DECISION-FOCUSED LEARNING UNDER JOINT UNCERTAINTY: LOCAL GEOMETRY AND COMPLEMENTARITY STRUCTURE

Anonymous authors
Paper under double-blind review

ABSTRACT

Decision-focused learning trains predictors for decisions in downstream optimisation problems. While objective uncertainty leaves the feasible set unchanged, constraint uncertainty can reshape the feasible set, render the represented decision infeasible for the true problem, and alter the active set governing local sensitivity. To study these effects, we focus on strongly convex quadratic programs and introduce `SIGNAL` by formalising a principled training loss function integrating objective-induced decision divergence, true-constraint violation, and active-set disagreement. We establish finite-sample bounds linking the `SIGNAL` loss to operational regret, true-violation frequency and magnitude, and dormant-bound prediction error. We then characterise the local learning information available through the represented optimiser, i.e., identify decision-sensitivity components alongside decision-null bound directions that remain observable through the primal–dual complementarity structure. We further show that `SIGNAL` is regionally convex quadratic in the prediction errors, with an affine output-space gradient and explicit active-set mismatch offsets. We finally demonstrate the effectiveness of `SIGNAL` on IEEE power-grid benchmarks, with particular emphasis on constraint shifts and their consequences for decision quality and true feasibility.

1 INTRODUCTION

Decision-focused learning (DFL) trains predictors for decisions in downstream optimisation problems—such as power dispatch, supply-chain planning, portfolio allocation, and healthcare resource allocation—rather than prediction accuracy (Donti et al., 2017; Wilder et al., 2019; Elmachtoub & Grigas, 2022; Mandi et al., 2024; Sadana et al., 2025; Lahoud et al., 2025). Yet, many decision problems involve uncertainty not only in the objective, but also in the constraints that define what is feasible. This distinction is fundamental. Predicting a cost changes which feasible decision is preferred; predicting a limit can change the feasible set itself, activate or deactivate constraints, and thereby alter the local sensitivity of the optimiser. It also breaks a basic intuition behind regret-based training: an infeasible decision may attain a lower true objective value than the true feasible optimum, so raw objective regret can become negative precisely when the decision is unacceptable. DFL under joint objective–constraint uncertainty therefore requires learning signals that distinguish downstream decision quality from true feasibility.

Only perfect prediction would erase this distinction. In practice, learning must contend with prediction error. These errors propagate anisotropically through the optimiser, as reflected in its local Karush-Kuhn-Tucker (KKT) structure: large parameter errors can lie in directions to which the optimiser is locally insensitive, while much smaller errors can produce substantial decision changes when they align with sensitive directions. Constraint uncertainty also introduces a blind spot. Within a regular active-set region, perturbing a represented-inactive bound does not move the optimiser at all, so its local decision Jacobian vanishes in that direction. Such a bound is nevertheless not irrelevant: the same constraint can become active under a nearby operating condition, at which point its prediction error directly affects the decision. Accordingly, a useful training objective must distinguish parameter errors that move the current decision from those that encode latent regime information not visible through the current decision map.

Contributions. To study these effects analytically while retaining nontrivial active-set changes, we focus on strongly convex quadratic programs (QPs) of the form $\text{minimise}_z \, \frac{1}{2} z^\top H z + c^\top z$ subject

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.

to $Az = b$ and $Gz \leq h$, with fixed H, A, G and jointly learned objective and constraint parameters (c, b, h). Our main contributions are as follows:

(i) We introduce SIGNAL (**s**ensitivity-**i**nformed **g**eometric and **n**ullspace-**a**ware **l**earning) by formalising a principled training loss function that integrates three complementary learning signals. The first term is an objective-induced Bregman divergence, which removes the signed first-order term from raw regret. For our quadratic objective, this divergence reduces exactly to the nonnegative objective-curvature discrepancy. The second term directly penalises true-constraint violations, capturing finite-error feasibility failures that vanish as the decision approaches the true optimum. The third term introduces explicit regime supervision for constraint-activity information that is not reliably exposed by local decision sensitivity. We instantiate this channel with a signed primal-dual margin built from primal slacks and dual multipliers. When the true margins on the supervised inequalities are bounded away from zero, sufficiently small regime loss guarantees recovery of their true active/inactive status. We further derive finite-sample bounds linking the combined loss to operational regret, true-violation frequency and magnitude, and dormant-bound prediction error.

(ii) Adapting parametric sensitivity to joint objective-constraint learning (Bemporad et al., 2002; Bonnans & Shapiro, 2000), we characterise the local learning information available through the represented optimiser. We show that, within a shared regular critical region, objective errors and active right-hand-side (RHS) errors affect the downstream decision through geometrically distinct tangent and H-normal directions, which are orthogonal in the objective-curvature metric. This explains what can be recovered through local optimiser sensitivity: the induced decision displacement carries structured information about both objective and active-constraint prediction errors. Yet, decision displacement alone does not reveal whether that movement remains inside the true feasible set. True-feasibility shaping supplies this missing information by identifying finite errors that cross feasibility boundaries. Finally, we show that the bound directions associated with represented-inactive inequalities span a kernel of the local decision Jacobian: their prediction errors produce no local decision motion, even though proximity to activation remains encoded in the primal-dual complementarity structure.

(iii) We further characterise the regional structure of the SIGNAL gradient: within each regular represented active-set region and fixed true-violation pattern, the loss is convex quadratic in the prediction errors and its output-space gradient is affine in $(\delta c, \delta b, \delta h)$, with active-set mismatch appearing through explicit regional offsets. This structure admits efficient evaluation through a single active-set adjoint, while dormant bounds contribute direct derivatives. We additionally characterise weakly active constraints through one-sided sensitivity, clarifying how the geometry extends at active-set boundaries where the local derivative is no longer unique. We evaluate SIGNAL on commonly used IEEE power-grid benchmarks (Christie, 1999; Zimmerman et al., 2011), with jointly uncertain power generation costs and directional transmission limits, including monitored line limits that are dormant in the nominal training regime but become decision-relevant under constraint shift. Across the compared methods, SIGNAL outperforms the adapted state-of-the-art baselines under joint uncertainty, delivering the lowest operational regret, true violation frequency, and true violation magnitude.

2 Related Work

Decision-focused learning and loss function design. DFL and related predict-then-optimise (PtO) methods train predictive models for downstream decision quality (Bertsimas & Kallus, 2020; Donti et al., 2017; Wilder et al., 2019; Elmachtoub & Grigas, 2022). In particular, Elmachtoub & Grigas (2022) introduce the smart predict-then-optimise loss and its convex surrogate SPO_+ for objective-parameter prediction with a fixed feasible set. Surrogate, smoothing, and perturbation methods provide alternative learning signals when direct differentiation through the downstream problem is difficult (Mandi & Guns, 2020; Berthet et al., 2020; Vlastelica Pogančić et al., 2020; Mulamba et al., 2021). Most of this literature emphasizes objective uncertainty or fixed feasible sets. Closest to our geometric perspective is PEAR, which interprets regret gradients under locally stable active constraints through tangent-space projection (Lee et al., 2026). Our setting instead considers joint objective and constraint uncertainty, where decision motion alone is insufficient.

Differentiable optimisation and local sensitivity. A complementary line of work develops mechanisms for differentiating through optimisation problems. OptNet established an influential differ-

entiable QP layer (Amos & Kolter, 2017). Later, differentiable cone and convex-program layers generalised this approach (Agrawal et al., 2019b;a), and modular implicit differentiation separates derivative rules from solver implementations (Blondel et al., 2022). More recent systems emphasise scalable or solver-agnostic QP backward passes, including `BPQP` (Pan et al., 2024), `dQP-DFL` (Magoon et al., 2025), and `DiffAPQP` (Xu & Teng, 2026). We instead use active-set and parametric-sensitivity structure (Bemporad et al., 2002; Bonnans & Shapiro, 2000) to characterise which components of joint objective-constraint prediction error are locally visible through the optimiser map and which remain hidden, and exploit that structure systematically in the design of our framework.

Learning with uncertain constraints and true infeasibility. Prior work also studies predict-and-optimize settings in which uncertain parameters enter the constraints. Hu et al. (2023) consider packing and covering LPs with unknown parameters in both the objective and constraints, introducing the notion of post-hoc regret that adds a penalty to the true objective gap in order to account for infeasible predicted decisions. Gupta & Zhang (2023) formulate DFL with unknown constraint parameters using bilevel inverse optimisation and feasibility safeguards, while Mandi et al. (2025) develops feasibility-aware losses for generic constrained optimisation and explicitly studies the trade-off between decision quality and infeasibility. Our analysis is complementary: we characterise the distinct learning signals generated by joint objective and constraint uncertainty.

Departure from prior work. To our knowledge, this is the first DFL analysis to provide a unified geometric account of joint objective-constraint uncertainty. Existing work either accommodates joint uncertainty without exposing its learning structure (Hu et al., 2023), or illuminates only one side of the problem—objective-side decision geometry (Lee et al., 2026) or constraint-side feasibility-aware learning (Mandi et al., 2025). We bring these perspectives together by characterising how divergence, violation, and disagreement encode distinct but coupled learning information, whose interaction through the optimiser is not revealed when objective and constraint uncertainty are analysed in isolation. Thus, methods centred on only one of these channels can leave consequential learning information unaccounted for, increasing the risk of true infeasibility or downstream suboptimality; `SIGNAL` is designed to expose and integrate all three.

3 PRELIMINARIES

Learning and optimisation setup. In DFL, the learner observes a context x available before the decision is made, predicts the unknown parameters of a downstream optimisation problem, and then solves that problem using the predictions. Formally, a sample $(x, \eta) \sim \mathcal{D}$ is drawn from the joint data distribution $\mathcal{D}$, where x is the observed context and $\eta = (c, b, h)$ contains the true optimisation parameters. A predictor ψ_θ, with trainable parameters θ, maps x to $\hat{\eta}_\theta(x) = (\hat{c}_\theta(x), \hat{b}_\theta(x), \hat{h}_\theta(x))$, abbreviated as $\hat{\eta} = (\hat{c}, \hat{b}, \hat{h})$. Throughout, unhatted quantities refer to the true problem, whereas hatted quantities refer to the represented problem. The true and represented downstream decisions, $z^\star(\eta) \in \mathbb{R}^n$ and $\hat{z}_\theta(x) \in \mathbb{R}^n$, abbreviated as $z^\star$ and $\hat{z}$, are defined by

$$z^\star(\eta) := \operatorname*{argmin}_{z \in \mathbb{R}^n} \Big\{ \frac{1}{2} z^\top H z + c^\top z \;:\; Az = b,\; Gz \le h \Big\}, \tag{1}$$

$$\hat{z}_\theta(x) := z^\star(\hat{\eta}_\theta(x)) = \operatorname*{argmin}_{z \in \mathbb{R}^n} \Big\{ \frac{1}{2} z^\top H z + \hat{c}_\theta(x)^\top z \;:\; Az = \hat{b}_\theta(x),\; Gz \le \hat{h}_\theta(x) \Big\}. \tag{2}$$

(3)

Here, $H \in \mathbb{S}^n_{++}$, $A \in \mathbb{R}^{p \times n}$, and $G \in \mathbb{R}^{m \times n}$ are fixed across samples. Since $H \succ 0$, each feasible instance has a unique primal optimiser. Standard regularity assumptions used in the theoretical results are stated explicitly where needed and collected in Appendix B. As only a subset of the constraint RHSs may be uncertain, we define the index sets of the predicted inequality and equality RHSs as $\mathcal{U} := \{i : h_i \text{ is predicted}\}$ and $\mathcal{E} := \{j : b_j \text{ is predicted}\}$. For the remaining rows, $\hat{h}_i = h_i$ for $i \notin \mathcal{U}$ and $\hat{b}_j = b_j$ for $j \notin \mathcal{E}$. All constraints remain enforced in the represented optimisation problem. For any constraint-row index set $\mathcal{S}$, $v_\mathcal{S}$ denotes the subvector of v indexed by $\mathcal{S}$, while $V_\mathcal{S}$ denotes the submatrix formed by the rows of V indexed by $\mathcal{S}$.

KKT conditions and active/inactive sets. Let $\lambda^\star \in \mathbb{R}^p$ and $\mu^\star \in \mathbb{R}^m_+$ be equality and inequality dual multipliers associated with $z^\star$. The KKT conditions are

$$Hz^\star + c + A^\top \lambda^\star + G^\top \mu^\star = 0, \qquad Az^\star = b, \qquad Gz^\star \le h, \tag{4}$$

$$\mu^\star \ge 0, \qquad \mu_i^\star (G_i z^\star - h_i) = 0, \qquad i = 1, \dots, m. \tag{5}$$

For each inequality row, define the true active and inactive sets

$$\mathcal{A}^\star := \mathcal{A}(\eta) = \{i \in \{1, \dots, m\} \,:\, G_i z^\star(\eta) = h_i\}, \tag{6}$$

$$\mathcal{I}^\star := \mathcal{I}(\eta) = \{1, \dots, m\} \setminus \mathcal{A}^\star. \tag{7}$$

The represented problem analogously has active and inactive sets $\hat{\mathcal{A}} := \mathcal{A}(\hat{\eta})$ and $\hat{\mathcal{I}} := \mathcal{I}(\hat{\eta})$.

True versus represented regime. The distinction between $(\mathcal{A}^\star, \mathcal{I}^\star)$ and $(\hat{\mathcal{A}}, \hat{\mathcal{I}})$ is essential. Each inequality row falls into one of four categories: correctly inactive ($i \in \mathcal{I}^\star \cap \hat{\mathcal{I}}$), correctly active ($i \in \mathcal{A}^\star \cap \hat{\mathcal{A}}$), missed active ($i \in \mathcal{A}^\star \cap \hat{\mathcal{I}}$), or falsely active ($i \in \mathcal{I}^\star \cap \hat{\mathcal{A}}$). The missed-active case is particularly important because the represented decision is ultimately applied to the true system: a constraint treated as inactive in the represented problem is in fact binding in the true problem. Within a regular critical region, $\mathcal{A}^\star$ is constant. We refer to this combinatorial active/inactive pattern as the active-set regime; a regime change occurs when this pattern changes.

4 Divergence, violation, and disagreement

Why raw regret fails when constraints are predicted. Let $f(z) := \frac{1}{2} z^\top H z + c^\top z$ denote the true objective. The raw regret is defined as

$$R(\hat{z}) := f(\hat{z}) - f(z^\star). \tag{8}$$

With a fixed feasible set, $\hat{z}$ and $z^\star$ are both feasible for the true problem, so optimality gives $R(\hat{z}) \geq 0$. However, once constraints are predicted, $\hat{z}$ may violate the true constraints and may satisfy $f(\hat{z}) < f(z^\star)$. Raw regret can therefore be negative and may reward infeasibility.

From raw regret to an objective-induced Bregman divergence. A replacement for raw regret should remain tied to the downstream objective while removing the signed first-order contribution whose sign need not be controlled when $\hat{z}$ is infeasible for the true problem. The quadratic objective makes this separation explicit. Let $\delta z := \hat{z} - z^\star$. Since $\nabla^2 f = H$, the exact second-order expansion of $f(\hat{z})$ around $z^\star$ gives $f(\hat{z}) = f(z^\star) + \nabla f(z^\star)^\top \delta z + \frac{1}{2} \delta z^\top H \delta z$. This decomposition suggests removing the first-order contribution while retaining the part of the objective discrepancy induced by curvature. More generally, this is precisely the construction of the Bregman divergence induced by the true downstream objective,

$$L_{\mathrm{dec}} := D_f(\hat{z}, z^\star) = f(\hat{z}) - f(z^\star) - \nabla f(z^\star)^\top (\hat{z} - z^\star). \tag{9}$$

For the strongly convex quadratic objective, the objective-induced Bregman divergence admits the exact form $L_{\mathrm{dec}} = \frac{1}{2} \|\hat{z} - z^\star\|_H^2$, i.e., it coincides with an H-weighted decision discrepancy. The metric comes directly from the objective Hessian rather than an arbitrary Euclidean choice; since $H \succ 0$, it is nonnegative and vanishes only at $\hat{z} = z^\star$. Proposition 1 explains how L_{dec} differs from raw regret at a constrained optimum.

Proposition 1 (Dual-residual correction for raw regret)**.** *Let $(z^\star, \lambda^\star, \mu^\star)$ satisfy the true KKT conditions and let $\mathcal{A}^\star$ denote the true active inequality set. For any represented decision $\hat{z}$,*

$$L_{\mathrm{dec}} = R(\hat{z}) + \lambda^{\star\top} (A\hat{z} - b) + \mu_{\mathcal{A}^\star}^{\star\top} (G_{\mathcal{A}^\star} \hat{z} - h_{\mathcal{A}^\star}). \tag{10}$$

No feasibility assumption on $\hat{z}$ is required. If $\hat{z}$ satisfies the true equalities and the true active inequalities at equality, the dual-residual correction vanishes and $L_{\mathrm{dec}} = R(\hat{z})$.

