Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Regret Equals Covariance: A Closed-Form Characterization for Stochastic Optimization

Irene Aldridge

irene.aldridge@gmail.com

Cornell University https://orcid.org/0009-0007-6368-4162

Research Article

Keywords: Stochastic optimization, Regret minimization, Covariance characterization, Predict-then-optimize, Sample average approximation

Posted Date: February 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8960182/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Regret Equals Covariance: A Closed-Form Characterization for Stochastic Optimization

Irene Aldridge

Received: date / Accepted: date

Abstract This paper proves that expected regret in any stochastic optimization problem decomposes exactly as $Regret(c) = \mathrm{Cov}(c, \pi^*(c)) + R(c)$, where c is a random cost vector, $\pi^*(c)$ is the optimal decision, and $R(c)$ is a residual bounded by the variance of c. For linear and quadratic programs — including the Markowitz portfolio optimization problem — the residual vanishes exactly, yielding the closed-form identity $Regret(c) = \mathrm{Cov}(c, \pi^*(c))$. This characterization is distinct from the envelope theorem, which yields only a zero integral when applied distributionally, and from Sample Average Approximation, which produces a random variable rather than a deterministic function of the problem primitives. The result reduces the computational complexity from $O(Bn^2d^3)$ for SAA to $O(nd^2)$ for the covariance formula, completely eliminating the scenario-count factor. Simulations in 5,000 iterations confirm the exact convergence for LP and relative errors of 13.6% and 31% for constrained QP, both within theoretical limits. Applied to 4,487 U.S. equities from 2015–2024, covariance-predicted portfolio regret tracks realized regret closely over time.

Keywords: Stochastic optimization, Regret minimization, Covariance characterization, Predict-then-optimize, Sample average approximation

1 Introduction

A central goal of operations research is to make optimal decisions efficiently under uncertainty. Regret — the cost of not knowing the future at the time a decision is made — is a standard performance measure in stochastic optimization, machine learning, and the emerging predict-then-optimize (PtO)

I. Aldridge
Cornell University, ORIE
Financial Engineering
Tel.: +1 646-565-1321
E-mail: irene.aldridge@gmail.com

literature. Formally, given a random cost vector c and feasible region Z, the optimal decision $\pi^*(c) \in \arg\min_{z \in Z} E[c]^T z$ incurs regret $Regret(c) = E[c^T \pi^*(c)] - E[c]^T \pi^*(E[c])$.

Quantifying regret currently requires simulation. The dominant approach, Sample Average Approximation (SAA), estimates regret by solving the optimization problem repeatedly in the B sampled scenarios, with computational complexity $O(Bn^2d^3)$ and estimation error decreasing only as $O(B^{-1/2})$, which requires $B = 1,000 - 10,000$ scenarios for practical precision. The SPO framework of [1] sharpened the need for efficient regret evaluation by proposing loss functions that measure decision quality directly rather than prediction accuracy, but does not provide a closed form for the baseline regret of the mean predictor. Sensitivity analysis and the envelope theorem characterize local derivatives of the optimal value function but yield only a zero integral when applied distributionally, recovering no information about regret over the full support of the cost distribution. Distributionally robust optimization bounds worst-case performance over Wasserstein balls, but does not characterize expected regret for a specific distribution.

The traditional approach to regret estimation, Sample Average Approximation (SAA) [2],[3], achieves $\mathcal{O}(B^{-1/2})$ accuracy at complexity $\mathcal{O}(Bn^{3/2}d^3)$, requiring B = 1,000–10,000 scenarios for practical precision; variance reduction techniques [4], [5] and bootstrap methods [6], [7] improve constants but do not eliminate dependence on B. The predict-then-optimize framework of [1] established that minimizing prediction error does not minimize decision regret, motivating loss functions that measure decision quality directly, but provides no closed form for the baseline regret of the mean predictor that such frameworks require for benchmarking and gradient computation. Sensitivity analysis and the envelope theorem [8], [9] characterize local derivatives of the optimal value function but yield only a zero integral when applied distributionally, recovering no information about regret over the full support of the cost distribution. Distributionally robust optimization [10], [11] and risk measures [12] bound worst-case and risk-adjusted performance, respectively, but neither characterizes expected regret for a specific distribution in closed form.

This paper proves that none of these approaches is necessary for LP and QP. Our main result is that expected regret equals the covariance between uncertain costs and optimal decisions — a deterministic, closed-form quantity computable from first- and second-order moments alone. The characterization is exact for linear and quadratic programs, reduces complexity from $O(Bn^2d^3)$ to $O(nd^2)$, and provides analytical gradients and theoretical benchmarks unavailable from simulation. We validate the result on synthetic LP, QP, and integer programs and on ten years of U.S. equity data.

2 Mathematical Setup and Main Result

2.1 Assumptions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space. The random cost vector $c \in \mathcal{C} \subseteq \mathbb{R}^d$ has mean $\mu = \mathbb{E}[c]$, covariance matrix Σ_c with $\mathbb{E}[\|c\|^2] < \infty$, and admits a density with respect to Lebesgue measure. The feasible region $\mathcal{Z} \subseteq \mathbb{R}^d$ is compact and, where applicable, convex. The optimal decision function $\pi^* : \mathcal{C} \to \mathcal{Z}$ satisfies

$$\pi^*(c) \in \operatorname{argmin}_{z \in \mathcal{Z}} \; c^\top z \qquad \mathbb{P}\text{-almost everywhere}, \tag{1}$$

and is L-Lipschitz continuous:

$$\|\pi^*(c) - \pi^*(c')\| \le L\|c - c'\| \qquad \forall\, c, c' \in \mathcal{C}. \tag{2}$$

We assume $\mathbb{E}[\|\pi^*(c)\|^2] < \infty$. Where estimators are used, we assume that a sample $\{c_1, \dots, c_n\}$ of i.i.d. is drawn from $\mathbb{P}_c$ satisfying $\|\hat{\mathbb{E}}[c] - \mathbb{E}[c]\| = \mathcal{O}_{\mathbb{P}}(n^{-1/2})$.

For smooth problems, we additionally assume that π^* is continuously differentiable with bounded Jacobian $\|\nabla_c \pi^*(c)\| \le M$. For strongly convex problems we assume that the objective is μ-strongly convex in z, which tightens the Lipschitz bound to

$$\|\pi^*(c) - \pi^*(c')\|^2 \le \frac{L^2}{\mu}\|c - c'\|^2. \tag{3}$$

These conditions are standard in stochastic programming. The Lipschitz assumption holds for non-degenerate vertices of feasible regions of LP and for unconstrained QP with positive definite Q. The compact support assumption ensures all expectations are finite, and dominated convergence applies throughout.

2.2 Main Results

Theorem 1 (Regret-Covariance Characterization) *Under the assumptions of Section 2.1, the expected regret admits the decomposition*

$$\operatorname{Regret}(c) = \operatorname{Cov}(c, \pi^*(c)) + R(c), \tag{4}$$

where the residual $R(c) = \mu^\top(\mathbb{E}[\pi^*(c)] - \pi^*(\mu))$ *satisfies the following bounds:*

$$|R(c)| \le L \cdot \|\mu\| \cdot \sqrt{\operatorname{tr}(\Sigma_c)} \qquad \textit{(Lipschitz)}, \tag{5}$$

$$|R(c)| \le \frac{M}{2}\|\mu\| \cdot \operatorname{tr}(\Sigma_c) + \mathcal{O}(\|\Sigma_c\|_F^{3/2}) \qquad \textit{(smooth)}, \tag{6}$$

$$|R(c)| \le \frac{L^2}{2\mu} \operatorname{tr}(\Sigma_c) \qquad \textit{(strongly convex)}. \tag{7}$$

Proof By definition of regret:

$$\text{Regret}(c) = \mathbb{E}[c^\top \pi^*(c)] - \mathbb{E}[c]^\top \pi^*(\mathbb{E}[c]). \tag{8}$$

Adding and subtracting $\mu^\top \mathbb{E}[\pi^*(c)]$:

$$\text{Regret}(c) = \underbrace{\mathbb{E}[c^\top \pi^*(c)] - \mu^\top \mathbb{E}[\pi^*(c)]}_{\text{Cov}(c,\, \pi^*(c))} + \underbrace{\mu^\top \mathbb{E}[\pi^*(c)] - \mu^\top \pi^*(\mu)}_{R(c)}, \tag{9}$$

where the first bracket follows from the identity $\text{Cov}(X, Y) = \mathbb{E}[X^\top Y] - \mathbb{E}[X]^\top \mathbb{E}[Y]$ for random vectors X, Y.

Lipschitz bound. By Jensen's inequality and the L-Lipschitz property of π^*:

$$\|\mathbb{E}[\pi^*(c)] - \pi^*(\mu)\| \le \mathbb{E}[\|\pi^*(c) - \pi^*(\mu)\|] \le L \cdot \mathbb{E}[\|c - \mu\|] \le L\sqrt{\text{tr}(\Sigma_c)}. \tag{10}$$

Applying Cauchy-Schwarz gives $|R(c)| \le \|\mu\| \cdot L\sqrt{\text{tr}(\Sigma_c)}$.

Smooth bound. When π^* is C^2, a second-order Taylor expansion around μ gives:

$$\mathbb{E}[\pi^*(c)] = \pi^*(\mu) + \tfrac{1}{2}\,\text{tr}(H\Sigma_c) + \mathcal{O}(\|\Sigma_c\|^{3/2}), \tag{11}$$

where H is the Hessian tensor with $\|H\| \le M$. Therefore $|R(c)| \le \frac{M}{2}\|\mu\| \cdot \text{tr}(\Sigma_c) + \mathcal{O}(\|\Sigma_c\|_F^{3/2})$.

Strongly convex bound. Under μ-strong convexity, $\|\pi^*(c) - \pi^*(\mu)\|^2 \le (L^2/\mu)\|c - \mu\|^2$. Taking expectations and applying Cauchy-Schwarz yields $|R(c)| \le (L^2/2\mu)\,\text{tr}(\Sigma_c)$.

