ := (\Delta u)_+, \qquad \Delta u_- := (\Delta u)_-.$$

For $g := \nabla_u \bar{\theta}^\star$, the identity

$$g^\top \Delta u = \left(g_+^\top \Delta u_+ + g_-^\top \Delta u_-\right) - \left(g_+^\top \Delta u_- + g_-^\top \Delta u_+\right)$$

together with Hölder's inequality gives

$$|g^\top \Delta u| \le (M+\varepsilon)\left(\|\Delta u_+\|_\infty + \|\Delta u_-\|_\infty\right)$$

almost everywhere in $\mathcal{D}_{s,r}$. The standard $W^{1,\infty}$ estimate on this convex set now gives

$$|\bar{\theta}^\star(\mathcal{I}^{(s)}) - \bar{\theta}^\star((\mathcal{I}')^{(s)})| \le (M+\varepsilon)\left(\|\Delta m\|_\infty + \|\Delta u_+\|_\infty + \|\Delta u_-\|_\infty\right) = (M+\varepsilon)\|\mathcal{I}' - \mathcal{I}\|_{\mathrm{blk}}.$$

Letting $\varepsilon \downarrow 0$ yields (10) for the inflated endpoints.

It remains to remove the inflation. For $\widetilde{\mathcal{I}} \in \{\mathcal{I}, \mathcal{I}'\}$, if $h_v(\widetilde{\mathcal{I}}) > 0$ for every v, the closed forms show directly that $\bar{\theta}^\star(\widetilde{\mathcal{I}}^{(s)}) \to \bar{\theta}^\star(\widetilde{\mathcal{I}})$ as $s \downarrow 0$. Otherwise, choose $k \in \mathcal{V}$ with $h_k(\widetilde{\mathcal{I}}) = 0$. The finite epigraph characterization and (23) give

$$|\bar{\theta}^\star(\widetilde{\mathcal{I}}^{(s)}) - \mu_k(\widetilde{\mathcal{I}}^{(s)})| \le h_k(\widetilde{\mathcal{I}}^{(s)})\bar{\alpha}^\star(\widetilde{\mathcal{I}}^{(s)}) \le M h_k(\widetilde{\mathcal{I}}^{(s)}).$$

At $\widetilde{\mathcal{I}}$, the same characterization gives $\bar{\theta}^\star(\widetilde{\mathcal{I}}) = \mu_k(\widetilde{\mathcal{I}})$. The inflation identities give $\mu_k(\widetilde{\mathcal{I}}^{(s)}) = \mu_k(\widetilde{\mathcal{I}})$ and $h_k(\widetilde{\mathcal{I}}^{(s)}) \to 0$. Therefore, $\bar{\theta}^\star(\widetilde{\mathcal{I}}^{(s)}) \to \bar{\theta}^\star(\widetilde{\mathcal{I}})$. Letting $s \downarrow 0$ proves (10). □

G Proofs and examples for the statistical analysis

Throughout this section, let

$$r := \max\{C/c, \rho\}.$$

Since $C/c \ge 1$, $\rho = O(1)$, and $L = O(1)$, we have $r = O(C/c)$, and hence the assumed rate $\delta(C/c)^{1+L} \to 0$ implies $\delta r^{1+L} \to 0$.

Lemma G.1 (Containment in the least-element interval)**.** *Let $(\mathcal{V}, \preceq)$ be a finite meet-semilattice with least element $\perp$. If the minimax inefficiency of an input $\mathcal{I} = (m, u)$ is finite, then*

$$\bar{\theta}^\star(\mathcal{I}) \in I_\perp = [m_\perp - u_\perp, m_\perp + u_\perp].$$

Proof. For $y \ge 1$, define the feasible interval

$$C_v(y) := [\mu_v + \delta_v^+ - h_v y,\ \mu_v - \delta_v^+ + h_v y].$$

The identities in Theorem 4.1 give

$$C_v(1) = [R_v - u_v, L_v + u_v].$$

This interval is nonempty because $R_v - L_v \le r_v - \ell_v = 2u_v$. Because $\perp \preceq v$, the definitions of L_v and R_v imply $L_v \ge m_\perp - u_\perp$ and $R_v \le m_\perp + u_\perp$. Consequently, every $C_v(1)$, and hence every $C_v(y)$ with $y \ge 1$, intersects $I_\perp$.

Set $y^\star = \bar{\alpha}^\star(\mathcal{I})$. At $y = y^\star$, the intervals $\{C_v(y^\star) : v \in \mathcal{V}\}$ intersect pairwise, and each also intersects $I_\perp$. Hence the finite family $\{I_\perp\} \cup \{C_v(y^\star) : v \in \mathcal{V}\}$ is pairwise intersecting. Every finite pairwise-intersecting family of intervals on the real line has a common point. Thus

$$I_\perp \cap \bigcap_{v \in \mathcal{V}} C_v(y^\star) \ne \varnothing.$$

By Theorem 4.1, $\bigcap_{v\in\mathcal{V}} C_v(y^\star) = \{\bar{\theta}^\star(\mathcal{I})\}$, which proves the claim. □

G.1 Proof of the quantile error bound

Proof of Proposition 5.1. For this proof, set

$$\theta_t^\circ := \bar{\theta}^\star(\widehat{m}, u(t)), \qquad R(t) := \min_{v\in\mathcal{V}}\{b_v + u_v(t)\},$$

and define the scale-approximation event

$$\mathcal{G}_\epsilon := \left\{ \max_{v\in\mathcal{V}} |\ln \widehat{\sigma}_v - \ln \sigma_v| \le \epsilon \right\}.$$

The nonnegativity of $q(t)$ and $z(t)$, together with condition i), implies

$$\begin{aligned} u_{v\wedge w}(t) &= \sigma_{v\wedge w} q(t) + \gamma_{v\wedge w} z(t) \\ &\le \kappa\{u_v(t) + u_w(t)\}, \qquad v, w \in \mathcal{V}. \end{aligned}$$

