
Beyond IID: data-driven decision-making in heterogeneous environments

Omar Besbes* Will Ma Omar Mouchtaki
Decision, Risk, and Operations Division
Columbia University
{ob2105,wm2428,om2316}@gsb.columbia.edu

Abstract

In this work, we study data-driven decision-making and depart from the classical identically and independently distributed (i.i.d.) assumption. We present a new framework in which historical samples are generated from unknown and different distributions, which we dub *heterogeneous environments*. These distributions are assumed to lie in a heterogeneity ball with known radius and centered around the (also) unknown future (out-of-sample) distribution on which the performance of a decision will be evaluated. We quantify the asymptotic worst-case regret that is achievable by central data-driven policies such as Sample Average Approximation, but also by rate-optimal ones, as a function of the radius of the heterogeneity ball. Our work shows that the type of achievable performance varies considerably across different combinations of problem classes and notions of heterogeneity. We demonstrate the versatility of our framework by comparing achievable guarantees for the heterogeneous version of widely studied data-driven problems such as pricing, ski-rental, and newsvendor. En route, we establish a new connection between data-driven decision-making and distributionally robust optimization.

1 Introduction

In optimization under uncertainty, the desirability of a decision (e.g., inventory) depends on an unknown future outcome (e.g., demand). Typically, past data is collected to be indicative of the future, and hence inform our decision. However, ideal data-driven decision-making requires postulating beliefs about the reliability of the past data, and importantly, whether the future may *deviate* from it. In practice, past data may depend on contexts, some of which can be controlled for, and some of which cannot (unobserved confounders). This may introduce data heterogeneity that is not "correctable."

In this paper, we consider a framework for modeling how the future may deviate from past data. At a high level, it accomplishes three goals: *i.*) capture different forms of future deviation, including no deviation, under the same umbrella; *ii.*) understand the performance of a central policy, which makes the decision that optimizes the *average* objective value over past data, under different forms of unexpected deviation; *iii.*) if the central policy performs poorly, then derive modifications that are robust to different forms of anticipated deviation, and quantify the achievable performance.

This framework is general, capturing different problems, and our analysis leads to insightful problem-specific conclusions. The only assumption is that the policy defined in *ii.*) above can be computed, which requires perfect counterfactual evaluation of any decision on all data points, as opposed to settings where past data are affected by previous actions [5, 16, 33, 41, 42, 43]. The policy described in *ii.*) is widely studied across different fields, known as *Sample Average Approximation (SAA)* [37, 36, 39], although we emphasize that we also derive *new policies beyond SAA* when it falters.

* Authors ordered alphabetically.

1.1 Framework description

Framework (Section 2.1). Our first contribution lies in studying a framework which models the future outcome as being drawn from an unknown distribution, and past data as being drawn independently from (also unknown) "nearby" distributions. In this way, past data is indicative of the future, and we emphasize that the samples of data can be drawn from different nearby distributions. In accord, we call this setting a *heterogeneous environment*. Using varying definitions of "nearby distribution" (e.g. based on Kolmogorov/Wasserstein distances; see Section 1.3), we analyze different forms of deviation possible between the past and future, and let a radius ϵ bound the allowed deviation. When $\epsilon = 0$, our framework captures the classical independent and identically distributed (i.i.d.) setting, in which past data are drawn from the same distribution as the future outcome [37, 45, 4]. This framework presents another possible parametrization to interpolate between the i.i.d. setting and the adversarial one and shares conceptually similar goals to previous approaches considering algorithm analysis beyond worst-case [51, 7, 53, 25, 10, 24].

Performance measure (Section 2.2). We consider the performance of different data-driven policies, which map past data into a decision for a given problem; SAA defines a feasible data-driven policy for any problem. Regret is measured as the difference in objective between the policy's decision and the optimal decision knowing the future distribution, taking an expectation over the draws of past data (which affect the policy's decision), any intrinsic randomness in the policy, and the outcome realization (which affects the objective of both the policy's decision and the optimal decision). We then take a worst case regret over all possible distributions that could have been chosen for the future outcome and the data points, to evaluate the performance of a fixed data-driven policy. We note that this performance depends on the problem (see Section 1.3 for examples), the number of data points n , the definition of "nearby distribution," and the radius ϵ .