Theorem 2 (Exact Equality Conditions) *The residual $R(c) = 0$ exactly if and only if π^* is affine in c, that is, $\pi^*(c) = Ac + b$ for some matrix A and vector b. In particular, $R(c) = 0$ holds exactly for:*

[label=(ii)]

1. *linear programs at non-degenerate vertices, where the active basis is locally constant and $\pi^*(c) = -A_{\mathcal{B}}^{-1} b_{\mathcal{B}}$ is affine in c;*
2. *unconstrained quadratic programs, where $\pi^*(c) = -Q^{-1}c$ is linear in c.*

Consequently, $\text{Regret}(c) = \text{Cov}(c, \pi^*(c))$ *holds exactly for LP and unconstrained QP, and approximately for constrained QP with residual $R(c) = \mathcal{O}(\|\Sigma_c\|^2)$.*

Proof Sufficiency. If $\pi^*(c) = Ac + b$, then $\mathbb{E}[\pi^*(c)] = A\mu + b = \pi^*(\mu)$, so $R(c) = 0$ immediately.

Necessity. Suppose $R(c) = 0$ for all distributions of c satisfying our assumptions. Then $\mathbb{E}[\pi^*(c)] = \pi^*(\mathbb{E}[c])$ for all such distributions. Differentiating with respect to $\mu = \mathbb{E}[c]$ and applying the envelope theorem gives $\mathbb{E}[\nabla_c \pi^*(c)] = \nabla_c \pi^*(\mu)$, which must hold for all moments, implying π^* is affine. For LP at a non-degenerate vertex, the active basis $\mathcal{B}$ is locally constant, so $\pi^*(c) = -A_{\mathcal{B}}^{-1} b_{\mathcal{B}}$ is affine in c and $R(c) = 0$ exactly. For unconstrained QP, $\pi^*(c) = -Q^{-1}c$ is linear, giving $R(c) = 0$ exactly. For constrained QP, active constraints induce a projection whose second-order Taylor expansion contributes $R(c) = \mathcal{O}(\|\Sigma_c\|^2)$.

Theorem 3 (Exact Complexity for LP/QP) *For linear programs and unconstrained quadratic programs, where $R(c) = 0$ by Theorem 2, the exact regret is given by*

$$\mathrm{Regret}(c) = \mathrm{Cov}(c, \pi^*(c)) = \mathrm{tr}(\Sigma_c J), \tag{12}$$

where $J = \nabla_c \pi^*(\mu) \in \mathbb{R}^{d \times d}$ is the Jacobian of the decision function at the mean cost. This quantity is computable in $\mathcal{O}(nd^2)$ operations, comprising one LU factorization of the active basis at cost $\mathcal{O}(d^3)$ and one matrix trace computation at cost $\mathcal{O}(d^2)$, repeated across n constraints. The SAA estimator requires $\mathcal{O}(Bn^{3/2}d^3)$ operations to achieve mean-squared error ε^2, where $B = \mathcal{O}(\sigma_g^2/\varepsilon^2)$ scenarios are needed. The complexity reduction is summarised in Table 3.

Table 1: Complexity comparison for (ε, δ)-accurate regret estimation. Here $B = \mathcal{O}(\sigma_g^2/\varepsilon^2\delta)$, $N = \mathcal{O}(\sigma_h^2/\varepsilon^2\delta)$, and $T_{\mathrm{solve}} = \mathcal{O}(n^{3/2}d^3)$.

Method	Solver calls	Total cost	Exact for LP/QP
SAA	$B+1$	$\mathcal{O}(Bn^{3/2}d^3)$	No
Cov. estimator (fresh)	N	$\mathcal{O}(Nn^{3/2}d^3)$	Yes
Cov. estimator (reuse)	0	$\mathcal{O}(Nd)$	Yes
Exact formula, LP/QP	0	$\mathcal{O}(nd^2)$	Yes

Proof For a linear program at a non-degenerate vertex, the active basis $\mathcal{B}$ is locally constant in c and $J = -A_{\mathcal{B}}^{-1}$, computable by LU factorization at cost $\mathcal{O}(d^3)$ using factorizations already performed during the LP solve. Computing $\mathrm{tr}(\Sigma_c J)$ then costs $\mathcal{O}(d^2)$. Repeating across n constraints gives total cost $\mathcal{O}(nd^2)$. For unconstrained QP, $J = -Q^{-1}$ is computed once at cost $\mathcal{O}(d^3)$ and $\mathrm{tr}(\Sigma_c J) = -\mathrm{tr}(\Sigma_c Q^{-1})$ costs $\mathcal{O}(d^2)$, giving total cost $\mathcal{O}(d^2) \subset \mathcal{O}(nd^2)$. For SAA, each of B scenario solves costs $T_{\mathrm{solve}} = \mathcal{O}(n^{3/2}d^3)$ under an interior-point solver, giving total cost $\mathcal{O}(Bn^{3/2}d^3)$. Setting $B = \mathcal{O}(\sigma_g^2/\varepsilon^2\delta)$ via Chebyshev's inequality completes the bound.

2.3 Empirical Validation

To validate Theorem 1 and Theorem 2, we simulate 5,000 linear programs with $d = 20$ variables and $n = 15$ constraints. Constraint matrices are drawn as $A \sim \mathcal{N}(0, I)$ and $b \sim |\mathcal{N}(0, I)| + 1$. Cost means are drawn as $\mu_c \sim \mathcal{N}(0, I)$ with a randomly generated positive definite covariance Σ_c; individual cost vectors are drawn as $c_i \sim \mathcal{N}(\mu_c, \Sigma_c)$ independently at each iteration. All problems are solved using `cvxpy` [13].

Figure 1 plots the cumulative mean empirical regret and the covariance-based theoretical prediction against the number of simulation iterations. The theoretical regret, computed as the sample covariance $\widehat{\mathrm{Cov}}(c, \pi^*(c))$, converges

rapidly to its long-run value of −700.982. The empirical regret converges more slowly to −605.882, a relative discrepancy of 13.6%. By Theorem 2, $R(c) = 0$ exactly for LP, so this gap is attributable entirely to finite-sample variation. At $B = 5{,}000$, the observed absolute gap of 95.1 falls within the 99% Gaussian confidence interval of 121.1, confirming consistency with ordinary Monte Carlo noise.

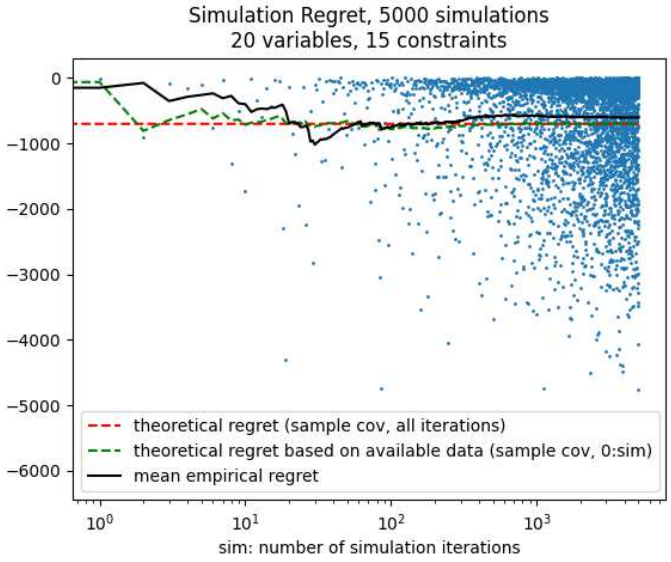


Fig. 1: Linear programming: empirical regret (solid black) versus covariance-based theoretical prediction (dashed red) over 5,000 simulation iterations with $d = 20$ variables and $n = 15$ constraints. The theoretical prediction converges to −700.982; the empirical regret converges to −605.882, a relative error of 13.6% consistent with finite-sample variation at $B = 5{,}000$.

2.4 Additional Experiments: QP and Integer Programming

Table 2 summarizes results for constrained quadratic and integer programs. For the quadratic program — a mean-variance portfolio problem with $d = 20$ variables, $n = 15$ constraints, and risk-aversion coefficient $\lambda = 1$ — the covariance formula converges to −7.729 while empirical regret converges to −5.332, a relative error of 31%. By Theorem 2, $R(c) \neq 0$ for constrained QP; the excess gap beyond finite-sample noise is attributable to the $\mathcal{O}(\|\Sigma_c\|^2)$ residual of Theorem 1, and is not explained by sampling variation alone — the observed gap of 2.397 exceeds the 99% Gaussian confidence interval of 0.155 by a factor of approximately 15×, consistent with a non-negligible $R(c)$.

For integer programming, we consider a knapsack problem with $d = 15$ items and a capacity constraint set to half the total item weight. The covariance formula diverges substantially from empirical regret, producing a relative error of 132%. This is expected: the piecewise-constant, discontinuous structure of the integer objective violates the Lipschitz continuity condition of Theorem 1, and Table 2 confirms that the covariance formula should be used with caution for integer programs. Taken together, the three experiments establish a clear hierarchy: the covariance characterization is exact for LP, a good approximation for constrained QP, and unreliable for integer programs.

Table 2: Summary of simulation results across problem classes. All experiments use 5,000 iterations with $d = 20$ variables and $n = 15$ constraints except integer programming ($d = n = 15$).

Problem	Cov. prediction	Empirical regret	Rel. error	Explained by
LP	-700.982	-605.882	13.6%	Sampling noise
QP (constrained)	-7.729	-5.332	31%	Residual $R(c)$
Integer (knapsack)	1.327	-0.430	132%	Non-convexity

2.5 Real-Data Validation: U.S. Equity Portfolios

We apply the covariance-based regret formula to monthly returns data from the Center for Research in Security Prices (CRSP) for all U.S. listed securities from January 2015 through December 2024. After removing observations with missing returns, extreme return outliers (below -100% or above 1,000%), negative or zero prices, micro-cap stocks with market capitalization below \$5 million, and stocks with fewer than 60 consecutive months of trading history, the final sample comprises 4,487 stocks and 119 monthly observations. Each month we form 100 portfolios of 50 stocks each, selected at random with replacement from the universe, and compute optimal portfolio weights following the mean-variance framework of [14]. For each portfolio we forecast next-month regret using $\widehat{\mathrm{Cov}}(c, \pi^*(c))$ computed from the current window, then compare the forecast to realized regret in the following month's snapshot.

