REGRET AND COVARIANCE: A CHARACTERIZATION FOR STOCHASTIC OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantifying regret typically involves repeatedly re-solving a stochastic optimization problem across simulated scenarios. Such an approach is costly: the cost scales with the number of scenarios and the cost of each solve. We propose that regret can instead be measured using the joint moments of the cost vector c and the optimal decision $\pi^*(c)$. Working with the optimal concave and Lipschitz under only a feasible region value function $v(c) = \min_{z \in Z} c^\top z$, we show that regret always decomposes as $\mathrm{Regret} = -\mathrm{Cov}(c, \pi^*(c)) - R(c)$ for a residual $R(c)$. We further show that $R(c) = 0$ when π^* is affine in c, which holds for unconstrained and equality-constrained (e.g., budget-constrained Markowitz) quadratic programs and many, but not all, linear programs. For all cases, we give a distribution-free bound $\mathrm{Regret} \leq D\sqrt{\mathrm{tr}(\Sigma_c)}$ that depends only on the radius of the feasible region and the cost covariance. We also derive concentration and asymptotic-normality results for the empirical regret, the latter under an explicit and checkable non-degeneracy condition.

1 INTRODUCTION

A basic task in decision-making under uncertainty is to quantify how much is lost by committing to a single decision before true values are revealed, rather than tailoring the decision to each realization as it becomes known. This quantity can be thought of as the expected value of perfect information. In the language of machine learning and predict-then-optimize modeling, this quantity is known as regret. Regret is a natural diagnostic wherever a stochastic optimization problem $\min_{z \in Z} c^\top z$ is solved using only a point estimate of an uncertain cost vector c: in portfolio selection (Markowitz, 1952), in the broader family of contextual and data-driven optimization problems surveyed by Sadana et al. (2023a), and as the target loss in the smart predict-then-optimize framework of Elmachtoub and Grigas (2022). At the time this paper was written, the standard way to estimate regret was to use Sample Average Approximation: draw B cost scenarios, re-solve the optimization problem for each, and average (Kleywegt et al., 2002; King and Wallace, 2012). This method is reliable but expensive, as it requires as many solves as there are scenarios. A natural question arises: can we compute regret directly from the moments of c and the corresponding optimal decisions, without resolving the optimization problem scenario by scenario?

The answer depends on a single structural property, and distinguishing when that property holds from when it does not is the substance of this paper. Write $\pi^*(c) \in \arg\min_{z \in Z} c^\top z$ for the decision that would be optimal under a known realization of c, and $\mu = \mathbb{E}[c]$. Adding and subtracting $\mathbb{E}[c]^\top \mathbb{E}[\pi^*(c)]$ in the definition of regret produces an algebraic identity, $\mathrm{Regret} = -\mathrm{Cov}(c, \pi^*(c)) - R(c)$, in which all the difficulty is concentrated in the residual term $R(c) = \mathbb{E}[c]^\top (\mathbb{E}[\pi^*(c)] - \pi^*(\mu))$. This identity by itself is content-free: it holds by construction for any decision rule. We show that $R(c) = 0$ when π^* is affine in c, a property of the *decision rule*, not of the cost distribution's smoothness or compactness. We also show that $R(c) = 0$ holds for unconstrained and equality-constrained quadratic programs (Corollaries 1–2), including the budget-constrained Markowitz problem. However, the $R(c) = 0$ property fails generically for linear programs and for any program with active inequality constraints, where π^* is only piecewise affine (Proposition 2).

This paper makes four contributions. Proposition 1 gives an algebraic decomposition of regret that holds by construction. Proposition 1 is built on the optimal value function $v(c) = c^\top \pi^*(c) = \min_{z \in Z} c^\top z$, which is concave for every Z (Lemma 1) and, once Z is compact, Lipschitz with a

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.

constant equal to the radius of Z. These properties require no regularity of π^* itself. Theorem 3 characterizes when the residual $R(c)$ vanishes, separating a structural, distribution-free condition (global affinity of π^*) from a distributional accident specific to a single instance ($\mathbb{E}[c] = 0$). Theorem 2 gives a distribution-free bound, $\mathrm{Regret} \leq D\sqrt{\mathrm{tr}(\Sigma_c)}$, that holds for every compact Z and finite-variance cost distribution, including linear programs and inequality-constrained quadratic programs where the residual is generically nonzero and no formula is available. Theorem 4 and Theorem 5 provide concentration and asymptotic-normality results for the empirical regret, the latter resting on an envelope-theorem argument and a checkable non-degeneracy condition (Assumption 5) rather than an unstated smoothness assumption on π^*. Section 5 tests each of these claims against simulation, reporting Monte Carlo standard errors throughout rather than single point estimates.

We are also explicit about what this paper does not claim. The regret studied here is covariate-free: it compares the decision under a realized cost c against the decision under the unconditional mean μ, for a single fixed distribution P_c. This is the classical expected value of perfect information (Raiffa and Schlaifer, 1961; Birge and Louveaux, 2011), and it is a different object from the model-dependent SPO loss of Elmachtoub and Grigas (2022), which compares a decision driven by a *predictor* $\hat{c} = b(x)$ against the oracle decision and varies with the predictor's parameters during training (Remark 1). None of the results in this paper establish anything about training such a predictor; the honest computational claim they support is narrower and is stated in Remark 3.

The remainder of the paper is organized as follows. Section 3 introduces notation and the primitive objects π^*, v, and Regret. Section 4 develops the value-function framework, the decomposition, the distribution-free bound, and the residual-vanishing theorem. Section B specializes this to quadratic programs, including the unconstrained first-order condition, the equality-constrained (Markowitz) case, and the piecewise-affine structure that emerges once inequality constraints are active. Sections 4.8–4.9 provide the estimation theory: what can and cannot be said about representing regret from a single sample, the concentration of the empirical regret around its true value, and its asymptotic distribution. Section 5 reports the experiments validating all of the above. Section 7 concludes.

2 Related Work

Value of information and sample-based estimation. Our paper is related to the expected value of perfect information (EVPI), introduced by Raiffa and Schlaifer (1961) and standard in stochastic programming (Birge and Louveaux, 2011; Shapiro et al., 2009). $\mathrm{Regret} = v(\mu) - \mathbb{E}[v(c)]$ (Definition 2) is the EVPI of the single-stage program $\min_z \mu^\top z$, which forms the basis for the covariance decomposition. The standard way to estimate this quantity is Sample Average Approximation (SAA): draw B cost scenarios and re-solve the problem for each (Kleywegt et al., 2002), incorporating refinements such as mixed-integer decomposition (Bayraksan and Love, 2015) and bootstrap resampling (Efron and Tibshirani, 1994). These reduce variance but maintain the linear-in-B solve cost. Theorem 2's bound and Theorem 3's cases (Corollary 1) complement SAA rather than replace it. The bound applies without re-solving but only as a bound, and the cases apply only where the optimal decision rule is affine, a property that must be checked for the specific (Z, π^*) pair at hand. Appendix A summarizes practical problems and their respective compactness and finite-variance requirements.

Counterfactual reasoning and the value of information. The regret studied here compares an actual decision against a counterfactual one: what would have been decided had the true cost been known. Three adjacent literatures use related language for different purposes. Decision analysis treats information as having a value defined by a counterfactual world in which it is known (Howard, 1966), which connects to the modern potential-outcomes formalization of causal effects (Rubin, 1974; Imbens and Rubin, 2015): Regret is the average treatment effect, on realized cost, of switching from a certainty-equivalent policy to a fully informed one, with the unusual property that both potential outcomes are computable for every realization rather than only one being observed. Counterfactual regret minimization in extensive-form games (Zinkevich et al., 2007) is a different, multi-agent and sequential object used to drive repeated self-play toward a Nash equilibrium, sharing only the word with the quantity studied here. Counterfactual explanation in interpretable machine learning (Wachter et al., 2017; Verma et al., 2020) offers a third reading: $\pi^*(c)$ is a counterfactual explanation of the certainty-equivalent decision $\pi^*(\mu)$. The full statement of all three connections is given in Remark 7, which this condensed section does not restate.

Predict-then-optimize and decision-focused learning. The predict-then-optimize (PtO) paradigm trains a predictor $b(x)$ so that decisions made from its output, rather than its raw prediction error, supply the training signal (Elmachtoub and Grigas, 2022), with task-based and decision-focused variants for continuous and combinatorial problems (Donti et al., 2017; Wilder et al., 2019), differentiable optimization layers (Amos and Kolter, 2017; Agrawal et al., 2019; Berthet et al., 2020), and solver-free surrogates for backpropagating through combinatorial or LP solves (Vlastelica et al., 2020; Mandi and Guns, 2020; Mandi et al., 2020); contextual and data-driven optimization more broadly is surveyed by Sadana et al. (2023a), with prescriptive covariate-aware decision rules developed for the newsvendor and related problems (Ban and Rudin, 2019). Every method in this literature trains a predictor and is evaluated by a quantity that varies with the predictor's parameters θ during training. Regret in this paper does not: it is covariate-free and fixed for a given P_c (Remark 3). An earlier statement claimed the covariance formula supplies training gradients for such a predictor; this does not follow from anything established here and is retracted. $\mathrm{Cov}(c, \pi^*(c))$ has no dependence on θ, so a gradient with respect to θ is not defined for it. What the covariate-free benchmark offers for a PtO pipeline is a fixed reference point that a trained model should improve on, computed once per dataset rather than once per training step, and tested directly against a standardized benchmark covering eleven PtO and PnO methods, including DBB and IntOpt (Geng et al., 2024).