Proof. See Appendix E. Proposition 1 shows that L_{dec} removes the signed first-order contribution induced by constraint residuals while retaining the objective-curvature discrepancy between $\hat{z}$ and $z^\star$.

True feasibility shaping. While L_{dec} measures discrepancy from the true optimal decision, it does not explicitly distinguish decision errors according to their effect on true constraint satisfaction. To measure these effects on a common scale, for each nonzero equality row A_j, $j \in \mathcal{E}$, and inequality row G_i, $i \in \mathcal{U}$, define the canonical geometric scales $d_i := \sqrt{G_i H^{-1} G_i^\top} > 0$ and $e_j := \sqrt{A_j H^{-1} A_j^\top} > 0$. We then introduce a row-rescaling-invariant feasibility-shaping loss

$$L_{\mathrm{feas}} = \frac{1}{2} \sum_{j \in \mathcal{E}} \Big(\frac{A_j \hat{z} - b_j}{e_j} \Big)^2 + \frac{1}{2} \sum_{i \in \mathcal{U}} \Big[\frac{G_i \hat{z} - h_i}{d_i} \Big]_+^2, \tag{11}$$

Exact feasibility of the represented solution implies $A\hat{z} = \hat{b}$ and $G_i\hat{z} = \hat{h}_i$ for $i \in \hat{\mathcal{A}}$, so the corresponding true-constraint residuals equal the RHS prediction errors, i.e., $A_j\hat{z} - b_j = \hat{b}_j - b_j$ for $j \in \mathcal{E}$, and $G_i\hat{z} - h_i = \hat{h}_i - h_i$ for $i \in \mathcal{U} \cap \hat{\mathcal{A}}$. This implies that the equality terms in L_{feas} penalise prediction errors in both directions, whereas the inequality terms penalise only bound overestimation. Proposition 2 shows that L_{feas} is uniformly controlled by the decision discrepancy, without requiring the true and represented problems to share an active set.

Proposition 2 (Uniform feasibility control)**.** *Let* $W_b := \operatorname{diag}(e_j^{-2}\mathbf{1}\{j \in \mathcal{E}\})_{j=1}^{p}$ *and* $W_h := \operatorname{diag}\big(d_i^{-2}\mathbf{1}\{i \in \mathcal{U}\}\big)_{i=1}^{m}$. *For every* $\hat{z} \in \mathbb{R}^n$,

$$L_{\text{feas}} \le \lambda_{\max}\Big(H^{-1/2}\big(A^\top W_b A + G^\top W_h G\big)H^{-1/2}\Big)L_{\text{dec}}. \tag{12}$$

The result requires neither a common active set for the true and represented problems nor local differentiability of the solution map.

Proof. See Appendix E. Proposition 2 shows that L_{feas} is globally controlled at the same quadratic order as L_{dec}: it preserves the same asymptotic target while reshaping finite decision errors according to their true-feasibility consequences.

Dormant-bound blind spot and need for regime supervision. For a represented-inactive row $i \in \mathcal{U} \cap \hat{\mathcal{I}}$, fixed-regime sensitivity gives $\partial\hat{z}/\partial\hat{h}_i = 0$. Both L_{dec} and L_{feas} depend on $\hat{h}_i$ only through $\hat{z}$, while the true targets are fixed. Hence, within a fixed regular represented regime, $\partial L_{\text{dec}}/\partial\hat{h}_i = 0$ and $\partial L_{\text{feas}}/\partial\hat{h}_i = 0$. This is the dormant-bound blind spot: a bound can be locally invisible to both decision and feasibility gradients, yet the very constraint may become active as operating conditions or problem parameters shift. A separate signal is therefore needed to train such bound predictions before the corresponding constraints enter the represented active set and the bounds become visible to the decision and feasibility gradients.

Complementarity structure. KKT complementarity suggests a single signed scalar representation that unifies the two sides of constraint activity. The true and represented slacks are $s_i^\star = h_i - G_i z^\star$ and $\hat{s}_i = \hat{h}_i - G_i\hat{z}$. The normalised slack s_i/d_i is the H-metric distance from the optimiser to the constraint hyperplane $G_i z = h_i$, while the scaled multiplier $d_i\mu_i$ provides the dual quantity on the same canonical scale. Define the signed primal-dual margins $q_i^\star := \frac{s_i^\star}{d_i} - d_i\mu_i^\star$ and $\hat{q}_i := \frac{\hat{s}_i}{d_i} - d_i\hat{\mu}_i$. At an exact strictly complementary KKT point, complementary slackness ensures that only one of the two terms is nonzero:

$$q_i^\star = \begin{cases} s_i^\star/d_i > 0, & i \in \mathcal{I}^\star, \\ -d_i\mu_i^\star < 0, & i \in \mathcal{A}^\star. \end{cases} \tag{13}$$

Thus, a two-variable complementarity pair $(s_i^\star, \mu_i^\star)$ collapses to a single signed margin: its sign identifies the constraint-activity branch, while its magnitude quantifies how strongly the row lies on that branch. The activity-switching boundary is represented by $q_i^\star = 0$. We then define a row-rescaling-invariant regime-supervision loss

$$L_{\text{reg}} = \frac{1}{2k}\sum_{i \in \mathcal{U}}(\hat{q}_i - q_i^\star)^2, \tag{14}$$

where $k = |\mathcal{U}|$. Because the sign of q_i identifies the active/inactive branch, L_{reg} can directly penalise constraint-regime disagreements: a missed-active row $i \in \mathcal{A}^\star \cap \hat{\mathcal{I}}$ or a falsely active row $i \in \mathcal{I}^\star \cap \hat{\mathcal{A}}$ places the true and represented signed margins on opposite sides of zero and therefore incurs a nonzero margin error. The sum in L_{reg} is restricted to $i \in \mathcal{U}$ because otherwise fixed rows $i \notin \mathcal{U}$ may contribute nonzero signed-margin errors; the canonical scaling d_i removes the need for an additional primal-dual weighting parameter. Proposition 3 provides an active-set certificate.

Proposition 3 (Active-set certificate)**.** *Assume exact primal-dual KKT pairs for the true and represented problems, and define the true margin* $\gamma_i := |q_i^\star|$. *Then,*

$$\sum_{i \in \mathcal{U} \cap (\mathcal{A}^\star \triangle \hat{\mathcal{A}})} \gamma_i^2 \le 2kL_{\text{reg}}. \tag{15}$$

Consequently, under strict complementarity on the supervised rows, with $\gamma_{\min,\mathcal{U}} := \min_{i\in\mathcal{U}} \gamma_i > 0$,

$$L_{\text{reg}} < \frac{\gamma_{\min,\mathcal{U}}^2}{2k} \;\Rightarrow\; \hat{\mathcal{A}} \cap \mathcal{U} = \mathcal{A}^\star \cap \mathcal{U}. \tag{16}$$

The result is not restricted to a shared critical region.

Proof. See Appendix E. Proposition 3 elevates L_{reg} from an auxiliary penalty to a quantitative active-set certificate. It states that every supervised constraint-regime disagreement incurs at least the squared true margin of that row; once L_{reg} falls below the minimum-margin threshold, exact active-set recovery on the supervised rows is guaranteed. Near weak activity, γ_i can vanish, so exact recovery cannot be inferred from a positive global margin; Appendix D characterises the corresponding directional sensitivity.

Training loss and empirical downstream control. Using L_{dec} as the reference scale, the resulting per-sample loss is then constructed as

$$L_{\text{signal}} = L_{\text{dec}} + \lambda_{\text{feas}} L_{\text{feas}} + \lambda_{\text{reg}} L_{\text{reg}}, \qquad \lambda_{\text{feas}}, \lambda_{\text{reg}} \geq 0. \tag{17}$$

The three terms, as discussed, play deliberately distinct structural roles in training. Evaluation, however, requires operational criteria that reflect the actual cost and dissatisfaction incurred by the true system. For test instance n, let $\hat{z}_n$ and $z_n^\star$ denote the represented and true-optimal decisions, respectively. Over N test instances, define the empirical operational regret as

$$R_{\text{op}} := \frac{1}{N} \sum_{n=1}^{N} \Big(f_n(\hat{z}_n) - f_n(z_n^\star) + \sum_{j\in\mathcal{E}} C_{b,j} |A_j \hat{z}_n - b_{n,j}| + \sum_{i\in\mathcal{U}} C_{h,i} [G_i \hat{z}_n - h_{n,i}]_+ \Big), \tag{18}$$

where $C_{b,j}, C_{h,i} \geq 0$ specify the application-dependent cost per unit of true-constraint violation. For test instance n, let $v_{n,i} := [G_i \hat{z}_n - h_{n,i}]_+$, for $i \in \mathcal{U}$, denote the true violation of supervised inequality i. Over N test instances, define the empirical true-violation frequency and magnitude as

$$V_{\text{freq}} := \frac{1}{N} \sum_{n=1}^{N} \mathbb{1}\big\{ \max_{i\in\mathcal{U}} v_{n,i} > \epsilon \big\}, \qquad V_{\text{mag}} := \frac{1}{kN} \sum_{n=1}^{N} \sum_{i\in\mathcal{U}} v_{n,i}, \tag{19}$$

where $\epsilon > 0$ is the numerical violation tolerance. Finally, when a supervised bound $\ell \in \mathcal{U}$ is inactive in both the true and represented problems for every test instance, define the mean normalised dormant-bound prediction error $E_{\text{dorm}} := \frac{1}{N} \sum_{n=1}^{N} |\delta h_\ell^{(n)}| / d_\ell$ as a theoretical control quantity. Note that while R_{op}, V_{freq}, and V_{mag} are the ultimate evaluation criteria, L_{signal} is designed to control the mechanisms that determine them. Theorem 1 formalises this concept.

Theorem 1 (Empirical downstream control)**.** *Consider N test instances and suppose that every represented decision satisfies all fixed constraints. Let $\lambda_{\text{feas}} > 0$ and $\lambda_{\text{reg}} > 0$, and define $\bar{L}_{\text{signal}} := \frac{1}{N} \sum_{n=1}^{N} L_{\text{signal}}^{(n)}$. For each instance n, let*

$$\Xi_n{}^\star := \left(\sum_{j\in\mathcal{E}} e_j^2 (C_{b,j} + |\lambda_{n,j}^\star|)^2 + \sum_{i\in\mathcal{U}} d_i^2 [C_{h,i} - \mu_{n,i}^\star]_+^2 \right)^{1/2}, \tag{20}$$

$$\Gamma_n^\star := \max_{\|u\|_H = 1} \sum_{i\in\mathcal{A}_n^\star} \mu_{n,i}^\star [-G_i u]_+, \qquad d_{\max} := \max_{i\in\mathcal{U}} d_i, \qquad d_{\mathcal{U}} = \big(\sum_{i\in\mathcal{U}} d_i^2 \big)^{1/2}. \tag{21}$$

Then, the empirical operational regret satisfies

$$R_{\text{op}} \leq \bar{L}_{\text{signal}} + \left(\Big(\frac{2}{N} \sum_{n=1}^{N} \Big(\Gamma_n^{\star\,2} + \frac{\Xi_n^{\star\,2}}{\lambda_{\text{feas}}} \Big) \Big) \bar{L}_{\text{signal}} \right)^{1/2}. \tag{22}$$

Moreover, the empirical true-violation frequency, true-violation magnitude, and mean normalised dormant-bound prediction error satisfy

$$V_{\text{freq}} \leq \frac{2 d_{\max}^2}{\lambda_{\text{feas}} \epsilon^2} \bar{L}_{\text{signal}}; \quad V_{\text{mag}} \leq \frac{d_{\mathcal{U}}}{k} \left(\frac{2 \bar{L}_{\text{signal}}}{\lambda_{\text{feas}}} \right)^{1/2}; \quad E_{\text{dorm}} \leq \left(\Big(2 + \frac{2k}{\lambda_{\text{reg}}} \Big) \bar{L}_{\text{signal}} \right)^{1/2}. \tag{23}$$

Proof. See Appendix E. Theorem 1 provides finite-sample control of downstream performance. Analogous bounds for uncertain equality residuals follow from the same feasibility-loss argument.

5 SENSITIVITY AND DECISION GEOMETRY

This section characterises the geometric structure of joint prediction errors, distinguishing directions that move the decision from those that are locally decision-null. Fix a true parameter instance η with regular solution $z^\star(\eta)$ and active inequality set $\mathcal{A}^\star$. Define $J = [A^\top\ G_{\mathcal{A}^\star}^\top]^\top$, $r = [b^\top\ h_{\mathcal{A}^\star}^\top]^\top$, and $y^\star = [\lambda^{\star\top}\ \mu^{\star\top}_{\mathcal{A}^\star}]^\top$, where r is the active-constraint RHS vector, and assume that J has full row rank. Define

$$S = JH^{-1}J^\top, \qquad P = H^{-1} - H^{-1}J^\top S^{-1}JH^{-1}, \qquad K = H^{-1}J^\top S^{-1}. \tag{24}$$

The KKT system reduces to

$$Hz^\star + c + J^\top y^\star = 0, \qquad Jz^\star - r = 0. \tag{25}$$

Define the prediction error $\delta\eta := \hat{\eta} - \eta = (\delta c, \delta b, \delta h)$, where $\delta c := \hat{c} - c$, $\delta b := \hat{b} - b$, and $\delta h := \hat{h} - h$. For the represented problem, let $\widehat{\mathcal{A}}$ denote its active set and analogously define $\hat{J} = [A^\top\ G_{\hat{\mathcal{A}}}^\top]^\top$, $\hat{r} = [\hat{b}^\top\ \hat{h}_{\hat{\mathcal{A}}}^\top]^\top$, and $\hat{y} = [\hat{\lambda}^\top\ \hat{\mu}_{\hat{\mathcal{A}}}^\top]^\top$. Whenever $\hat{\mathcal{A}} = \mathcal{A}^\star$, we have $\hat{J} = J$ and define the active-constraint RHS error $\delta r := \hat{r} - r$. Theorem 2 characterises the exact fixed-region geometry induced by joint objective and active-constraint RHS perturbations.

Theorem 2 (Fixed-region decision and feasibility geometry)**.** *Suppose the line segment joining the true and represented parameter instances remains in a common regular critical region. Then, $\hat{\mathcal{A}} = \mathcal{A}^\star$ and the decision displacement satisfies*

$$\delta z = -P\delta c + K\delta r =: \delta z_{\tan} + \delta z_{\mathrm{nor}}, \qquad \delta z_{\tan} := -P\delta c, \quad \delta z_{\mathrm{nor}} := K\delta r, \tag{26}$$

where $\delta z_{\tan} \in \ker J$, $\delta z_{\mathrm{nor}} \in \mathrm{range}(H^{-1}J^\top) = (\ker J)^{\perp_H}$, and $\delta z_{\tan}^\top H \delta z_{\mathrm{nor}} = 0$. Hence, objective error induces tangent motion along the current active manifold, whereas active-constraint RHS error induces the H-normal motion required to track displacement of that manifold. Furthermore, the induced decision loss and true-feasibility loss decompose exactly as

$$L_{\mathrm{dec}} = \frac{1}{2}\,\delta c^\top P\,\delta c + \frac{1}{2}\,\delta r^\top S^{-1}\delta r, \tag{27}$$

$$L_{\mathrm{feas}} = \frac{1}{2}\sum_{j\in\mathcal{E}}\left(\frac{\delta b_j}{e_j}\right)^2 + \frac{1}{2}\sum_{i\in\mathcal{U}\cap\mathcal{A}^\star}\left[\frac{\delta h_i}{d_i}\right]_+^2 + \frac{1}{2}\sum_{i\in\mathcal{U}\cap\mathcal{I}^\star}\left[\frac{-G_iP\delta c + G_iK\delta r - s_i^\star}{d_i}\right]_+^2. \tag{28}$$

Proof. See Appendix E. Theorem 2 states that objective and active-RHS errors generate orthogonal tangent and H-normal decision motion, while objective perturbations in $\mathrm{range}(J^\top)$ and inactive-bound perturbations are locally decision-null. The latter can nevertheless alter future active-set status, motivating Proposition 4.

Proposition 4 (Decision-null dormant-bound subspace)**.** *Let $\hat{\eta}$ lie in a regular represented critical region and define $J_h := \partial\hat{z}/\partial\hat{h}_{\mathcal{U}}$. For each $i \in \mathcal{U}$, let u_i denote the standard basis vector of $\mathbb{R}^{|\mathcal{U}|}$ associated with bound coordinate i. Then,*

$$\ker J_h = \mathrm{span}\{u_i : i \in \mathcal{U}\cap\hat{\mathcal{I}}\}. \tag{29}$$

Hence, perturbations of represented-inactive predicted bounds are locally invisible to the represented decision map. Moreover, for any perturbation $\delta\hat{h}$ supported on $\mathcal{U}\cap\hat{\mathcal{I}}$ such that $\hat{h} + \delta\hat{h}$ remains in the same represented critical region, the regime loss satisfies

$$L_{\mathrm{reg}}(\hat{h} + \delta\hat{h}) = L_{\mathrm{reg}}(\hat{h}) + \frac{1}{k}\sum_{i\in\mathcal{U}\cap\hat{\mathcal{I}}}\left(\hat{q}_i - q_i^\star\right)\frac{\delta\hat{h}_i}{d_i} + \frac{1}{2k}\sum_{i\in\mathcal{U}\cap\hat{\mathcal{I}}}\left(\frac{\delta\hat{h}_i}{d_i}\right)^2. \tag{30}$$

Thus, while the represented decision is locally constant along dormant-bound directions, L_{reg} retains an explicit direct derivative and strictly positive quadratic curvature along those same directions.