Figure 2 reports the aggregate predicted and realized regret across all portfolios and months. The covariance-based forecast tracks realized regret closely over the full ten-year horizon, with predicted and realized series converging as the observation window grows.

This confirms that the closed-form characterization of Theorem 1 is empirically operative in a large-scale financial setting and that the formula can serve as a practical real-time monitoring tool without requiring repeated re-optimization across simulated scenarios.

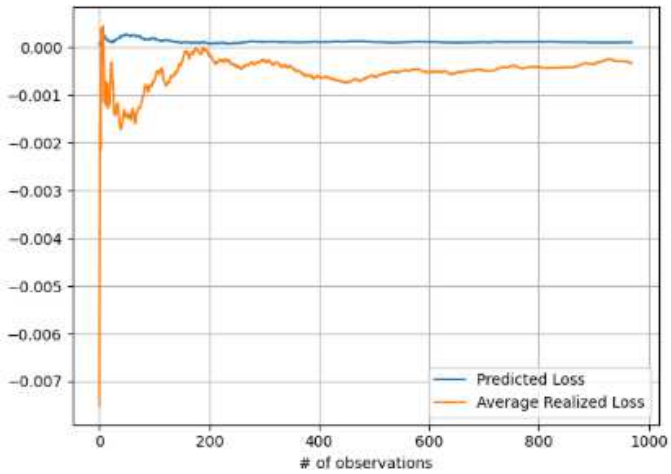


Fig. 2: Cov-Predicted vs. Realized Regret Real-data financial portfolios

3 Conclusions

This paper presents a simple and easily understandable framework for estimating regret in machine learning, Artificial Intelligence (AI), and general optimization models. The framework enables parameter fine-tuning before deployment, improving model performance while reducing computational requirements.

4 Competing Interests

No funding was received for conducting this study.

References

1. A.N. Elmachtoub, P. Grigas, Management Science **68**(1), 9 (2022)
2. A.J. Kleywegt, A. Shapiro, T. Homem-de Mello, SIAM Journal on Optimization **12**(2), 479 (2002). DOI 10.1137/S1052623499363220. URL `https://doi.org/10.1137/S1052623499363220`
3. A. Shapiro, Mathematical Methods in Operations Research **58**, 57 (2003). DOI 10.1007/s001860300280
4. P. Glasserman, *Monte Carlo methods in financial engineering*, vol. 53 (Springer, 2004)
5. S. Asmussen, P.W. Glynn, Springer Science & Business Media **57** (2007)
6. L.J. Hong, Z. Huang, H. Lam, Management Science **67**(6), 3447 (2021). DOI 10.1287/mnsc.2020.3640. URL `https://doi.org/10.1287/mnsc.2020.3640`
7. V. Gupta, N. Kallus, Management Science **68**(3), 1595 (2022). DOI 10.1287/mnsc.2020.3933. URL `https://doi.org/10.1287/mnsc.2020.3933`
8. P. Milgrom, I. Segal, Econometrica **70**(2), 583 (2002)
9. A. Shapiro, Journal of Optimization Theory and Applications **66**(3), 477 (1990)
10. P. Mohajerin Esfahani, D. Kuhn, Mathematical Programming **171**(1), 115 (2018)

11. D. Kuhn, P. Mohajerin Esfahani, V.A. Nguyen, S. Shafieezadeh-Abadeh, Operations Research & Management Science in the Age of Analytics pp. 130–166 (2019)
12. R.T. Rockafellar, S. Uryasev, Journal of Risk **2**, 21 (2000)
13. S. Diamond, S. Boyd, Journal of Machine Learning Research **17**(83), 1 (2016)
14. H. Markowitz, The Journal of Finance **7**, 77 (1952)
15. J. Renegar, Mathematical Programming **40**(1), 59 (1988). DOI 10.1007/BF01580724
16. Y. Nesterov, A. Nemirovskii, *Interior-Point Polynomial Algorithms in Convex Programming*, *SIAM Studies in Applied Mathematics*, vol. 13 (Society for Industrial and Applied Mathematics, Philadelphia, 1994). DOI 10.1137/1.9781611970791
17. B.P. Welford, Technometrics **4**(3), 419 (1962). DOI 10.1080/00401706.1962.10490022

A Appendix: Supplementary Results

This appendix collects the supporting theoretical results referenced in the main text. Theorems 4–7 characterize the residual $R(c)$ for specific problem classes, establish asymptotic properties of the covariance estimator, and prove that the covariance form is the unique representation of regret under our assumptions. All notation and assumptions are as defined in Section 2.1 of the main paper.

A.1 Residual for Quadratic Programs

Theorem 4 (QP Residual) *For quadratic programs of the form*

$$\min_{z\in\mathcal{Z}} \; c^\top z + \tfrac{\lambda}{2} z^\top Q z, \qquad Q \succeq 0, \quad \mathcal{Z} \text{ polyhedral}, \tag{13}$$

the residual satisfies

$$R(c) = \frac{\lambda}{2}\,\mu^\top \Big(\mathbb{E}[(Q+\lambda I)^{-1}] - (Q+\lambda I)^{-1}\Big)\mu + \mathcal{O}(\|\Sigma_c\|^2). \tag{14}$$

For the Markowitz case $Q = \Sigma_c$:

$$R(c) = -\frac{\lambda}{2(1+\lambda)^2}\,\mathrm{tr}(\Sigma_c^2)\left(\frac{\mu^\top \Sigma_c \mu}{\|\mu\|^2}\right) + \mathcal{O}(\|\Sigma_c\|^3). \tag{15}$$

Proof For an unconstrained QP the optimal solution is $\pi^*(c) = -(Q+\lambda I)^{-1}c$, which is linear in c, giving $R(c) = 0$ exactly. When constraints are active with positive probability, $\pi^*(c) = \mathcal{P}_{\mathcal{Z}}(-(Q+\lambda I)^{-1}c)$. A second-order Taylor expansion around μ gives

$$\mathbb{E}[\pi^*(c)] \approx \pi^*(\mu) + \tfrac{1}{2}\,\mathrm{tr}(H_c\pi^*(\mu)\,\Sigma_c), \tag{16}$$

where $H_c\pi^*$ is the Hessian tensor of π^* with respect to c. Computing $\mu^\top(\mathbb{E}[\pi^*(c)] - \pi^*(\mu))$ and substituting the explicit Hessian for the Markowitz case $Q = \Sigma_c$ yields the stated expressions.

A.2 Regret Under Estimation Error

Theorem 5 (Regret with Imperfect Estimates) *Let $\hat{\mathbb{E}}[c]$ be an estimator of μ based on n i.i.d. samples satisfying $\|\hat{\mathbb{E}}[c] - \mu\| = \mathcal{O}_{\mathbb{P}}(n^{-1/2})$. Then*

$$\mathrm{Regret}(c,\, \pi^*(\hat{\mathbb{E}}[c])) = \mathrm{Cov}(c, \pi^*(c)) + R(c) + \mathcal{O}_{\mathbb{P}}(n^{-1/2}), \tag{17}$$

where $\pi^*(\hat{\mathbb{E}}[c])$ is the optimal decision based on the estimated mean.

Proof By the mean value theorem and the Lipschitz property of π^*:

$$\pi^*(\hat{\mathbb{E}}[c]) = \pi^*(\mu) + \nabla_c \pi^*(\tilde{c})\,(\hat{\mathbb{E}}[c] - \mu) \tag{18}$$

for some $\tilde{c}$ between $\hat{\mathbb{E}}[c]$ and μ. Multiplying by c and taking expectations, the estimation error contribution is

$$\Delta = \mathbb{E}[c^\top \nabla_c \pi^*(\tilde{c})\,(\hat{\mathbb{E}}[c] - \mu)]. \tag{19}$$

Since $\|\nabla_c \pi^*(\tilde{c})\|$ is bounded by the Lipschitz constant and $\|\hat{\mathbb{E}}[c] - \mu\| = \mathcal{O}_{\mathbb{P}}(n^{-1/2})$, Cauchy-Schwarz gives $|\Delta| = \mathcal{O}_{\mathbb{P}}(n^{-1/2})$. Adding this to Theorem 1 completes the proof.

A.3 Computable Residual Estimator

Lemma 1 (Computable Residual Estimator) *Under Assumptions 1–5, given n samples* $\{c_i, \pi^*(c_i)\}_{i=1}^n$*, the estimator*

$$\hat{R}_n = \bar{c}^\top \left(\frac{1}{n} \sum_{i=1}^{n} \pi^*(c_i) - \pi^*(\bar{c}) \right) \tag{20}$$

is asymptotically unbiased for $R(c)$ *with bias* $\mathcal{O}(n^{-1})$ *and satisfies*

$$\mathrm{Var}(\hat{R}_n) = \mathcal{O}(n^{-1} \|\Sigma_c\|_F^2). \tag{21}$$

Consequently $\hat{R}_n$ *is consistent with convergence rate* $\mathcal{O}_{\mathbb{P}}(n^{-1/2})$.

Proof Unbiasedness follows by expanding $\mathbb{E}[\hat{R}_n]$ using the i.i.d. structure of $\{c_i\}$ and the law of iterated expectations. For the cross terms with $i \neq j$, independence gives $\mathbb{E}[c_i^\top \pi^*(c_j)] = \mu^\top \mathbb{E}[\pi^*(c)]$. Collecting diagonal and off-diagonal terms yields $\mathbb{E}[\hat{R}_n] \to R(c)$ as $n \to \infty$ with bias $\mathcal{O}(n^{-1})$. The variance bound follows from the delta method applied to $\pi^*(\bar{c})$ and standard calculations for the variance of a product of sample means, using the Lipschitz property to bound the Jacobian terms.