1.2 General results

Our second contribution lies in developing a set of general results for this framework. A central result we present in Section 3 is a general reduction we develop in Theorem 1 in which we establish the following: asymptotically, as the number of samples goes to ∞ , the worst-case regret of a policy (in a broad subclass we call sample-size-agnostic) is upper-bounded by that of an alternative problem with two major simplifications: *i.*) the worst-case is now taken over only a single distribution for historical environments as opposed to a sequence of heterogeneous distributions; and *ii.*) the decision-maker (DM) has access to the exact historical distribution. This upper bound can be seen as a uniform distributionally robust optimization problem under a regret performance metric, and offers a systematic way to derive upper bounds for policies of interest. This result is closely related to equation (7) in [46] and we discuss in detail the connections to their result after stating Theorem 1.

We develop problem-independent results that hold whenever the objective function is Lipschitz as a function of the future outcome that materializes (Section 4). In that case, we establish in Theorem 2 that SAA scales linearly in ϵ for both the Kolmogorov and Wasserstein forms of deviation. We also show that for this class of problems, the Lipschitz constant characterizes the hardness of the instance.

1.3 Problem-specific, deviation-specific results

Our third contribution lies in deriving problem-specific conclusions in our general heterogeneous environments. In particular, we consider three classical problems in the context of our framework—newsvendor, pricing, and ski-rental—that have received significant attention in the literature across Operations Research, Economics, and Computer Science in the known-prior, data-driven, distributionally robust, adversarial, and advice-augmented settings. Problem-specific results are obtained by using: *i.*) our general reduction and results to upper-bound their asymptotic worst-case regret, and *ii.*) a combination of algorithm design and the development of impossibility results (lower bounds) on the worst-case regret. We analyze the regret for the SAA policy as well as alternative policies.

Problem descriptions. First, we consider the newsvendor problem (Section 4.1), in which a DM must decide how much capacity or inventory to plan in the face of unknown demand. The objective is to minimize the sum of *underage cost*, paid for each unit by which capacity falls short of realized demand, and *overage cost*, paid for each unit of capacity in excess of demand that needs to sit idle. This is a foundational problem in Operations Research, with early focus on the setting where the demand distribution is known [see 47], and since then studied in various distributionally robust

[54, 22, 49] settings but also in data-driven ones [40, 15, 6]. Newsvendor captures the fundamental tradeoff in guessing an unknown value when there are asymmetric consequences for going under/over.

Second, we consider the single-item pricing problem (Section 5.1), in which a DM must find the best price to offer to a customer with unknown willingness to pay (wtp). The objective is to maximize revenue. This is a foundational problem in Economics, with early focus on the setting with known prior [see 48, 52], and more recently in data-driven [21, 29, 2, 1, 17] and distributionally robust [12, 23] settings. Pricing captures the fundamental tradeoff between margin and volume.

Finally, we consider the ski-rental problem (Section 5.2), in which a DM must decide each day whether to buy skis or keep renting, without knowing the length of their trip. The objective is to minimize total cost, where buying eliminates the rental cost going forward. Although typically presented as an online problem that requires a decision to be made every day, any (randomized) buy-or-rent policy can be described by a (randomized) duration for which to rent skis, before committing to buying (if the trip has not ended by then). This is a foundational problem in Computer Science, critical to the development of competitive analysis [see 11] which considers the *adversarial* setting where nothing is known, and recently also important to the development of *advice-augmented algorithms* which receive a prediction on the length of the ski trip [50, 19]. Ski rental captures the fundamental tradeoff between proceeding with a suboptimal process vs. paying the cost of refactoring.

Definitions of "nearby distribution." Our framework generally models past data to be drawn from distributions that lie within an ϵ -distance of the future distribution, for *any* notion of distance between probability distributions. For our problems of interest, we focus on two specific distance metrics: Kolmogorov and Wasserstein (see Section 2.1 for formal definitions). This dichotomy suffices to demonstrate the importance of distance metric in determining performance in our problems of interest, and moreover, represents contrasting inductive philosophies about how the future might deviate from the past: Wasserstein allows *every* data point to be erroneous by ϵ , whereas Kolmogorov allows an ϵ fraction of data points to be *arbitrarily* erroneous. We note that similar notions of "nearby distributions" were considered in [51] in the online learning setting and in [7, Example 5] for the expert selection problem. Our work differs from these two papers because we consider an offline setting and our benchmark is stronger since we focus on the oracle knowing the ground-truth out-of-sample distribution, as opposed to the "best action in hindsight" benchmark.