Proof. See Appendix E. Proposition 4 identifies active-set regime information that decision-based learning signals cannot recover locally.

6 TRAINING AND DIFFERENTIATION

This section translates the preceding geometric structure into an efficient training rule for `SIGNAL`. Training minimises the empirical average of L_{signal} over the training samples. Algorithm 1 in Appendix A summarises one training step.

Gradient evaluation. Consider a regular represented critical region with active set $\hat{\mathcal{A}}$ and inactive set $\hat{\mathcal{I}}$. Under the regularity assumptions in Appendix B, the represented QP is feasible, $\hat{J} = [A; G_{\hat{\mathcal{A}}}]$ has full row rank, the active set is locally constant, and the primal-dual solution map is differentiable. Let $k := |\mathcal{U}|$; when $k = 0$, the regime channel is omitted. For $k > 0$, define the masked margin errors $w_i := (\hat{q}_i - q_i^\star)/k$ for $i \in \mathcal{U}$ and $w_i := 0$ otherwise, and let $w_i^{\mathrm{I}} := w_i \mathbb{1}\{i \in \hat{\mathcal{I}}\}$. Let $D = \mathrm{diag}(d_1, \ldots, d_m)$, assigning arbitrary positive entries on unsupervised rows, whose weights are zero. We use the canonical feasibility weights W_b, W_h defined in Proposition 2. We differentiate the branch-wise margin representation: $\hat{q}_i = -d_i \hat{\mu}_i$ on $\hat{\mathcal{A}}$ and $\hat{q}_i = (\hat{h}_i - G_i \hat{z})/d_i$ on $\hat{\mathcal{I}}$. Its differential is $dL_{\mathrm{reg}} = (w^{\mathrm{I}})^\top D^{-1}(d\hat{h} - G\, d\hat{z}) - (Dw)^\top_{\hat{\mathcal{A}}} d\hat{\mu}_{\hat{\mathcal{A}}}$. The upstream derivatives of the combined loss are therefore $q_z = H(\hat{z} - z^\star) + \lambda_{\mathrm{feas}} T - \lambda_{\mathrm{reg}} G^\top D^{-1} w^{\mathrm{I}}$ and $q_y = -\lambda_{\mathrm{reg}}[0_p; (Dw)_{\hat{\mathcal{A}}}]$, where $T = [A^\top W_b(A\hat{z} - b) + G^\top W_h[G\hat{z} - h]_+]$. Here, $[u; v]$ denotes vertical concatenation, and $\hat{y} = [\hat{\lambda}; \hat{\mu}_{\hat{\mathcal{A}}}]$. The partial derivatives $q_z = \partial_{\hat{z}} L_{\mathrm{signal}}$ and $q_y = \partial_{\hat{y}} L_{\mathrm{signal}}$ treat the primal-dual state and predicted RHSs as independent inputs to the branch-wise loss before differentiating the KKT reconstruction. The dual contribution q_y transmits regime supervision through the active multipliers and must be retained.

Define the symmetric represented KKT matrix $\hat{M} = \left[\begin{smallmatrix} H & \hat{J}^\top \\ \hat{J} & 0 \end{smallmatrix}\right]$. All optimiser-mediated contributions share the adjoint solve $\hat{M}[a; v] = [q_z; q_y]$ and $v = [v_b; v_{h,\hat{\mathcal{A}}}]$. Equivalently, eliminating the primal adjoint yields $v = \hat{S}^{-1}(\hat{J}H^{-1}q_z - q_y)$ and $a = H^{-1}(q_z - \hat{J}^\top v)$. The Schur-complement system has $p + |\hat{\mathcal{A}}|$ unknowns, compared with $n + p + |\hat{\mathcal{A}}|$ in the saddle-point system, and can reduce computational cost when the active set is small and solves with the fixed H can be reused efficiently. Accordingly, the gradients with respect to the predicted outputs are

$$g_c := \nabla_{\hat{c}} L_{\mathrm{signal}} = -a, \qquad g_{b,\mathcal{E}} := \nabla_{\hat{b}_{\mathcal{E}}} L_{\mathrm{signal}} = (v_b)_{\mathcal{E}}, \tag{31}$$

$$g_{h,i} := \frac{\partial L_{\mathrm{signal}}}{\partial \hat{h}_i} = \begin{cases} v_{h,i}, & i \in \mathcal{U} \cap \hat{\mathcal{A}}, \\ \lambda_{\mathrm{reg}} w_i / d_i, & i \in \mathcal{U} \cap \hat{\mathcal{I}}. \end{cases} \tag{32}$$

Regional error-to-gradient structure. Fix a regular represented critical region with active set $\hat{\mathcal{A}}$. Within this region, the KKT reconstruction gives affine error representations $\delta z = Z\delta\eta + z_0$ and $\delta q_{\mathcal{U}} = Q\delta\eta + q_0$, where $\delta q_{\mathcal{U}} := \hat{q}_{\mathcal{U}} - q^\star_{\mathcal{U}}$, Z and Q are regionally constant, and the offsets z_0 and q_0 arise from true-represented active-set mismatch through the associated true slacks and multipliers. Their explicit construction is given in Appendix C. Proposition 5 characterises how prediction errors map to the `SIGNAL` gradient.

Proposition 5 (Regional affine structure)**.** *Fix a regular represented active-set region with active set $\hat{\mathcal{A}}$, and let Z, Q, z_0, q_0 denote the regional matrices and offsets. On any subregion characterised by the true-violation pattern $\chi_i := \mathbb{1}\{G_i \delta z - s_i^\star > 0\}$, with $C_\chi := \mathrm{diag}(\chi_1, \ldots, \chi_m)$, the output-space gradient satisfies*

$$\nabla_{\hat{\eta}} L_{\mathrm{signal}} = \Omega_{\hat{\mathcal{A}},\chi} \delta\eta + \zeta_{\hat{\mathcal{A}},\chi}, \tag{33}$$

where

$$\Omega_{\hat{\mathcal{A}},\chi} = Z^\top H Z + \lambda_{\mathrm{feas}} Z^\top \Big(A^\top W_b A + G^\top W_h C_\chi G \Big) Z + \frac{\lambda_{\mathrm{reg}}}{k} Q^\top Q \succeq 0, \tag{34}$$

$$\zeta_{\hat{\mathcal{A}},\chi} = Z^\top H z_0 + \lambda_{\mathrm{feas}} Z^\top \Big(A^\top W_b A z_0 + G^\top W_h C_\chi \big(G z_0 - s^\star \big) \Big) + \frac{\lambda_{\mathrm{reg}}}{k} Q^\top q_0, \tag{35}$$

for $k > 0$. Thus, within each fixed regional regime, L_{signal} is convex quadratic in the prediction errors, and its output-space gradient is an affine function of $(\delta c, \delta b, \delta h)$.

Proof. See Appendix E. Proposition 5 exposes the anisotropy of DFL under joint uncertainty: prediction errors of comparable magnitude can generate different `SIGNAL` gradients depending on their direction and the local active-set structure.

Predictor update and boundary behaviour. The chain rule maps the assembled output-space gradients to the predictor parameters: $\nabla_\theta L_{\mathrm{signal}} = (\frac{\partial(\hat{c}, \hat{b}_{\mathcal{E}}, \hat{h}_{\mathcal{U}})}{\partial \theta})^\top [-a;\ (v_b)_{\mathcal{E}};\ g_{h,\mathcal{U}}]$. This gradient is used by a standard first-order optimiser. Note that the preceding derivatives apply within a regular represented critical region. At an active-set transition, a classical Jacobian may fail to exist, and a fixed-active-set adjoint alone does not characterise all directional responses. Appendix D gives the corresponding one-sided sensitivity under its stated assumptions.

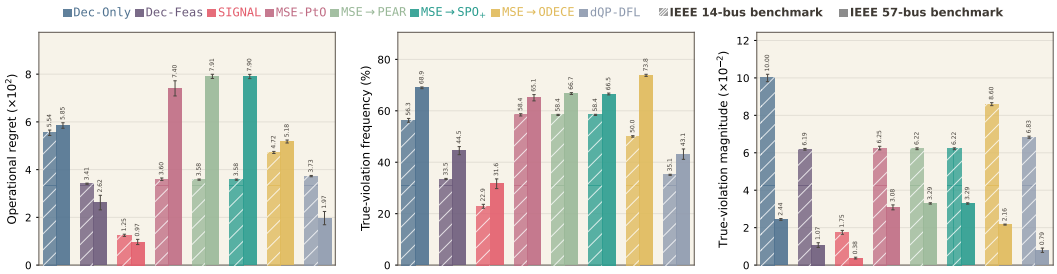

Figure 1: Performance on IEEE power-grid benchmarks under constraint shift and observation noise: operational regret (left), true-violation frequency (middle), and true-violation magnitude (right).

7 Experiments

Experimental setup and evaluation criteria. We evaluate SIGNAL on the IEEE 14-bus and IEEE 57-bus benchmarks (Christie, 1999; Zimmerman et al., 2011), jointly predicting power generation costs and directional transmission-limit inequalities under constraint shift and observation noise. We compare eight methods using the same predictor class and paired datasets: Dec-Only, Dec-Feas, SIGNAL, MSE-PtO, dQP-DFL, MSE → PEAR, MSE → SPO_+, and MSE → ODECE. The first three form a nested ablation sequence that isolates the contributions of feasibility shaping and regime supervision in SIGNAL; MSE-PtO follows the standard predict-then-optimise approach, while dQP-DFL (Magoon et al., 2025) serves as a differentiable-optimisation baseline that propagates gradients through the represented problem. For the adapted sequential baselines, an MSE-trained constraint predictor is frozen before training the objective predictor with PEAR (Lee et al., 2026) or SPO_+ (Elmachtoub & Grigas, 2022); conversely, an MSE-trained objective predictor is frozen before training the constraints with ODECE (Mandi et al., 2025). Appendix F provides experimental details, including the underlying optimisation problems, data generation and shift, predictors, baseline adaptations, calibration protocol, and repeated-run evaluation. We evaluate performance using three criteria: operational regret R_{op}, true-violation frequency V_{freq}, and true-violation magnitude V_{mag}.

Results and interpretation. Under constraint shift at observation-noise level $\sigma_{\text{obs}} = 0.10$ and with training sample size $N_{\text{train}} = 1500$, Figure 1 shows that SIGNAL achieves the lowest mean value on all three criteria. In particular, for the IEEE 57-bus benchmark, the left panel shows an operational regret of approximately 97 for SIGNAL, compared with 197 for the strongest external comparator, dQP-DFL, a reduction of approximately 51%. The nested ablation sequence shows a marked progression: operational regret decreases from approximately 585 for Dec-Only to 262 for Dec-Feas and further to 97 for SIGNAL. The middle panel shows that these improvements accompany fewer infeasible decisions: true-violation frequency decreases from 68.9% for Dec-Only to 44.5% for Dec-Feas and further to 31.6% for SIGNAL, below the 43.1% attained by dQP-DFL. The right panel shows that SIGNAL reduces average violation magnitude by approximately 64% relative to Dec-Feas, 52% relative to dQP-DFL, and 88% relative to both MSE → PEAR and MSE → SPO_+. Thus, the reduction in operational regret is accompanied by lower frequency and magnitude of true-constraint violations. The results for the IEEE 14-bus benchmark follow the same pattern. Overall, the results indicate that SIGNAL generalises well across both benchmarks under constraint shift; Appendix F reports additional results.

8 Conclusion

DFL under joint uncertainty is not a homogeneous learning problem. Objective errors, active-constraint RHS errors, and dormant bound errors interact with the optimiser through different geometric channels and therefore carry different downstream semantics. SIGNAL makes this structure explicit. Empirically, SIGNAL outperforms competing DFL methods and exhibits stronger operational robustness. The central lesson is that effective DFL under joint uncertainty must reason not only about how prediction errors move the current optimiser, but also about which errors threaten true feasibility and which remain hidden from its local sensitivity. The present study is limited to strongly convex QPs with fixed H, A, and G, and two IEEE power-grid benchmarks, so the empirical conclusions should not be interpreted as establishing universal scalability or robustness beyond this setting. Extending this perspective beyond fixed-geometry QPs to broader convex programs, uncertain Hessians and constraint matrices, and nonstationary settings is a natural next step.

Reproducibility Statement

The assumptions underlying the theoretical results are stated in Appendix B, and complete proofs are provided in Appendix E. Appendix F documents the experimental protocol, benchmark construction, data generation and constraint shift, predictor and baseline configurations, calibration procedure, and repeated-run evaluation. The supplementary material contains the implementation codes used for all experiments. Together, these materials provide the information needed to reproduce the theoretical and empirical results reported in the paper.

AI Use Disclosure

Generative AI tools were used to assist with cleaning and reformatting datasets, to support qualitative and thematic analysis of experimental materials, and to edit the final manuscript for clarity and readability. All AI-assisted outputs were reviewed and verified by the authors, who take full responsibility for the final content, analyses, claims, and presentation of this work.

References

Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J. Zico Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, volume 32, pp. 9558–9570, 2019a. URL https://proceedings.neurips.cc/paper/2019/hash/9ce3c52fc54362e22053399d3181c638-Abstract.html.

Akshay Agrawal, Shane Barratt, Stephen Boyd, Enzo Busseti, and Walaa M. Moursi. Differentiating through a cone program. *Journal of Applied and Numerical Optimization*, 1(2):107–115, 2019b. URL https://doi.org/10.23952/jano.1.2019.2.02.

Brandon Amos and J. Zico Kolter. OptNet: Differentiable optimization as a layer in neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 136–145. PMLR, 2017. URL https://proceedings.mlr.press/v70/amos17a.html.

Alberto Bemporad, Manfred Morari, Vivek Dua, and Efstratios N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002. URL https://doi.org/10.1016/S0005-1098(01)00174-1.

Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. In *Advances in Neural Information Processing Systems*, volume 33, pp. 9508–9519, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/6bb56208f672af0dd65451f869fedfd9-Abstract.html.

Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020. URL https://doi.org/10.1287/mnsc.2018.3253.

Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and modular implicit differentiation. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/228b9279ecf9bbafe582406850c57115-Abstract-Conference.html.

J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, New York, 2000. URL https://doi.org/10.1007/978-1-4612-1394-9.

Richard D. Christie. Power systems test case archive. University of Washington, Electrical Engineering, 1999. URL https://labs.ece.uw.edu/pstca/. Accessed 2026.

Priya L. Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5484–5494, 2017. URL https://proceedings.neurips.cc/paper/2017/hash/3fc2c60b5782f641f76bcefc39fb2392-Abstract.html.

Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1): 9–26, 2022. URL https://doi.org/10.1287/mnsc.2020.3922.

Rishabh Gupta and Qi Zhang. An inverse optimization approach to decision-focused learning. In *Computer Aided Chemical Engineering*, volume 52, pp. 1359–1365. Elsevier, 2023. URL https://doi.org/10.1016/B978-0-443-15274-0.50216-X.

Xinyi Hu, Jasper C. H. Lee, and Jimmy H. M. Lee. Predict+optimize for packing and covering LPs with unknown parameters in constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 3987–3995, 2023. URL https://doi.org/10.1609/aaai.v37i4.25513.

Alan Anis Lahoud, Ahmad Saeed Khan, Erik Schaffernicht, Marco Trincavelli, and Johannes Andreas Stork. Predict-and-optimize techniques for data-driven optimization problems: A review. *Neural Processing Letters*, 57(2):40, 2025. URL https://doi.org/10.1007/s11063-025-11746-w.

Junhyeong Lee, Sangjin Jin, and Yongjae Lee. Decision-focused learning via tangent-space projection of prediction error. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026. URL https://icml.cc/virtual/2026/poster/63194.

Connor W. Magoon, Fengyu Yang, Noam Aigerman, and Shahar Z. Kovalsky. Differentiation through black-box quadratic programming solvers. In *Advances in Neural Information Processing Systems*, volume 38, pp. 120073–120104, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/9bd9aa8c2326b67321eb1f5f281ba618-Abstract-Conference.html.

Jayanta Mandi and Tias Guns. Interior point solving for LP-based prediction+optimisation. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/51311013e51adebc3c34d2cc591fefee-Abstract.html.

Jayanta Mandi, James Kotary, Senne Berden, Maxime Mulamba, Victor Bucarey, Tias Guns, and Ferdinando Fioretto. Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*, 80:1623–1701, 2024. URL https://doi.org/10.1613/JAIR.1.15320.

Jayanta Mandi, Marianne Defresne, Senne Berden, and Tias Guns. Feasibility-aware decision-focused learning for predicting parameters in the constraints. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/ef158a0d4f5364741a571b8d1d44fb1b-Abstract-Conference.html.