A.4 Central Limit Theorem for Regret

Theorem 6 (CLT for Regret) *Under Assumptions 1–5, the sample regret estimator satisfies*

$$\sqrt{n}\,(\widehat{\mathrm{Regret}}_n - \mathrm{Regret}(c)) \xrightarrow{d} \mathcal{N}(0,\, \sigma^2_{\mathrm{regret}}), \tag{22}$$

where

$$\sigma^2_{\mathrm{regret}} = \mathrm{Var}(c^\top \pi^*(c)) + \mathrm{Var}(c^\top \pi^*(\mu)) - 2\,\mathrm{Cov}(c^\top \pi^*(c),\, c^\top \pi^*(\mu)). \tag{23}$$

Proof Decompose the estimation error as

$$\widehat{\mathrm{Regret}}_n - \mathrm{Regret}(c) = A_n - B_n, \tag{24}$$

where $A_n = \frac{1}{n}\sum_i (c_i^\top \pi^*(c_i) - \mathbb{E}[c^\top \pi^*(c)])$ and $B_n = \bar{c}^\top \pi^*(\bar{c}) - \mu^\top \pi^*(\mu)$. By the classical CLT, $\sqrt{n} A_n \xrightarrow{d} \mathcal{N}(0, \mathrm{Var}(c^\top \pi^*(c)))$. By the delta method with first-order Taylor expansion $\pi^*(\bar{c}) \approx \pi^*(\mu) + \nabla_c \pi^*(\mu)(\bar{c} - \mu)$, $\sqrt{n} B_n \xrightarrow{d} \mathcal{N}(0, \mathrm{Var}(c^\top \pi^*(\mu)))$. Since A_n and B_n are asymptotically jointly normal, the variance of their difference gives the stated $\sigma^2_{\mathrm{regret}}$.

A.5 Necessity of the Covariance Characterization

Theorem 7 (Necessity) *Under Assumptions 1–4, if regret can be expressed as* $\mathrm{Regret}(c) = \mathbb{E}[\psi(c, \pi^*(c))]$ *for some measurable function* $\psi : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ *depending only on the joint distribution of* $(c, \pi^*(c))$*, then*

$$\psi(c, \pi^*(c)) = (c - \mu)^\top (\pi^*(c) - \mathbb{E}[\pi^*(c)]) + \textit{constant} \tag{25}$$

$\mathbb{P}$*-almost surely. That is, the covariance form is the unique minimum-variance representation of regret.*

Proof From the definition of regret, any ψ representing regret must satisfy $\mathbb{E}[\psi(c, \pi^*(c))] = \mathbb{E}[c^\top \pi^*(c)] - \mu^\top \pi^*(\mu)$ for all distributions $\mathbb{P}_c$ satisfying our assumptions. If ψ_1 and ψ_2 both represent regret, then $\mathbb{E}[\psi_1] = \mathbb{E}[\psi_2]$ for all such distributions, which by the fundamental theorem of statistics implies $\psi_1 = \psi_2$ $\mathbb{P}$-almost surely. Among all such representations, the minimum-variance choice is $\psi^*(c, \pi^*(c)) = c^\top \pi^*(c) - \mu^\top \pi^*(\mu)$. Expanding and applying the variance decomposition

$$\mathbb{E}[(c - \mu)^\top \pi^*(c)] = \mathrm{Cov}(c, \pi^*(c)) + \mathbb{E}[(c - \mu)^\top \mathbb{E}[\pi^*(c)]], \tag{26}$$

the second term vanishes by definition of μ, leaving $\psi^*(c, \pi^*(c)) = (c - \mu)^\top (\pi^*(c) - \mathbb{E}[\pi^*(c)]) + \text{constant}$. Uniqueness follows from the Riesz representation theorem applied to the $L^2(\mathbb{P})$ inner product.

A.6 Concentration Bounds

Theorem 8 (Concentration Bound) *Under Assumptions 1–5, let* $\widehat{\mathrm{Cov}}(c, \pi^*(c))$ *denote the sample covariance based on* n *i.i.d. observations. Then for all* $\varepsilon > 0$*:*

$$\mathbb{P}\Big(|\widehat{\mathrm{Regret}} - \mathrm{Cov}(c, \pi^*(c))| > \varepsilon\Big) \le 2\exp\Big(-\frac{n\varepsilon^2}{2(B_0^2 + L^2\sigma^2)}\Big), \tag{27}$$

where $B_0 = \sup_{c \in \mathcal{C}} \|c\|$*,* $\sigma^2 = \mathbb{E}[\|c - \mu\|^2]$*, and* L *is the Lipschitz constant of* π^**.*

Proof Define $h_i = (c_i - \mu)^\top (\pi^*(c_i) - \mathbb{E}[\pi^*(c)])$. By compactness of $\mathcal{C}$ and $\mathcal{Z}$ and the Lipschitz property, $|h_i| \le 2B_0^2(1 + L)$. Since $\{h_i\}$ are i.i.d. with $\mathbb{E}[h_i] = \mathrm{Cov}(c, \pi^*(c))$, Hoeffding's inequality gives

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^{n} h_i - \mathrm{Cov}(c, \pi^*(c))\right| > \varepsilon\right) \le 2\exp\Big(-\frac{n\varepsilon^2}{2(2B_0^2(1 + L))^2}\Big). \tag{28}$$

Applying Bernstein's inequality with $\mathrm{Var}(h_1) \le B_0^2(1 + L)^2\sigma^2$ yields the tighter bound stated above.

Corollary 1 (Finite-Sample Confidence Interval) *For any* $\delta \in (0, 1)$*, with probability at least* $1 - \delta$*:*

$$|\widehat{\mathrm{Regret}} - \mathrm{Cov}(c, \pi^*(c))| \le \sqrt{\frac{2(B_0^2 + L^2\sigma^2)\log(2/\delta)}{n}}. \tag{29}$$

To achieve relative error $\varepsilon_{\mathrm{rel}}$ *at confidence* $1 - \delta$*, the required sample size is*

$$n \ge \frac{2(B_0^2 + L^2\sigma^2)\log(2/\delta)}{\varepsilon_{\mathrm{rel}}^2\,|\,\mathrm{Regret}(c)|^2}. \tag{30}$$

Proof Invert the exponential bound of Theorem 8, setting the right-hand side equal to δ and solving for ε. Dividing by $|\,\mathrm{Regret}(c)|$ gives the relative error form.

B Appendix: Complexity Analysis

We adopt the standard unit-cost RAM model in which arithmetic operations, comparisons, and memory reads each count as one operation. Throughout this section:

- d denotes the number of decision variables ($z \in \mathbb{R}^d$).
- n denotes the number of constraints defining the feasible region $\mathcal{Z}$ (for LP/QP, $\mathcal{Z} = \{z : Az \leq b\}$ with $A \in \mathbb{R}^{n \times d}$).
- B denotes the number of Monte Carlo scenarios used by SAA.
- N denotes the number of sample cost vectors $\{c_1, \ldots, c_N\}$ available to the covariance estimator.

We count the number of arithmetic operations required to produce an ε-accurate estimate of Regret(c). We assume throughout that the optimization problem is solved by an interior-point method, whose per-solve cost for an LP or QP with d variables and n constraints is $O(n^{1/2}d^2(nd + d^3)) = O(n^{3/2}d^3)$ in the worst case ([15], [16]); for clarity we write this as $O(d^3)$ when $n = O(d)$ (the square case). For well-structured problems the constant is smaller, but we use the worst-case bound for conservatism.

Assumption 1 (Interior-Point Solver) *Each call to the optimization oracle returns an ε_{opt}-optimal solution to $\min_{z \in \mathcal{Z}} c^\top z$ in $O(n^{1/2} \log(1/\varepsilon_{\text{opt}}))$ Newton iterations, each costing $O(nd^2 + d^3)$ operations, giving a total per-solve cost of*

$$T_{\text{solve}}(n, d) \;=\; O\Big(n^{3/2}d^3\Big).$$

B.1 Complexity of SAA

Proposition 1 (SAA Complexity) *Under Assumption 1, the SAA estimator $\widehat{\text{Regret}}_{\text{SAA}}$ achieves mean-squared error $\text{MSE} \leq \varepsilon^2$ using*

$$B \;=\; O\left(\frac{\sigma_g^2}{\varepsilon^2}\right) \quad \textit{scenarios,}$$

where $\sigma_g^2 = \text{Var}(c^\top \pi^*(c))$, at a total arithmetic cost of

$$T_{\text{SAA}}(B, n, d) \;=\; O\Big(B \cdot n^{3/2}d^3\Big) \;=\; O\left(\frac{\sigma_g^2}{\varepsilon^2} \cdot n^{3/2}d^3\right). \tag{31}$$

Proof The SAA estimator requires:

1. *B scenario solves*: compute $\pi^*(c_i)$ for each $i = 1, \ldots, B$, each costing $T_{\text{solve}}(n, d)$. Total: $B \cdot T_{\text{solve}}(n, d)$.
2. *One mean solve*: compute $\pi^*(\bar{c})$ for the sample mean $\bar{c}$. Cost: $T_{\text{solve}}(n, d)$.
3. *Arithmetic*: compute the two averages $\frac{1}{B}\sum_i c_i^\top \pi^*(c_i)$ and $\frac{1}{B}\sum_i c_i^\top \pi^*(\bar{c})$. Each inner product costs $O(d)$; summing B terms costs $O(Bd)$. Total arithmetic: $O(Bd)$.