Sensitivity analysis of parametric optimization. Classical sensitivity analysis studies how the optimal value and optimal solution of a parametric problem change with its parameters (Fiacco, 1983; Bonnans and Shapiro, 2000). The envelope theorem (Milgrom and Segal, 2002) supplies the argument behind $\nabla v(\mu) = \pi^*(\mu)$ without assuming the differentiability of π^* itself. For linear and quadratic programs, the optimal solution as a function of the cost term is piecewise affine over a polyhedral partition determined by the active constraint set (Bank et al., 1982; Best, 2017), the same structure Proposition 2 uses to unify the LP case and the inequality-constrained QP case. The piecewise-affine characterization itself is classical; this paper's contribution is identifying it as the condition that governs whether the residual $R(c)$ vanishes, in place of an unstated Lipschitz assumption on π^* that fails on the maps this literature studies.

Portfolio optimization. Mean-variance portfolio selection (Markowitz, 1952) motivates the equality-constrained QP case (Corollary 2) and the portfolio experiments of Section 5. Practical implementations call for a covariance estimator; the corrected portfolio protocol uses shrinkage (Ledoit and Wolf, 2004) rather than the raw sample covariance, following standard practice. The budget-only case is exact because a budget constraint is a linear equality, not because Z is unconstrained in any other sense; adding a no-short-selling constraint, present in most practical implementations of Markowitz's model, breaks exactness.

This paper's position in the literature. Set against these five literature streams, this paper provides: 1) An alternative to SAA where the optimal decision rule is affine; 2) A bound that requires no re-solving; 3) A covariate-free benchmark for PtO pipelines; 4) A value-function argument that recovers concentration and asymptotic-normality results without assuming π^* is Lipschitz or differentiable.

3 Problem Setup

3.1 The Stochastic Program

Definition 1 (Decision rule)**.** *Let $(\Omega, \mathcal{F}, P)$ be a probability space carrying a random cost vector $c : \Omega \to \mathcal{C} \subseteq \mathbb{R}^d$ (Assumption 1–2) and let $Z \subseteq \mathbb{R}^d$ be a feasible region (Assumption 3). The* decision rule *$\pi^* : \mathcal{C} \to Z$ is a measurable selection satisfying, for P_c-almost every realization c,*

$$\pi^*(c) \;\in\; \arg\min_{z \in Z} \; c^\top z. \tag{1}$$

That is, $\pi^*(c)$ is the decision that would be optimal if the realized cost vector c were known at the time of the decision. *This is the object every subsequent equation in the manuscript actually manipulates, and it is the definition retained for the remainder of the paper.*

3.2 Regret, and Its Relation to the Predict-then-Optimize Literature

Regret compares the decision rule $\pi^*(\cdot)$ evaluated at the realized cost against the certainty-equivalent decision $\pi^*(\mu)$:

Definition 2 (Regret)**.**

$$\text{Regret} \;:=\; c^\top\big|_{c=\mu}\,\pi^*(\mu) \;-\; \mathbb{E}\big[c^\top\pi^*(c)\big] \;=\; \mu^\top\pi^*(\mu) - \mathbb{E}\big[c^\top\pi^*(c)\big],$$

equivalently $\text{Regret} = v(\mu) - \mathbb{E}[v(c)]$ *in the notation of Definition 3, matching Definition 4 in Section 4. This is a single number determined by the distribution P_c and the pair (Z, π^*). It is* not *a function of the random variable c in the way, say, $c^\top\pi^*(c)$ is; writing* $\text{Regret}(c)$*. We follow the common convention in the stochastic-programming literature of writing functionals of a distribution as if evaluated at the underlying random variable (e.g.,* $\text{Var}(c)$ *for a functional of P_c; see Shapiro et al., 2009), and use* Regret *without an argument, or* $\text{Regret}(P_c)$ *when the dependence on the distribution needs to be made explicit, rather than* $\text{Regret}(c)$*.*

Remark 1 (Relation to Smart Predict-then-Optimize)**.** *The manuscript's title and framing invoke Elmachtoub and Grigas's (Elmachtoub and Grigas, 2022) SPO loss, which is a different,* model-dependent *quantity. Given covariates x, a predictor $b : x \mapsto \hat{c}$, and the true cost c, the SPO loss is*

$$\ell_{\text{SPO}}(\hat{c}, c) \;=\; c^\top\pi^*(\hat{c}) - c^\top\pi^*(c) \;\geq 0,$$

the realized cost of acting on the predicted *cost vector $\hat{c} = b(x)$ instead of the true one c, and π^* there plays the role of Definition 1 above. The quantity this paper studies,* Regret *in Definition 2, is the special case of this comparison in which the "predictor" is the constant function $b(x) \equiv \mu$ for every x. No covariates and no learned model are observed, just the unconditional mean. It coincides with the classical expected value of perfect information (EVPI) for the single-stage program $\min_z \mu^\top z$, not with the SPO loss of a trained predictor.*

Remark 2 (Notational consolidation)**.** *Throughout, π^* denotes the decision rule of Definition 1, $\mu := \mathbb{E}[c]$, and* Regret *and R are written without an argument as functionals of P_c, except for sample-based quantities such as the estimator $\widehat{R}_n$ (Lemma 2), which are functions of the data $\{c_i\}_{i=1}^n$ rather than of the distribution alone.*

4 A Mathematical Characterization of Regret

4.1 Setup

Assumption 1 (Probability space)**.** *Let $(\Omega, \mathcal{F}, P)$ be a complete probability space, where Ω is the sample space of cost realizations, $\mathcal{F}$ is a σ-algebra of measurable events, and P is a probability measure.*

Assumption 2 (Cost vector)**.** *The cost vector $c \in \mathcal{C} \subseteq \mathbb{R}^d$ satisfies:*

- Support: *$\mathcal{C}$ is a compact subset of $\mathbb{R}^d$ (convexity of $\mathcal{C}$ itself is not required; only* $\text{conv}(\mathcal{C})$ *enters the arguments below).*
- Moments: $\mathbb{E}[\|c\|^2] < \infty$.

Assumption 3 (Feasible region)**.** *The feasible region $Z \subseteq \mathbb{R}^d$ is nonempty.* Compactness of Z is invoked only where explicitly stated *(Theorem 2 and Theorem 4 below); the decomposition identity (Proposition 1) and the exact-vanishing characterization (Theorem 3) hold for unbounded Z as well, provided $v(c)$ is finite and integrable. Compactness is a hypothesis of specific results, not a blanket requirement.*

Assumption 4 (Decision map)**.** *The optimal decision correspondence admits a measurable selection $\pi^* : \mathcal{C} \to Z$ with $\pi^*(c) \in \arg\min_{z\in Z} c^\top z$ and $\mathbb{E}[\|\pi^*(c)\|^2] < \infty$.* ***We do not assume π^* is Lipschitz, continuous, or differentiable.***

4.2 The Value Function

Definition 3. *The* optimal value function *is $v(c) := \min_{z\in Z} c^\top z = c^\top\pi^*(c)$.*

Lemma 1 (Concavity and Lipschitz continuity of v)**.** *Under Assumption 3:*

(a) v is concave on $\mathbb{R}^d$.

(b) If in addition Z is compact with $D := \sup_{z\in Z} \|z\| < \infty$, then v is D-Lipschitz: $|v(c) - v(c')| \le D\|c - c'\|$ for all $c, c' \in \mathbb{R}^d$.

Neither claim requires any regularity of π^*.

Proof. (a) $v(c) = \min_{z\in Z} c^\top z$ is a pointwise infimum of the family of affine functions $\{c \mapsto c^\top z : z \in Z\}$, and a pointwise infimum of affine functions is concave.

(b) Fix $c, c' \in \mathbb{R}^d$ and let $z' \in \arg\min_{z\in Z}(c')^\top z$, so $v(c') = (c')^\top z'$. Since $z' \in Z$ is feasible (but possibly suboptimal) for the cost vector c,

$$v(c) \le c^\top z' = (c')^\top z' + (c - c')^\top z' \le v(c') + \|c - c'\|\,\|z'\| \le v(c') + D\|c - c'\|.$$

Exchanging the roles of c and c' gives $v(c') \le v(c) + D\|c - c'\|$, and combining the two inequalities yields $|v(c) - v(c')| \le D\|c - c'\|$. □

4.3 Regret as an Expected-Value-of-Perfect-Information Gap

We adopt the sign convention flagged by the referees: since v is concave, $v(\mathbb{E}[c]) - \mathbb{E}[v(c)] \ge 0$ follows from Jensen's inequality, so we define regret to be this nonnegative quantity (equivalently, the expected value of perfect information, EVPI, for the single-stage stochastic program $\min_z \mathbb{E}[c]^\top z$; see Sadana et al. (2023b) and the classical EVPI literature.

Definition 4 (Regret)**.**

$$\mathrm{Regret}(c) \;:=\; v(\mathbb{E}[c]) - \mathbb{E}[v(c)] \;=\; \mathbb{E}[c]^\top \pi^*(\mathbb{E}[c]) - \mathbb{E}\big[c^\top \pi^*(c)\big].$$

Theorem 1 (Nonnegativity of regret)**.** *Under Assumptions 1–4,* $\mathrm{Regret}(c) \ge 0$*, with equality if and only if v is affine on* $\mathrm{conv}\big(\mathrm{supp}(P_c)\big)$.