On $\mathcal{G}_\epsilon$, $e^{-\epsilon}\sigma_v \le \widehat{\sigma}_v \le e^{\epsilon}\sigma_v$ for every v, and hence

$$\widehat{\sigma}_{v\wedge w} \le e^{\epsilon}\sigma_{v\wedge w} \le \kappa e^{\epsilon}(\sigma_v + \sigma_w) \le \kappa e^{2\epsilon}(\widehat{\sigma}_v + \widehat{\sigma}_w).$$

Together with the κ-subadditivity of γ, this shows that, on this event, $u(t)$ and $\widehat{u}(t)$ both belong to the convex cone $\mathbb{S}^{\mathcal{V}}_{\kappa e^{2\epsilon}}$, and the entire line segment joining them is contained in the same cone. Moreover, $\kappa e^{2\epsilon} \ge 1/2$, so Theorem 4.2 applies to this cone. The same coordinatewise bounds also imply

$$\|(\widehat{\sigma} - \sigma)_+\|_\infty \le (e^{\epsilon} - 1)\sigma_{\max}, \qquad \|(\widehat{\sigma} - \sigma)_-\|_\infty \le (1 - e^{-\epsilon})\sigma_{\max}.$$

On the simultaneous coverage event, the instance $\mathfrak{s} = (\theta^\star, \eta, b)$ is compatible with the input $(\widehat{m}, u(t))$, because

$$\eta_v \in [\widehat{m}_v - u_v(t), \widehat{m}_v + u_v(t)], \qquad |\eta_v - \theta^\star| \le b_v, \qquad v \in \mathcal{V},$$

and b is isotonic. Hence Theorem 4.2 and the definition of the minimax inefficiency give

$$|\theta^\circ_t - \theta^\star| \le 2\kappa R(t).$$

On $\mathcal{G}_\epsilon$, Theorem 4.3 and Remark 4.2 also give

$$\begin{aligned} |\widehat{\theta}_t - \theta^\circ_t| &\le 2\kappa e^{2\epsilon} q(t) \left\{\|(\widehat{\sigma} - \sigma)_+\|_\infty + \|(\widehat{\sigma} - \sigma)_-\|_\infty\right\} \\ &\le 2\kappa\sigma_{\max} q(t)(e^{3\epsilon} - e^{\epsilon}) \\ &\le 2\kappa\rho(e^{3\epsilon} - e^{\epsilon})R(t). \end{aligned}$$

For $a_\epsilon := e^{3\epsilon} - e^{\epsilon}$,

$$\begin{aligned} e^{-2\epsilon}(1 + a_\epsilon) &= e^{-2\epsilon} + 2\sinh\epsilon \\ &\le \cosh(2\epsilon) \le e^{2\epsilon^2}. \end{aligned}$$

Here, the first inequality follows from

$$\cosh(2\epsilon) - e^{-2\epsilon} - 2\sinh\epsilon = 2\sinh\epsilon(\cosh\epsilon - 1) \ge 0,$$

and the second is the standard bound $\cosh x \le e^{x^2/2}$. Therefore, by the generalized Bernoulli inequality and $\rho \ge 1$,

$$1 + \rho a_\epsilon \le (1 + a_\epsilon)^\rho \le e^{2\rho\epsilon(1+\epsilon)}.$$

The triangle inequality and the union bound now yield

$$\Pr\left\{|\widehat{\theta}_t - \theta^\star| \le 2\kappa e^{2\rho\epsilon(1+\epsilon)} R(t)\right\} \ge 1 - t - \delta. \tag{24}$$

Finally, $\kappa = O(1)$, $\rho = O(1)$, and $\epsilon \to 0$ give

$$2\kappa e^{2\rho\epsilon(1+\epsilon)} = 2\kappa + o(1),$$

which proves the claim. □

G.2 Proof of the mean absolute error bound

Proof of Proposition 5.2. For $y, w \ge 0$, define

$$g(y, w) := \min_{v\in\mathcal{V}}(b_v + \sigma_v y + \gamma_v w),$$

$$J_q(t) := \int_0^t \{q(x) - q(t)\}\,dx, \qquad J_z(t) := \int_0^t \{z(x) - z(t)\}\,dx,$$

and write $\theta^\circ_s := \bar{\theta}^\star(\widehat{m}, u(s))$ for $s \in (0, 1)$. Fix $x \in (0, t]$. Since $q(x) \ge q(t)$ and $z(x) \ge z(t)$, we have $u(x) - u(t) \ge 0$ coordinatewise. Applying Theorem 4.3 and Remark 4.2 with $\Delta m = 0$ and $\Delta u = u(x) - u(t) \ge 0$ therefore gives

$$\begin{aligned} |\theta^\circ_t - \theta^\circ_x| &\le 2\kappa\|u(x) - u(t)\|_\infty \\ &\le 2\kappa\left[\sigma_{\max}\{q(x) - q(t)\} + \gamma_{\max}\{z(x) - z(t)\}\right]. \end{aligned}$$

The simultaneous coverage condition and the deterministic oracle bound at level x therefore show that

$$\Pr\left\{\frac{|\theta_t^\circ - \theta^\star|}{2\kappa} \le g(q(x), z(x)) + \sigma_{\max}\{q(x) - q(t)\} + \gamma_{\max}\{z(x) - z(t)\}\right\} \ge 1 - x,$$

yielding an upper bound on the $(1-x)$-quantile of $|\theta_t^\circ - \theta^\star|/(2\kappa)$. Integrating this bound with respect to $x \in (0, t]$ and combining with the oracle bound at level t gives

$$\begin{aligned}\frac{\mathbb{E}|\theta_t^\circ - \theta^\star|}{2\kappa} &\le (1-t)g(q(t), z(t)) + \int_0^t g(q(x), z(x))\, dx \\ &\quad + \sigma_{\max} J_q(t) + \gamma_{\max} J_z(t) \\ &\le g(q(t) + J_q(t), z(t) + J_z(t)) + \sigma_{\max} J_q(t) + \gamma_{\max} J_z(t) \\ &\le \min_{v \in \mathcal{V}}\{b_v + \bar{u}_v(t)\}. \end{aligned} \tag{25}$$

Here, the penultimate inequality uses the joint concavity of g and Jensen's inequality. The final inequality follows from $J_q(t), J_z(t) \ge 0$, $\rho\sigma_v \ge \sigma_{\max}$, and $\rho\gamma_v \ge \gamma_{\max}$ for every v.