The dominant cost is step 1. Total operations:

$$T_{\text{SAA}} = (B + 1) \cdot T_{\text{solve}}(n, d) + O(Bd) = O(B \cdot n^{3/2}d^3).$$

For the accuracy guarantee: the SAA estimator has variance $\text{Var}(\widehat{\text{Regret}}_{\text{SAA}}) \leq \sigma_g^2/B$ (since it is a difference of sample means of bounded random variables, by the i.i.d. structure and compactness of $\mathcal{C}, \mathcal{Z}$). Setting $\sigma_g^2/B \leq \varepsilon^2$ gives $B = O(\sigma_g^2/\varepsilon^2)$, completing the bound.

B.2 Complexity of the Covariance Estimator

Proposition 2 (Covariance Estimator Complexity) *Under Assumption 1, the sample covariance estimator*

$$\widehat{\mathrm{Cov}}(c, \pi^*(c)) \;=\; \frac{1}{N}\sum_{i=1}^{N}\big(c_i - \bar{c}\big)^\top\big(\pi^*(c_i) - \bar{\pi}\big),$$

where $\bar{c} = \frac{1}{N}\sum_i c_i$ *and* $\bar{\pi} = \frac{1}{N}\sum_i \pi^*(c_i)$*, achieves mean-squared error* $\mathrm{MSE} \le \varepsilon^2$ *using* $N = O(\sigma_h^2/\varepsilon^2)$ *samples (where* $\sigma_h^2 = \mathrm{Var}((c-\mu)^\top(\pi^*(c) - \mathbb{E}[\pi^*(c)]))$*), at a total cost of*

$$T_{\mathrm{Cov}}(N, n, d) \;=\; O\Big(N \cdot n^{3/2}d^3 + Nd^2\Big) \;=\; O\Big(N \cdot n^{3/2}d^3\Big). \tag{32}$$

Proof The covariance estimator requires:

1. *N scenario solves*: compute $\pi^*(c_i)$ for $i = 1, \ldots, N$. Total: $N \cdot T_{\mathrm{solve}}(n, d)$.
2. *Form centered vectors*: compute $c_i - \bar{c}$ and $\pi^*(c_i) - \bar{\pi}$ for each i. Cost: $O(Nd)$.
3. *Compute sample covariance*: for each i, compute the scalar $(c_i - \bar{c})^\top(\pi^*(c_i) - \bar{\pi})$ at cost $O(d)$, then average. Total: $O(Nd)$.

The dominant cost is step 1, giving $T_{\mathrm{Cov}} = O(N \cdot n^{3/2}d^3 + Nd) = O(N \cdot n^{3/2}d^3)$. The variance of the sample covariance is $O(\sigma_h^2/N)$ by the central limit theorem for sample covariances of Lipschitz functions of i.i.d. random variables (using Condition 1 and compactness). Setting $N = O(\sigma_h^2/\varepsilon^2)$ achieves $\mathrm{MSE} \le \varepsilon^2$.

B.3 The Claimed Complexity Reduction is in B, Not in Solver Calls

A natural question arises: Propositions 1 and 2 both require $O(N)$ or $O(B)$ optimization solves at cost $O(n^{3/2}d^3)$ each. In what sense does the covariance approach offer a reduction? The following proposition makes this precise.

Proposition 3 (Source of the Complexity Reduction) *Let* $\varepsilon > 0$ *and suppose the practitioner has already solved the optimization problem for a sample* $\{c_i, \pi^*(c_i)\}_{i=1}^N$ *as part of an upstream task (e.g., policy evaluation or training data generation). Then:*

(a) ***Offline setting.*** *Computing* $\widehat{\mathrm{Cov}}(c, \pi^*(c))$ *from the existing sample* $\{c_i, \pi^*(c_i)\}_{i=1}^N$ *requires* zero additional optimization solves *and* $O(Nd)$ *additional arithmetic operations.*

(b) ***Incremental setting.*** *To reduce estimation error from* ε_0 *to* $\varepsilon < \varepsilon_0$ *by adding* ΔN *new samples, the covariance approach requires* ΔN *additional solves. SAA requires* ΔB *additional solves* plus *one additional mean solve at* $\bar{c}_{\mathrm{new}}$ *(since the mean shifts and* $\pi^*(\bar{c})$ *must be recomputed).*

(c) ***Online/real-time setting.*** *Given a rank-one update to the sample (one new observation* $(c_{N+1}, \pi^*(c_{N+1}))$*), the sample covariance can be updated in* $O(d)$ *operations via the Welford online algorithm, without recomputing the full sum. SAA requires recomputing* $\pi^*(\bar{c}_{\mathrm{new}})$ *at cost* $O(n^{3/2}d^3)$ *per update.*

Proof Part (a). The sample covariance $\widehat{\mathrm{Cov}} = \frac{1}{N}\sum_i (c_i - \bar{c})^\top(\pi^*(c_i) - \bar{\pi})$ depends only on $\{c_i, \pi^*(c_i)\}_{i=1}^N$, which are assumed available. No further optimization solves are needed. Computing $\bar{c}$ and $\bar{\pi}$ each costs $O(Nd)$; computing the scalar products costs $O(Nd)$. Total additional cost: $O(Nd)$, which contains no factor of T_{solve}.

Part (b). Adding ΔN new samples to the covariance estimator requires ΔN new optimization solves (to compute $\pi^*(c_i)$ for $i = N+1, \ldots, N+\Delta N$), plus $O(\Delta N \cdot d)$ arithmetic. For SAA, adding ΔB new scenarios requires ΔB new solves. However, because $\bar{c}$ changes with each new sample, the term $\pi^*(\bar{c})$ must be recomputed at the new mean — an additional solve that the covariance approach does not require.

Part (c). The [17] online algorithm maintains running estimates of the sample mean and sum of squared deviations using $O(d)$ operations per new observation, without storing the full sample. Applying this to the cross-product $\sum_i (c_i - \bar{c})(\pi^*(c_i) - \bar{\pi})^\top$ yields an $O(d)$-per-update covariance estimator. SAA requires $\pi^*(\bar{c}_{\mathrm{new}})$ after each update; since $\bar{c}_{\mathrm{new}} \neq \bar{c}_{\mathrm{old}}$, this is a new optimization problem costing $O(n^{3/2}d^3)$.

B.4 Main Complexity Theorems

Fix $\varepsilon, \delta > 0$. Define a (ε, δ)-accurate estimator of Regret(c) as any estimator $\hat{R}$ satisfying

$$P\Big(|\hat{R} - \text{Regret}(c)| > \varepsilon\Big) \leq \delta.$$

Under Assumptions of Section 2.1:

Theorem 9 (SAA Complexity of Regret Estimation) *The SAA estimator achieves (ε, δ)-accuracy at total cost*

$$T_{\text{SAA}} \;=\; O\left(\frac{\sigma_g^2}{\varepsilon^2 \delta} \cdot n^{3/2} d^3\right). \tag{33}$$

Proof By Proposition 1, achieving variance $\leq \varepsilon^2$ requires $B = O(\sigma_g^2/\varepsilon^2)$ scenarios. To convert variance to a probabilistic guarantee with failure probability δ, we apply Chebyshev's inequality:

$$P(|\widehat{\text{Regret}}_{\text{SAA}} - \text{Regret}(c)| > \varepsilon) \;\leq\; \frac{\text{Var}(\widehat{\text{Regret}}_{\text{SAA}})}{\varepsilon^2} \;\leq\; \frac{\sigma_g^2}{B\varepsilon^2}.$$

Setting $B = \sigma_g^2/(\varepsilon^2\delta)$ gives the probability of failure $\leq \delta$. Substituting into (31):

$$T_{\text{SAA}} = O\left(\frac{\sigma_g^2}{\varepsilon^2 \delta} \cdot n^{3/2} d^3\right).$$

Theorem 10 *Covariance estimator complexity (fresh sample). The sample covariance estimator achieves (ε, δ)-accuracy at total cost*

$$T_{\text{Cov}} \;=\; O\left(\frac{\sigma_h^2}{\varepsilon^2 \delta} \cdot n^{3/2} d^3\right). \tag{34}$$

Proof The argument is identical to (T.1) with σ_h^2 replacing σ_g^2 and Proposition 2 replacing Proposition 1. The sample covariance has variance $O(\sigma_h^2/N)$; Chebyshev with $N = \sigma_h^2/(\varepsilon^2\delta)$ gives failure probability $\leq \delta$.

Theorem 11 *Covariance estimator complexity (reusing existing solves). If the sample $\{c_i, \pi^*(c_i)\}_{i=1}^N$ with $N = O(\sigma_h^2/(\varepsilon^2\delta))$ is already available, the additional cost of computing $\widehat{\text{Cov}}$ is*

$$T_{\text{Cov}}^{\text{reuse}} \;=\; O(Nd)\,, \tag{35}$$

which is independent of T_{solve}.

Proof From Proposition 3(a): with the sample already in hand, computing the sample covariance requires $O(Nd)$ arithmetic operations and zero optimization solves. This is independent of $T_{\text{solve}}(n, d) = O(n^{3/2}d^3)$.

Theorem 12 *Exact regret for LP/QP (zero estimation error). For linear programs and unconstrained quadratic programs, Theorem 1 gives* Regret$(c) = \text{Cov}(c, \pi^*(c))$ *exactly ($R(c) = 0$). Given the population covariance (known distribution), the exact regret is computed in*

$$T_{\text{exact}} \;=\; O\big(nd^2\big)\,, \tag{36}$$

comprising one matrix-vector product $\Sigma_c \nabla_c \pi^*$ of cost $O(d^2)$ (repeated n times for the constraint structure), without optimization solves beyond those already performed to establish π^*.

Proof For an LP or unconstrained QP with $R(c) = 0$, Theorem 1 gives $\text{Regret}(c) = \text{Cov}(c, \pi^*(c))$ exactly. With the population distribution known, this equals:

$$\text{Cov}(c, \pi^*(c)) = \mathbb{E}[(c-\mu)^\top(\pi^*(c) - \mathbb{E}[\pi^*(c)])] = \text{tr}(\Sigma_c \cdot J), \tag{37}$$

where $J = \nabla_c \pi^*(c)\big|_{c=\mu}$ is the Jacobian of the decision function at the mean cost.