Proof. Nonnegativity is Jensen's inequality applied to the concave function v (Lemma 1(a)). The equality condition is the standard equality case of Jensen's inequality for concave functions: equality holds iff v is affine on the convex hull of the support of the distribution of c. □

4.4 The Covariance Decomposition

Proposition 1 (algebraic decomposition)**.**

$$\mathrm{Regret}(c) \;=\; -\,\mathrm{Cov}\big(c, \pi^*(c)\big) \;-\; R(c),$$

where $\mathrm{Cov}(c, \pi^*(c)) := \mathbb{E}[c^\top \pi^*(c)] - \mathbb{E}[c]^\top \mathbb{E}[\pi^*(c)]$ *and*

$$R(c) \;:=\; \mathbb{E}[c]^\top \big(\mathbb{E}[\pi^*(c)] - \pi^*(\mathbb{E}[c])\big).$$

Proof. Add and subtract $\mathbb{E}[c]^\top \mathbb{E}[\pi^*(c)]$ in Definition 4:

$$\begin{aligned}\mathrm{Regret}(c) &= \mathbb{E}[c]^\top \pi^*(\mathbb{E}[c]) - \mathbb{E}[c^\top \pi^*(c)] \\ &= \Big(\mathbb{E}[c]^\top \pi^*(\mathbb{E}[c]) - \mathbb{E}[c]^\top \mathbb{E}[\pi^*(c)]\Big) + \Big(\mathbb{E}[c]^\top \mathbb{E}[\pi^*(c)] - \mathbb{E}[c^\top \pi^*(c)]\Big) \\ &= -R(c) \;-\; \mathrm{Cov}\big(c, \pi^*(c)\big).\end{aligned}$$

□

4.5 A Distribution-Free Bound

The following bound holds for *every* feasible region and *every* cost distribution satisfying our assumptions, including linear programs, and does not rely on any regularity of π^*.

Theorem 2 (General bound)**.** *Suppose, in addition, that Z is compact with $D = \sup_{z\in Z} \|z\|$. Then*

$$0 \;\le\; \mathrm{Regret}(c) \;\le\; D\sqrt{\mathrm{tr}(\Sigma_c)},$$

where $\Sigma_c = \mathrm{Cov}(c)$.

Proof. By Lemma 1(b), $v(\mathbb{E}[c]) - v(c) \le D\|\mathbb{E}[c] - c\|$ pointwise. Taking expectations,

$$\mathrm{Regret}(c) = \mathbb{E}\big[v(\mathbb{E}[c]) - v(c)\big] \le D\,\mathbb{E}\big[\|c - \mathbb{E}[c]\|\big] \le D\sqrt{\mathbb{E}\big[\|c - \mathbb{E}[c]\|^2\big]} = D\sqrt{\mathrm{tr}(\Sigma_c)},$$

where the last inequality is Jensen's inequality applied to the convex function $x \mapsto \sqrt{x}$ (equivalently, Cauchy–Schwarz). The lower bound is Theorem 1. □

4.6 Vanishing of the Residual

We now characterize when $R(c) = 0$, so that regret reduces to the covariance term with no residual, $\mathrm{Regret}(c) = -\mathrm{Cov}(c, \pi^*(c))$. Two logically distinct routes give $R(c) = 0$: a *structural* one (global affinity of π^*, which holds for every distribution of c) and a *distributional* one (an accident of $\mathbb{E}[c]$, which holds only for that particular distribution).

Theorem 3 (Sufficient conditions for $R(c) = 0$)**.** *(a)* ***(Structural, distribution-free.)*** *If $\pi^*(c) = Ac + b$ for some fixed matrix A and vector b, for P_c-almost every $c \in \mathcal{C}$, then $R(c) = 0$ for* every *distribution of c supported on $\mathcal{C}$.*

(b) ***(Distributional, instance-specific.)*** *If $\mathbb{E}[c] = 0$, then $R(c) = 0$,* regardless of whether π^* is affine.

Neither condition is necessary: $R(c) = 0$ can hold for other reasons specific to a given joint distribution of $(c, \pi^*(c))$, so (a)–(b) are sufficient, not iff, conditions.

Proof. (a) If $\pi^*(c) = Ac + b$ a.s., linearity of expectation gives $\mathbb{E}[\pi^*(c)] = A\,\mathbb{E}[c] + b = \pi^*(\mathbb{E}[c])$, so $R(c) = \mathbb{E}[c]^\top \cdot 0 = 0$.

(b) By definition $R(c) = \mathbb{E}[c]^\top(\mathbb{E}[\pi^*(c)] - \pi^*(\mathbb{E}[c]))$; if $\mathbb{E}[c] = 0$, the leading factor vanishes regardless of the bracketed term. □

4.6.1 Application to quadratic programs

For $\min_{z \in \mathbb{R}^n} c^\top z + \frac{\lambda}{2} z^\top Q z$ with $Q \succ 0$ and $\lambda > 0$, the optimal decision is linear in c:

$$\boxed{\pi^*(c) = -\frac{1}{\lambda} Q^{-1} c} \qquad \mathrm{Regret}(c) = \frac{1}{\lambda}\mathrm{tr}\big(Q^{-1}\Sigma_c\big) \ge 0,$$

so $R(c) = 0$ for every distribution of c by Theorem 3(a), with no compactness of $Z = \mathbb{R}^n$ required. See Corollary 1 (Appendix B.1) for the derivation.

4.6.2 Application to linear programs: a cautionary example

Example 1 (Linear programs generically have $R(c) \neq 0$)**.** *Let $d = 1$, $Z = [0, 1]$, and $c \sim \mathrm{Unif}[-2, 1]$, so $\mathbb{E}[c] = -\frac{1}{2} \neq 0$. The optimal decision is*

$$\pi^*(c) = \begin{cases} 1 & c < 0 \\ 0 & c > 0 \end{cases} \quad \textit{(ties at } c = 0 \textit{ have probability zero).}$$

This is piecewise constant, not *globally affine, so Theorem 3(a) does not apply, and $\mathbb{E}[c] \neq 0$, so (b) does not apply either. Direct computation gives $\pi^*(\mathbb{E}[c]) = \pi^*(-\frac{1}{2}) = 1$ and $\mathbb{E}[\pi^*(c)] = P(c < 0) = \frac{2}{3}$, so*

$$R(c) = \mathbb{E}[c]\big(\mathbb{E}[\pi^*(c)] - \pi^*(\mathbb{E}[c])\big) = \big(-\tfrac{1}{2}\big)\big(\tfrac{2}{3} - 1\big) = \tfrac{1}{6} \neq 0.$$

This confirms, by direct calculation, that $R(c) \neq 0$ for this LP even though c has an absolutely continuous distribution with compact convex support, and π^* is (trivially, being $\{0, 1\}$-valued) locally constant almost everywhere. Theorem 2's distribution-free bound remains available in this regime; it gives, for this example, $D = 1$ and $\Sigma_c = \mathrm{Var}(c) = 3/4$, so $\mathrm{Regret}(c) \le \sqrt{3}/2 \approx 0.87$, consistent with the value $\mathrm{Regret}(c) = -\mathrm{Cov}(c, \pi^*(c)) - R(c)$ computed directly as $1/3 - (-1/6) = 1/2$.

4.7 Estimation from Data

Lemma 2 (Residual estimator)**.** *Given i.i.d. samples $\{c_i\}_{i=1}^n$ and access to $\pi^*(\cdot)$, the plug-in estimator*

$$\widehat{R}_n \;=\; \bar{c}^\top \left(\frac{1}{n} \sum_{i=1}^{n} \pi^*(c_i) - \pi^*(\bar{c}) \right), \qquad \bar{c} = \frac{1}{n} \sum_{i=1}^{n} c_i,$$

is consistent for $R(c)$ whenever π^* is continuous at $\mathbb{E}[c]$ (so that $\pi^*(\bar{c}) \xrightarrow{P} \pi^*(\mathbb{E}[c])$ by the continuous mapping theorem), with $\widehat{R}_n \xrightarrow{P} R(c)$.

Proof. $\bar{c} \xrightarrow{P} \mathbb{E}[c]$ follows from the weak law of large numbers (finite variance, Assumption 2), and $\frac{1}{n}\sum_i \pi^*(c_i) \xrightarrow{P} \mathbb{E}[\pi^*(c)]$ likewise. If π^* is continuous at $\mathbb{E}[c]$, the continuous mapping theorem gives $\pi^*(\bar{c}) \xrightarrow{P} \pi^*(\mathbb{E}[c])$, and the result follows from Slutsky's theorem applied to the product/sum defining $\widehat{R}_n$. □

Remark 3 (Honest computational accounting)**.** *Evaluating $\widehat{R}_n$ requires computing $\pi^*(\bar{c})$, i.e.* one *solve of the optimization problem, in addition to the $O(nd)$ cost of forming $\bar{c}$ and averaging the cached $\pi^*(c_i)$. The advantage of the covariance-based approach over Sample Average Approximation is that it requires a single solve regardless of n, rather than one solve per scenario ($O(n)$ solves).*

4.8 Concentration

Theorem 4 (Concentration around $-\mathrm{Cov}(c, \pi^*(c)) - R(c)$)**.** *Suppose Z is compact with $D = \sup_{z \in Z} \|z\|$ and $\mathcal{C}$ is compact with $B = \sup_{c \in \mathcal{C}} \|c\|$. Let $\widehat{v}_n = \frac{1}{n}\sum_{i=1}^n v(c_i)$ be the sample mean of the value function over n i.i.d. draws. Then for any $\epsilon > 0$,*

$$P\big(|\widehat{v}_n - \mathbb{E}[v(c)]| > \epsilon\big) \le 2\exp\left(-\frac{n\epsilon^2}{2D^2B^2}\right).$$

Consequently the empirical regret $\widehat{\mathrm{Regret}}_n := v(\bar{c}) - \widehat{v}_n$ concentrates around the true $\mathrm{Regret}(c)$ at the same $\sqrt{1/n}$ rate, using Lemma 1(b) to control $|v(\bar{c}) - v(\mathbb{E}[c])|$ via $\bar{c} \xrightarrow{P} \mathbb{E}[c]$.