Set

$$B(t) := \min_{v \in \mathcal{V}}\{b_v + \bar{u}_v(t)\},$$

and let $\mathcal{G}_\epsilon$ be the scale-approximation event defined in the proof of Proposition 5.1. On $\mathcal{G}_\epsilon$, the same stability calculation as in that proof, together with $B(t) \ge \sigma_{\min} q(t)$, gives

$$|\widehat{\theta}_t - \theta_t^\circ| \le 2\kappa\rho(e^{3\epsilon} - e^{\epsilon})B(t).$$

Condition i) ensures that the input with radii $u(t)$ has finite minimax inefficiency. Condition iii) implies that $\widehat{\sigma}$ is $\max\{1, C/(2c)\}$-subadditive almost surely. Combined with the κ-subadditivity of γ and the nonnegativity of $q(t)$ and $z(t)$, this property allows Theorem 4.2 to be applied, ensuring that the input with radii $\widehat{u}(t)$ also has finite minimax inefficiency almost surely. On $\mathcal{G}_\epsilon^c$, apply Lemma G.1 to the inputs with radii $u(t)$ and $\widehat{u}(t)$. Both least-element intervals are centered at $\widehat{m}_\perp$, so condition iii) gives

$$\begin{aligned}|\widehat{\theta}_t - \theta_t^\circ| &\le \widehat{u}_\perp(t) + u_\perp(t) \\ &\le 2Cq(t) + 2\gamma_\perp z(t) \\ &\le 2r\{cq(t) + \gamma_{\min} z(t)\} \\ &\le 2rB(t). \end{aligned}$$

Here, the final inequality uses $\bar{u}_v(t) \ge u_v(t) \ge cq(t) + \gamma_{\min} z(t)$ for every v. Consequently, the triangle inequality, (25), and $\Pr(\mathcal{G}_\epsilon^c) \le \delta$ yield

$$\begin{aligned}\mathbb{E}|\widehat{\theta}_t - \theta^\star| &\le \mathbb{E}|\theta_t^\circ - \theta^\star| + \mathbb{E}\left[|\widehat{\theta}_t - \theta_t^\circ|\mathbf{1}_{\mathcal{G}_\epsilon}\right] + \mathbb{E}\left[|\widehat{\theta}_t - \theta_t^\circ|\mathbf{1}_{\mathcal{G}_\epsilon^c}\right] \\ &\le \left\{2\kappa\left[1 + \rho(e^{3\epsilon} - e^{\epsilon})\right] + 2r\delta\right\} B(t). \end{aligned}$$

Applying the exponential bound used in the preceding proof gives the explicit nonasymptotic inequality

$$\mathbb{E}|\widehat{\theta}_t - \theta^\star| \le \left\{2\kappa e^{2\rho\epsilon(1+\epsilon)} + 2r\delta\right\} B(t). \tag{26}$$

Since $r \ge 1$ and $L > 0$, the assumption $\delta r^{1+L} \to 0$ implies $r\delta \to 0$. Together with $\kappa = O(1)$, $\rho = O(1)$, and $\epsilon \to 0$, this gives

$$2\kappa e^{2\rho\epsilon(1+\epsilon)} + 2r\delta = 2\kappa + o(1),$$

which proves (12).

Finally, since u_v is differentiable and integrable, $\bar{u}_v$ is differentiable and, for every $t \in (0, 1)$,

$$\bar{u}_v'(t) = \{1 - (1+\rho)t\}u_v'(t).$$

Thus $\bar{u}_v$ is nonincreasing on $(0, t^\star]$ and nondecreasing on $[t^\star, 1)$. Hence each $\bar{u}_v$ is minimized at $t = t^\star$, as is the right-hand side of (12). Moreover, since $1 + \rho = 1/t^\star$,

$$\begin{aligned}\bar{u}_v(t^\star) &= u_v(t^\star) + \frac{1}{t^\star}\int_0^{t^\star} \{u_v(x) - u_v(t^\star)\}\, dx \\ &= \frac{1}{t^\star}\int_0^{t^\star} u_v(x)\, dx.\end{aligned}$$

It remains to prove the plug-in bound. Set

$$B^\star := \min_{v\in\mathcal{V}}\{b_v + \bar{u}_v(t^\star)\}.$$

On $\mathcal{G}_\epsilon$, the coordinatewise scale bounds give

$$e^{-2\epsilon}\frac{\sigma_{\max}}{\sigma_{\min}} \le \frac{\widehat{\sigma}_{\max}}{\widehat{\sigma}_{\min}} \le e^{2\epsilon}\frac{\sigma_{\max}}{\sigma_{\min}}.$$

Taking the maximum with the fixed ratio $\gamma_{\max}/\gamma_{\min}$ gives

$$e^{-2\epsilon}\rho \le \widehat{\rho} \le e^{2\epsilon}\rho.$$

Consequently,

$$e^{-2\epsilon} \le \frac{\hat{t}}{t^\star} = \frac{1+\rho}{1+\widehat{\rho}} \le e^{2\epsilon}, \qquad |\ln\hat{t} - \ln t^\star| \le 2\epsilon. \tag{27}$$