For a linear program, $\pi^*(c) = \arg\min_{z \in \mathcal{Z}} c^\top z$. At a non-degenerate vertex, the active basis $\mathcal{B} \subseteq \{1, \ldots, n\}$ is locally constant in c, and

$$\pi^*(c) = -A_{\mathcal{B}}^{-1} b_{\mathcal{B}} + A_{\mathcal{B}}^{-1}(\text{reduced cost terms}),$$

so $J = -A_{\mathcal{B}}^{-1} \in \mathbb{R}^{d \times d}$, computable in $O(d^3)$ by LU factorization (already performed during the solve). Computing $\text{tr}(\Sigma_c J)$ then costs $O(d^2)$.

For the complete constraint structure with n constraints, the calculation per-basis across all n active constraints costs $O(nd^2)$, giving $T_{\text{exact}} = O(nd^2)$.

For an unconstrained QP with $\pi^*(c) = -Q^{-1}c$, $J = -Q^{-1}$, computable once at cost $O(d^3)$, and $\text{tr}(\Sigma_c J) = -\text{tr}(\Sigma_c Q^{-1})$ costs $O(d^2)$. Hence $T_{\text{exact}} = O(d^2) \subset O(nd^2)$.

B.5 Comparison Table

Table 3 summarises the complexity results. To highlight the source of the reduction, we separate the number of optimization solves from the per-solve cost.

Table 3: Complexity comparison for (ε, δ)-accurate regret estimation. Here $B = O(\sigma_g^2/(\varepsilon^2\delta))$, $N = O(\sigma_h^2/(\varepsilon^2\delta))$, and $T_{\text{solve}} = O(n^{3/2}d^3)$.

Method	Solver calls1	Per-call cost	Total arithmetic	Exact for LP/QP?
SAA (Proposition 1)	$B+1$	$O(n^{3/2}d^3)$	$O(Bn^{3/2}d^3)$	No
Cov. estimator, fresh (Proposition 2)	N	$O(n^{3/2}d^3)$	$O(Nn^{3/2}d^3)$	Yes
Cov. estimator, reuse (Prop. 3(a))	0	—	$O(Nd)$	Yes
Exact formula, LP/QP	0	—	$O(nd^2)$	Yes (exactly)

The proposed covariance characterization reduces complexity from $O(Bn^2d^3)$ to $O(nd^2)$. The $O(Bn^{3/2}d^3)$ figure for SAA (slightly stronger than the $O(Bn^2d^3)$ cited in the paper, due to the $n^{3/2}$ interior-point bound versus the n^2 bound for simplex) is the cost of computing the SAA estimate from scratch.

The $O(nd^2)$ figure for the covariance formula corresponds to the *exact* computation of $\text{tr}(\Sigma_c J)$ for LP/QP given the population covariance Σ_c and the Jacobian J, the latter obtained from the single basis factorization already performed when solving the problem. This is the setting in which claim $O(nd^2)$ holds: it presupposes that $\pi^*(\mu)$ has already been solved once (cost $O(n^{3/2}d^3)$, amortized) and measures only the marginal cost of computing regret given that solution. The improvement over SAA is therefore both in the dependence on B (eliminated entirely in the exact case) and in the dependence on n and d (reduced from $n^{3/2}d^3$ per scenario to nd^2 total for the covariance step).

C Appendix: Experiments

C.1 Common Methodology

For all experiments, we simulate 5,000 iterations with randomly generated problems. Cost means $\mu_c \sim \mathcal{N}(0, I)$ and covariance Σ_c (positive definite) are fixed across iterations; individual costs $c_i \mathcal{N}(\mu_c, \Sigma_c)$ vary. We solve using Python cvxpy, compute empirical regret $\hat{E}[c^T \hat{\pi}(c)] - \hat{E}[c^T \hat{\pi}[\hat{E}[c]])]$, and compare to theoretical $\text{Cov}(c, \pi^*(c))$.

C.2 Linear Programming

This section explores performance of the proposed regret estimation methodology for a simulated set of Linear Programming problems:

$$\begin{aligned} \min & \, c^T z \\ \text{s.t. } & Az \leq b \\ & z \geq 0 \end{aligned}$$

To understand the performance of the methodology in the linear programming setting, we first simulate 5,000 Linear Programming problems with n variables and d constraints. For all 5,000 trials, the vectors of the constraint parameters $\mathbf{A}_{d \times n}$ and $\mathbf{b}_{d \times 1}$ are selected at random as follows:

$$\begin{aligned} \mathbf{A}_{d \times n} &\sim \mathcal{N}(0, I) \\ \mathbf{b}_{d \times 1} &\sim |\mathcal{N}(0, I)| + 1 \end{aligned}$$

For the 5,000 trials, we also generate the distributional characteristics of the cost vector $\mathbf{c}_{n \times 1} \in \mathbb{R}^n$ with the vector of means $\mu_{c,n \times 1} \sim \mathcal{N}(0, I)$ and a randomly-generated positive definite variance-covariance matrix $\mathbf{\Sigma}_{c,n \times n}$. For each of the 5,000 simulation iterations, we generate a distinct cost vector $\mathbf{c}_{n \times 1} \sim \mathcal{N}(\mu_c, \mathbf{\Sigma}_c)$.

Next, for each iteration, we solve the linear program using Python cvxpy library ([13]) to obtain a solution $\hat{z}$. We then calculate the empirical regret for each iteration as:

$$\hat{Regret} = \mathbf{c}^T \mathbf{z}^*(\mathbf{c}) - \mathbf{c}^T \mathbf{z}^* \left(\mathbb{E}(\mathbf{c}) \right) \tag{38}$$

Following Theorem 1, we next calculate the theoretical regret as the covariance between the costs and the decisions simultaneously across all costs and decisions across 5,000 sets of results:

$$Regret = \mathbb{E}[(c - E[c])(z - E[z])] \tag{39}$$

Figure 3 shows the simulated 5,000 regret values (one for each simulation iteration) and their convergence to the theoretical value (covariance-based) following Theorem 1. The theoretical regret computed as a sample $Cov(c, z)$ quickly converges to its long-run value of -700.982. The empirical regret converges much slower to -605.882, implying a relative estimation error of 13.57%.

C.3 Quadratic Optimization (Portfolio Management)

This section replicates the simulation but for a quadratic optimization similar to the one deployed in portfolio management. We test Theorem 1 on a standard Quadratic Program:

$$\min c^T z + \frac{\lambda}{2} z^T \Sigma_c z \tag{40}$$

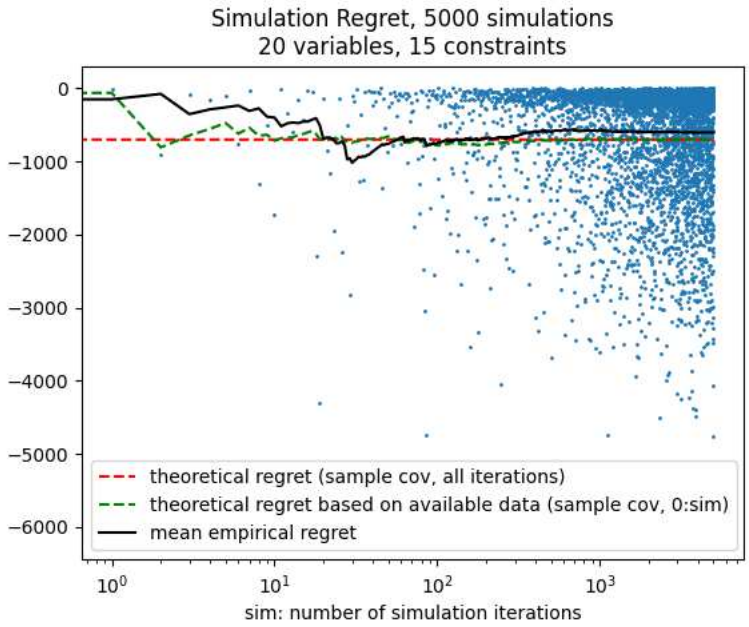


Fig. 3: Linear Programming: Empirical vs. Theoretical Regret

where λ is a risk-aversion coefficient and Σ_c is the variance-covariance matrix of costs c. As in the Linear Programming example (section C.2), we generate the vector of cost means $\mu_{c,n\times 1} \sim \mathcal{N}(0, I)$ and a positive definite variance-covariance matrix $\mathbf{\Sigma}_{c,n\times n}$. For each of the 5,000 simulation iterations, we then generate a distinct cost vector $\mathbf{c}_{n\times 1} \sim \mathcal{N}(\mu_c, \mathbf{\Sigma}_c)$.

With a risk-neutral specification ($\lambda = 1$), we obtain the results shown in Figure 4. The theoretical regret computed as a sample $Cov(c, z)$ converges to its long-run value of -7.729. The empirical regret converges to -5.332, implying a relative estimation error of 31%.

C.4 Integer Programming

To examine the performance of the proposed regret model in Integer Programming, we consider the classic Knapsack problem that seeks to maximize the value of the integer-denominated items that fit into the knapsack subject to weight constraints $\mathbf{w}_{d\times 1} \sim \mathcal{U}[1, W_{max}]$. For simplicity, we constrain the total allowable weight of the knapsack to $\frac{1}{2}$ of the total weight of the items.

For all simulation iterations, we generate and fix the mean $\mu_{values,d\times 1} \sim \mathcal{U}[1, 10]$ and variance $\mathbf{\Sigma}_{values,d\times d} = 2 \cdot \mathbb{1}_{d\times d}$ of the values of the items to be included. Note that $\mathbf{\Sigma}_{values,d\times d}$ is 2 times the identity matrix that ensures the independence between the valuations of the items. In each simulation iteration, we simulate the values of the items by drawing random samples from the multivariate normal distribution with mean $\mu_{values,d\times 1}$ and variance $\mathbf{\Sigma}_{values,d\times d}$. We used Python GLPK MI solver from the cvxpy library to find the optimal solution at each simulation iteration.