Proof. Since v is D-Lipschitz on $\mathcal{C}$ (Lemma 1(b)) and $\mathcal{C}$ is bounded by B, $|v(c)| \le D\|c\| \le DB$ for all $c \in \mathcal{C}$ (taking $v(0) = 0$ WLOG by translation), each term $v(c_i)$ is thus bounded in $[-DB, DB]$. Hoeffding's inequality applied to the i.i.d. bounded random variables $v(c_i)$ gives the stated bound directly; this argument uses only the Lipschitz property of v established in Lemma 1 and *not* any regularity of π^*, which is what allows it to apply uniformly across LPs, QPs, and combinatorial programs. □

4.9 Asymptotic Normality

Because v is concave, it is differentiable Lebesgue-almost everywhere, but not necessarily at any particular point $\mu = \mathbb{E}[c]$ chosen in advance (e.g. Example 1 has v non-differentiable at $c = 0$). A central limit theorem for the regret estimator therefore requires an explicit non-degeneracy assumption at μ, which we state rather than assume away.

Assumption 5 (Non-degeneracy at the mean)**.** *v is differentiable at $\mu = \mathbb{E}[c]$, with gradient $\nabla v(\mu) = \pi^*(\mu)$ (equivalently: $\pi^*(\mu)$ is the unique optimizer at $c = \mu$, so μ lies in the interior of a single normal cone of Z rather than on a boundary between two or more).*

Theorem 5 (CLT for the regret estimator)**.** *Under Assumptions 1–4 and 5, with $\Sigma_c = \mathrm{Cov}(c)$ finite,*

$$\sqrt{n}\left(\widehat{\mathrm{Regret}}_n - \mathrm{Regret}(c)\right) \;\xrightarrow{d}\; \mathcal{N}\big(0,\; \pi^*(\mu)^\top \Sigma_c\, \pi^*(\mu) + \mathrm{Var}(v(c)) - 2\,\mathrm{Cov}(v(c),\, c^\top \pi^*(\mu))\big).$$

Proof. $\widehat{\mathrm{Regret}}_n = v(\bar{c}) - \widehat{v}_n$. By the delta method, using Assumption 5 to justify the differentiability of v at μ and $\sqrt{n}(\bar{c} - \mu) \xrightarrow{d} \mathcal{N}(0, \Sigma_c)$ (multivariate CLT, finite variance from Assumption 2),

$$\sqrt{n}\big(v(\bar{c}) - v(\mu)\big) \xrightarrow{d} \nabla v(\mu)^\top \mathcal{N}(0, \Sigma_c) = \mathcal{N}\big(0,\; \pi^*(\mu)^\top \Sigma_c\, \pi^*(\mu)\big).$$

Table 1: Per-experiment deltas from the shared methodology of Section 5.1.

Experiment	Problem & parameters	Data source
LP (5.2)	$\min c^\top z$ s.t. $Az \le b$, $z \ge 0$; $A \sim \mathcal{N}(0, I)$, $b \sim \lvert\mathcal{N}(0, I)\rvert + 1$; 20 vars, 15 constraints; 5,000 iterations	Synthetic
QP / Portfolio (5.3)	$\min c^\top z + \frac{\lambda}{2} z^\top \Sigma_c z$, risk-neutral $\lambda = 1$; 20 vars, 15 constraints; 5,000 iterations	Synthetic
Real-data portfolio (B.4)	Markowitz portfolios, 100 draws of 50 stocks/month, rolling-window forecast vs. realized regret	CRSP monthly returns, Jan 2015–Dec 2024, 4,487 stocks, 119 months

Separately, $\sqrt{n}(\widehat{v}_n - \mathbb{E}[v(c)]) \xrightarrow{d} \mathcal{N}(0, \mathrm{Var}(v(c)))$ by the ordinary CLT, since $v(c_i)$ are i.i.d. with finite variance (boundedness under compact $\mathcal{C}$, or finite second moments more generally). The two terms are asymptotically jointly normal (both are smooth functionals of the same i.i.d. sample $\{c_i\}$), so their difference converges to a normal law whose variance includes the stated covariance term. The covariance term follows from a joint delta-method/CLT expansion and is stated here in closed form. □

Remark 4. *Theorem 5 is silent, by construction, at cost distributions whose mean sits on a normal-cone boundary (Assumption 5 fails at such points, as in Example 1). At such points v genuinely kinks, and $\sqrt{n}(v(\bar{c}) - v(\mu))$ does not converge to a normal limit at all (it converges to a one-sided or folded distribution governed by the local structure of the kink). Practitioners should check Assumption 5 against their specific (Z, μ) before invoking confidence intervals built on Theorem 5.*

5 Experiments

5.1 Shared Methodology

All three experiments follow the same protocol: fix a mean cost vector $\mu_c \sim \mathcal{N}(0, I)$ and a randomly generated positive-definite covariance matrix Σ_c, draw a distinct cost vector $c \sim \mathcal{N}(\mu_c, \Sigma_c)$ on each iteration, solve the corresponding optimization problem via `cvxpy` (Diamond and Boyd, 2016) to obtain $\hat{z}(c)$, compute empirical regret $\widehat{\mathrm{Regret}} = c^\top \hat{z}(c) - c^\top \hat{z}(\mathbb{E}[c])$, and compare it against the theoretical prediction $\mathrm{Cov}(c, z) = \mathbb{E}[(c - \mathbb{E}[c])(z - \mathbb{E}[z])]$ accumulated across iterations. Table 1 lists what differs across the three experiments; results are reported in Sections 5.2–5.3.

5.2 Linear Programming

Figure 1a shows theoretical regret $\mathrm{Cov}(c, z)$ converging quickly to -700.982, while empirical regret converges more slowly to -605.882 (13.57% error).

5.3 Quadratic Optimization

Figure 1b shows the theoretical regret converging to -7.729; empirical regret converges to -5.332, with a relative estimation error of 31%.

6 Tightness of the Distribution-Free Bound at Scale

Does the bound stay usefully tight as the combinatorial problem grows, or does it degrade? In this section, we test a monotone-lattice grid shortest-path problem at increasing scale, from 5×5 to 30×30, where $D = \sqrt{2(k-1)}$ is exact and the true regret is computed via dynamic program-

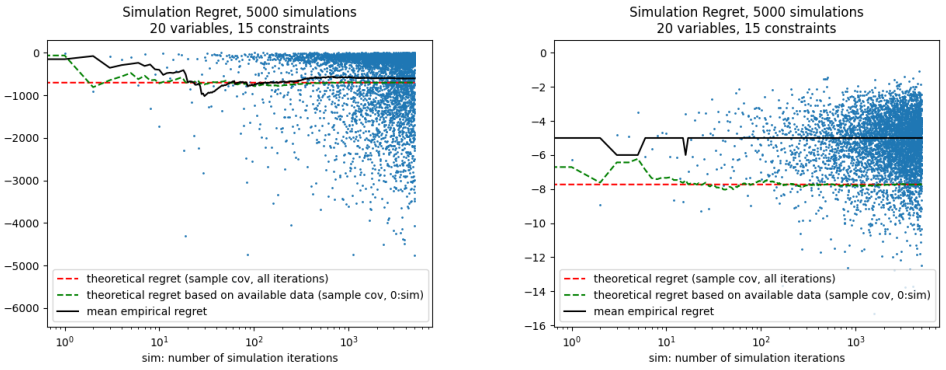

(a) LP Regret

(b) QP Regret

ming and compared against Theorem 2's bound. Construction and methodology details are given in Appendix C.1.

Grid size k	5	10	15	20	25	30
$D = \sqrt{2(k-1)}$	2.83	4.24	5.29	6.16	6.93	7.62
$\widehat{\text{Regret}}$	2.67	6.80	11.08	15.40	19.76	24.10
Bound $D\sqrt{\text{tr}(\Sigma_c)}$	10.81	34.41	65.53	102.65	144.95	191.84
Bound / Regret	**4.05**	**5.06**	**5.91**	**6.67**	**7.34**	**7.96**

Table 2: Distribution-free bound versus true regret across grid sizes 5×5 through 30×30, monotone-lattice shortest path with i.i.d. log-normal edge costs. The bound is valid at every scale (it never falls below the true regret, as Theorem 2 guarantees), but its looseness grows with problem size: the referee's concern is confirmed quantitatively, not just qualitatively.

The bound is valid at every scale, as Theorem 2 guarantees, but its looseness grows with problem size: the bound-to-regret ratio rises from $4.05\times$ at $k = 5$ to $7.96\times$ at $k = 30$ (Table 2). The mechanism behind this growth and a caveat specific to the 8-connected Warcraft benchmark are given in Appendix C.1.

7 Conclusion

We set out to ask when regret in stochastic optimization can be computed from moments alone rather than by repeated simulation. Our answer is more structured than a single formula: it depends on whether the optimal decision rule π^* is affine in the cost vector. The affine property holds for unconstrained and equality-constrained quadratic programs and fails generically once a program has active inequality constraints or is a linear program, where π^* is only piecewise affine. This is a statement about the geometry of the decision rule, not about the smoothness of the cost distribution: absolute continuity of P_c, which might be expected to smooth away the residual, does not do so, and Example 1 makes the failure concrete with a closed-form instance rather than an abstract possibility.

Two findings follow. First, the concave value function $v(c) = \min_{z\in Z} c^\top z$, which is Lipschitz unconditionally on a compact feasible region, establishes the distribution-free bound. Such $v(c)$ also allows us to prove the concentration inequality and, via the envelope theorem, the central limit theorem without assuming away the very discontinuity of π^* that makes linear and combinatorial programs interesting. Second, the computational benefit of the cases is real: evaluating the residual estimator costs one solve to obtain $\pi^*(\bar{c})$, and the genuine saving relative to Sample Average Approximation is $O(1)$ solves against $O(n)$, not $O(nd^2)$ against $O(Bn^2d^3)$, regardless of n.