Maxime Mulamba, Jayanta Mandi, Michelangelo Diligenti, Michele Lombardi, Victor Bucarey, and Tias Guns. Contrastive losses and solution caching for predict-and-optimize. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pp. 2833–2840, 2021. URL https://doi.org/10.24963/ijcai.2021/390.

Jianming Pan, Zeqi Ye, Xiao Yang, Xu Yang, Weiqing Liu, Lewen Wang, and Jiang Bian. BPQP: A differentiable convex optimization framework for efficient end-to-end learning. In *Advances in Neural Information Processing Systems*, volume 37, pp. 77468–77493, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/8db12f7214d3a1a0c450ba751163e0fd-Abstract-Conference.html.

Utsav Sadana, Abhilash Reddy Chenreddy, Erick Delage, Alexandre Forel, Emma Frejinger, and Thibaut Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025. URL https://doi.org/10.1016/j.ejor.2024.03.020.

Marin Vlastelica Pogančić, Anselm Paulus, Vít Musil, Georg Martius, and Michal Rolínek. Differentiation of blackbox combinatorial solvers. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=BkevoJSYPB.

Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1658–1665, 2019. URL https://doi.org/10.1609/aaai.v33i01.33011658.

Wangkun Xu and Fei Teng. Structured differentiable optimization for efficient decision-focused learning in power systems, 2026. URL https://arxiv.org/abs/2608.04189.

Ray D. Zimmerman, Carlos E. Murillo-Sánchez, and Robert J. Thomas. MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on Power Systems*, 26(1):12–19, 2011. URL https://doi.org/10.1109/TPWRS.2010.2051168.

A ALGORITHM

Algorithm 1 summarises an adjoint form of one training step of SIGNAL: a represented-QP solve identifies the active set, a primal-dual KKT reconstruction provides the quantities required to evaluate the loss, and a single adjoint solve yields the output-space gradient.

Algorithm 1 Adjoint form of one training step of SIGNAL

Require: Context x; predictor parameters θ; fixed H, A, G; true RHSs (b, h); targets $z^\star, \{q_i^\star\}_{i\in\mathcal{U}}$; positive scales e_j, d_i; loss weights $\lambda_{\text{feas}}, \lambda_{\text{reg}}$

Forward pass

1: Predict $\hat{\eta} \leftarrow \psi_\theta(x)$, where $\hat{\eta} = [\hat{c}; \hat{b}_\mathcal{E}; \hat{h}_\mathcal{U}]$
2: Set $\hat{b}_j \leftarrow b_j$ for $j \notin \mathcal{E}$ and $\hat{h}_i \leftarrow h_i$ for $i \notin \mathcal{U}$
3: Solve the represented QP outside autodiff to obtain $\tilde{z}$; identify its active set $\hat{\mathcal{A}}$
4: Form $\hat{J} \leftarrow [A; G_{\hat{\mathcal{A}}}]$, $\hat{r} \leftarrow [\hat{b}; \hat{h}_{\hat{\mathcal{A}}}]$, and $\hat{M} \leftarrow \begin{bmatrix} H & \hat{J}^\top \\ \hat{J} & 0 \end{bmatrix}$
5: Factorise $\hat{M}$ and reconstruct $(\hat{z}, \hat{y})$ by solving $\hat{M}[\hat{z}; \hat{y}] = [-\hat{c}; \hat{r}]$, where $\hat{y} = [\hat{\lambda}; \hat{\mu}_{\hat{\mathcal{A}}}]$
6: For each $i \in \mathcal{U}$, set $\hat{q}_i \leftarrow -d_i\hat{\mu}_i$ if $i \in \hat{\mathcal{A}}$, and $\hat{q}_i \leftarrow (\hat{h}_i - G_i\hat{z})/d_i$ otherwise
7: Evaluate $L_{\text{dec}}, L_{\text{feas}}, L_{\text{reg}}$ using the true targets; set $L_{\text{signal}} \leftarrow L_{\text{dec}} + \lambda_{\text{feas}}L_{\text{feas}} + \lambda_{\text{reg}}L_{\text{reg}}$

Backward pass (active set held fixed)

8: Compute $q_z \leftarrow \partial_{\hat{z}} L_{\text{signal}}$, $q_y \leftarrow \partial_{\hat{y}} L_{\text{signal}}$, and $t_i \leftarrow \lambda_{\text{reg}}\, \partial_{\hat{q}_i} L_{\text{reg}}$ for $i \in \mathcal{U}$
9: Reuse the factorisation to solve $\hat{M}[a; v] = [q_z; q_y]$, where $v = [v_b; v_h]$
10: Set $g_c \leftarrow -a$ and $g_{b,j} \leftarrow v_{b,j}$ for $j \in \mathcal{E}$
11: For each $i \in \mathcal{U}$, set $g_{h,i} \leftarrow v_{h,i}$ if $i \in \hat{\mathcal{A}}$, and $g_{h,i} \leftarrow t_i/d_i$ otherwise
12: Assemble $g_\eta \leftarrow [g_c; g_{b,\mathcal{E}}; g_{h,\mathcal{U}}]$
13: Compute $g_\theta \leftarrow (D_\theta\hat{\eta})^\top g_\eta$; update θ using g_θ ▷ standard backprop through ψ_θ

Equivalent explicit-gradient form. The explicit fixed-active-set expressions for the gradient of L_{signal} derived in Section 6 provide an equivalent alternative to the adjoint construction in Algorithm 1. In particular, one may evaluate $\nabla_{\hat{\eta}} L_{\text{signal}}$ directly from those expressions and then apply standard backpropagation through the predictor, $\nabla_\theta L_{\text{signal}} = (D_\theta\hat{\eta})^\top \nabla_{\hat{\eta}} L_{\text{signal}}$. Both forms are understood within the same fixed-active-set regularity regime; weakly active boundaries are treated by the directional construction described in Appendix D.

B STANDING ASSUMPTIONS

We collect here the assumptions used by the regular-region results. For either the true or represented QP under consideration, let $\mathcal{A}$ and $\mathcal{I}$ denote its active and inactive inequality sets, respectively, and let $J := [A; G_\mathcal{A}]$ denote the corresponding equality-plus-active-inequality constraint matrix. Thus, for the true problem, we have $(\mathcal{A}, \mathcal{I}, J) = (\mathcal{A}^\star, \mathcal{I}^\star, J^\star)$, whereas for the represented problem we have $(\mathcal{A}, \mathcal{I}, J) = (\hat{\mathcal{A}}, \hat{\mathcal{I}}, \hat{J})$.

Regular-region assumptions. Unless a result explicitly concerns a weakly active boundary, the local analysis is stated at a feasible QP instance satisfying: (A1) $H \succ 0$; (A2) the considered QP is feasible; (A3) the linear independence constraint qualification (LICQ) holds, i.e., J has full row rank; (A4) strict complementarity holds on the active inequalities, $\mu_i > 0$ for all $i \in \mathcal{A}$; and (A5) inactive inequalities have strictly positive slack, $s_i := h_i - G_i z > 0$ for all $i \in \mathcal{I}$. These are local regularity conditions for the smooth critical-region analysis rather than guarantees for every possible predictor output. Together, they ensure the nondegenerate regular regime used in the local analysis, with a locally stable active set and differentiable primal-dual solution map. Appendix D relaxes strict complementarity on the weakly active subset while retaining LICQ for all geometrically active constraints.

C REGIONAL AFFINE ERRORS

Fix a regular represented critical region with active set $\hat{\mathcal{A}}$ and inactive set $\hat{\mathcal{I}}$. Define $\hat{J} := [A; G_{\hat{\mathcal{A}}}]$, $\hat{S} := \hat{J}H^{-1}\hat{J}^\top$, $\hat{P} := H^{-1} - H^{-1}\hat{J}^\top\hat{S}^{-1}\hat{J}H^{-1}$, and $\hat{K} := H^{-1}\hat{J}^\top\hat{S}^{-1}$. By regularity, $\hat{J}$ has full row rank, so $\hat{S}$ is nonsingular. Within the region, $\hat{\mathcal{A}}$ and hence $\hat{J}, \hat{P}, \hat{K}$ are fixed.

Decision error. Let $\delta y := [\hat{\lambda} - \lambda^\star; \hat{\mu}_{\hat{\mathcal{A}}} - \mu^\star_{\hat{\mathcal{A}}}]$. Subtracting the true KKT stationarity condition from the represented active-set KKT system gives

$$\begin{bmatrix} H & \hat{J}^\top \\ \hat{J} & 0 \end{bmatrix} \begin{bmatrix} \delta z \\ \delta y \end{bmatrix} = \begin{bmatrix} -\delta c + G_{\hat{\mathcal{I}}}^\top \mu^\star_{\hat{\mathcal{I}}} \\ \delta b \\ \delta h_{\hat{\mathcal{A}}} + s^\star_{\hat{\mathcal{A}}} \end{bmatrix}, \tag{36}$$

where $s^\star := h - Gz^\star$ is the true slack vector. Indeed, $A\delta z = \delta b$, while for $i \in \hat{\mathcal{A}}$, $G_i\delta z = \hat{h}_i - G_i z^\star = \delta h_i + s_i^\star$. The stationarity offset $G_{\hat{\mathcal{I}}}^\top \mu^\star_{\hat{\mathcal{I}}}$ arises because true multipliers associated with rows omitted from the represented active set remain present in the true KKT system.

Solving (36) for the primal component yields

$$\delta z = -\hat{P}\delta c + \hat{K}\begin{bmatrix} \delta b \\ \delta h_{\hat{\mathcal{A}}} \end{bmatrix} + \hat{P}G_{\hat{\mathcal{I}}}^\top \mu^\star_{\hat{\mathcal{I}}} + \hat{K}\begin{bmatrix} 0 \\ s^\star_{\hat{\mathcal{A}}} \end{bmatrix}. \tag{37}$$

Thus, defining the regional linear map Z and offset z_0 by

$$Z\delta\eta := -\hat{P}\,\delta c + \hat{K}\begin{bmatrix} \delta b \\ \delta h_{\hat{\mathcal{A}}} \end{bmatrix}, \qquad z_0 := \hat{P}G_{\hat{\mathcal{I}}}^\top \mu^\star_{\hat{\mathcal{I}}} + \hat{K}\begin{bmatrix} 0 \\ s^\star_{\hat{\mathcal{A}}} \end{bmatrix}, \tag{38}$$

we obtain

$$\delta z = Z\delta\eta + z_0. \tag{39}$$

The two components of z_0 correspond to the two forms of active-set mismatch: true multiplier forces from rows that are inactive in the represented problem, and true slacks of rows that are active in the represented problem.

Multiplier error. The dual component of (36) satisfies

$$\delta y = -\hat{S}^{-1}\left(\hat{J}H^{-1}\delta c + \begin{bmatrix} \delta b \\ \delta h_{\hat{\mathcal{A}}} \end{bmatrix}\right) + \hat{S}^{-1}\left(\hat{J}H^{-1}G_{\hat{\mathcal{I}}}^\top \mu^\star_{\hat{\mathcal{I}}} - \begin{bmatrix} 0 \\ s^\star_{\hat{\mathcal{A}}} \end{bmatrix}\right). \tag{40}$$

Accordingly, write the full inequality-multiplier error as

$$\delta\mu := \hat{\mu} - \mu^\star = U\delta\eta + u_0. \tag{41}$$

On $\hat{\mathcal{A}}$, $U\delta\eta$ and u_0 are the inequality-multiplier components of the first and second terms in (40), respectively. On $\hat{\mathcal{I}}$, $(U\delta\eta)_{\hat{\mathcal{I}}} = 0$ and $(u_0)_{\hat{\mathcal{I}}} = -\mu^\star_{\hat{\mathcal{I}}}$, because $\hat{\mu}_{\hat{\mathcal{I}}} = 0$ throughout the represented region.

Margin error. Let $D_{\mathcal{U}} := \mathrm{diag}(d_i)_{i\in\mathcal{U}}$. From the signed-margin definition on the supervised rows,

$$q_{\mathcal{U}} = D_{\mathcal{U}}^{-1}(h_{\mathcal{U}} - G_{\mathcal{U}}z) - D_{\mathcal{U}}\mu_{\mathcal{U}}. \tag{42}$$

Hence,

$$\delta q_{\mathcal{U}} := \hat{q}_{\mathcal{U}} - q^\star_{\mathcal{U}} = D_{\mathcal{U}}^{-1}\delta h_{\mathcal{U}} - D_{\mathcal{U}}^{-1}G_{\mathcal{U}}\delta z - D_{\mathcal{U}}\delta\mu_{\mathcal{U}}. \tag{43}$$

Substituting (39) and (41) gives

$$\delta q_{\mathcal{U}} = Q\delta\eta + q_0, \tag{44}$$

where the regional linear map Q is defined by

$$Q\delta\eta := D_{\mathcal{U}}^{-1}\delta h_{\mathcal{U}} - D_{\mathcal{U}}^{-1}G_{\mathcal{U}}(Z\delta\eta) - D_{\mathcal{U}}(U\delta\eta)_{\mathcal{U}}, \qquad q_0 := -D_{\mathcal{U}}^{-1}G_{\mathcal{U}}z_0 - D_{\mathcal{U}}(u_0)_{\mathcal{U}}. \tag{45}$$

Therefore, throughout a fixed regular represented active-set region, the decision error and supervised primal-dual margin error are affine in the prediction error:

$$\delta z = Z\delta\eta + z_0, \qquad \delta q_{\mathcal{U}} = Q\delta\eta + q_0. \tag{46}$$

When $\hat{\mathcal{A}} = \mathcal{A}^\star$, complementary slackness gives $\mu^\star_{\hat{\mathcal{I}}} = 0$ and $s^\star_{\hat{\mathcal{A}}} = 0$. Hence, $z_0 = 0$, $u_0 = 0$, $q_0 = 0$. Thus, in a shared regular critical region, both error maps are linear in $\delta\eta$; true-represented active-set mismatch generates the affine offsets in the general case.

D WEAK ACTIVITY

Let $\mathcal{A}_+^\star := \{i \in \mathcal{A}^\star : \mu_i^\star > 0\}$ and $\mathcal{A}_0^\star := \{i \in \mathcal{A}^\star : \mu_i^\star = 0\}$, so that $\mathcal{A}^\star = \mathcal{A}_+^\star \dot\cup \mathcal{A}_0^\star$. Thus, $\mathcal{A}_0^\star$ contains the weakly active inequalities, while strictly inactive inequalities remain in $\mathcal{I}^\star$. Theorem 3 provides the directional derivative at weak activity.

Theorem 3 (Directional derivative at weak activity)**.** *Assume $H \succ 0$, fixed H, A, G, and LICQ for all geometrically active constraints, i.e., $\left[\begin{smallmatrix} A \\ G_{\mathcal{A}^\star} \end{smallmatrix}\right]$ has full row rank. For a directional perturbation $\Delta\eta := (\Delta c, \Delta b, \Delta h)$, the one-sided directional derivative*

$$d := z^{\star\prime}(\eta; \Delta\eta) = \lim_{t \downarrow 0} \frac{z^\star(\eta + t\Delta\eta) - z^\star(\eta)}{t} \tag{47}$$

exists and is the unique solution of

$$\begin{aligned} \underset{d}{\text{minimise}} \quad & \frac{1}{2} d^\top H d + \Delta c^\top d \\ \text{subject to} \quad & Ad = \Delta b, \\ & G_i d = \Delta h_i, \qquad i \in \mathcal{A}_+^\star, \\ & G_i d \le \Delta h_i, \qquad i \in \mathcal{A}_0^\star. \end{aligned} \tag{48}$$

Under the uncertain-RHS convention, $\Delta b_j = 0$ for $j \notin \mathcal{E}$ and $\Delta h_i = 0$ for $i \notin \mathcal{U}$; fixed active constraints nevertheless remain in the directional system above.

Proof. Partition the true active inequalities into $\mathcal{A}_+^\star$ and $\mathcal{A}_0^\star$. For the one-sided expansion

$$z^\star(\eta + t\Delta\eta) = z^\star + td + o(t), \qquad t \downarrow 0, \tag{49}$$

linearised feasibility gives $Ad = \Delta b$ and $G_i d \le \Delta h_i$ for $i \in \mathcal{A}^\star$. For $i \in \mathcal{A}_+^\star$, positivity of the base multiplier together with complementarity forces $G_i d = \Delta h_i$. Strictly inactive inequalities have positive slack at the base point and therefore impose no first-order restriction.

The directional KKT stationarity condition is

$$Hd + \Delta c + A^\top \dot\lambda + G_{\mathcal{A}_+^\star}^\top \dot\mu_+ + G_{\mathcal{A}_0^\star}^\top \nu = 0, \tag{50}$$

where $\nu \ge 0$ are the first-order multipliers associated with the weakly active inequalities. Their directional complementarity conditions are

$$\nu_i \big(G_i d - \Delta h_i\big) = 0, \qquad i \in \mathcal{A}_0^\star. \tag{51}$$

Together with $Ad = \Delta b$, $G_i d = \Delta h_i$ for $i \in \mathcal{A}_+^\star$, and $G_i d \le \Delta h_i$ for $i \in \mathcal{A}_0^\star$, these are exactly the KKT conditions of the stated auxiliary QP.