Both $t^\star$ and $\hat{t}$ belong to $(0, 1/2]$. Condition iv) implies, for $f \in \{q, z\}$ and $s, t \in (0, 1/2]$,

$$e^{-L|\ln s - \ln t|} f(t) \le f(s) \le e^{L|\ln s - \ln t|} f(t).$$

Indeed, differentiation shows that $t^L f(t)$ is nondecreasing and $t^{-L} f(t)$ is nonincreasing. This argument also covers a profile that is identically zero. Hence

$$\begin{aligned}&|q(\hat{t}) - q(t^\star)| \le a_L(\epsilon) q(t^\star), \\ &|z(\hat{t}) - z(t^\star)| \le a_L(\epsilon) z(t^\star), \qquad\qquad a_L(\epsilon) := e^{2L\epsilon} - 1.\end{aligned}$$

As in the proof of Proposition 5.1, the estimated scale vector is $\kappa e^{2\epsilon}$-subadditive on $\mathcal{G}_\epsilon$. Together with the κ-subadditivity of γ, this implies that, on the same event, the radius vectors along the line segment joining $\widehat{u}(\hat{t})$ and $\widehat{u}(t^\star)$ belong to $\mathbb{S}^{\mathcal{V}}_{\kappa e^{2\epsilon}}$. Therefore, the results in Theorem 4.3 and Remark 4.2 give

$$\begin{aligned}|\widehat{\theta}_{\hat{t}} - \widehat{\theta}_{t^\star}| &\le 2\kappa e^{2\epsilon}\left\{\widehat{\sigma}_{\max}|q(\hat{t}) - q(t^\star)| + \gamma_{\max}|z(\hat{t}) - z(t^\star)|\right\} &(28)\\ &\le 2\kappa e^{2\epsilon}\left\{e^{\epsilon}\rho a_L(\epsilon) + \rho a_L(\epsilon)\right\} B^\star. &(29)\end{aligned}$$

Here, we used

$$B^\star \ge \min_{v\in\mathcal{V}} \bar{u}_v(t^\star) \ge \sigma_{\min} q(t^\star) + \gamma_{\min} z(t^\star).$$

Moreover, $\kappa = O(1)$, $\rho = O(1)$, $L = O(1)$, and $\epsilon \to 0$ imply $a_L(\epsilon) = o(1)$. Thus the coefficient on the right-hand side of (29) is $o(1)$.

It remains to control $\mathcal{G}_\epsilon^c$. Set

$$t_0 := \frac{1}{1+r}.$$

Condition iii) and the definitions of ρ, $\widehat{\rho}$, and r give

$$t_0 \le t^\star \le \frac{1}{2}, \qquad t_0 \le \hat{t} \le \frac{1}{2}.$$

Apply Lemma G.1 to the two inputs with radii $\widehat{u}(t^\star)$ and $\widehat{u}(\hat{t})$. Since their least-element intervals are centered at the same point, monotonicity of q and z gives

$$|\widehat{\theta}_{\hat{t}} - \widehat{\theta}_{t^\star}| \le 2\{Cq(t_0) + \gamma_\perp z(t_0)\}.$$

Table G.1: **Common confidence profiles and their tail-calibrated radii**. Here $K := |\mathcal{V}|$ and $B := 2K(1+\rho)$. The notation Φ^{-1} denotes the standard normal quantile function, and n denotes the underlying sample size. The scale parameters $\sigma_v > 0$ in the Wald and Bernstein profiles are usually estimated from data, whereas $\gamma_v > 0$ is a known coefficient in the Hoeffding and Bernstein profiles. The order bounds in the third column have universal implicit constants, and the last column gives a uniform choice of L in condition iv), rounded up to two decimal places; see Section G.3.

Profile	$u_v(t)$	Upper bound on $\bar{u}_v(t^\star)$	L
Hoeffding	$\frac{\gamma_v}{\sqrt{n}}\sqrt{2\log\frac{2K}{t}}$	$\lesssim \frac{\gamma_v}{\sqrt{n}}\sqrt{\log B}$	0.37
Wald	$\frac{\sigma_v}{\sqrt{n}}\Phi^{-1}\left(1-\frac{t}{2K}\right)$	$\lesssim \frac{\sigma_v}{\sqrt{n}}\sqrt{\log B}$	1.17
Bernstein	$\frac{\sigma_v}{\sqrt{n}}\sqrt{2\log\frac{2K}{t}} + \frac{2\gamma_v}{3n}\log\frac{2K}{t}$	$\lesssim \frac{\sigma_v}{\sqrt{n}}\sqrt{\log B} + \frac{2\gamma_v}{3n}\log B$	0.73

The profile bound above and $t^\star/t_0 = (1+r)/(1+\rho) \le 1+r$ yield

$$q(t_0) \le (1+r)^L q(t^\star), \qquad z(t_0) \le (1+r)^L z(t^\star).$$

Since $Cq(t^\star) \le rB^\star$ and $\gamma_\perp z(t^\star) \le \rho B^\star$, we obtain

$$\mathbb{E}\left[|\widehat{\theta}_{\hat{t}} - \widehat{\theta}_{t^\star}|\mathbf{1}_{\mathcal{G}_\epsilon^c}\right] \le 2\delta\left\{r(1+r)^L + \rho(1+r)^L\right\}B^\star \tag{30}$$

$$\le 4\,2^L\delta r^{1+L}B^\star \tag{31}$$

$$= o(B^\star). \tag{32}$$

Here, we used $r \ge \rho \ge 1$ and $1+r \le 2r$. The final equality follows from $L = O(1)$ and $\delta r^{1+L} \to 0$.