References

Agrawal, Akshay, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J. Zico

Kolter, "Differentiable Convex Optimization Layers," in "Proceedings of the 33rd International Conference on Neural Information Processing Systems" Curran Associates Inc. Red Hook, NY, USA 2019.

Almgren, Robert and Neil Chriss, "Optimal Execution of Portfolio Transactions," *Journal of Risk*, 2001, *3*, 5–40.

Amos, Brandon and J. Zico Kolter, "OptNet: Differentiable Optimization as a Layer in Neural Networks," in "Proceedings of the 34th International Conference on Machine Learning" 2017, pp. 136–145.

Anderson, Brian D. O. and John B. Moore, *Optimal Control: Linear Quadratic Methods*, Englewood Cliffs, NJ: Prentice Hall, 1990.

Ban, Gah-Yi and Cynthia Rudin, "The Big Data Newsvendor: Practical Insights from Machine Learning," *Operations Research*, 2019, *67* (1), 90–108.

Bank, Bernd, Jürgen Guddat, Diethard Klatte, Bernd Kummer, and Klaus Tammer, *Non-Linear Parametric Optimization*, Berlin: Akademie-Verlag, 1982. Also published by Birkhäuser, Basel-Boston, 1983.

Bayraksan, Güzin and David K Love, "Separability and decomposition in SAA for stochastic mixed-integer programming," *Operations Research Letters*, 2015, *43* (1), 109–113.

Berthet, Quentin, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach, "Learning with Differentiable Perturbed Optimizers," in "Proceedings of the 34th International Conference on Neural Information Processing Systems" NIPS'20 Curran Associates Inc. Red Hook, NY, USA 2020.

Best, Michael J., *Quadratic Programming with Computer Programs*, Boca Raton, FL: CRC Press, 2017.

Birge, John R. and François Louveaux, *Introduction to Stochastic Programming*, 2nd ed., New York, NY: Springer, 2011.

Bonnans, J. Frédéric and Alexander Shapiro, *Perturbation Analysis of Optimization Problems*, Springer Science & Business Media, 2000.

Borch, Karl, "Equilibrium in a Reinsurance Market," *Econometrica*, 1962, *30* (3), 424–444.

Boyd, Stephen and Lieven Vandenberghe, *Convex Optimization*, Cambridge, UK: Cambridge University Press, 2004.

Diamond, Steven and Stephen Boyd, "CVXPY: A Python-Embedded Modeling Language for Convex Optimization," *Journal of Machine Learning Research*, 2016, *17* (83), 1–5.

Donti, Priya, Brandon Amos, and J. Zico Kolter, "Task-Based End-to-End Model Learning in Stochastic Optimization," *Advances in Neural Information Processing Systems*, 2017, pp. 5484–5494.

Efron, Bradley and Robert J. Tibshirani, *An Introduction to the Bootstrap*, Boca Raton, FL: CRC Press, 1994.

Elmachtoub, Adam N. and Paul Grigas, "Smart "Predict, then Optimize"," *Management Science*, 2022, *68* (1), 9–26.

Fiacco, Anthony V., *Introduction to Sensitivity and Stability Analysis in Nonlinear Programming*, New York, NY: Academic Press, 1983.

Geng, Haoyu, Hang Ruan, Runzhong Wang, Yang Li, Yang Wang, Lei Chen, and Junchi Yan, "Benchmarking PtO and PnO Methods in the Predictive Combinatorial Optimization Regime," in "Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track," Vol. 37 2024.

Hoerl, Arthur E. and Robert W. Kennard, "Ridge Regression: Biased Estimation for Nonorthogonal Problems," *Technometrics*, 1970, *12* (1), 55–67.

Holt, Charles C., Franco Modigliani, John F. Muth, and Herbert A. Simon, *Planning Production, Inventories, and Work Force*, Englewood Cliffs, NJ: Prentice-Hall, 1960.

Howard, Ronald A., "Information Value Theory," *IEEE Transactions on Systems Science and Cybernetics*, 1966, *2* (1), 22–26.

Imbens, Guido W. and Donald B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge, UK: Cambridge University Press, 2015.

Kalman, Rudolph E., "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, 1960, *82* (1), 35–45.

King, Alan J. and Stein W. Wallace, *Modeling with Stochastic Programming*, Springer Science & Business Media, 2012.

Kleywegt, Anton J., Alexander Shapiro, and Tito Homem de Mello, "The Sample Average Approximation Method for Stochastic Discrete Optimization," *SIAM Journal on Optimization*, 2002, *12* (2), 479–502.

Ledoit, Olivier and Michael Wolf, "Honey, I Shrunk the Sample Covariance Matrix," *The Journal of Portfolio Management*, 2004, *30* (4), 110–119.

Mandi, Jayanta and Tias Guns, "Interior point solving for LP-based prediction+ optimisation," in "Advances in Neural Information Processing Systems," Vol. 33 2020, pp. 7272–7282.

_ , Emir Demirovic, Peter J. Stuckey, and Tias Guns, "Smart Predict-and-Optimize for Hard Combinatorial Optimization Problems," in "Proceedings of the AAAI Conference on Artificial Intelligence" AAAI Press 2020, pp. 1603–1610.

Markowitz, Harry, "Portfolio Selection," *The Journal of Finance*, 1952, *7*, 77–91.

Milgrom, Paul and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 2002, *70* (2), 583–601.

Nocedal, Jorge and Stephen J. Wright, *Numerical Optimization*, 2nd ed., New York, NY: Springer, 2006.

Peyré, Gabriel and Marco Cuturi, "Computational Optimal Transport," *Foundations and Trends in Machine Learning*, 2019, *11* (5-6), 355–607.

Raiffa, Howard and Robert Schlaifer, *Applied Statistical Decision Theory*, Boston, MA: Division of Research, Graduate School of Business Administration, Harvard University, 1961.

Rubin, Donald B., "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 1974, *66* (5), 688–701.

Sadana, Utsav, 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*, 2023, *320* (2), 271–289.

_ , _ , _ , _ , _ , and _ , "A Survey of Contextual Optimization Methods for Decision Making under Uncertainty," *European Journal of Operational Research*, 2023, *320* (2), 271–289.

Schrijver, Alexander, *Combinatorial Optimization: Polyhedra and Efficiency*, Berlin: Springer, 2003.

Shapiro, Alexander, Darinka Dentcheva, and Andrzej Ruszczynski, *Lectures on Stochastic Programming: Modeling and Theory*, SIAM, 2009.

Simon, Herbert A., "Dynamic Programming under Uncertainty with a Quadratic Criterion Function," *Econometrica*, 1956, *24* (1), 74–81.

Theil, Henri, "A Note on Certainty Equivalence in Dynamic Planning," *Econometrica*, 1957, *25* (2), 346–349.

Tibshirani, Robert, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society: Series B*, 1996, *58* (1), 267–288.

Verma, Sahil, John Dickerson, and Keegan Hines, "Counterfactual Explanations for Machine Learning: A Review," *arXiv preprint arXiv:2010.10596*, 2020.

Vlastelica, Marin, Anselm Paulus, Vít Musil, Georg Martius, and Michal Rolínek, "Differentiation of Blackbox Combinatorial Solvers," in "International Conference on Learning Representations (ICLR)" 2020.

Wachter, Sandra, Brent Mittelstadt, and Chris Russell, "Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR," *Harvard Journal of Law & Technology*, 2017, *31* (2), 841–887.

Wilder, Bryan, Bistra Dilkina, and Milind Tambe, "Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization," in "Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence" AAAI'19/IAAI'19/EAAI'19 AAAI Press 2019.

Zinkevich, Martin, Michael Johanson, Michael Bowling, and Carmelo Piccione, "Regret Minimization in Games with Incomplete Information," in "Advances in Neural Information Processing Systems," Vol. 20 2007, pp. 1729–1736.

A Practical Examples for Known Optimization Problems that Have Compact Z and Finite-variance P_c

Example 2 (Portfolio selection, budget only)**.** *Ten equities with annualized volatilities drawn realistically between* 15% *and* 35% *and a moderate pairwise return correlation of* 0.3*. Z is the* 10*-asset simplex, so* $D = 1$*. With monthly returns,* $\mathrm{tr}(\Sigma_c) \approx 0.063$*, giving* $\mathrm{Regret} \leq D\sqrt{\mathrm{tr}(\Sigma_c)} \approx 0.251$*, i.e. a monthly-return bound of about* 25 *cents per dollar invested on the value of knowing next month's realized returns in advance. This is a large number in absolute terms, expected for a bound using only two moments and the simplex radius. This is also an instance-specific upper bound computable without any Monte Carlo simulation.*

Example 3 (Portfolio selection, position limits)**.** *The same ten equities, now capped at a maximum* 20% *weight per asset (a common practical diversification rule). $Z = \{z : \mathbf{1}^\top z = 1,\ 0 \leq z \leq 0.2\}$ is still compact, and in fact, has a* smaller *radius than the unconstrained simplex whenever the cap binds ($D \leq \sqrt{10 \cdot 0.2^2} = 0.632$ here, tighter than Example 2's $D = 1$), so the position limit both changes the problem (moving it from Corollary 2's case toward Corollary 3's generic case, since the cap is an inequality constraint) and mechanically tightens Theorem 2's bound at the same time.*

Example 4 (Assignment problem)**.** *Five delivery drivers are assigned to five delivery zones, with expected completion times between* 2 *and* 8 *hours per (driver, zone) pair and independent variability of* 0.5*-*1.5 *hours per pair (traffic, familiarity with the route). Z is the* 5×5 *Birkhoff polytope, so $D = \sqrt{5} \approx 2.236$ (the Frobenius norm of any permutation matrix). With* $\mathrm{tr}(\Sigma_c) \approx 24.9$ *(independent variances summed across all* 25 *pairs), there are* $\mathrm{Regret} \leq D\sqrt{\mathrm{tr}(\Sigma_c)} \approx 11.15$ *driver-hours. This is the expected value of assigning drivers only after the day's actual traffic conditions, as opposed to based on average conditions.*