Since $H \succ 0$, the auxiliary QP is strongly convex and therefore has a unique primal solution. Under LICQ and the stated regularity conditions, the standard directional-sensitivity result applies, so this unique solution is the one-sided directional derivative $z^{\star\prime}(\eta; \Delta\eta)$. □

E PROOFS OF THE MAIN RESULTS

E.1 PROOF OF PROPOSITION 1

Let $\delta z := \hat z - z^\star$. Since $f(z) = \frac{1}{2} z^\top H z + c^\top z$, the exact quadratic expansion around $z^\star$ gives

$$R(\hat z) = f(\hat z) - f(z^\star) = (Hz^\star + c)^\top \delta z + \frac{1}{2} \delta z^\top H \delta z. \tag{52}$$

Using the true KKT stationarity condition $Hz^\star + c + A^\top \lambda^\star + G_{\mathcal{A}^\star}^\top \mu_{\mathcal{A}^\star}^\star = 0$, we obtain

$$R(\hat z) = \frac{1}{2} \delta z^\top H \delta z - \lambda^{\star\top} A \delta z - \mu_{\mathcal{A}^\star}^{\star\top} G_{\mathcal{A}^\star} \delta z. \tag{53}$$

Since $Az^\star = b$ and $G_{\mathcal{A}^\star} z^\star = h_{\mathcal{A}^\star}$, we can write

$$A\delta z = A\hat z - b, \qquad G_{\mathcal{A}^\star} \delta z = G_{\mathcal{A}^\star} \hat z - h_{\mathcal{A}^\star}. \tag{54}$$

Moreover, by definition of the quadratic Bregman divergence, $L_{\text{dec}} = \frac{1}{2}\delta z^\top H \delta z$. Therefore,

$$L_{\text{dec}} = R(\hat{z}) + \lambda^{\star\top}(A\hat{z} - b) + \mu_{\mathcal{A}^\star}^{\star\top}(G_{\mathcal{A}^\star}\hat{z} - h_{\mathcal{A}^\star}), \tag{55}$$

which proves (10). No feasibility assumption on $\hat{z}$ is required. If the true equality residuals and true active-inequality residuals vanish, the dual correction is zero and $L_{\text{dec}} = R(\hat{z})$. This completes the proof of Proposition 1.

E.2 Proof of Proposition 2

Let $\delta z := \hat{z} - z^\star$. Since $z^\star$ is feasible for the true problem, $Az^\star = b$ and $Gz^\star \le h$. Hence, $A\hat{z} - b = A\delta z$. Moreover, for each $i \in \mathcal{U}$,

$$G_i\hat{z} - h_i = G_i\delta z + (G_i z^\star - h_i) \le G_i\delta z, \tag{56}$$

and therefore

$$[G_i\hat{z} - h_i]_+ \le [G_i\delta z]_+ \le |G_i\delta z|. \tag{57}$$

Using the definitions of W_b and W_h,

$$2L_{\text{feas}} = (A\delta z)^\top W_b(A\delta z) + \sum_{i\in\mathcal{U}} \frac{[G_i\hat{z} - h_i]_+^2}{d_i^2} \tag{58}$$

$$\le \delta z^\top \big(A^\top W_b A + G^\top W_h G\big)\delta z. \tag{59}$$

Define $M_{\text{feas}} := A^\top W_b A + G^\top W_h G$. Since $H \succ 0$,

$$\delta z^\top M_{\text{feas}}\delta z = (H^{1/2}\delta z)^\top \big(H^{-1/2}M_{\text{feas}}H^{-1/2}\big)(H^{1/2}\delta z) \tag{60}$$

$$\le \lambda_{\max}\big(H^{-1/2}M_{\text{feas}}H^{-1/2}\big)\delta z^\top H\delta z. \tag{61}$$

Therefore,

$$2L_{\text{feas}} \le \lambda_{\max}\big(H^{-1/2}(A^\top W_b A + G^\top W_h G)H^{-1/2}\big)\delta z^\top H\delta z. \tag{62}$$

Using $L_{\text{dec}} = \frac{1}{2}\delta z^\top H\delta z$ gives

$$L_{\text{feas}} \le \lambda_{\max}\Big(H^{-1/2}(A^\top W_b A + G^\top W_h G)H^{-1/2}\Big)L_{\text{dec}}, \tag{63}$$

which proves (12). The argument uses only feasibility of $z^\star$ for the true problem and $H \succ 0$; it requires neither a common active set nor local differentiability of the solution map. This completes the proof of Proposition 2.

E.3 Proof of Proposition 3

Consider any $i \in \mathcal{U} \cap (\mathcal{A}^\star \triangle \hat{\mathcal{A}})$. There are two possible forms of active-set disagreement. If $i \in \mathcal{I}^\star \cap \hat{\mathcal{A}}$, then the true row is inactive, so $q_i^\star = \gamma_i > 0$, while represented activity implies $\hat{q}_i \le 0$. Hence,

$$|\hat{q}_i - q_i^\star| \ge q_i^\star = \gamma_i. \tag{64}$$

If instead $i \in \mathcal{A}^\star \cap \hat{\mathcal{I}}$, then $q_i^\star = -\gamma_i \le 0$, while represented inactivity implies $\hat{q}_i \ge 0$, and therefore

$$|\hat{q}_i - q_i^\star| \ge |q_i^\star| = \gamma_i. \tag{65}$$

The weakly active case $q_i^\star = 0$ gives $\gamma_i = 0$, for which the same bound holds trivially. Thus, every supervised active-set disagreement satisfies

$$(\hat{q}_i - q_i^\star)^2 \ge \gamma_i^2. \tag{66}$$

Summing over the disagreement set and then enlarging the sum to all supervised rows gives

$$\sum_{i\in\mathcal{U}\cap(\mathcal{A}^\star\triangle\hat{\mathcal{A}})} \gamma_i^2 \le \sum_{i\in\mathcal{U}\cap(\mathcal{A}^\star\triangle\hat{\mathcal{A}})} (\hat{q}_i - q_i^\star)^2 \le \sum_{i\in\mathcal{U}} (\hat{q}_i - q_i^\star)^2 = 2kL_{\text{reg}}, \tag{67}$$

which proves (15). Now suppose $\gamma_{\min,\mathcal{U}} > 0$. If $\hat{\mathcal{A}} \cap \mathcal{U} \ne \mathcal{A}^\star \cap \mathcal{U}$, then the disagreement set contains at least one supervised row, and therefore $2kL_{\text{reg}} \ge \gamma_{\min,\mathcal{U}}^2$. Taking the contrapositive yields

$$L_{\text{reg}} < \frac{\gamma_{\min,\mathcal{U}}^2}{2k} \;\Rightarrow\; \hat{\mathcal{A}} \cap \mathcal{U} = \mathcal{A}^\star \cap \mathcal{U}, \tag{68}$$

which proves (16). This completes the proof of Proposition 3.

E.4 Proof of Theorem 1

Step 1. For test instance n, define $r_{b,n} := A\hat{z}_n - b_n$, $r_{h,n} := G\hat{z}_n - h_n$, and $v_{n,i} := [r_{h,n,i}]_+$. Applying Proposition 1 to instance n gives

$$f_n(\hat{z}_n) - f_n(z_n^\star) = L_{\mathrm{dec}}^{(n)} - \lambda_n^{\star\top} r_{b,n} - \mu_{n,\mathcal{A}_n^\star}^{\star\top} r_{h,n,\mathcal{A}_n^\star}. \tag{69}$$

Since every represented decision satisfies all fixed constraints, the equality residuals vanish on $j \notin \mathcal{E}$, while fixed inequality rows satisfy $r_{h,n,i} \le 0$. For the equality terms,

$$C_{b,j}|r_{b,n,j}| - \lambda_{n,j}^\star r_{b,n,j} \le (C_{b,j} + |\lambda_{n,j}^\star|)|r_{b,n,j}|. \tag{70}$$

For the inequality terms, decompose each residual as $r_{h,n,i} = [r_{h,n,i}]_+ - [-r_{h,n,i}]_+$. We distinguish supervised active, supervised inactive, and fixed active rows. For $i \in \mathcal{U} \cap \mathcal{A}_n^\star$, we have

$$\begin{aligned} C_{h,i}[r_{h,n,i}]_+ - \mu_{n,i}^\star r_{h,n,i} &= (C_{h,i} - \mu_{n,i}^\star)[r_{h,n,i}]_+ + \mu_{n,i}^\star[-r_{h,n,i}]_+ \\ &\le [C_{h,i} - \mu_{n,i}^\star]_+ \, v_{n,i} + \mu_{n,i}^\star[-r_{h,n,i}]_+. \end{aligned} \tag{71}$$

If $i \in \mathcal{U} \cap \mathcal{I}_n^\star$, then $\mu_{n,i}^\star = 0$, so the corresponding contribution is $C_{h,i}v_{n,i}$. For fixed active rows $i \in \mathcal{A}_n^\star \setminus \mathcal{U}$, we have $r_{h,n,i} \le 0$, and hence their contribution is $\mu_{n,i}^\star[-r_{h,n,i}]_+$. Therefore, the nth summand in the definition of R_{op} satisfies

$$\begin{aligned} &f_n(\hat{z}_n) - f_n(z_n^\star) + \sum_{j\in\mathcal{E}} C_{b,j}|r_{b,n,j}| + \sum_{i\in\mathcal{U}} C_{h,i}v_{n,i} \\ &\le L_{\mathrm{dec}}^{(n)} + \sum_{j\in\mathcal{E}} (C_{b,j} + |\lambda_{n,j}^\star|)|r_{b,n,j}| + \sum_{i\in\mathcal{U}} [C_{h,i} - \mu_{n,i}^\star]_+ \, v_{n,i} + \sum_{i\in\mathcal{A}_n^\star} \mu_{n,i}^\star[-r_{h,n,i}]_+. \end{aligned} \tag{72}$$

By weighted Cauchy-Schwarz and the definitions of $\Xi_n^\star$ and $L_{\mathrm{feas}}^{(n)}$, we can write

$$\begin{aligned} &\sum_{j\in\mathcal{E}} (C_{b,j} + |\lambda_{n,j}^\star|)|r_{b,n,j}| + \sum_{i\in\mathcal{U}} [C_{h,i} - \mu_{n,i}^\star]_+ \, v_{n,i} \\ &\le \Xi_n^\star \Bigg(\sum_{j\in\mathcal{E}} \Big(\frac{r_{b,n,j}}{e_j}\Big)^2 + \sum_{i\in\mathcal{U}} \Big(\frac{v_{n,i}}{d_i}\Big)^2 \Bigg)^{1/2} = \Xi_n^\star \sqrt{2L_{\mathrm{feas}}^{(n)}}. \end{aligned} \tag{73}$$

Moreover, for $i \in \mathcal{A}_n^\star$, true activity implies $r_{h,n,i} = G_i(\hat{z}_n - z_n^\star) = G_i \delta z_n$. Hence, by the definitions of $\Gamma_n^\star$ and L_{dec}, we get

$$\sum_{i\in\mathcal{A}_n^\star} \mu_{n,i}^\star[-r_{h,n,i}]_+ \le \Gamma_n^\star \|\delta z_n\|_H = \Gamma_n^\star \sqrt{2L_{\mathrm{dec}}^{(n)}}. \tag{74}$$

Combining (72)-(74) yields

$$\begin{aligned} &f_n(\hat{z}_n) - f_n(z_n^\star) + \sum_{j\in\mathcal{E}} C_{b,j}|r_{b,n,j}| + \sum_{i\in\mathcal{U}} C_{h,i}v_{n,i} \\ &\le L_{\mathrm{dec}}^{(n)} + \sqrt{2}\Gamma_n^\star \sqrt{L_{\mathrm{dec}}^{(n)}} + \sqrt{2}\Xi_n^\star \sqrt{L_{\mathrm{feas}}^{(n)}}. \end{aligned} \tag{75}$$

Since $\lambda_{\mathrm{feas}} > 0$, we deduce that

$$\begin{aligned} \Gamma_n^\star \sqrt{L_{\mathrm{dec}}^{(n)}} + \Xi_n^\star \sqrt{L_{\mathrm{feas}}^{(n)}} &= \Gamma_n^\star \sqrt{L_{\mathrm{dec}}^{(n)}} + \frac{\Xi_n^\star}{\sqrt{\lambda_{\mathrm{feas}}}} \sqrt{\lambda_{\mathrm{feas}} L_{\mathrm{feas}}^{(n)}} \\ &\le \Big({\Gamma_n^\star}^2 + \frac{{\Xi_n^\star}^2}{\lambda_{\mathrm{feas}}} \Big)^{1/2} \Big(L_{\mathrm{dec}}^{(n)} + \lambda_{\mathrm{feas}} L_{\mathrm{feas}}^{(n)} \Big)^{1/2}. \end{aligned} \tag{76}$$

Because $L_{\mathrm{dec}}^{(n)} \le L_{\mathrm{signal}}^{(n)}$ and $L_{\mathrm{dec}}^{(n)} + \lambda_{\mathrm{feas}} L_{\mathrm{feas}}^{(n)} \le L_{\mathrm{signal}}^{(n)}$, the nth summand in R_{op} is bounded by

$$\begin{aligned} &f_n(\hat{z}_n) - f_n(z_n^\star) + \sum_{j\in\mathcal{E}} C_{b,j}|r_{b,n,j}| + \sum_{i\in\mathcal{U}} C_{h,i}v_{n,i} \\ &\le L_{\mathrm{signal}}^{(n)} + \sqrt{2\Big({\Gamma_n^\star}^2 + \frac{{\Xi_n^\star}^2}{\lambda_{\mathrm{feas}}}\Big) L_{\mathrm{signal}}^{(n)}}. \end{aligned} \tag{77}$$

Averaging over $n = 1, \ldots, N$ gives

$$R_{\text{op}} \leq \bar{L}_{\text{signal}} + \frac{\sqrt{2}}{N} \sum_{n=1}^{N} \sqrt{\left(\Gamma_n^{\star\,2} + \frac{\Xi_n^{\star\,2}}{\lambda_{\text{feas}}}\right) L_{\text{signal}}^{(n)}}. \tag{78}$$

Applying Cauchy-Schwarz across the N instances,

$$\begin{aligned}
&\frac{1}{N} \sum_{n=1}^{N} \left(\left(\Gamma_n^{\star\,2} + \frac{\Xi_n^{\star\,2}}{\lambda_{\text{feas}}}\right) L_{\text{signal}}^{(n)} \right)^{1/2} \\
&\leq \left(\frac{1}{N} \sum_{n=1}^{N} \left(\Gamma_n^{\star\,2} + \frac{\Xi_n^{\star\,2}}{\lambda_{\text{feas}}} \right) \right)^{1/2} \left[\frac{1}{N} \sum_{n=1}^{N} L_{\text{signal}}^{(n)} \right]^{1/2} \\
&= \left(\frac{1}{N} \sum_{n=1}^{N} \left(\Gamma_n^{\star\,2} + \frac{\Xi_n^{\star\,2}}{\lambda_{\text{feas}}} \right) \right)^{1/2} \bar{L}_{\text{signal}}^{1/2}.
\end{aligned} \tag{79}$$

This proves the bound in (22).

Step 2. For each test instance n, define $I_n := \mathbb{1}\{\max_{i \in \mathcal{U}} v_{n,i} > \epsilon\}$. We first show that

$$I_n \leq \frac{d_{\max}^2}{\epsilon^2} \sum_{i \in \mathcal{U}} \frac{v_{n,i}^2}{d_i^2}. \tag{80}$$

If $I_n = 0$, then the inequality holds trivially because the right-hand side is nonnegative. If $I_n = 1$, then there exists some $i^\star \in \mathcal{U}$ such that $v_{n,i^\star} > \epsilon$. Since $d_{i^\star} \leq d_{\max}$, we get

$$\frac{v_{n,i^\star}^2}{d_{i^\star}^2} > \frac{\epsilon^2}{d_{\max}^2}. \tag{81}$$

Moreover, because every term in the sum is nonnegative, we obtain

$$\sum_{i \in \mathcal{U}} \frac{v_{n,i}^2}{d_i^2} \geq \frac{v_{n,i^\star}^2}{d_{i^\star}^2} > \frac{\epsilon^2}{d_{\max}^2}. \tag{82}$$

Hence,

$$\frac{d_{\max}^2}{\epsilon^2} \sum_{i \in \mathcal{U}} \frac{v_{n,i}^2}{d_i^2} > 1 = I_n, \tag{83}$$

which proves (80). Next, by the definition of $L_{\text{feas}}^{(n)}$, we find that

$$\frac{1}{2} \sum_{i \in \mathcal{U}} \frac{v_{n,i}^2}{d_i^2} \leq L_{\text{feas}}^{(n)}. \tag{84}$$

Since $L_{\text{signal}}^{(n)} = L_{\text{dec}}^{(n)} + \lambda_{\text{feas}} L_{\text{feas}}^{(n)} + \lambda_{\text{reg}} L_{\text{reg}}^{(n)}$, with all three losses nonnegative and $\lambda_{\text{feas}} > 0$, we know that $\lambda_{\text{feas}} L_{\text{feas}}^{(n)} \leq L_{\text{signal}}^{(n)}$. Therefore,

$$\sum_{i \in \mathcal{U}} \frac{v_{n,i}^2}{d_i^2} \leq \frac{2 L_{\text{signal}}^{(n)}}{\lambda_{\text{feas}}}. \tag{85}$$

Substituting this into (80) gives

$$\mathbb{1}\Big\{ \max_{i \in \mathcal{U}} v_{n,i} > \epsilon \Big\} \leq \frac{2 d_{\max}^2}{\lambda_{\text{feas}} \epsilon^2} L_{\text{signal}}^{(n)}. \tag{86}$$

Finally, averaging over $n = 1, \ldots, N$ yields

$$\begin{aligned}
V_{\text{freq}} &= \frac{1}{N} \sum_{n=1}^{N} \mathbb{1}\Big\{ \max_{i \in \mathcal{U}} v_{n,i} > \epsilon \Big\} \\
&\leq \frac{2 d_{\max}^2}{\lambda_{\text{feas}} \epsilon^2} \frac{1}{N} \sum_{n=1}^{N} L_{\text{signal}}^{(n)} = \frac{2 d_{\max}^2}{\lambda_{\text{feas}} \epsilon^2} \bar{L}_{\text{signal}},
\end{aligned} \tag{87}$$

This proves the first bound in (23).