SAA vs. robustified policies. Finally, performance is demarcated by the policy used. We first provide tight asymptotic performance bounds for the SAA policy, which we note is agnostic to the heterogeneity that might arise. We then analyze the best-possible asymptotic performance bounds that can be achieved by any policy when the notion of distance and the radius ϵ are known.

In Table 1, we present a high level summary of our results across problem classes. We assume that the support of the unknown environment is in $[0, M]$ for some positive real number M which M can be interpreted as parametrizing the "extent" of uncertainty one faces (e.g., maximal values of demand in newsvendor) We track the dependence on both the heterogeneity radius ϵ , as well as M . These results are for the asymptotic setting in which $n \rightarrow \infty$.

			Kolmogorov	Wasserstein
Lipschitz	Newsvendor (§4.1)	SAA	$\Theta(M\epsilon)$	$\Theta(\epsilon)$
		best policy	$\Theta(M\epsilon)$	$\Theta(\epsilon)$
Beyond Lipschitz	Pricing (§5.1)	SAA	$\Theta(M\epsilon)$	$\Omega(M)$
		best policy	$\Theta(M\epsilon)$	$\Theta(\sqrt{M\epsilon})$
	Ski-rental (§5.2)	SAA	$\Theta(M\epsilon)$	$\Theta(1)$
		best policy	$\tilde{\Theta}(\epsilon)^*$	$\Theta(\sqrt{\epsilon})$

Table 1: **SAA and achievable performance for different problem classes and heterogeneity balls.**

*(Here $\tilde{\Theta}$ provides rate order up to logarithmic factors.)

For newsvendor problems, through a combination of the general result for Lipschitz-continuous problems and lower bounds, we show that SAA actually achieves the optimal rate for asymptotic worst-case regret both as a function of ϵ and M and that this rate is linear.

We then turn to problems whose objective is not Lipschitz. For the pricing problem, we establish that while under Kolmogorov distance, SAA is near-optimal and the optimal performance is again

of order $M\epsilon$, the picture is starkly different under the Wasserstein distance. In this case, SAA has arbitrarily poor performance. We then show that a robustification procedure that deflates SAA by an appropriate factor can yield the best achievable asymptotic worst-case regret that shrinks to zero at rate $\sqrt{M}\epsilon$. To our understanding, analyzing these non-SAA policies (which are required for good performance on pricing) deviates from standard analyses used in learning theory (see Remark 2).

A second prototypical problem that is not Lipschitz is the ski-rental problem. While SAA was near-optimal for the Kolmogorov distance for pricing and newsvendor, we establish the surprising fact that SAA, even under the Kolmogorov distance, has arbitrarily poor performance. Indeed, its asymptotic worst-case regret scales as $\Theta(M\epsilon)$, while the optimal policy achieves a performance of $\Theta(\epsilon)$. As a consequence, the performance of SAA can be arbitrarily worse than the optimal one as the support grows to infinity. Meanwhile, for the Wasserstein distance, we show that the performance of SAA is also suboptimal. Indeed, SAA incurs a regret of $\Theta(1)$; but by inflating SAA appropriately, one can construct a policy that achieves order $\sqrt{\epsilon}$ performance, which is the best dependence on ϵ . These examples demonstrate that SAA fails for both heterogeneity types for ski-rental, but also more broadly the intricate interplay between problem class and type of heterogeneity.

Connection to offline data corruption models. Different notions of offline data corruptions and robustness have been previously studied in statistics [30, 26], learning theory [55, 28, 35, 34, 55, 38, 18, 8, 20] and, more recently in revenue management [13, 12, 14, 23]. Most approaches consider ϵ -contamination settings [30] in which the adversary can shift a small proportion of data arbitrarily far from the true distribution. They consider either oblivious adversaries who fix a *single* corrupted distribution from which samples are generated or adaptive one who observe samples before corrupting them. [8] establishes that adaptive adversaries yield equivalent performance to oblivious ones in many settings. Closest to us are [12, 23] who study optimal auctions. They assume the distribution is common across all past observations; the former focuses on regret and the latter on a ratio metric when data is generated from MHR distributions close in Kolmogorov distance. The framework we develop is anchored around a setting that allows for *heterogeneous* "nearby" distributions for each past observation, and is general in that it allows us to unify a variety of problems/metrics and highlight how these affect the levels of achievable performance.