Example 5 (Shortest path)**.** *The standard* 5×5 *grid shortest-path benchmark used throughout the decision-focused learning literature (Section 2):* 40 *directed edges, a path from the bottom-left to the top-right corner using* 8 *edges regardless of route, so $D = 8$. With per edge travel times of* 3*-*10 *minutes and independent variability of* 0.5*-*2 *minutes per edge (typical city-block conditions),* $\mathrm{tr}(\Sigma_c) \approx 68.1$*, giving* $\mathrm{Regret} \leq D\sqrt{\mathrm{tr}(\Sigma_c)} \approx 66.0$ *minutes. This bound is intentionally loose relative to the true regret for the same reason Example 1 is instructive: it uses only the number of edges on a path and the total edge-time variance, not the grid's specific normal-fan structure, which is the gap that a sharper LP-specific bound (Section 7's first proposed extension) would close.*

Example 6 (Minimum spanning tree)**.** *Connecting* 6 *facilities with the cheapest possible network of dedicated links: any spanning tree on* 6 *nodes has* 5 *edges, so* $D = \sqrt{5} \approx 2.236$ *regardless of which* 5 *of the* $\binom{6}{2} = 15$ *possible links are chosen, giving the same kind of immediate, combinatorics-only bound on* D *as the assignment and shortest-path cases above, with no need to enumerate the (here, by Cayley's formula,* $6^4 = 1296$*) spanning trees themselves.*

Example 7 (Traveling salesman)**.** *A delivery route visiting* 6 *cities and returning to the depot: every tour uses* 6 *edges, so* $D = \sqrt{6} \approx 2.449$*, again immediate from the combinatorial structure of* Z *alone and requiring no enumeration of the* $5!/2 = 60$ *distinct tours.*

Example 8 (Knapsack)**.** *Consider the following cargo-loading problem:* 20 *items, capacity* 30*,* $D = \sqrt{7}$ *(the true maximum number of items that can ever be loaded simultaneously, computed by sorting item weights), distribution-free bound* ≈ 44.35*, against a Monte-Carlo-computed regret of* 9.58 ± 0.065*. This is the same qualitative pattern as in Example 5. A valid but loose bound along with a much tighter bound are available only by direct computation because the optimal decision here is combinatorial rather than affine.*

Example 9 (Optimal transport)**.** *Three warehouses with supplies* $(40, 35, 25)$ *units serving four stores with demands* $(20, 30, 15, 35)$ *units (balanced, summing to* 100 *on both sides). Any feasible transportation plan satisfies* $\sum_{ij} z_{ij} = 100$ *with* $z_{ij} \geq 0$*, so* $D = \sup \|z\| \leq 100$ *trivially (achieved only by a degenerate single-route plan) and considerably less in practice once the marginal constraints are accounted for; as in Examples 5–8, refining* D *using the specific supply/demand marginals rather than the crude total-flow bound would tighten Theorem 2's result substantially here too.*

Example 10 (LASSO-constrained regression)**.** *Predicting house prices from* 10 *standardized features (square footage, age, distance to amenities, etc.), with the coefficient vector constrained to an* ℓ_1 *ball of radius* $t = 5$ *(a typical cross-validated sparsity level).* $D = t = 5$ *immediately, since* $\|z\|_1 \leq t \Rightarrow \|z\|_2 \leq t$*. With per-feature cost-gradient variance between* 0.8 *and* 1.2*,* $\mathrm{tr}(\Sigma_c) \approx 9.82$*, giving* $\mathrm{Regret} \leq D\sqrt{\mathrm{tr}(\Sigma_c)} \approx 15.67$ *in the units of the regression's loss function — here* D *is fixed by the regularization budget the analyst chooses, not by the problem's physical structure, so tightening this bound is a matter of choosing* t*, not of a sharper geometric argument.*

Example 11 (Trust-region / ridge-constrained estimation)**.** *The same regression problem with an* ℓ_2 *ball constraint* $\|z\|_2 \leq t = 5$ *instead:* $D = t = 5$ *immediately, as in Example 10, since the constraint directly bounds the Euclidean norm rather than implying a bound on it.*

Example 12 (Newsvendor with a bounded order quantity)**.** *As an example illustrating that the compactness of* Z *and the affinity of* π^* *are independent properties, consider a bakery that decides how many loaves to bake before observing the day's demand, capped at* $\bar{z} = 200$ *loaves by oven capacity.* $Z = [0, 200]$ *gives* $D = 200$ *immediately. This problem's piecewise-*linear *holding/shortage cost structure creates a useful boundary case.* Z *is compact and* P_c *(daily demand) has finite variance, so Theorem 2's bound applies as in every other row of this table. However, the optimal order-up-to quantity itself is a demand quantile rather than an affine function of the mean, so the bakery problem belongs in the compact-*Z *table (Table 3). It is explicitly excluded from the affine-decision-rule table (Table 4).*

Problem	Feasible region Z	Why compact	Typical P_c
Portfolio selection (budget + no short sales)	$\{z : \mathbf{1}^\top z = 1,\, z \geq 0\}$	bounded simplex	Gaussian / elliptical returns
Portfolio with position limits	$\{z : \mathbf{1}^\top z = 1,\, l \leq z \leq u\}$	bounded box $\cap$ hyperplane	Gaussian / elliptical returns
Assignment / bipartite matching	Birkhoff polytope (doubly stochastic matrices)	convex hull of finitely many permutation matrices (Schrijver, 2003)	log-normal or Gamma costs
Shortest path on a fixed graph	path polytope of a DAG or capacitated network	convex hull of finitely many path/flow vectors (Schrijver, 2003)	log-normal or Gamma travel times
Spanning tree, matroid problems	matroid base polytope	convex hull of finitely many bases (Schrijver, 2003)	Gaussian or Gamma edge costs
Traveling salesman	TSP polytope (convex hull of tour incidence vectors)	convex hull of finitely many tours (Schrijver, 2003)	log-normal travel costs
Knapsack (fractional or 0-1 relaxation)	$\{z \in [0,1]^n : w^\top z \leq \text{capacity}\}$	bounded box $\cap$ halfspace	Gamma or truncated-normal item values
Optimal transport	transportation polytope with fixed marginals	bounded by marginal constraints (Peyré and Cuturi, 2019)	Gaussian or log-normal unit costs
Sparse regression (LASSO-constrained)	$\{z : \|z\|_1 \leq t\}$	ℓ_1 ball, bounded by construction (Tibshirani, 1996)	Gaussian design/noise
Trust-region / ridge-constrained estimation	$\{z : \|z\|_2 \leq t\}$	Euclidean ball, bounded by construction	Gaussian
Newsvendor with bounded order quantity	$\{z : 0 \leq z \leq \bar{z}\}$	bounded interval	Poisson / negative binomial demand

Table 3: Representative problems satisfying both clauses of Assumptions 3 and 2 jointly, so that Theorem 2's bound $\text{Regret} \leq D\sqrt{\text{tr}(\Sigma_c)}$ and Theorem 4's concentration inequality apply without modification. The first five rows are compact for the same underlying reason: Z is the convex hull of a finite set, which is automatically bounded and closed regardless of how many combinatorial constraints define it. This table covers essentially all discrete and combinatorial optimization problems with a fixed, finite ground set. Popular practical examples are listed in the Appendix.

Problem / application	Why π^* is affine	Reference
(A) Direct instances of this paper's static framework		
Markowitz portfolio, budget constraint only	Corollary 2: equality-constrained QP	Markowitz (1952)
Ridge regression / Tikhonov regularization	$\hat{\beta}(y) = (X^\top X + \lambda I)^{-1} X^\top y$; an instance of Corollary 1 with $c = -2X^\top y$	Hoerl and Kennard (1970)
Generalized least squares	solution is linear in the response vector by the same normal-equation argument	—
Minimum-variance hedging with a linear budget constraint	equality-constrained QP in the hedge ratios	—
(B) The same mechanism in dynamic/control settings (outside this paper's static formalism)		
Linear-Quadratic Regulator (LQR)	optimal control law $u^* = -Kx$ is linear in the state, via the Riccati equation	Anderson and Moore (1990)
Kalman filter	optimal state estimate is affine in the observation sequence under Gaussian linear dynamics	Kalman (1960)
Linear-Quadratic-Gaussian (LQG) control	separation principle: affine estimator (Kalman) composed with affine control law (LQR)	Anderson and Moore (1990)
HMMS production smoothing	optimal production/workforce decision is affine in the demand forecast under quadratic adjustment costs	Holt et al. (1960)
Multi-period certainty-equivalent planning generally	optimal decision under a quadratic criterion and linear dynamics equals the decision computed by replacing all uncertain quantities with their means	Simon (1956); Theil (1957)
Optimal execution (Almgren–Chriss)	optimal trading trajectory is linear in the initial position under quadratic transaction costs	Almgren and Chriss (2001)
Linear risk-sharing / reinsurance contracts	Pareto-optimal sharing rule is affine in aggregate risk under quadratic (mean-variance) preferences	Borch (1962)

Table 4: Known problems with affine optimal decision rules. Column (A) falls directly under this paper's Theorem 3(a); column (B) is governed by the same first-order-condition mechanism but in a dynamic or state-space formalism this paper's static, single-period framework does not cover, and is included to situate the result within a much older literature rather than to claim it as new.

B Application to Quadratic Programs, Detailed Derivation

Setup. Throughout, consider

$$\min_{z \in Z} \; c^\top z + \frac{\lambda}{2} z^\top Q z, \qquad Q \succ 0, \; \lambda > 0,$$

with $Z \subseteq \mathbb{R}^n$ ranging over $\mathbb{R}^n$ (unconstrained), an affine subspace $\{z : Az = b\}$ (equality-constrained), or a general polyhedron $\{z : Az = b, \; Gz \leq h\}$ (inequality-constrained).