Step 3. For each test instance n, we can write

$$\frac{1}{k}\sum_{i\in\mathcal{U}} v_{n,i} = \frac{1}{k}\sum_{i\in\mathcal{U}} d_i \frac{v_{n,i}}{d_i} \le \frac{1}{k}\Big(\sum_{i\in\mathcal{U}} d_i^2\Big)^{1/2}\Big(\sum_{i\in\mathcal{U}} \frac{v_{n,i}^2}{d_i^2}\Big)^{1/2} = \frac{d_{\mathcal{U}}}{k}\Big(\sum_{i\in\mathcal{U}} \frac{v_{n,i}^2}{d_i^2}\Big)^{1/2}. \tag{88}$$

By the definition of $L_{\mathrm{feas}}^{(n)}$, we have

$$\sum_{i\in\mathcal{U}} \frac{v_{n,i}^2}{d_i^2} \le 2L_{\mathrm{feas}}^{(n)}, \tag{89}$$

and therefore

$$\frac{1}{k}\sum_{i\in\mathcal{U}} v_{n,i} \le \frac{d_{\mathcal{U}}}{k}\sqrt{2L_{\mathrm{feas}}^{(n)}}. \tag{90}$$

Averaging over $n = 1, \dots, N$ gives

$$V_{\mathrm{mag}} \le \frac{d_{\mathcal{U}}}{k}\frac{1}{N}\sum_{n=1}^{N}\sqrt{2L_{\mathrm{feas}}^{(n)}}. \tag{91}$$

Applying Cauchy-Schwarz across the N instances, we get

$$\frac{1}{N}\sum_{n=1}^{N}\sqrt{2L_{\mathrm{feas}}^{(n)}} \le \left(\frac{2}{N}\sum_{n=1}^{N} L_{\mathrm{feas}}^{(n)}\right)^{1/2}. \tag{92}$$

Finally, since $\lambda_{\mathrm{feas}} L_{\mathrm{feas}}^{(n)} \le L_{\mathrm{signal}}^{(n)}$, we have

$$\frac{1}{N}\sum_{n=1}^{N} L_{\mathrm{feas}}^{(n)} \le \frac{\bar{L}_{\mathrm{signal}}}{\lambda_{\mathrm{feas}}}. \tag{93}$$

Hence,

$$V_{\mathrm{mag}} \le \frac{d_{\mathcal{U}}}{k}\left(\frac{2\bar{L}_{\mathrm{signal}}}{\lambda_{\mathrm{feas}}}\right)^{1/2}, \tag{94}$$

which proves the second bound in (23).

Step 4. Finally, consider the auxiliary dormant-bound diagnostic E_{dorm}. By its definition, the supervised bound ℓ is inactive in both the true and represented problems for every test instance, so $\mu_{\ell,n}^{\star} = \hat{\mu}_{\ell,n} = 0$. By the definition of the signed margin, $\hat{q}_{\ell,n} - q_{\ell,n}^{\star} = (\delta h_{\ell}^{(n)} - G_{\ell}\delta z_n)/d_{\ell}$, and therefore

$$\frac{|\delta h_{\ell}^{(n)}|}{d_{\ell}} \le |\hat{q}_{\ell,n} - q_{\ell,n}^{\star}| + \frac{|G_{\ell}\delta z_n|}{d_{\ell}}. \tag{95}$$

By the definition of $L_{\mathrm{reg}}^{(n)}$, $|\hat{q}_{\ell,n} - q_{\ell,n}^{\star}| \le \sqrt{2kL_{\mathrm{reg}}^{(n)}}$. Moreover, since $d_{\ell}^2 = G_{\ell}H^{-1}G_{\ell}^{\top}$, Cauchy-Schwarz in the H metric gives

$$\frac{|G_{\ell}\delta z_n|}{d_{\ell}} \le \|\delta z_n\|_H. \tag{96}$$

By the definition of L_{dec}, we have $\|\delta z\|_H = \sqrt{2L_{\mathrm{dec}}}$. Thus, we can write

$$\frac{|\delta h_{\ell}^{(n)}|}{d_{\ell}} \le \sqrt{2kL_{\mathrm{reg}}^{(n)}} + \sqrt{2L_{\mathrm{dec}}^{(n)}}. \tag{97}$$

Since $\lambda_{\mathrm{reg}} > 0$, we deduce that

$$\begin{aligned}\sqrt{kL_{\mathrm{reg}}^{(n)}} + \sqrt{L_{\mathrm{dec}}^{(n)}} &= \sqrt{\frac{k}{\lambda_{\mathrm{reg}}}}\sqrt{\lambda_{\mathrm{reg}}L_{\mathrm{reg}}^{(n)}} + \sqrt{L_{\mathrm{dec}}^{(n)}} \\ &\le \sqrt{1+\frac{k}{\lambda_{\mathrm{reg}}}}\sqrt{L_{\mathrm{dec}}^{(n)} + \lambda_{\mathrm{reg}}L_{\mathrm{reg}}^{(n)}} \\ &\le \sqrt{1+\frac{k}{\lambda_{\mathrm{reg}}}}\sqrt{L_{\mathrm{signal}}^{(n)}}.\end{aligned} \tag{98}$$

This implies that

$$\frac{|\delta h_\ell^{(n)}|}{d_\ell} \le \left(\left(2 + \frac{2k}{\lambda_{\text{reg}}}\right) L_{\text{signal}}^{(n)} \right)^{1/2}. \tag{99}$$

Averaging over n and applying Cauchy-Schwarz gives

$$\begin{aligned} E_{\text{dorm}} &\le \left(\left(2 + \frac{2k}{\lambda_{\text{reg}}}\right) \right)^{1/2} \frac{1}{N} \sum_{n=1}^{N} \sqrt{L_{\text{signal}}^{(n)}} \\ &\le \left(\left(2 + \frac{2k}{\lambda_{\text{reg}}}\right) \right)^{1/2} \left(\frac{1}{N} \sum_{n=1}^{N} L_{\text{signal}}^{(n)} \right)^{1/2} \\ &= \left(\left(2 + \frac{2k}{\lambda_{\text{reg}}}\right) \bar{L}_{\text{signal}} \right)^{1/2}, \end{aligned} \tag{100}$$

which proves the final bound in (23) and completes the proof of Theorem 1.

E.5 Proof of Theorem 2

Step 1. Because the true and represented parameter instances lie in the same regular critical region, they share the active set $\hat{\mathcal{A}} = \mathcal{A}^\star$, and hence the same reduced constraint matrix J. Their reduced KKT systems are

$$Hz^\star + c + J^\top y^\star = 0, \qquad Jz^\star - r = 0, \tag{101}$$

$$H\hat{z} + \hat{c} + J^\top \hat{y} = 0, \qquad J\hat{z} - \hat{r} = 0. \tag{102}$$

Subtracting the true system from the represented system gives

$$H\delta z + \delta c + J^\top \delta y = 0, \qquad J\delta z - \delta r = 0, \tag{103}$$

where $\delta y := \hat{y} - y^\star$. Equivalently, we can write

$$\begin{bmatrix} H & J^\top \\ J & 0 \end{bmatrix} \begin{bmatrix} \delta z \\ \delta y \end{bmatrix} = \begin{bmatrix} -\delta c \\ \delta r \end{bmatrix}. \tag{104}$$

From the first row, $\delta z = -H^{-1}\delta c - H^{-1}J^\top \delta y$. Substituting into the second row yields $-JH^{-1}\delta c - JH^{-1}J^\top \delta y = \delta r$. Using $S = JH^{-1}J^\top$, we obtain $\delta y = -S^{-1}(\delta r + JH^{-1}\delta c)$. Therefore,

$$\delta z = -H^{-1}\delta c + H^{-1}J^\top S^{-1}\big(\delta r + JH^{-1}\delta c\big) \tag{105}$$

$$= -\Big(H^{-1} - H^{-1}J^\top S^{-1} JH^{-1}\Big)\delta c + H^{-1}J^\top S^{-1}\delta r \tag{106}$$

$$= -P\delta c + K\delta r. \tag{107}$$

By definition $\delta z_{\tan} := -P\delta c$ and $\delta z_{\text{nor}} := K\delta r$. Hence, $\delta z = \delta z_{\tan} + \delta z_{\text{nor}}$.

To establish the geometric property, first note that

$$JP = JH^{-1} - JH^{-1}J^\top S^{-1} JH^{-1} = 0. \tag{108}$$

Hence, $\delta z_{\tan} \in \ker J$. Moreover, $\text{range}(K) = \text{range}(H^{-1}J^\top)$, because S^{-1} is nonsingular. For any $v \in \ker J$ and any $w \in \text{range}(H^{-1}J^\top)$, there exists $\alpha \in \mathbb{R}^{p+|\mathcal{A}^\star|}$ such that $w = H^{-1}J^\top \alpha$. Hence,

$$v^\top H w = v^\top J^\top \alpha = (Jv)^\top \alpha = 0. \tag{109}$$

Since $\ker J$ and $\text{range}(H^{-1}J^\top)$ have complementary dimensions, $\text{range}(H^{-1}J^\top) = (\ker J)^{\perp_H}$. Therefore,

$$\delta z_{\tan} \in \ker J, \qquad \delta z_{\text{nor}} \in (\ker J)^{\perp_H}, \qquad \delta z_{\tan}^\top H \delta z_{\text{nor}} = 0, \tag{110}$$

which establishes the tangent-H-normal decomposition.

Step 2. Since $L_{\mathrm{dec}} = \frac{1}{2}\delta z^\top H \delta z = \frac{1}{2}(\delta z_{\tan} + \delta z_{\mathrm{nor}})^\top H (\delta z_{\tan} + \delta z_{\mathrm{nor}})$, the H-orthogonality gives

$$2L_{\mathrm{dec}} = \delta z_{\tan}^\top H \delta z_{\tan} + \delta z_{\mathrm{nor}}^\top H \delta z_{\mathrm{nor}}. \tag{111}$$

The identities $PHP = P$ and $K^\top H K = S^{-1}$ follow directly from the definitions:

$$PHP = \left(I - H^{-1}J^\top S^{-1} J\right) P = P, \tag{112}$$

$$K^\top H K = S^{-1} J H^{-1} J^\top S^{-1} = S^{-1}. \tag{113}$$

Using $\delta z_{\tan} = -P\delta c$, $\delta z_{\mathrm{nor}} = K\delta r$, and the symmetry of P, we obtain

$$\delta z_{\tan}^\top H \delta z_{\tan} = \delta c^\top P \delta c, \tag{114}$$

$$\delta z_{\mathrm{nor}}^\top H \delta z_{\mathrm{nor}} = \delta r^\top S^{-1} \delta r. \tag{115}$$

Hence,

$$L_{\mathrm{dec}} = \frac{1}{2}\delta c^\top P \delta c + \frac{1}{2}\delta r^\top S^{-1} \delta r, \tag{116}$$

which proves (27). This is the first decomposition.

Now, it remains to prove the second decomposition. For every predicted equality row $j \in \mathcal{E}$, represented feasibility gives $A_j \hat{z} = \hat{b}_j$, and hence

$$A_j \hat{z} - b_j = \hat{b}_j - b_j = \delta b_j. \tag{117}$$

For every predicted inequality $i \in \mathcal{U} \cap \mathcal{A}^\star$, the common active set implies $G_i \hat{z} = \hat{h}_i$, so

$$G_i \hat{z} - h_i = \hat{h}_i - h_i = \delta h_i. \tag{118}$$

Finally, for $i \in \mathcal{U} \cap \mathcal{I}^\star$, using $s_i^\star := h_i - G_i z^\star > 0$ and $\delta z = -P\delta c + K\delta r$ gives

$$G_i \hat{z} - h_i = G_i(z^\star + \delta z) - h_i \tag{119}$$

$$= G_i \delta z - s_i^\star \tag{120}$$

$$= -G_i P \delta c + G_i K \delta r - s_i^\star. \tag{121}$$

Substituting these three identities into the definition of L_{feas} yields

$$L_{\mathrm{feas}} = \frac{1}{2}\sum_{j\in\mathcal{E}} \left(\frac{\delta b_j}{e_j}\right)^2 + \frac{1}{2}\sum_{i\in\mathcal{U}\cap\mathcal{A}^\star} \left[\frac{\delta h_i}{d_i}\right]_+^2 + \frac{1}{2}\sum_{i\in\mathcal{U}\cap\mathcal{I}^\star} \left[\frac{-G_i P\delta c + G_i K \delta r - s_i^\star}{d_i}\right]_+^2, \tag{122}$$

which proves (28). This completes the proof of Theorem 2.

E.6 Proof of Proposition 4

Let $\hat{J} := [A; G_{\hat{\mathcal{A}}}]$, $\hat{S} := \hat{J} H^{-1} \hat{J}^\top$, and $\hat{K} := H^{-1}\hat{J}^\top \hat{S}^{-1}$. By regularity, $\hat{J}$ has full row rank, so $\hat{S}$ is nonsingular and $\hat{J}\hat{K} = I$. For $i \in \mathcal{U} \cap \hat{\mathcal{A}}$, let $\hat{k}_i$ denote the column of $\hat{K}$ associated with the active inequality row G_i. Fixed-active-set sensitivity gives

$$J_h u_i = \frac{\partial \hat{z}}{\partial \hat{h}_i} = \begin{cases} \hat{k}_i, & i \in \mathcal{U} \cap \hat{\mathcal{A}}, \\ 0, & i \in \mathcal{U} \cap \hat{\mathcal{I}}. \end{cases} \tag{123}$$

Hence, every u_i with $i \in \mathcal{U} \cap \hat{\mathcal{I}}$ belongs to $\ker J_h$, so $\mathrm{span}\{u_i : i \in \mathcal{U} \cap \hat{\mathcal{I}}\} \subseteq \ker J_h$. For the reverse inclusion, let $x \in \ker J_h$. Using (123), we get

$$J_h x = \sum_{i\in\mathcal{U}\cap\hat{\mathcal{A}}} x_i \hat{k}_i = 0. \tag{124}$$

The columns $\{\hat{k}_i : i \in \mathcal{U} \cap \hat{\mathcal{A}}\}$ are linearly independent: indeed, they form a subset of the columns of $\hat{K}$, and $\hat{J}\hat{K} = I$ implies that $\hat{K}$ has full column rank. Therefore, $x_i = 0$ and $i \in \mathcal{U} \cap \hat{\mathcal{A}}$. Thus, x is supported only on $\mathcal{U} \cap \hat{\mathcal{I}}$, proving

$$\ker J_h = \operatorname{span}\{u_i : i \in \mathcal{U} \cap \hat{\mathcal{I}}\}. \tag{125}$$

Now let $\delta\hat{h}$ be supported on $\mathcal{U} \cap \hat{\mathcal{I}}$ and suppose that $\hat{h} + \delta\hat{h}$ remains in the same represented critical region. Since only represented-inactive bounds are perturbed, the active-set KKT system is unchanged. Hence, $\hat{z}$ and the active multipliers remain unchanged, while the represented-inactive multipliers remain zero. For every $i \in \mathcal{U} \cap \hat{\mathcal{I}}$, the branch-wise signed margin is $\hat{q}_i = (\hat{h}_i - G_i\hat{z})/d_i$. Therefore, under the perturbation,

$$\hat{q}_i(\hat{h} + \delta\hat{h}) = \hat{q}_i(\hat{h}) + \frac{\delta\hat{h}_i}{d_i}, \qquad i \in \mathcal{U} \cap \hat{\mathcal{I}}, \tag{126}$$

whereas the signed margins on all other supervised rows are unchanged. Following the definition L_{reg}, we consequently find

$$L_{\mathrm{reg}}(\hat{h} + \delta\hat{h}) - L_{\mathrm{reg}}(\hat{h}) = \frac{1}{2k} \sum_{i \in \mathcal{U} \cap \hat{\mathcal{I}}} \left(\left(\hat{q}_i - q_i^\star + \frac{\delta\hat{h}_i}{d_i} \right)^2 - (\hat{q}_i - q_i^\star)^2 \right) \tag{127}$$

$$= \frac{1}{k} \sum_{i \in \mathcal{U} \cap \hat{\mathcal{I}}} (\hat{q}_i - q_i^\star) \frac{\delta\hat{h}_i}{d_i} + \frac{1}{2k} \sum_{i \in \mathcal{U} \cap \hat{\mathcal{I}}} \left(\frac{\delta\hat{h}_i}{d_i} \right)^2, \tag{128}$$

which proves (30). The quadratic term is strictly positive for every nonzero perturbation supported on $\mathcal{U} \cap \hat{\mathcal{I}}$, since $k > 0$ and $d_i > 0$. This completes the proof of Proposition 4.