Combining (29) and (32) gives

$$\mathbb{E}|\widehat{\theta}_{\hat{t}} - \widehat{\theta}_{t^\star}| = o(B^\star).$$

Finally, the triangle inequality and (12) at $t = t^\star$ give

$$\begin{aligned}\mathbb{E}|\widehat{\theta}_{\hat{t}} - \theta^\star| &\le \mathbb{E}|\widehat{\theta}_{t^\star} - \theta^\star| + \mathbb{E}|\widehat{\theta}_{\hat{t}} - \widehat{\theta}_{t^\star}| \\ &\le \{2\kappa + o(1)\}B^\star + o(B^\star) \\ &= \{2\kappa + o(1)\}B^\star.\end{aligned}$$

This proves the plug-in statement in Proposition 5.2. □

G.3 Confidence-profile evaluations

We first verify $u_v(t)$ and $\bar{u}_v(t^\star)$ for the three common confidence profiles in Table G.1. The calculation of L is deferred to the end of this section.

Hoeffding profile. Define the Hoeffding profile

$$z_{\mathrm{H},n}(t) := \frac{1}{\sqrt{n}}\sqrt{2\log\frac{2K}{t}}.$$

This profile fits the general radius form by taking $q = 0$, $z = z_{\mathrm{H},n}$, and $\sigma = \widehat{\sigma} = \mathbf{1}$. Thus $u_v(t) = \gamma_v z_{\mathrm{H},n}(t)$ and $\rho = \gamma_{\max}/\gamma_{\min}$. Here and below, $K = |\mathcal{V}|$ and $B = 2K/t^\star$, as in Table G.1. Since $K \ge 1$ and $t^\star \le 1/2$, we have $B \ge 4$. The change of variables $x = t^\star y$ gives

$$\frac{1}{t^\star}\int_0^{t^\star} z_{\mathrm{H},n}(x)\,dx = \frac{1}{\sqrt{n}}\int_0^1 \sqrt{2\log\frac{B}{y}}\,dy.$$

Monotonicity of the integrand and Jensen's inequality give

$$\frac{\sqrt{2\log B}}{\sqrt{n}} \le \frac{1}{t^\star}\int_0^{t^\star} z_{\mathrm{H},n}(x)\,dx \le \frac{\sqrt{2(\log B + 1)}}{\sqrt{n}}.$$

Therefore,

$$\bar{u}_v(t^\star) = \frac{\gamma_v}{t^\star}\int_0^{t^\star} z_{\mathrm{H},n}(x)\,dx \asymp \frac{\gamma_v}{\sqrt{n}}\sqrt{\log B}, \qquad v \in \mathcal{V},$$

with universal implicit constants. This proves the Hoeffding row of Table G.1.

Wald profile. For the Wald setting, define

$$q_{\mathrm{W},n}(t) := \frac{1}{\sqrt{n}}\Phi^{-1}\left(1 - \frac{t}{2K}\right).$$

Let $a_B := \Phi^{-1}(1 - 1/B)$, and let $\phi = \Phi'$ denote the standard normal density. The substitution $a = \Phi^{-1}(1 - x/(2K))$ gives

$$x = 2K\{1 - \Phi(a)\}, \qquad dx = -2K\phi(a)\,da.$$

As x increases from 0 to $t^\star$, a decreases from ∞ to a_B. Using $\phi'(a) = -a\phi(a)$, we obtain

$$\begin{aligned}\frac{1}{t^\star}\int_0^{t^\star} q_{\mathrm{W},n}(x)\,dx &= \frac{2K}{t^\star\sqrt{n}}\int_{a_B}^{\infty} a\phi(a)\,da \\ &= \frac{2K}{t^\star\sqrt{n}}\phi(a_B) \\ &= \frac{1}{\sqrt{n}}\frac{\phi(a_B)}{1 - \Phi(a_B)},\end{aligned}$$

where the last equality uses $1 - \Phi(a_B) = 1/B = t^\star/(2K)$. Since $B \geq 4$, we have $a_B \geq a_0 := \Phi^{-1}(3/4) > 0$. The standard Mills-ratio bounds give

$$a_B \leq \frac{\phi(a_B)}{1 - \Phi(a_B)} \leq a_B + \frac{1}{a_B} \leq (1 + a_0^{-2})a_B.$$

To compare a_B with $\sqrt{\log B}$, note that for $a \geq 0$, monotonicity of ϕ and the Gaussian Chernoff bound give

$$\phi(a+1) \leq \int_a^{a+1}\phi(s)\,ds \leq 1 - \Phi(a) \leq e^{-a^2/2}.$$

Applying these inequalities at $a = a_B$ and taking negative logarithms yields

$$\frac{a_B^2}{2} \leq \log B \leq \frac{(a_B+1)^2}{2} + \frac{1}{2}\log(2\pi) \leq Ca_B^2$$

because $a_B \geq a_0 > 0$, where C is a universal constant. Combining these comparisons with the integral identity above yields

$$\begin{aligned}\bar{u}_v(t^\star) &= \frac{\sigma_v}{t^\star}\int_0^{t^\star} q_{\mathrm{W},n}(x)\,dx \\ &= \frac{\sigma_v}{\sqrt{n}}\frac{\phi(a_B)}{1 - \Phi(a_B)} \\ &\asymp \frac{\sigma_v}{\sqrt{n}}a_B \asymp \frac{\sigma_v}{\sqrt{n}}\sqrt{\log B}, \qquad v \in \mathcal{V}.\end{aligned}$$

Both comparisons hold with universal implicit constants for all $B \geq 4$. This proves the Wald row of Table G.1.