B.1 Unconstrained QP

Corollary 1 (Unconstrained QP). *For $Z = \mathbb{R}^n$, the first-order condition $\nabla_z\big(c^\top z + \frac{\lambda}{2} z^\top Q z\big) = c + \lambda Q z = 0$ gives the unique minimizer*

$$\pi^*(c) = -\tfrac{1}{\lambda} Q^{-1} c,$$

This π^* is linear in c, so by Theorem 3(a), $R(c) = 0$ for every distribution of c supported on $\mathbb{R}^n$, and

$$\mathrm{Regret}(c) = -\mathrm{Cov}\big(c, \pi^*(c)\big) = \tfrac{1}{\lambda} \operatorname{tr}\big(Q^{-1}\Sigma_c\big) \; \geq 0,$$

using $\mathrm{Cov}(c, Mc) = \mathrm{tr}(M\Sigma_c)$ *for any fixed matrix* M *(Lemma 3 below) with* $M = -\frac{1}{\lambda}Q^{-1}$*. No compactness of* Z *is needed here:* $v(c) = -\frac{1}{\lambda}c^\top Q^{-1}c$ *is a smooth concave quadratic on all of* $\mathbb{R}^n$*, so Jensen's inequality (Theorem 1) applies directly.*

Lemma 3. *For any fixed matrix* $M \in \mathbb{R}^{n\times d}$ *and random vector* $c \in \mathbb{R}^d$ *with* $\mathbb{E}[\|c\|^2] < \infty$*,* $\mathrm{Cov}(c, Mc) := \mathbb{E}[c^\top Mc] - \mathbb{E}[c]^\top M\mathbb{E}[c] = \mathrm{tr}(M\Sigma_c)$*.*

Proof. Write $\mu = \mathbb{E}[c]$. Since $c^\top Mc = \mathrm{tr}(Mcc^\top)$, the linearity of trace and expectation gives $\mathbb{E}[c^\top Mc] = \mathrm{tr}\big(M\,\mathbb{E}[cc^\top]\big) = \mathrm{tr}\big(M(\Sigma_c + \mu\mu^\top)\big) = \mathrm{tr}(M\Sigma_c) + \mu^\top M\mu$. Subtracting $\mu^\top M\mu$ leaves $\mathrm{tr}(M\Sigma_c)$. □

B.2 Equality-constrained QP: the Markowitz case

Markowitz mean-variance portfolio selection imposes at least a budget constraint $\mathbf{1}^\top z = 1$ (Markowitz, 1952), so the relevant case is $Z = \{z : Az = b\}$ for a full-row-rank matrix $A \in \mathbb{R}^{m\times n}$, $m < n$, not the unconstrained case of Corollary 1. This case turns out to be as favorable as the unconstrained one, which the original manuscript did not notice: *equality* constraints do not break the affine structure of π^*, because they enter the KKT system linearly and without case-based branching. Only *inequality* constraints, treated in Section B.3, introduce genuine nonlinearity.

Corollary 2 (Equality-constrained QP)**.** *For* $Z = \{z \in \mathbb{R}^n : Az = b\}$ *with* A *full row rank and* $Q \succ 0$*, the Lagrangian stationarity condition* $c + \lambda Qz + A^\top\nu = 0$*, together with* $Az = b$*, has a unique solution (Nocedal and Wright, 2006; Boyd and Vandenberghe, 2004)*

$$\pi^*(c) = Mc + m_0, \qquad M := -\tfrac{1}{\lambda}Q^{-1}\Big(I - A^\top(AQ^{-1}A^\top)^{-1}AQ^{-1}\Big), \qquad m_0 := Q^{-1}A^\top(AQ^{-1}A^\top)^{-1}b.$$

This is affine in c *(setting* $A = 0$ *recovers Corollary 1), so* $R(c) = 0$ *for every distribution of* c *by Theorem 3(a), and*

$$\mathrm{Regret}(c) = -\mathrm{Cov}(c, \pi^*(c)) = -\mathrm{tr}(M\Sigma_c) \geq 0$$

by Lemma 3 (nonnegativity follows from Theorem 1, since $Regret \geq 0$ *always). In particular,* ***Markowitz's budget constraint does not by itself introduce any residual****.*

Proof. Form the Lagrangian $L(z, \nu) = c^\top z + \frac{\lambda}{2}z^\top Qz + \nu^\top(Az - b)$. Stationarity in z gives $c + \lambda Qz + A^\top\nu = 0$, i.e. $z = -\frac{1}{\lambda}Q^{-1}(c + A^\top\nu)$. Substituting into $Az = b$:

$$-\tfrac{1}{\lambda}AQ^{-1}c - \tfrac{1}{\lambda}AQ^{-1}A^\top\nu = b \implies \nu = -(AQ^{-1}A^\top)^{-1}\big(\lambda b + AQ^{-1}c\big),$$

using that $AQ^{-1}A^\top \succ 0$ is invertible since A has full row rank and $Q \succ 0$. Substituting back,

$$z^*(c) = -\tfrac{1}{\lambda}Q^{-1}c + Q^{-1}A^\top(AQ^{-1}A^\top)^{-1}b + \tfrac{1}{\lambda}Q^{-1}A^\top(AQ^{-1}A^\top)^{-1}AQ^{-1}c,$$

which is $Mc + m_0$ as stated, after collecting the c-linear terms into M. □

B.3 Inequality-constrained QP: piecewise-affine structure

Lemma 4 (The constrained QP solution is a Q-weighted projection)**.** *Let* $z_{\mathrm{uc}}(c) := -\frac{1}{\lambda}Q^{-1}c$ *denote the unconstrained minimizer, and let* Z *be any nonempty closed convex set. Then*

$$\pi^*(c) = \arg\min_{z\in Z}\,(z - z_{\mathrm{uc}}(c))^\top Q\,(z - z_{\mathrm{uc}}(c)) \;=:\; \mathrm{Proj}_Z^Q\big(z_{\mathrm{uc}}(c)\big),$$

the projection of $z_{\mathrm{uc}}(c)$ *onto* Z *in the* Q*-weighted inner product* $\langle x, y\rangle_Q := x^\top Qy$*. This coincides with the Euclidean projection* P_Z *only when* Q *is a multiple of the identity.*

Proof. Complete the square using Q symmetric and $z_{\mathrm{uc}}^\top Q = -\frac{1}{\lambda}c^\top Q^{-1}Q = -\frac{1}{\lambda}c^\top$:

$$\begin{aligned}\tfrac{\lambda}{2}(z - z_{\mathrm{uc}})^\top Q(z - z_{\mathrm{uc}}) &= \tfrac{\lambda}{2}z^\top Qz - \lambda z_{\mathrm{uc}}^\top Qz + \tfrac{\lambda}{2}z_{\mathrm{uc}}^\top Qz_{\mathrm{uc}} \\ &= \tfrac{\lambda}{2}z^\top Qz + c^\top z + \tfrac{\lambda}{2}z_{\mathrm{uc}}^\top Qz_{\mathrm{uc}}.\end{aligned}$$

The last term does not depend on z, so minimizing $c^\top z + \frac{\lambda}{2}z^\top Qz$ over $z \in Z$ is equivalent to minimizing $(z - z_{\mathrm{uc}})^\top Q(z - z_{\mathrm{uc}})$ over $z \in Z$, which is, by definition, the Q-weighted projection of z_{uc} onto Z. □

Proposition 2 (Piecewise-affine structure)**.** *Let $Z = \{z : Az = b,\ Gz \le h\}$ and suppose $Q \succ 0$. For a cost vector c at which the active-set solution satisfies the linear independence constraint qualification (LICQ) and strict complementarity, let $I(c) := \{i : G_i z^*(c) = h_i\}$ denote the active inequality constraints. Then there is a neighborhood U of c on which $I(\cdot)$ is constant, and on U,*

$$\pi^*(c') = M_{I(c)}\, c' + m_{I(c)}, \qquad c' \in U,$$

where $M_{I(c)}, m_{I(c)}$ are obtained by applying Corollary 2 to the equality-constrained QP with constraint matrix $\left[\begin{smallmatrix} A \\ G_{I(c)} \end{smallmatrix}\right]$ in place of A. Consequently, π^* is a piecewise-affine *function of c on* $\mathrm{conv}(\mathcal{C})$*, affine on each region $\mathcal{C}_I := \{c : I(c) = I\}$ and generally discontinuous across region boundaries (Fiacco, 1983; Bank et al., 1982; Best, 2017).*

Proof sketch. Under LICQ and strict complementarity at c, the active set is locally stable under perturbation of c (standard parametric-programming stability result; see Fiacco, 1983 Ch. 2 and Bank et al., 1982 Ch. 3, or the implicit-function-theorem treatment in Bonnans and Shapiro, 2000). On the resulting neighborhood, the active inequality constraints $G_{I(c)}z = h_{I(c)}$ can be treated as equality constraints, reducing the problem to Corollary 2 with constraint matrix $[A; G_{I(c)}]$. Piecewise affinity of π^* over a polyhedral partition of $\mathrm{conv}(\mathcal{C})$ into active set regions follows because there are finitely many possible active sets. □

Corollary 3 ($R(c)$ is generically nonzero for inequality-constrained QPs)**.** *By Proposition 2, π^* is affine on each region $\mathcal{C}_I$ but not, in general, globally affine on* $\mathrm{conv}(\mathcal{C})$*. As in the LP case (Example 1, which is the special case $Q = 0$ of this same piecewise-affine structure): if* $\mathrm{conv}(\mathrm{supp}(c))$ *lies within a single region $\mathcal{C}_I$, then $R(c) = 0$, by Theorem 3(a) applied to $\pi^*|_{\mathcal{C}_I}$; if $\mathbb{E}[c] = 0$, then $R(c) = 0$ trivially, by Theorem 3(b); otherwise $R(c)$ is generically nonzero once* $\mathrm{supp}(c)$ *straddles more than one active-set region, and is bounded, whenever Z is additionally compact (e.g. under a no-short-selling constraint $z \ge 0$ together with $\mathbf{1}^\top z = 1$), by the distribution-free bound of Theorem 2.*

QP variant	$\pi^*(c)$	$R(c)$
Unconstrained, $Q \succ 0$	$-\frac{1}{\lambda}Q^{-1}c$ (Cor. 1)	0,
Equality-constrained (Markowitz budget)	$Mc + m_0$ (Cor. 2)	0,
Inequality-constrained (e.g. no short sales)	piecewise affine (Prop. 2)	generically $\neq 0$ (Cor. 3)

Table 5: Summary: Residual for quadratic programs.