E.7 Proof of Proposition 5

Within the fixed represented active-set region, $\delta z = Z\delta\eta + z_0$ and $\delta q_{\mathcal{U}} = Q\delta\eta + q_0$. Since the true parameters are fixed, $\nabla_{\hat{\eta}} L_{\mathrm{signal}} = \nabla_{\delta\eta} L_{\mathrm{signal}}$. We consider the three terms of L_{signal} separately. First, for the decision discrepancy loss, we have

$$L_{\mathrm{dec}} = \frac{1}{2}\delta z^\top H \delta z = \frac{1}{2}(Z\delta\eta + z_0)^\top H (Z\delta\eta + z_0). \tag{129}$$

True feasibility implies $Az^\star = b$ and, with $s^\star := h - Gz^\star$, we get $A\hat{z} - b = A\delta z$ and $G\hat{z} - h = G\delta z - s^\star$. On any region with fixed true-violation pattern $\chi_i = \mathbb{1}\{G_i\delta z - s_i^\star > 0\}$, the positive-part terms have a fixed branch, and the feasibility-shaping loss becomes

$$2L_{\mathrm{feas}} = \delta z^\top A^\top W_b A\, \delta z + (G\delta z - s^\star)^\top W_h C_\chi (G\delta z - s^\star). \tag{130}$$

Substituting $\delta z = Z\delta\eta + z_0$ gives

$$\begin{aligned} 2L_{\mathrm{feas}} = {} & (Z\delta\eta + z_0)^\top A^\top W_b A (Z\delta\eta + z_0) \\ & + \Big(G(Z\delta\eta + z_0) - s^\star \Big)^\top W_h C_\chi \Big(G(Z\delta\eta + z_0) - s^\star \Big). \end{aligned} \tag{131}$$

Finally, by definition of the supervised margin error $\delta q_{\mathcal{U}} := \hat{q}_{\mathcal{U}} - q^\star_{\mathcal{U}}$, the regime supervision loss satisfies

$$L_{\mathrm{reg}} = \frac{1}{2k}\|\delta q_{\mathcal{U}}\|_2^2 = \frac{1}{2k}\|Q\delta\eta + q_0\|_2^2. \tag{132}$$

Differentiating $L_{\mathrm{signal}} = L_{\mathrm{dec}} + \lambda_{\mathrm{feas}} L_{\mathrm{feas}} + \lambda_{\mathrm{reg}} L_{\mathrm{reg}}$ with respect to $\delta\eta$ yields

$$\begin{aligned} \nabla_{\hat{\eta}} L_{\mathrm{signal}} = {} & Z^\top H (Z\delta\eta + z_0) \\ & + \lambda_{\mathrm{feas}} Z^\top \Big(A^\top W_b A (Z\delta\eta + z_0) + G^\top W_h C_\chi \big(G(Z\delta\eta + z_0) - s^\star \big) \Big) \\ & + \frac{\lambda_{\mathrm{reg}}}{k} Q^\top (Q\delta\eta + q_0). \end{aligned} \tag{133}$$

Collecting the terms that multiply $\delta\eta$ and the terms independent of $\delta\eta$ gives

$$\nabla_{\hat{\eta}} L_{\text{signal}} = \Omega_{\hat{\mathcal{A}},\chi}\delta\eta + \zeta_{\hat{\mathcal{A}},\chi}, \tag{134}$$

where

$$\Omega_{\hat{\mathcal{A}},\chi} = Z^\top HZ + \lambda_{\text{feas}} Z^\top \Big(A^\top W_b A + G^\top W_h C_\chi G\Big) Z + \frac{\lambda_{\text{reg}}}{k} Q^\top Q, \tag{135}$$

$$\zeta_{\hat{\mathcal{A}},\chi} = Z^\top Hz_0 + \lambda_{\text{feas}} Z^\top \Big(A^\top W_b A z_0 + G^\top W_h C_\chi (Gz_0 - s^\star)\Big) + \frac{\lambda_{\text{reg}}}{k} Q^\top q_0. \tag{136}$$

These are exactly (34) and (35).

It remains to establish positive semidefiniteness. For any vector x,

$$\begin{aligned} x^\top \Omega_{\hat{\mathcal{A}},\chi} x = (Zx)^\top H(Zx) + \lambda_{\text{feas}}(AZx)^\top W_b(AZx) \\ + \lambda_{\text{feas}}(GZx)^\top W_h C_\chi (GZx) + \frac{\lambda_{\text{reg}}}{k}\|Qx\|_2^2 \geq 0, \end{aligned} \tag{137}$$

because $H \succ 0$, $W_b \succeq 0$, $W_h C_\chi \succeq 0$, and $\lambda_{\text{feas}}, \lambda_{\text{reg}} \geq 0$. Hence, $\Omega_{\hat{\mathcal{A}},\chi} \succeq 0$. This completes the proof of Proposition 5.

F Experimental details

F.1 IEEE 14-bus benchmark

Optimisation problem. The IEEE 14-bus benchmark has 20 physical branches and three controllable generators at buses $\{1, 2, 3\}$, with p_g denoting the power output of generator g. The generator at bus 3 (generator 3 in our indexing) is eliminated through power balance, i.e., $p_3 = D - p_1 - p_2$, where D denotes the total system demand for the instance, yielding a two-dimensional strongly convex reduced QP with decision variable $z = (p_1, p_2) \in \mathbb{R}^2$. The represented problem contains the generator-bound inequalities together with two uncertain directional transmission constraints. Because the eliminated generator has a quadratic generation cost, the reduction also induces a known demand-dependent contribution to the reduced linear objective. In particular, if the full generation cost is $\sum_{g=1}^{3}(a_g p_g^2 + \ell_g p_g)$, where a_g and ℓ_g denote the quadratic and linear generation-cost coefficients, then $c_j = \ell_j - \ell_3 - 2a_3 D$, for $j = 1, 2$, while the reduced Hessian contains the corresponding contribution from the quadratic cost of generator 3.

Joint uncertainty. The uncertain directional transmission limits correspond to branches 5-6 and 2-4. For both uncertain lines, the true limit is generated by perturbing a nominal rating through a nonlinear function of the latent context:

$$h_j(W, \varepsilon_j) = \bar{h}_j + \phi_j(W) + \varepsilon_j, \qquad j = 1, 2, \tag{138}$$

where $W = (w_1, w_2, w_3) \in \mathbb{R}^3$ denotes the latent context, ε_j are independent line-rating perturbations with $\varepsilon_j \sim \text{Unif}[-0.12, 0.12]$, and the nominal ratings are $\bar{h}_1 = 41$ and $\bar{h}_2 = 49$. The nonlinear context maps are

$$\phi_1(W) = 0.30w_1 - 0.15w_2 + 0.50\sin(\pi(w_1 + 0.3w_2)) + 0.18w_1 w_2, \tag{139}$$

$$\phi_2(W) = 0.25w_2 + 0.10w_1 + 0.45\sin(\pi(w_2 - 0.2w_1)) + 0.15w_1 w_2. \tag{140}$$

The predictor ranges are $h_1 \in [39.5, 42.5]$ and $h_2 \in [48.0, 50.5]$.

The three full linear generation-cost coefficients are also nonlinear functions of W:

$$\ell_1(W) = 20 + 0.55w_1 + 0.20w_2 + 0.18\sin(\pi w_1) + 0.10w_1 w_2, \tag{141}$$

$$\ell_2(W) = 20 + 0.50w_2 - 0.18w_1 + 0.16\sin(\pi w_2) + 0.08w_1 w_2, \tag{142}$$

$$\ell_3(W) = 40 + 1.20w_3 + 0.40(w_3^2 - (1/3)) + 0.22\sin(\pi w_3) + 0.10w_1 w_3. \tag{143}$$

After eliminating generator 3, the two uncertain full-cost coefficients induce the two predicted coordinates of the reduced linear objective, corresponding to $\ell_1(W) - \ell_3(W)$ and $\ell_2(W) - \ell_3(W)$. Their predictor ranges are $[-23, -17]$. The reduced linear objective additionally contains a known demand-dependent contribution induced by the quadratic cost of the eliminated generator. Thus, the IEEE 14-bus benchmark jointly predicts two objective and two constraint parameters using a predictor with 16 trainable parameters. The branch 2-4 limit is activation-exposed in the nominal regime, whereas the branch 5-6 limit is dormant over the nominal training sample and can become decision-relevant under the high-load shift.

F.2 IEEE 57-BUS BENCHMARK

Optimisation problem. The IEEE 57-bus benchmark has 80 physical branches and seven controllable generators at buses $\{1, 2, 3, 6, 8, 9, 12\}$. The generator at bus 12 (generator 7 in our indexing) is eliminated through power balance, i.e., $p_7 = D - \sum_{j=1}^{6} p_j$, yielding a six-dimensional strongly convex reduced QP with decision variable $z = (p_1, \ldots, p_6) \in \mathbb{R}^6$, corresponding respectively to the generators at buses $\{1, 2, 3, 6, 8, 9\}$. The represented problem retains the complete reduced inequality system: 14 generator-bound rows and 160 directional branch-flow inequalities, for a total of 174 inequality rows. All represented inequalities are checked when accepting a forward solution. Because the eliminated generator has a quadratic generation cost, the reduction also induces a known demand-dependent contribution to the reduced linear objective.

Joint uncertainty. Four physical line ratings are uncertain, corresponding to branches 8-9, 24-26, 28-29, and 57-56. As in the IEEE 14-bus benchmark, each rating is obtained by perturbing a nominal value through a nonlinear function of the latent context. Here, the perturbation is expressed relative to the nominal rating:

$$h_j(W) = \bar{h}_j + \bar{h}_j \alpha_j \phi_j(W) = \bar{h}_j \big[1 + \alpha_j \phi_j(W)\big], \qquad j = 1, \ldots, 4, \tag{144}$$

where $(\bar{h}_1, \bar{h}_2, \bar{h}_3, \bar{h}_4) = (206.026, 19.206, 32.550, 4.132)$ and the relative perturbation amplitudes are $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (0.030, 0.030, 0.035, 0.035)$. The nonlinear context maps are

$$\phi_1(W) = 0.60w_1 - 0.25w_2 + 0.35\sin(\pi(w_1 + 0.25w_2)) + 0.15w_1w_2 + 0.08w_3, \tag{145}$$

$$\phi_2(W) = 0.50w_2 + 0.20w_1 + 0.35\sin(\pi(w_2 - 0.20w_1)) + 0.15w_1w_2 - 0.08w_3, \tag{146}$$

$$\phi_3(W) = 0.48w_3 - 0.18w_1 + 0.32\sin(\pi(w_3 + 0.20w_1)) + 0.13w_1w_3 + 0.06w_2, \tag{147}$$

$$\phi_4(W) = 0.42w_1 - 0.30w_2 + 0.24w_3 + 0.30\sin(\pi(0.55w_1 - 0.35w_3)) + 0.12w_2w_3. \tag{148}$$

The predictor ranges are $h_j \in [0.95\bar{h}_j, 1.05\bar{h}_j]$, for $j = 1, \ldots, 4$. Each predicted physical rating is shared by the two directional branch-flow inequalities associated with that line, so the four uncertain ratings induce eight uncertain inequality rows.

Two full linear generation-cost coefficients are uncertain: those of the generators at buses 1 and 2,

$$\ell_1(W) = 20 + 0.55w_1 + 0.20w_2 + 0.18\sin(\pi w_1) + 0.10w_1w_2, \tag{149}$$

$$\ell_2(W) = 40 + 1.20w_3 + 0.40(w_3^2 - (1/3)) + 0.22\sin(\pi w_3) + 0.10w_1w_3. \tag{150}$$

All remaining full linear generation-cost coefficients are fixed. After elimination of the generator at bus 12, these two uncertain full-cost coefficients enter the first two coordinates of the reduced linear objective. All six reduced linear coordinates additionally contain the known demand-dependent term $-2a_7D$ induced by the quadratic cost of the eliminated generator. Thus, although the reduced linear objective varies with demand in every coordinate, only two objective coefficients are predicted; the demand-dependent contribution is known for each instance. Their full-cost predictor ranges are $\ell_1 \in [16, 24]$ and $\ell_2 \in [32, 48]$. Therefore, the IEEE 57-bus benchmark jointly predicts two objective coefficients and four physical line-rating parameters, which induce eight uncertain directional constraint rows, using a predictor with 24 trainable parameters. The limits on branches 8-9 and 24-26 are exposed and already decision-relevant over the nominal load range. The limits on branches 28-29 and 57-56 are dormant during nominal training and can become relevant under the high-load regime.

F.3 EXPERIMENTAL PROTOCOL

Methods and training. We compare the eight methods introduced in Section 7: `Dec-Only`, `Dec-Feas`, `SIGNAL`, `MSE-PtO`, `dQP-DFL`, `MSE→PEAR`, `MSE→SPO+`, and `MSE→ODECE`. All methods use the same case-specific bounded affine-sigmoid predictor with separate objective and constraint heads and are trained and evaluated on the same paired datasets. Because several external baselines were originally developed for one-sided uncertainty—either objective or constraint prediction—we adapt them to the joint-uncertainty setting through sequential training, retaining the corresponding objective-side or constraint-side learning criterion. Specifically, `MSE→PEAR` and `MSE→SPO+` freeze an MSE-trained constraint predictor before objective-side training, whereas `MSE→ODECE` freezes an MSE-trained objective predictor before constraint-side training. Decision-stage training uses 20 epochs, batch size 100, and global gradient-norm clipping at 5. MSE-stage training uses a maximum of 100 epochs, with validation-based early stopping when enabled by the benchmark protocol.

Observation model and constraint shift. The latent context is $W \sim \mathrm{Unif}([-1,1]^3)$, and the predictor observes $X = W + \sigma_{\mathrm{obs}}\varepsilon$, where $\varepsilon \sim \mathcal{N}(0, I_3)$, at observation-noise levels $\sigma_{\mathrm{obs}} \in \{0, 0.05, 0.10, 0.15, 0.20\}$. The true objective and constraint parameters are nonlinear functions of W, inducing deliberate model misspecification relative to the affine-sigmoid predictor. All predicted uncertain constraints are inequality constraints. The in-distribution (ID) operating regime uses a global demand scale in $[0.82, 0.90]$, whereas out-of-distribution (OOD) evaluation uses the high-load regime $[0.95, 0.98]$. Bus-level demands are further perturbed around the globally scaled nominal profile by small independent multiplicative fluctuations. The shift therefore increases congestion pressure while leaving the network topology, the mappings from W to the uncertain objective and constraint parameters, and the structure of the downstream optimisation problem unchanged.

Hyperparameter calibration. Both benchmarks use the same calibration protocol, applied independently at each training-set size. Screening uses seeds $\{99501, 99502, 99503\}$ and confirmation uses seeds $\{99501, \ldots, 99510\}$ over the five observation-noise levels. Decision-stage configurations are compared using the weighted validation score $S = 0.8R_{\mathrm{dev-ID}} + 0.2R_{\mathrm{dev-OOD}}$, followed by stability screening, confirmation, and a one-standard-error eligibility rule. For `Dec-Feas` and `SIGNAL`, calibration additionally searches λ_{feas}, while `SIGNAL` also searches λ_{reg}. Stage-1 learning rates of sequential baselines are selected using prediction MSE, since these stages produce a predictor that is subsequently frozen for Stage 2. For `MSE`$\rightarrow$`ODECE`, the method-specific hyperparameters are additionally calibrated. A common learning-rate grid is used for all methods.