Bernstein profile. For the Bernstein radii, take $q_{\mathrm{B},n} = z_{\mathrm{H},n}$ and

$$z_{\mathrm{B},n}(t) := \frac{2}{3n}\log\frac{2K}{t}.$$

Direct integration gives

$$\frac{1}{t^\star}\int_0^{t^\star} z_{\mathrm{B},n}(x)\,dx = \frac{2}{3n}\left(\log B + 1\right).$$

Since $B \geq 4$, $\log B + 1 \lesssim \log B$. Combining this observation with the Hoeffding integral bound gives

$$\bar{u}_v(t^\star) \lesssim \frac{\sigma_v}{\sqrt{n}}\sqrt{\log B} + \frac{2\gamma_v}{3n}\log B,$$

which proves the Bernstein row of Table G.1. In all three cases, Proposition 5.2 incorporates the displayed upper bound on $\bar{u}_v(t^\star)$ into the MAE bound at $t = t^\star$.

Constants L in condition iv). For the Hoeffding and Bernstein profiles, direct differentiation gives

$$\left|\frac{d\ln z_{\mathrm{H},n}(t)}{d\ln t}\right| = \left|\frac{d\ln q_{\mathrm{B},n}(t)}{d\ln t}\right| = \frac{1}{2\ln(2K/t)} \le \frac{1}{2\ln 4}, \qquad \left|\frac{d\ln z_{\mathrm{B},n}(t)}{d\ln t}\right| = \frac{1}{\ln(2K/t)} \le \frac{1}{\ln 4}$$

for $t \in (0, 1/2]$. Thus we have $L = 1/(2\ln 4)$ for Hoeffding and $L = 1/(\ln 4)$ for Bernstein.

For the Wald profile, let $a_0 := \Phi^{-1}(3/4)$ and $a(t) := \Phi^{-1}(1 - t/(2K))$. Then $a(t) \ge a_0$. Define

$$h(a) := \frac{1-\Phi(a)}{a\phi(a)}.$$

The standard Mills-ratio bound $1 - \Phi(a) < \phi(a)/a$ gives

$$\frac{d}{da}\ln h(a) = a - \frac{1}{a} - \frac{\phi(a)}{1-\Phi(a)} < -\frac{1}{a} < 0, \qquad a > 0.$$

Consequently,

$$\left|\frac{d\ln q_{\mathrm{W},n}(t)}{d\ln t}\right| = h(a(t)) \le h(a_0) = \frac{1}{4a_0\phi(a_0)}.$$

Thus all three choices of L displayed in Table G.1 are independent of n and K.

G.4 Example: clipped empirical standard deviations for the Bernstein bound

Let $X_1, \dots, X_n$ be i.i.d. random vectors indexed by $\mathcal{V}$, with

$$\mathbb{E}X_{i,v} = \eta_v, \qquad \mathrm{Var}(X_{i,v}) = \sigma_v^2 > 0, \qquad |X_{i,v} - \eta_v| \le B_v$$

almost surely for every $v \in \mathcal{V}$, where $B \in \mathbb{R}^{\mathcal{V}}_{>0}$ is known. Let $\gamma \in \mathbb{R}^{\mathcal{V}}_{>0}$ be a known coefficient vector satisfying

$$\gamma_v \ge B_v, \qquad v \in \mathcal{V}.$$

Because $\mathcal{V}$ is finite and σ and γ are fixed positive vectors, both are κ-subadditive for some fixed $\kappa < \infty$. Thus condition i) holds with $\kappa = O(1)$, and ρ is fixed and finite. For each $v \in \mathcal{V}$, let $\widehat{m}_v$ and $S_{n,v}$ be the empirical mean and standard deviation, respectively, and write $S_n = (S_{n,v})_{v\in\mathcal{V}}$. Define

$$\widehat{\sigma}_v := S_{n,v} \vee \frac{B_v}{\sqrt{n}}, \quad v \in \mathcal{V},$$

This enforces the positive lower bound $\widehat{\sigma}_v \ge B_v/\sqrt{n} > 0$, and hence $\widehat{\sigma}_v^2 \ge B_v^2/n$. We verify that this estimator satisfies the assumptions of the plug-in statement in Proposition 5.2. Boundedness of the observations gives a constant $C_0 < \infty$ such that $\widehat{\sigma}_{\max} \le C_0$ almost surely. After increasing C_0 if necessary, we may also take it to be an upper bound for the clipping levels and all coordinates of the fixed vector σ. Since $\sigma_{\min} > 0$, condition iii) holds for all sufficiently large n with

$$c_n = \frac{B_{\min}}{\sqrt{n}}, \qquad C_n = C_0.$$

For condition ii), note that $\sigma_v \ge B_v/\sqrt{n}$ for every v once n is sufficiently large. Componentwise clipping therefore does not increase the distance from σ, and hence

$$\|\widehat{\sigma} - \sigma\|_2 \le \|S_n - \sigma\|_2.$$

Take $\epsilon_n = n^{-1/4}$. On the event

$$\|S_n - \sigma\|_2 \le \frac{\sigma_{\min}}{2}\epsilon_n,$$

the logarithm is $2/\sigma_{\min}$-Lipschitz on every interval with endpoints σ_v and $\widehat{\sigma}_v$, and therefore

$$\max_{v\in\mathcal{V}} |\ln\widehat{\sigma}_v - \ln\sigma_v| \le \epsilon_n.$$

Concentration of the bounded empirical first and second moments, together with the positive marginal variances, gives constants $c_2, C_2 > 0$ such that the complementary event has probability at most

$$\delta_n := C_2 e^{-c_2\sqrt{n}}.$$

Thus condition ii) holds, and

$$\delta_n(C_n/c_n)^{1+L} = C_2 e^{-c_2\sqrt{n}} \left(\frac{C_0\sqrt{n}}{B_{\min}}\right)^{1+L} \longrightarrow 0.$$

Here, L is the fixed constant in condition iv) verified above for the Bernstein profiles, so condition iv) and the required asymptotic rate both hold.