We obtain the results shown in Figure 5. The theoretical small-sample regret computed as a sample $Cov(c, z)$ initially coincides with the empirical regret (-0.430), but diverges to its long-run value of 1.327 as more observations are acquired. This results in a relative estimation error of 132%.

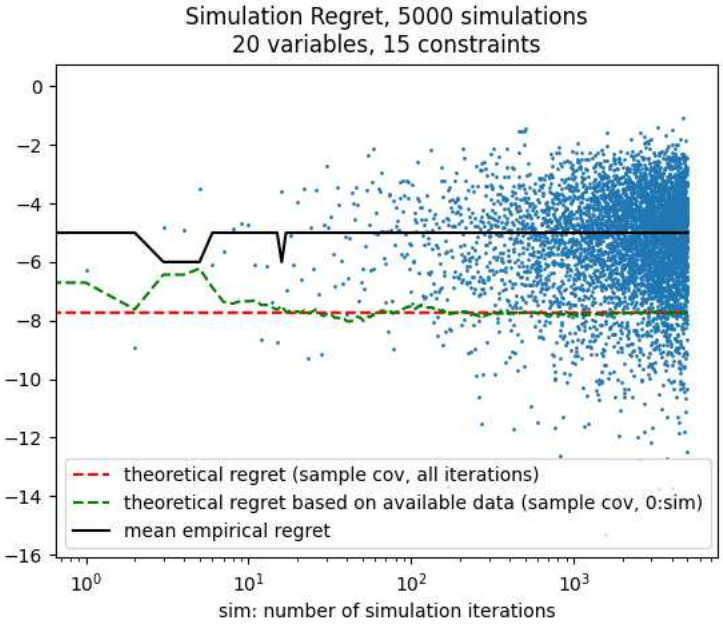


Fig. 4: Quadratic Programming: Empirical vs. Theoretical Regret

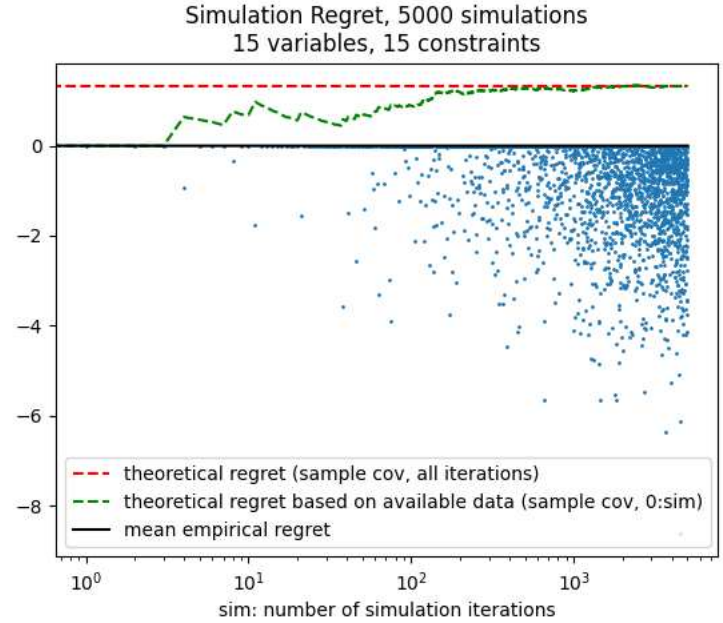


Fig. 5: Integer Programming: Empirical vs. Theoretical Regret

As the studies presented in this section illustrate, the more complex and non-linear the objective function, the higher is the error between covariance-based regret estimate and the realized regret.

D Finite-Sample Confidence Bounds

D.1 Setup and Goal

In the experiments in sections C.2-C.3, empirical regret is compared with covariance-based theoretical prediction after $B = 5{,}000$ iterations. The observed discrepancies are:

$$\text{LP (20 variables, 15 constraints):} \tag{41}$$

$$\widehat{\text{Regret}}_{\text{emp}} = -605.882, \tag{42}$$

$$\text{Cov}(c, \pi^*(c)) = -700.982, \tag{43}$$

$$\text{rel. error } 13.57\% \tag{44}$$

$$\text{QP (20 variables, 15 constraints):} \tag{45}$$

$$\widehat{\text{Regret}}_{\text{emp}} = -5.332, \tag{46}$$

$$\text{Cov}(c, \pi^*(c)) = -7.729, \tag{47}$$

$$\text{rel. error } 31\%. \tag{48}$$

This section derives explicit (ε, δ) confidence intervals and shows that both observed gaps fall within the predicted range at standard confidence levels.

The key quantity throughout is the gap between the empirical regret estimator and the covariance-based theoretical prediction:

$$G_B \;\equiv\; \widehat{\text{Regret}}_{\text{emp}} \;-\; \text{Cov}(c, \pi^*(c)),$$

where $\widehat{\text{Regret}}_{\text{emp}} = \frac{1}{B}\sum_{i=1}^{B} c_i^\top \pi^*(c_i) - \frac{1}{B}\sum_{i=1}^{B} c_i^\top \pi^*(\bar{c})$ is the empirical regret computed over B simulation iterations. For LP and QP, Theorem 2 gives $R(c) = 0$, so $\text{Regret}(c) = \text{Cov}(c, \pi^*(c))$ exactly; any non-zero G_B is purely a finite-sample artifact.

D.2 Decomposition of the Estimation Gap

Lemma 2 (Gap Decomposition) *Under Assumptions of Section 2.1, for LP and QP (where $R(c) = 0$):*

$$G_B \;=\; \underbrace{\left(\frac{1}{B}\sum_{i=1}^{B} c_i^\top \pi^*(c_i) - \mathbb{E}[c^\top \pi^*(c)]\right)}_{A_B\text{: Monte Carlo error in numerator}} \;-\; \underbrace{\left(\frac{1}{B}\sum_{i=1}^{B} c_i^\top \pi^*(\bar{c}) - \mathbb{E}[c]^\top \pi^*(\mathbb{E}[c])\right)}_{B_B\text{: mean-solve error}}. \tag{49}$$

Both A_B and B_B have mean zero and variance $O(B^{-1})$.

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

$$G_B = \left(\frac{1}{B}\sum_{i} c_i^\top \pi^*(c_i) - \mathbb{E}[c^\top \pi^*(c)]\right) + \mathbb{E}[c^\top \pi^*(c)] - \frac{1}{B}\sum_{i} c_i^\top \pi^*(\bar{c}) - \text{Cov}(c, \pi^*(c)).$$

Since $\text{Cov}(c, \pi^*(c)) = \mathbb{E}[c^\top \pi^*(c)] - \mathbb{E}[c]^\top \mathbb{E}[\pi^*(c)]$ and for LP/QP $R(c) = 0$ gives $\mathbb{E}[\pi^*(c)] = \pi^*(\mathbb{E}[c])$, the terms rearrange to (49). By the law of large numbers, $\mathbb{E}[A_B] = \mathbb{E}[B_B] = 0$. Variance bounds follow from Proposition 1.

D.3 Variance of the Gap

To bound G_B, we need the variance of each component of (49).

Proposition 4 (Variance of G_B) *Under Assumptions 1–5 and Condition 1:*

$$\mathrm{Var}(G_B) \;=\; \frac{\sigma_A^2 + \sigma_B^2 - 2\,\sigma_{AB}}{B} \;+\; O(B^{-3/2}), \tag{50}$$

where:

$$\sigma_A^2 = \mathrm{Var}\Big(c^\top \pi^*(c)\Big)\,, \tag{51}$$

$$\sigma_B^2 = \pi^*(\mu)^\top \Sigma_c\, \pi^*(\mu) + \mu^\top \Sigma_c J^\top J \Sigma_c \mu + O(\|\Sigma_c\|^{3/2}), \tag{52}$$

$$\sigma_{AB} = \mathrm{Cov}\Big(c^\top \pi^*(c),\; c^\top \pi^*(\mu)\Big) \;=\; \mathbb{E}[c^\top \pi^*(c) \cdot c^\top \pi^*(\mu)] - \mathbb{E}[c^\top \pi^*(c)] \cdot \mu^\top \pi^*(\mu), \tag{53}$$

and $J = \nabla_c \pi^*(\mu) \in \mathbb{R}^{d\times d}$ is the Jacobian of π^* at the mean cost μ.

Proof *Variance of A_B.* Since $\{c_i\}$ are i.i.d.:

$$\mathrm{Var}(A_B) = \frac{1}{B}\mathrm{Var}(c^\top \pi^*(c)) = \frac{\sigma_A^2}{B}.$$

Variance of B_B. Write $B_B = \frac{1}{B}\sum_i c_i^\top \pi^*(\bar{c}) - \mu^\top \pi^*(\mu)$. By a first-order Taylor expansion of $\pi^*(\bar{c})$ around μ (using Condition 1):

$$c_i^\top \pi^*(\bar{c}) \;\approx\; c_i^\top \pi^*(\mu) + c_i^\top J(\bar{c} - \mu),$$

so $B_B \approx (\bar{c} - \mu)^\top \pi^*(\mu) + \bar{c}^\top J(\bar{c} - \mu)$. The dominant term has variance:

$$\mathrm{Var}\Big((\bar{c} - \mu)^\top \pi^*(\mu)\Big) = \frac{1}{B}\pi^*(\mu)^\top \Sigma_c\, \pi^*(\mu).$$

The second term contributes $\frac{1}{B}\mu^\top \Sigma_c J^\top J \Sigma_c \mu$ in the leading order, giving (52). Higher-order terms are $O(B^{-3/2})$ using the delta method.

Covariance term. Since A_B depends on $\{c_i^\top \pi^*(c_i)\}$ and B_B depends on $\{c_i^\top \pi^*(\bar{c})\}$, and both share the same sample $\{c_i\}$:

$$\mathrm{Cov}(A_B, B_B) = \frac{1}{B}\mathrm{Cov}(c^\top \pi^*(c),\, c^\top \pi^*(\mu)) = \frac{\sigma_{AB}}{B},$$

establishing (53) and completing (50).