Connection to online and batch learning in non-stationary environments. Our work also relate broadly to papers studying learnability in changing environment. [27, 9] adopt a universal learning approach and analyze algorithms which are able to achieve vanishing long-term average loss for general data-generation processes. [44, 3, 46] derive generalization guarantees in a setting where the training set distribution may differ from the test set distribution.

2 Problem formulation: data-driven decisions in heterogeneous environments

We consider a general stochastic optimization problem in which a DM wants to optimize an objective in a stochastic environment. Formally, we denote by \mathcal{X} the decision space and by Ξ the space of environment realizations. The objective of the DM is represented by a function $g : \mathcal{X} \times \Xi \rightarrow \mathbb{R}$ which maps each pair of decision $x \in \mathcal{X}$ and environment realization $\xi \in \Xi$ to an objective value.

The DM faces a subset of possible probability measures $\mathcal{P} \subset \Delta(\Xi)$, where $\Delta(\Xi)$ denotes the space of probability measures supported on Ξ . For every probability measure $\mu \in \mathcal{P}$, the goal of the DM is to optimize the objective defined as

$$\mathcal{G}(\cdot, \mu) : \begin{cases} \mathcal{X} \rightarrow \mathbb{R} \\ x \mapsto \mathbb{E}_{\xi \sim \mu} [g(x, \xi)] \end{cases} \quad (1)$$

We extend the definition of $\mathcal{G}(\cdot, \mu)$ to allow for randomized decisions $\rho \in \Delta(\mathcal{X})$, in which case we define $\mathcal{G}(\rho, \mu) = \mathbb{E}_{x \sim \rho} [\mathcal{G}(x, \mu)]$ to also take an expectation over the randomized decision.

Remark that depending on the nature of the problem, the objective could represent for example a function to maximize (profit) or to minimize (cost). In turn, we formally define for every $\mu \in \mathcal{P}$, the optimal objective $\text{opt}(\mu)$ as the objective value achieved by an *optimal action*, given by an *oracle* operator ORACLE that is a mapping from $\Delta(\Xi)$ to \mathcal{X} .

2.1 Heterogeneous environments

In the absence of distributional knowledge on the underlying environment, a common approach consists in assuming that historical samples are independent and identically drawn from a fixed probability measure μ . We relax the assumption that the samples are identically distributed and introduce a generalization of this framework that we refer to as ϵ -heterogeneity.

Let $d : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$ be a metric on \mathcal{P} . Given a target probability measure $\mu \in \mathcal{P}$ and a level of heterogeneity ϵ , we define the heterogeneity ball anchored around μ and with radius ϵ as

$$\mathcal{U}_\epsilon(\mu) := \{\nu \in \mathcal{P}, \text{ s.t } d(\mu, \nu) \leq \epsilon\}.$$

The heterogeneity ball $\mathcal{U}_\epsilon(\mu)$ includes all distributions that are within a distance of ϵ from μ . Next, we formally define a data-driven decision-making problem in a heterogeneous environment.

Definition 1 (Data-driven decision-making problem in heterogeneous environment). *A data-driven decision-making problem in a heterogeneous environment is defined by the tuple: $(\mathcal{X}, \Xi, \mathcal{P}, d, g, \text{ORACLE})$, where \mathcal{X} is the set of possible actions, Ξ is the environment space, \mathcal{P} is the space of measures nature can select from, d is a metric on the space of measures, g is an objective function, and ORACLE is a mapping from environment measures to target actions.*

Focal metrics. The definition above is general, and the framework can be applied to a variety of metrics. In this work, we aim at developing an understanding of the effect of heterogeneity under various metrics and problem classes. We focus on two prototypical and central distances: *Kolmogorov* and *Wasserstein*. We formally define them in one dimension as follows.