Training and evaluation. The benchmarks use $N_{\mathrm{train}} \in \{500, 750, 1000, 1250, 1500\}$ training samples, depending on the specific experiment; both use 300 ID-validation and 300 development-OOD samples. We use the same 20 paired reporting seeds, $\{99601, \ldots, 99620\}$, for every method and condition. Both benchmarks are evaluated on 3000 ID and 3000 OOD test instances per reporting condition. The three evaluation criteria are those defined in Section 4: operational regret, true-violation frequency, and true-violation magnitude. The application-dependent cost per unit of true-constraint violation is fixed at $C_{h,i} = 3000$ for $i \in \mathcal{U}$, all methods, and both benchmarks. For each training size and observation-noise level, results are paired across the 20 reporting seeds. For each sweep, the overall value first averages the five evaluated conditions within each reporting seed and then averages across the 20 reporting seeds. Statistical comparisons use paired two-sided t-tests on the seed-level overall values, with Holm correction over the seven method-vs-`SIGNAL` comparisons separately within each sweep and evaluation criterion.

Numerical implementation details. The IEEE 14-bus benchmark is solved with the exact low-dimensional active-set enumerative solver, while the IEEE 57-bus benchmark uses the custom primal active-set backend with an SLSQP fallback (`ftol` $= 10^{-11}$, at most 1000 iterations). Active rows are identified using a multiplier threshold of 10^{-7}; weak activity is resolved by the one-sided directional QP using slack and multiplier thresholds of 10^{-7} and directional tolerance 10^{-9}. The implementation differentiates the reconstructed fixed-active-set KKT map with autodiff, which is algebraically equivalent to the adjoint formulas. For solutions obtained with the custom primal active-set backend, accepted solutions satisfy primal violation at most 2×10^{-7} and stationarity residual at most 2×10^{-7} after KKT polishing, or 2×10^{-6} for the SLSQP fallback. Bus-level demand is independently perturbed multiplicatively as $d_{n,b} = \bar{d}_b \rho_n (1 + \xi_{n,b})$, where $\xi_{n,b} \sim \mathrm{Unif}[-0.006, 0.006]$ and ρ_n denotes the sampled global demand scale for instance n. All computational jobs were run on the same CPU type (AMD EPYC 7713), with 1 GB of RAM allocated per job.

F.4 Additional numerical results

Robustness to observation noise. We begin with the in-distribution setting to isolate the effect of imperfect contextual information from distribution shift. Holding the IEEE 57-bus benchmark operating regime fixed with training sample size $N_{\mathrm{train}} = 1500$, we progressively increase the observation noise over the five levels $\sigma_{\mathrm{obs}} \in \{0, 0.05, 0.10, 0.15, 0.20\}$. This experiment provides the baseline stress test: it examines how the methods respond when prediction becomes increasingly difficult, while the train and test distributions remain aligned. Figure 2 shows the seed-averaged performance of all methods across the five observation-noise levels for operational regret, true-violation frequency, and true-violation magnitude.

Robustness under constraint shift. We next increase the difficulty by introducing the OOD high-load regime with training sample size $N_{\mathrm{train}} = 1500$ while retaining the same observation-noise levels. The shift increases congestion pressure and makes previously dormant transmission limits more decision-relevant, thereby testing the setting that most directly motivates joint objective-constraint learning. In contrast to the preceding ID experiment, performance now depends not only on robustness to noisy observations but also on how well the learned objective and constraint predictions transfer to a changed operating regime. Table 1 reports results across the five noise levels, while Figure 3 shows the corresponding seed-averaged performance of all methods across observation-noise levels for the three evaluation criteria.

Robustness under data scarcity. Finally, we retain the more challenging OOD constraint-shift setting and additionally vary the amount of training data. This experiment introduces a third source of difficulty—limited supervision—and tests whether the methods remain effective when the predictor must learn jointly uncertain objective and constraint parameters from fewer samples. We vary the training-set size over five levels while preserving the same OOD evaluation regime and observation-noise level $\sigma_{\mathrm{obs}} = 0.10$, so that changes in performance can be attributed directly to data availability. Table 2 reports the detailed results across training sizes, and Figure 4 shows the seed-averaged performance of all methods across training sample sizes for the three evaluation criteria.

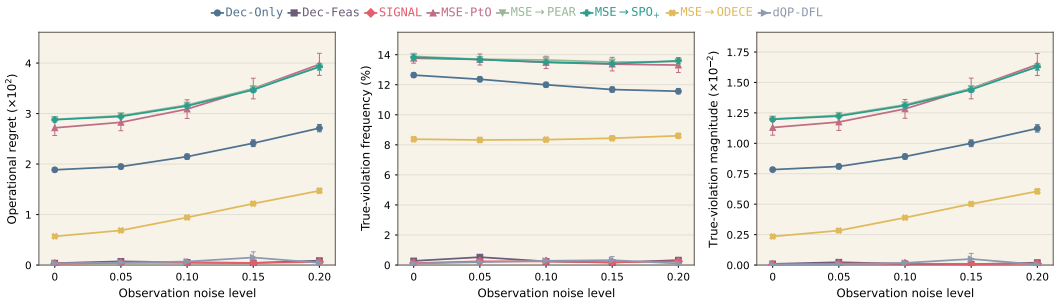

Figure 2: **Robustness to observation noise experiment:** in-distribution performance across observation-noise levels under the nominal operating regime with training sample size $N_{\text{train}} = 1500$. Curves report seed-averaged operational regret, true-violation frequency, and true-violation magnitude for all methods; error bars denote standard errors across reporting seeds. Lower is better in all panels. `SIGNAL`, together with `Dec-Feas` and `dQP-DFL`, maintains values close to zero across all three criteria, whereas several external baselines degrade noticeably as observation noise increases.

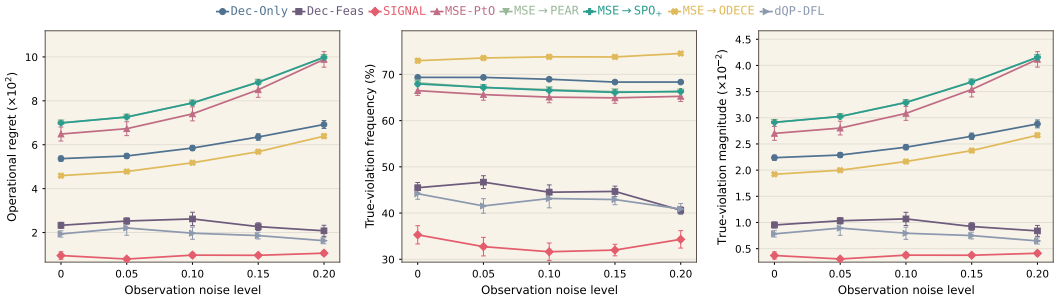

Figure 3: **Robustness under constraint shift experiment:** out-of-distribution performance across the five observation-noise levels under the high-load constraint shift with training sample size $N_{\text{train}} = 1500$. Curves report seed-averaged operational regret, true-violation frequency, and true-violation magnitude for all methods; error bars denote standard errors across reporting seeds. Lower is better in all panels. `SIGNAL` attains the lowest values across all three criteria throughout the noise range and remains comparatively stable as observation noise increases.

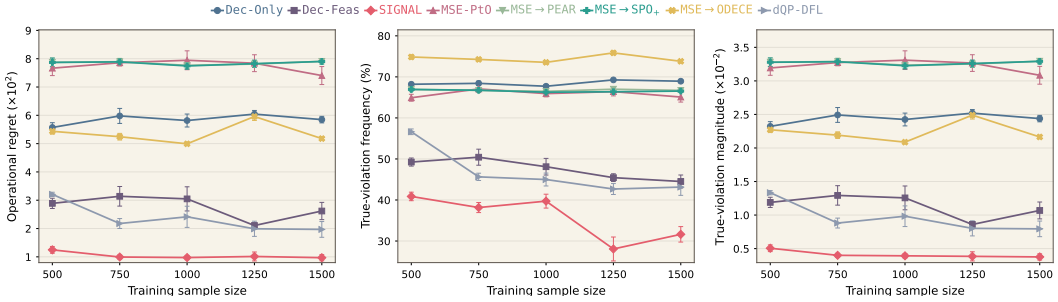

Figure 4: **Robustness under data scarcity experiment:** out-of-distribution performance across five training sample sizes under the high-load constraint shift at observation-noise level $\sigma_{\text{obs}} = 0.10$. Curves report seed-averaged operational regret, true-violation frequency, and true-violation magnitude for all methods; error bars denote standard errors across reporting seeds. Lower is better in all panels. `SIGNAL` maintains the lowest values across the evaluated training sizes, indicating robustness to reduced training-data availability.

Table 1: Performance across observation-noise levels under the OOD constraint shift with training sample size $N_{\text{train}} = 1500$, reported for operational regret, true-violation frequency, and true-violation magnitude. Entries are $\text{mean}\pm\text{SE}$ across reporting seeds. p_{Holm} denotes the Holm-adjusted paired two-sided t-test against SIGNAL, computed from seed-level overall values.

Criteria	Method	$\sigma_{\text{obs}} = 0$	0.05	0.10	0.15	0.20	Overall	p_{Holm}
Operational regret ($\times 10^2$)	Dec-Only	5.37±0.12	5.49±0.10	5.85±0.12	6.35±0.15	6.92±0.18	5.99±0.12	4.91e-19
	Dec-Feas	2.33±0.14	2.52±0.14	2.62±0.30	2.27±0.17	2.08±0.25	2.36±0.13	2.71e-08
	SIGNAL	**0.95±0.17**	**0.79±0.06**	**0.97±0.11**	**0.96±0.08**	**1.06±0.10**	**0.95±0.06**	–
	MSE-PtO	6.49±0.32	6.73±0.32	7.40±0.32	8.50±0.34	9.89±0.35	7.80±0.30	3.47e-15
	MSE→PEAR	7.00±0.07	7.26±0.08	7.91±0.09	8.85±0.10	9.98±0.11	8.20±0.09	1.64e-25
	MSE→SPO+	6.99±0.07	7.26±0.07	7.90±0.09	8.85±0.10	9.98±0.11	8.20±0.09	1.64e-25
	MSE→ODECE	4.59±0.05	4.78±0.05	5.18±0.06	5.68±0.08	6.39±0.10	5.32±0.06	7.13e-21
	dQP-DFL	1.93±0.14	2.21±0.33	1.97±0.28	1.86±0.15	1.63±0.14	1.92±0.12	1.14e-07
True-violation frequency (%)	Dec-Only	69.35±0.37	69.34±0.32	68.94±0.34	68.33±0.34	68.34±0.29	68.86±0.28	8.91e-19
	Dec-Feas	45.48±1.14	46.69±1.39	44.51±1.58	44.68±1.14	40.62±0.70	44.40±0.66	1.27e-09
	SIGNAL	**35.29±1.97**	**32.74±2.01**	**31.63±1.89**	**31.99±1.25**	**34.31±1.88**	**33.19±0.91**	–
	MSE-PtO	66.48±1.03	65.62±1.22	65.08±1.22	64.92±1.17	65.25±1.11	65.47±1.07	5.40e-17
	MSE→PEAR	68.16±0.29	67.17±0.31	66.71±0.37	66.25±0.34	66.20±0.32	66.90±0.30	1.57e-18
	MSE→SPO+	67.94±0.32	67.18±0.33	66.54±0.36	66.11±0.39	66.33±0.36	66.82±0.32	1.57e-18
	MSE→ODECE	72.96±0.47	73.54±0.45	73.79±0.41	73.77±0.38	74.52±0.39	73.72±0.38	6.44e-19
	dQP-DFL	44.18±1.21	41.54±1.57	43.15±2.02	42.92±1.11	40.87±1.19	42.53±0.93	4.00e-07
True-violation magnitude ($\times 10^{-2}$)	Dec-Only	2.24±0.05	2.29±0.04	2.44±0.05	2.65±0.06	2.88±0.07	2.50±0.05	3.45e-19
	Dec-Feas	0.95±0.06	1.03±0.06	1.07±0.13	0.92±0.07	0.84±0.10	0.96±0.05	1.75e-08
	SIGNAL	**0.37±0.07**	**0.30±0.03**	**0.38±0.05**	**0.37±0.03**	**0.41±0.04**	**0.37±0.03**	–
	MSE-PtO	2.70±0.13	2.80±0.13	3.08±0.13	3.54±0.14	4.12±0.15	3.25±0.13	2.97e-15
	MSE→PEAR	2.91±0.03	3.03±0.03	3.29±0.04	3.69±0.04	4.16±0.05	3.42±0.04	1.37e-25
	MSE→SPO+	2.91±0.03	3.02±0.03	3.29±0.04	3.68±0.04	4.16±0.05	3.41±0.04	1.37e-25
	MSE→ODECE	1.92±0.02	2.00±0.02	2.16±0.03	2.37±0.03	2.67±0.04	2.22±0.02	4.72e-21
	dQP-DFL	0.78±0.06	0.89±0.14	0.79±0.12	0.75±0.06	0.65±0.06	0.77±0.05	1.08e-07

Table 2: Performance across training sample sizes under the OOD constraint shift at observation-noise level $\sigma_{\text{obs}} = 0.10$, reported for operational regret, true-violation frequency, and true-violation magnitude. Entries are $\text{mean}\pm\text{SE}$ across reporting seeds. p_{Holm} denotes the Holm-adjusted paired two-sided t-test against SIGNAL, computed from seed-level overall values.

Criteria	Method	$N_{\text{train}} = 500$	750	1000	1250	1500	Overall	p_{Holm}
Operational regret ($\times 10^2$)	Dec-Only	5.57±0.17	5.98±0.26	5.82±0.23	6.04±0.13	5.85±0.12	5.85±0.13	6.98e-19
	Dec-Feas	2.89±0.18	3.14±0.35	3.05±0.43	2.11±0.09	2.62±0.30	2.76±0.12	1.22e-10
	SIGNAL	**1.25±0.12**	**0.99±0.07**	**0.97±0.08**	**1.01±0.16**	**0.97±0.11**	**1.04±0.04**	–
	MSE-PtO	7.67±0.26	7.86±0.12	7.94±0.33	7.84±0.30	7.40±0.32	7.74±0.11	1.66e-22
	MSE→PEAR	7.87±0.15	7.89±0.12	7.77±0.11	7.81±0.12	7.91±0.09	7.85±0.11	2.07e-22
	MSE→SPO+	7.87±0.16	7.89±0.12	7.74±0.11	7.83±0.13	7.90±0.09	7.85±0.11	2.07e-22
	MSE→ODECE	5.43±0.10	5.24±0.12	4.99±0.06	5.96±0.14	5.18±0.06	5.36±0.07	1.66e-22
	dQP-DFL	3.21±0.07	2.18±0.18	2.41±0.37	1.99±0.27	1.97±0.28	2.35±0.12	3.59e-09
True-violation frequency (%)	Dec-Only	68.20±0.45	68.42±0.53	67.72±0.41	69.26±0.34	68.94±0.34	68.51±0.24	4.56e-19
	Dec-Feas	49.25±0.98	50.43±1.95	48.11±2.03	45.44±0.95	44.51±1.58	47.55±0.82	5.27e-09
	SIGNAL	**40.87±1.04**	**38.16±1.23**	**39.72±1.69**	**28.06±2.91**	**31.63±1.89**	**35.69±0.80**	–
	MSE-PtO	64.92±0.84	67.06±0.35	65.92±0.79	66.41±1.04	65.08±1.22	65.88±0.43	2.23e-17
	MSE→PEAR	66.92±0.34	66.72±0.47	66.47±0.44	66.99±0.54	66.71±0.37	66.76±0.32	6.13e-18
	MSE→SPO+	66.95±0.38	66.74±0.47	66.26±0.41	66.37±0.56	66.54±0.36	66.57±0.36	7.43e-18
	MSE→ODECE	74.85±0.56	74.25±0.47	73.54±0.38	75.85±0.49	73.79±0.41	74.46±0.34	7.80e-21
	dQP-DFL	56.62±0.69	45.64±0.89	44.99±1.57	42.67±1.33	43.15±2.02	46.61±0.62	1.29e-09
True-violation magnitude ($\times 10^{-2}$)	Dec-Only	2.32±0.07	2.49±0.11	2.42±0.10	2.52±0.05	2.44±0.05	2.44±0.05	5.71e-19
	Dec-Feas	1.19±0.08	1.29±0.15	1.26±0.18	0.86±0.04	1.07±0.13	1.13±0.05	1.28e-10
	SIGNAL	**0.51±0.05**	**0.40±0.03**	**0.39±0.03**	**0.38±0.07**	**0.38±0.05**	**0.41±0.02**	–
	MSE-PtO	3.19±0.11	3.27±0.05	3.31±0.14	3.27±0.12	3.08±0.13	3.22±0.05	1.37e-22
	MSE→PEAR	3.28±0.06	3.29±0.05	3.23±0.05	3.25±0.05	3.29±0.04	3.27±0.04	1.96e-22
	MSE→SPO+	3.28±0.06	3.29±0.05	3.23±0.05	3.26±0.05	3.29±0.04	3.27±0.04	1.96e-22
	MSE→ODECE	2.27±0.04	2.19±0.05	2.09±0.03	2.49±0.06	2.16±0.03	2.24±0.03	1.37e-22
	dQP-DFL	1.33±0.03	0.88±0.07	0.98±0.15	0.80±0.11	0.79±0.12	0.96±0.05	3.67e-09