Finally, with $x = \log(2K/t)$, Bernstein's inequality gives

$$\Pr\left\{|\widehat{m}_v - \eta_v| > \sigma_v\sqrt{\frac{2x}{n}} + \frac{2B_v x}{3n}\right\} \le 2e^{-x}.$$

The assumption $\gamma_v \ge B_v$ and a union bound therefore give the simultaneous coverage condition for

$$u_v(t) = \frac{\sigma_v}{\sqrt{n}}\sqrt{2\log\frac{2K}{t}} + \frac{2\gamma_v}{3n}\log\frac{2K}{t}.$$

Consequently, the plug-in statement in Proposition 5.2 applies: the output using the clipped empirical standard deviations and the plug-in confidence parameter has the same first-order MAE upper bound as the output at $t^\star$.

H OACIS AND D-OACIS ALGORITHMS: IMPLEMENTATION AND PROOFS

For the DAG $G = (\mathcal{V}, E)$ introduced in Section 6, let $\mathrm{pred}(v) = \{w : (w, v) \in E\}$ denote the set of predecessors (in-neighbors) of $v \in \mathcal{V}$ in G. The only optimization subroutine required by OACIS acts on a finite family $\mathcal{F}$ of affine functions. Define the planar epigraph linear program

$$\mathrm{EpigraphLP}(\mathcal{F}) \in \operatorname*{arg\,min}_{(x,y)\in\mathbb{R}^2} \{y : y \ge f(x) \text{ for every } f \in \mathcal{F}\}.$$

Write $T_{\mathrm{LP}}(k)$ for the time complexity of $\mathrm{EpigraphLP}(\mathcal{F})$ with $|\mathcal{F}| = k$. Possible implementations are summarized in Table H.1. For the strongest worst-case guarantee, the proof below instantiates $\mathrm{EpigraphLP}$ with the prune-and-search algorithm.

Algorithms H.1 and H.2 compute OACIS and d-OACIS, respectively, when their minimax inefficiencies are finite; otherwise they return NONFINITE. Algorithm H.3 is a shared subroutine that evaluates the worst-case inefficiency at a prescribed point. For d-OACIS, ties between minimizing midpoint values are resolved using the fixed rule in Definition 3.3.

Algorithm H.1 OACIS

Require: Intervals $([\ell_v, r_v])_{v\in\mathcal{V}}$ and a DAG $G = (\mathcal{V}, E)$ representing $\preceq$
Ensure: NONFINITE if and only if $\bar{\alpha}^\star(\mathcal{I}) = \infty$; otherwise $(\bar{\theta}^\star(\mathcal{I}), \bar{\alpha}^\star(\mathcal{I}))$

Compute a topological ordering of G
for all $v \in \mathcal{V}$ in topological order **do**
    $L_v \leftarrow \max\big(\{\ell_v\} \cup \{L_w : w \in \mathrm{pred}(v)\}\big)$
    $R_v \leftarrow \min\big(\{r_v\} \cup \{R_w : w \in \mathrm{pred}(v)\}\big)$
    Compute $u_v, \mu_v, \delta_v^+, \delta_v^-$, and h_v as in Section 4
end for
$Z \leftarrow \{v \in \mathcal{V} : h_v = 0\}$
if $Z \neq \emptyset$ **then**
    Choose $z_0 \in Z$ and set $x^\star \leftarrow \mu_{z_0}$
    if $\mu_v \neq x^\star$ for some $v \in Z$ **then**
        return NONFINITE
    end if
    $y^\star \leftarrow \textsc{WorstCaseInefficiency}(\mathcal{I}, G, x^\star)$
    return $(x^\star, y^\star)$
end if
$\mathcal{F} \leftarrow \{x \mapsto 1\}$
for all $v \in \mathcal{V}$ **do**
    Add $x \mapsto (x - \mu_v + \delta_v^+)/h_v$ to $\mathcal{F}$
    Add $x \mapsto (-x + \mu_v + \delta_v^+)/h_v$ to $\mathcal{F}$
end for
return $\mathrm{EpigraphLP}(\mathcal{F})$

Algorithm H.2 Discretized OACIS

Require: Intervals $([\ell_v, r_v])_{v\in\mathcal{V}}$ and a DAG $G = (\mathcal{V}, E)$ representing $\preceq$
Ensure: NONFINITE if and only if $\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}) = \infty$; otherwise $(\bar{\theta}^\star_{\mathrm{d}}(\mathcal{I}), \bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}))$

$q \leftarrow \mathrm{OACIS}(\mathcal{I}, G)$
if $q =$ NONFINITE **then**
    return NONFINITE
end if
$(x^\star, y^\star) \leftarrow q$
$\mathcal{V}_- \leftarrow \arg\max_{v\in\mathcal{V}: m_v \le x^\star} m_v$
$\mathcal{V}_+ \leftarrow \arg\min_{v\in\mathcal{V}: m_v \ge x^\star} m_v$
$D \leftarrow \emptyset$
if $\mathcal{V}_- \neq \emptyset$ **then**
    Add any $v_- \in \mathcal{V}_-$ to D
end if
if $\mathcal{V}_+ \neq \emptyset$ **then**
    Add any $v_+ \in \mathcal{V}_+$ to D
end if
for all $v \in D$ **do**
    $a_v \leftarrow \textsc{WorstCaseInefficiency}(\mathcal{I}, G, m_v)$
end for
Pick $\hat{v} \in \arg\min_{v\in D} a_v$, applying the fixed rule to the midpoint values
return $(m_{\hat{v}}, a_{\hat{v}})$

Algorithm H.3 Worst-case inefficiency

Require: Intervals $([\ell_v, r_v])_{v\in\mathcal{V}}$, a DAG $G = (\mathcal{V}, E)$ representing $\preceq$, and a point $x \in \mathbb{R}$
Ensure: $\bar{\alpha}(x; \mathcal{I}) \in [1, \infty]$

Compute a topological ordering of G
$y \leftarrow 1$
for all $v \in \mathcal{V}$ in topological order **do**
  $L_v \leftarrow \max\big(\{\ell_v\} \cup \{L_w : w \in \mathrm{pred}(v)\}\big)$
  $R_v \leftarrow \min\big(\{r_v\} \cup \{R_w : w \in \mathrm{pred}(v)\}\big)$
  Compute $u_v, \mu_v, \delta_v^+, \delta_v^-$, and h_v as in Section 4
  $s_v \leftarrow |x - \mu_v| + \delta_v^+$
  if $h_v = 0$ and $s_v > 0$ **then**
    return ∞
  else if $h_v > 0$ **then**
    $y \leftarrow \max\{y, s_v/h_v\}$
  end if
end for
return y

Table H.1: Exact implementations of the planar epigraph LP subroutine.