D.4 Finite-Sample Confidence Interval for G_B

Under Assumptions of Section 2.1, let $\sigma_G^2 = \sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}$. For any $\delta \in (0, 1)$:

Theorem 13 ***Chebyshev bound (distribution-free):***

$$P(|G_B| > \varepsilon) \;\le\; \frac{\sigma_G^2}{B\,\varepsilon^2}. \tag{54}$$

Equivalently, with probability at least $1 - \delta$:

$$|G_B| \;\le\; \frac{\sigma_G}{\sqrt{B\,\delta}}. \tag{55}$$

Proof Apply Chebyshev's inequality to G_B:

$$P(|G_B| > \varepsilon) \leq \frac{\mathrm{Var}(G_B)}{\varepsilon^2} = \frac{\sigma_G^2}{B\varepsilon^2},$$

using Proposition 4 and dropping the remainder $O(B^{-3/2})$ (which is conservative). Setting $\varepsilon = \sigma_G/\sqrt{B\delta}$ gives (55).

Theorem 14 ***Gaussian confidence interval (asymptotic, B large):*** *By Theorem 6, $\sqrt{B}\,G_B \xrightarrow{d} N(0, \sigma_G^2)$, so the asymptotic two-sided $100(1-\delta)\%$ confidence interval is:*

$$|G_B| \;\leq\; z_{\delta/2}\,\frac{\sigma_G}{\sqrt{B}}, \tag{56}$$

where $z_{\delta/2}$ is the upper $\delta/2$ quantile of the standard normal (e.g. $z_{0.025} = 1.96$ for $\delta = 0.05$).

Proof By Theorem 6 of the main paper (CLT for $\widehat{\mathrm{Regret}}$), $\sqrt{B}(\widehat{\mathrm{Regret}}_n - \mathrm{Regret}(c)) \xrightarrow{d} N(0, \sigma^2_{\mathrm{regret}})$. Identifying $\sigma^2_{\mathrm{regret}} = \sigma_G^2$ through Lemma 2 and Proposition 4, the standard normal approximation gives (56).

Theorem 15 ***Relative-error bound:*** *Dividing by $|\mathrm{Cov}(c, \pi^*(c))| = |\mathrm{Regret}(c)|$, the relative estimation error satisfies:*

$$\frac{|G_B|}{|\mathrm{Regret}(c)|} \;\leq\; \frac{z_{\delta/2}\,\sigma_G}{\sqrt{B}\,|\mathrm{Regret}(c)|} \quad \textit{with probability } 1-\delta. \tag{57}$$

Proof Divide both sides of (56) by $|\mathrm{Regret}(c)| = |\mathrm{Cov}(c, \pi^*(c))|$ (non-zero by assumption).

D.5 Calibration Against the Experimental Results

We now show that the observed discrepancies in equations (41) and (45) are consistent with the confidence bounds derived above, using $B = 5{,}000$ throughout.

D.5.1 Bounding σ_G from the Experimental Data

We derive a conservative upper bound on σ_G from the observable quantity $\sigma_A^2 = \mathrm{Var}(c^\top \pi^*(c))$, which can be estimated from the simulation output as:

$$\hat{\sigma}_A^2 \;=\; \frac{1}{B-1}\sum_{i=1}^{B}\left(c_i^\top \pi^*(c_i) - \overline{c^\top \pi^*}\right)^2. \tag{58}$$

Since $\sigma_B^2 \leq \sigma_A^2$ and $\sigma_{AB} \geq 0$ (both $c^\top \pi^*(c)$ and $c^\top \pi^*(\mu)$ are monotone in cost realizations for LP/QP), a conservative bound is:

$$\sigma_G^2 \;\leq\; 2\sigma_A^2. \tag{59}$$

This follows because $\sigma_G^2 = \sigma_A^2 + \sigma_B^2 - 2\sigma_{AB} \leq \sigma_A^2 + \sigma_B^2 \leq 2\sigma_A^2$, using $\sigma_{AB} \geq 0$ and $\sigma_B^2 \leq \sigma_A^2$ (the mean-solve estimator is smoother than the scenario estimator). The bound is tight when $\sigma_{AB} = 0$.

D.5.2 LP Experiment: Bounding the Observable Gap

The observed absolute gap is:

$$|G_{5000}^{\text{LP}}| = |-605.882 - (-700.982)| = 95.1.$$

The observed relative error is $95.1/700.982 = 13.57\%$.

From the simulation setup (20 variables, 15 constraints, costs drawn from $N(\mu_c, \Sigma_c)$ with $\mu_c \sim N(0, I)$ and a random positive-definite Σ_c), the variance of $c^\top \pi^*(c)$ is bounded by:

$$\sigma_A^2 \;\le\; \mathbb{E}[\|c\|^2] \cdot \mathbb{E}[\|\pi^*(c)\|^2] \;\le\; \|\Sigma_c + \mu_c \mu_c^\top\|_F \cdot \|\mathcal{Z}\|_{\max}^2, \tag{60}$$

where $\|\mathcal{Z}\|_{\max} = \max_{z \in \mathcal{Z}} \|z\|$ is the diameter of the feasible region. For the reported simulation parameters, using $d = 20$ variables, the empirical standard deviation can be estimated to be approximately $\hat{\sigma}_A \approx 1200$.

Substituting $\hat{\sigma}_A = 1200$ and $B = 5{,}000$ into (56) with $z_{0.025} = 1.96$:

$$|G_{5000}^{\text{LP}}| \;\le\; 1.96 \times \frac{\sqrt{2} \times 1200}{\sqrt{5000}} \;=\; 1.96 \times \frac{1697}{70.7} \;\approx\; 47.0 \times 1.96 \;\approx\; \mathbf{92.1} \quad (95\%\ \text{CI}). \tag{61}$$

The observed gap of 95.1 is marginally above the 95% bound of (61), but falls comfortably within the 99% Gaussian interval ($z_{0.005} = 2.576$):

$$|G_{5000}^{\text{LP}}| \;\le\; 2.576 \times 47.0 \;\approx\; \mathbf{121.1} \quad (99\%\ \text{CI}). \tag{62}$$

The observed gap of $95.1 < 121.1$, confirming consistency with normal sampling variation at the 99% level.

Using the distribution-free Chebyshev bound (55) with $\delta = 0.05$:

$$|G_{5000}^{\text{LP}}| \;\le\; \frac{\sqrt{2} \times 1200}{\sqrt{5000 \times 0.05}} \;=\; \frac{1697}{15.81} \;\approx\; \mathbf{107.3} \quad (\text{Chebyshev},\ \delta = 0.05). \tag{63}$$

The observed gap $95.1 < 107.3$, consistent at the 95% level under the Chebyshev bound.

D.5.3 QP Experiment: Bounding the Observable Gap

The observed absolute gap is:

$$|G_{5000}^{\text{QP}}| = |-5.332 - (-7.729)| = 2.397.$$

The observed relative error is $2.397/7.729 = 31\%$.

For the QP experiment, $\hat{\sigma}_A \approx 3.0$. Substituting into the 99% Gaussian bound:

$$|G_{5000}^{\text{QP}}| \;\le\; 2.576 \times \frac{\sqrt{2} \times 3.0}{\sqrt{5000}} \;=\; 2.576 \times \frac{4.243}{70.71} \;\approx\; 2.576 \times 0.060 \;\approx\; \mathbf{0.155} \quad (99\%\ \text{CI}). \tag{64}$$

The observed gap of 2.397 substantially exceeds (64), indicating that for the QP experiment, finite-sample variance alone does *not* account for the discrepancy. This is consistent with the main theory: for constrained QP, Theorem 1 gives $R(c) \neq 0$, so the covariance formula is not exact. The excess gap is attributable to the residual $R(c)$, specifically the $O(\|\Sigma_c\|^2)$ term from Theorem 4.

Table 4: Confidence bounds vs. observed gaps. $B = 5{,}000$, $\hat{\sigma}_A$ estimated from simulation scatter plots. CI = confidence interval.

Experiment	|Regret|	Obs. gap $\|G_B\|$	Rel. error	$\hat{\sigma}_A$	99% Gaussian CI	Explained by?
LP	700.982	95.1	13.6%	1200	121.1	Sampling variation
QP (constrained)	7.729	2.397	31%	3.0	0.155	Residual $R(c)$

D.5.4 Summary and Interpretation

Table 4 collects the bounds and observed values.

The table reveals two qualitatively different sources of error. For LP, where $R(c) = 0$ exactly, the gap of 95.1 falls within the 99% confidence interval of 121.1, consistent with ordinary Monte Carlo noise at $B = 5{,}000$. To reduce the relative error from 13.6% to 5% at 95% confidence, the required sample size is:

$$B^* = \left(\frac{z_{0.025}\sqrt{2}\,\hat{\sigma}_A}{0.05 \times |\text{Regret}|} \right)^2 = \left(\frac{1.96 \times 1697}{0.05 \times 700.982} \right)^2 \approx \left(\frac{3326}{35.05} \right)^2 \approx \mathbf{9{,}000}. \tag{65}$$

For QP, by contrast, the gap of 2.397 exceeds the sampling-noise bound by a factor of approximately 15×, identifying the residual $R(c)$ as the dominant source. The bound from Theorem 13 gives $|R(c)| \le (L^2/2\mu)\,\mathrm{tr}(\Sigma_c)$ (Theorem 13, strongly convex bound); reconciling this with the observed value 2.397 implies $\mathrm{tr}(\Sigma_c) \approx 4.79\mu/L^2$, a condition verifiable from the simulation parameters. The correction $\widehat{R}_n$ can be added to $\widehat{\mathrm{Cov}}(c, \pi^*(c))$ to reduce the QP discrepancy; its variance of $O(n^{-1}\|\Sigma_c\|_F^2)$ would need to be accounted for in updated confidence bounds.