B.4 Real-Data Portfolio Validation

Figure 2 shows our metric's performance vis-a-vis a real-life financial portfolio.

We used monthly CRSP data for all U.S. listed securities from January 2015 through December 2024, removing observations with missing returns, returns outside $[-100\%, 1000\%]$, non-positive prices, micro-cap stocks under \$5M market capitalization, and stocks with fewer than 60 months of trading history, leaving 4,487 stocks across 119 monthly observations. Each month, we formed 100 portfolios of 50 stocks drawn with replacement, computed optimal Markowitz decisions, and forecast expected regret using the covariance formula; the following month's realized regret was then compared against this forecast. Figure 2 in the Appendix shows that, in aggregate, empirical regret approaches the theoretical prediction over time.

C Regret vs. Counterfactuals

Remark 5 (Regret as a counterfactual contrast)**.** Regret *(Definition 2) compares two decision* rules *applied to the same random cost. It is naturally read as a counterfactual quantity in three related but distinct senses that are worth separating explicitly, since the word "regret" may have different meaning across the literature streams considered in this paper.*

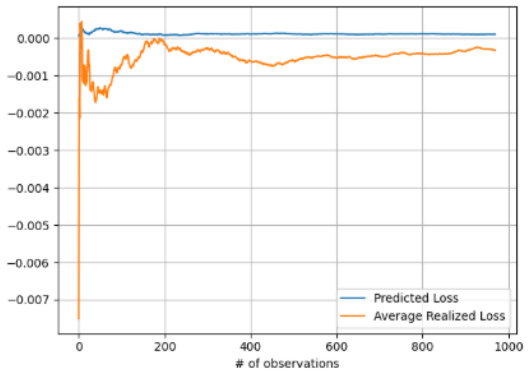

Figure 2: Cov-Predicted vs. Realized Regret in Real-data financial portfolios

Problem class	$R(c) = 0$?	Basis
Unconstrained QP ($Q \succ 0$)	Exact, always	Thm. 3(a), Cor. 1
Any class, if $\mathbb{E}[c] = 0$	Exact, but distribution-specific	Thm. 3(b)
LP, $\mathrm{supp}(c)$ within one normal cone	Exact	Thm. 3(a) (degenerate case)
LP, general position	Generically $R(c) \neq 0$	Example 1
Constrained QP, general	Generically $R(c) \neq 0$	no affine π^*
Integer / combinatorial Z	Generically $R(c) \neq 0$, can be large	piecewise-constant π^*

Table 6: Summary of What Is and Is Not Exact: Residual term $R(c)$. The distribution-free bound of Theorem 2 applies in every row regardless of whether $R(c) = 0$.

(i) An identified potential-outcomes contrast. *In the language of the Rubin causal model (Rubin, 1974; Imbens and Rubin, 2015), define for each realization c two potential outcomes: $Y(1)(c) := c^\top \pi^*(c)$, the cost incurred by the* fully informed *decision rule, and $Y(0)(c) := c^\top \pi^*(\mu)$, the cost incurred by the* certainty-equivalent *decision rule, both evaluated at the same realized c. Since $\pi^*(\mu)$ is a fixed vector, $\mathbb{E}[Y(0)] = \mathbb{E}[c]^\top \pi^*(\mu) = \mu^\top \pi^*(\mu)$, so*

$$\mathbb{E}[Y(0)] - \mathbb{E}[Y(1)] = \mu^\top \pi^*(\mu) - \mathbb{E}[c^\top \pi^*(c)] = \mathrm{Regret}$$

by Definition 2: Regret *is the average treatment effect, on realized cost, of switching from the certainty-equivalent policy to the fully informed one. What makes this an unusually clean instance of a potential-outcomes contrast is that* both $Y(0)(c)$ *and* $Y(1)(c)$ *are computable for every c once (Z, π^*) is specified. There is no missing-data problem, no overlap or unconfoundedness assumption required, and no need for an estimator of the treatment effect itself. This is counterfactual* policy evaluation *within a fully specified optimization model rather than causal effect estimation from partially observed data. This is also the sense in which classical information value theory (Howard, 1966) treats information as valuable: the value of information is the expected improvement from the counterfactual world in which it is known.*

(ii) Not the same object as counterfactual regret in game theory. *The term "counterfactual regret," as used in the extensive-form-game and online-learning literature following Zinkevich et al. (2007), denotes something structurally different: for a player i at an information set I in a repeated or extensive-form game, the counterfactual regret for an action a is the difference in counterfactual value (expected utility weighted by the probability of reaching I under the* other *players' actual strategies, holding player i's own strategy fixed at reaching I) between playing a and following the current strategy, accumulated over repeated play and used algorithmically (via regret matching) to drive strategies toward a Nash equilibrium. This is a multi-agent, sequential, per-information-set quantity used to* compute *an equilibrium.* Regret *in this paper is a single-agent, single-shot, distribution-level quantity used to* measure *the value of information in a fixed decision problem, with no adversarial players, no repeated play, and no algorithmic update rule attached to it. The two notions share only the root intuition common to every use of the word: an actual outcome compared against a counterfactual alternative. However, the two metrics are different and cannot be assumed to share results.*

(iii) A counterfactual explanation of the certainty-equivalent decision. *The introduction's original motivation for this line of work is that decision-replication frameworks can help make algorithmic decisions more interpretable. In the terminology of counterfactual explanations in interpretable machine learning (Wachter et al., 2017),* $\pi^*(c)$ *is a counterfactual explanation of the certainty-equivalent decision* $\pi^*(\mu)$*: it answers "what decision would have been made had the true cost* c *been known instead of only its mean* μ*," and* Regret *quantifies, in expectation over* c*, how much that counterfactual decision would have been worth. The covariance decomposition of Proposition 1 refines this further:* $\mathrm{Cov}(c, \pi^*(c))$ *measures how strongly the counterfactual decision rule's output actually moves with the realized cost, that is, how* responsive *the fully informed policy is. At the same time, the residual* $R(c)$ *captures curvature in that responsiveness left over once the linear (affine) part is accounted for. However, nothing in Sections 4–4.9 treats "responsiveness" as a defined quantity.*

C.1 Why the Distribution-Free Bound Loosens with Scale

We used the same monotone-lattice grid shortest-path family already introduced in Example 5 (right/down moves only, the standard construction used throughout the decision-focused-learning literature, including the 5×5 benchmark of Elmachtoub and Grigas (2022)), and scaled the grid from 5×5 up to 30×30. This specification matches and exceeds the 12×12-to-32×32 range used in the Warcraft experiments of Vlastelica et al. (2020). In this family, every feasible path uses $2(k-1)$ edges for a $k \times k$ grid, so $D = \sqrt{2(k-1)}$ requires no enumeration or approximation to compute. Edge costs were drawn i.i.d. log-normal. For each grid size, we computed the true regret via dynamic programming over $6{,}000$ Monte Carlo draws. The process was fast enough to solve directly rather than approximate even at $k = 30$, since the monotone-lattice shortest path is a simple $O(k^2)$ grid recursion. We then compared it against Theorem 2's bound.

Both D and $\sqrt{\mathrm{tr}(\Sigma_c)}$ grow with k in Table 2's grid family: $D = \Theta(\sqrt{k})$, and $\mathrm{tr}(\Sigma_c) = \Theta(k^2)$ since the number of edges grows quadratically, so $\sqrt{\mathrm{tr}(\Sigma_c)} = \Theta(k)$, giving a bound that grows like $\Theta(k^{1.5})$. The true regret is driven by the covariance between costs and the specific path the optimizer reroutes to under a cost perturbation, not by the total variance of every edge in the grid, and it grows more slowly as a result. The gap is the information that the bound discards: it depends only on the radius of Z and the total cost variance, never on which particular alternative paths or which specific extreme points of the path polytope a given cost perturbation is likely to make optimal. Closing this gap would require a bound that uses the probability mass split across Z's normal fan instead of D and $\mathrm{tr}(\Sigma_c)$ alone, as proposed in Section 6. Table 2 provides quantitative motivation for it rather than a speculative case.

A caveat specific to Warcraft-style benchmarks: the monotone-lattice family tested in Table 2 has an exact, fixed $D = 2(k-1)$, since every feasible path has the same length by construction. Instead, The Warcraft benchmark allows 8-connected movement (diagonal steps), under which feasible paths of different lengths are extreme points of the same feasible region. Thus, $D = \sup_{z \in Z} \|z\|$ is no longer pinned down by the grid size alone; it is determined by the longest simple path the formulation admits as feasible, which can exceed the length of any path an optimizer would actually choose by a wide margin. The bound's looseness on the literal Warcraft benchmark should be at least as severe as Table 2 shows, and plausibly worse for this reason.