Implementation	$T_{\mathrm{LP}}(k)$
Slope sorting (de Berg et al., 2008)	$O(k \log k)$ in the worst case
Randomized (Seidel, 1991)	expected $O(k)$
Prune-and-search (Megiddo, 1983)	$O(k)$ in the worst case

H.1 Proofs of correctness and computational complexity

We first prove the correctness of the common subroutine and establish its computational complexity.

Correctness of Algorithm H.3. Induction along a topological ordering gives $L_v = \max_{w\preceq v} \ell_w$ and $R_v = \min_{w\preceq v} r_w$. For a fixed x, Proposition 4.1 shows that a finite y is feasible exactly when $|x - \mu_v| + \delta_v^+ = 0$ for every v with $h_v = 0$. If any such constraint fails, the algorithm correctly returns ∞. Otherwise, the least feasible value is

$$\max\left\{1, \max_{v:h_v>0} \frac{|x - \mu_v| + \delta_v^+}{h_v}\right\},$$

which is exactly the value accumulated in y. Thus the returned value is $\bar{\alpha}(x; \mathcal{I})$.

Topological sorting and the recurrences for L_v and R_v scan every vertex and edge once, and the remaining updates take constant time per vertex. Hence the subroutine takes $O(|\mathcal{V}| + |E|)$ time under an adjacency-list representation of G. □

We next show the correctness and computational complexity of the OACIS and d-OACIS algorithms.

Proposition H.1 (Correctness and computational complexity of OACIS)**.** *Algorithm H.1 reports* NONFINITE *if and only if $\bar{\alpha}^\star(\mathcal{I}) = \infty$; otherwise it returns $(\bar{\theta}^\star(\mathcal{I}), \bar{\alpha}^\star(\mathcal{I}))$. Moreover, the algorithm runs in $O(|\mathcal{V}| + |E|)$ time.*

Proof. The same induction along a topological ordering guarantees the correct computation of L_v, R_v, μ_v, and $\delta_v^\pm$ used by Algorithm H.1. If $Z \neq \emptyset$, the constraints indexed by $v \in Z$ therefore force $x = \mu_v$. They are incompatible exactly when these centers differ, in which case Algorithm H.3 and Theorem 4.1 agree that $\bar{\alpha}^\star(\mathcal{I}) = \infty$. Otherwise they fix $x = x^\star$, and Algorithm H.3 returns the least feasible value $y^\star = \bar{\alpha}(x^\star; \mathcal{I})$.

If $Z = \emptyset$, splitting each absolute-value constraint in (6) into its two affine parts gives exactly the epigraph LP defined by $\mathcal{F}$. Hence the algorithm is correct.

The topological scan takes $O(|\mathcal{V}|+|E|)$ time. The LP has $2|\mathcal{V}|+1$ constraints, so prune-and-search takes $O(|\mathcal{V}|)$ time. In the case $Z \neq \emptyset$, the single call to Algorithm H.3 has the same linear-time bound. Therefore, the total running time is $O(|\mathcal{V}| + |E|)$. □

Proposition H.2 (Correctness and computational complexity of d-OACIS)**.** *Algorithm H.2 reports* NONFINITE *if and only if* $\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}) = \infty$*; otherwise it returns* $(\bar{\theta}^\star_{\mathrm{d}}(\mathcal{I}), \bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}))$*. Moreover, the algorithm runs in* $O(|\mathcal{V}| + |E|)$ *time.*

Proof. By Proposition 4.2, $\bar{\alpha}^\star(\mathcal{I})$ is finite if and only if $\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I})$ is finite. Therefore, propagating NONFINITE from the OACIS call reports exactly the inputs with $\bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}) = \infty$. Suppose henceforth that these values are finite. By Proposition 4.1, the function $x \mapsto \bar{\alpha}(x; \mathcal{I})$ is convex. Its unique minimizer $x^\star$ and convexity imply that it is strictly decreasing to the left and strictly increasing to the right wherever it is finite. Thus every midpoint attaining the finite discrete minimum is a closest midpoint to $x^\star$ from the left or right, including $x^\star$ itself when it is a candidate midpoint. For each such midpoint, the algorithm includes a representative candidate in D. Calls to Algorithm H.3 evaluate their exact objective values, so the set of minimizing midpoint values among candidates in D equals the full set of minimizers over Θ_{d}. Applying the same fixed rule to this set as in Definition 3.3 therefore returns $(\bar{\theta}^\star_{\mathrm{d}}(\mathcal{I}), \bar{\alpha}^\star_{\mathrm{d}}(\mathcal{I}))$.

The OACIS call takes $O(|\mathcal{V}| + |E|)$ time. A scan over the candidates finds the two bracketing midpoints and a representative for each in $O(|\mathcal{V}|)$ time, and at most two sequential calls to Algorithm H.3 each take $O(|\mathcal{V}| + |E|)$. The constant number of linear-time calls preserves the $O(|\mathcal{V}| + |E|)$ time bound. □