Definition 2 (Kolmogorov and Wasserstein distances). *Given an environment space $\Xi \subset \mathbb{R}$ and a subset of probability measures \mathcal{P} supported on Ξ , the Kolmogorov (resp. Wasserstein) distance d_K (resp. d_W) is defined for every $\mu_1, \mu_2 \in \mathcal{P}$ as*

$$d_K(\mu_1, \mu_2) = \sup_{\xi \in \Xi} |F_1(\xi) - F_2(\xi)| \quad \text{and} \quad d_W(\mu_1, \mu_2) = \int_{\xi \in \Xi} |F_1(\xi) - F_2(\xi)| d\xi,$$

where F_1 (resp. F_2) is the cumulative distribution associated with μ_1 (resp. μ_2).

Note that when Ξ is a compact interval of \mathbb{R} , both distances are related as for every $\mu_1, \mu_2 \in \mathcal{P}$,

$$d_W(\mu_1, \mu_2) \leq \text{diam } \Xi \cdot d_K(\mu_1, \mu_2). \quad (2)$$

2.2 Data-driven policies and performance

Data-driven policies. In most settings, the probability measure μ is unknown to the DM and the optimal objective $\text{opt}(\mu)$ is not achievable. In turn, the DM observes a sequence of historical samples $\xi_1, \xi_2, \dots, \xi_n$ representing the previous environment realizations. A data-driven policy is then a mapping from the past samples to a (randomized) decision. Define the empirical measure $\hat{\mu}_{\xi^n}$ as

$$\hat{\mu}_{\xi^n}(\xi) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{\xi = \xi_i\}, \quad \text{for every } \xi \in \Xi.$$

In a setting in which the order of the samples does not matter (as the one we will consider), to represent a policy taking as input $\xi^n := (\xi_1, \dots, \xi_n)$, one only needs to consider mappings that take as an input the number of samples and the empirical measure. We define a data-driven policy π as a mapping from $\mathbb{N}^* \times \mathcal{P}$ to $\Delta(\mathcal{X})$, the space of measures on the action space \mathcal{X} , that associates $(n, \hat{\mu})$ to a distribution of actions $\pi(n, \hat{\mu})$. We let Π denote the set of all such mappings. We assume that such a mapping can be made with knowledge of the metric d and radius ϵ . Moreover, we say that π is *sample-size-agnostic* if its decision does not depend on n , in which case we write $\pi(\hat{\mu})$ for brevity.

A widely-studied policy that achieves excellent performance across a vast range of data-driven decision-making problems is Sample Average Approximation (SAA). Given an empirical distribution $\hat{\mu}$, SAA (denoted by π^{SAA}) selects the target action ORACLE($\hat{\mu}$), the best action associated with $\hat{\mu}$.

Performance through worst-case regret. Given an out-of-sample distribution $\mu \in \mathcal{P}$ the best achievable objective is $\text{opt}(\mu)$. For a sequence of historical distributions ν_1, \dots, ν_n in $\mathcal{U}_\epsilon(\mu)$, the expected objective of a policy π is

$$\mathbb{E}_{\xi_i \sim \nu_i} [\mathcal{G}(\pi(n, \hat{\mu}_{\xi^n}), \mu)].$$

In turn, we define the expected (absolute) regret of a policy $\pi \in \Pi$ against an out-of-sample distribution $\mu \in \mathcal{P}$ and a sequence of n historical distributions $\nu_1, \dots, \nu_n \in \mathcal{U}_\epsilon(\mu)$ as

$$\mathcal{R}_n(\pi, \mu, \nu_1, \dots, \nu_n) := |\mathbb{E}_{\xi_i \sim \nu_i} [\mathcal{G}(\pi(n, \hat{\mu}_{\xi^n}), \mu)] - \text{opt}(\mu)|.$$

We note that we use the absolute value in this definition, as our framework applies to minimization and maximization problem. Once the type of problem is specified (i.e., the operator associated with ORACLE), this definition of regret coincides with the usual regret for minimization or maximization problems, as we will see in Sections 4 and 5, when describing specific applications.

In this work, we will assess the performance of a policy based on its *worst-case expected regret*. Given an instance $\mathcal{I} = (\mathcal{X}, \Xi, \mathcal{P}, d, g, \text{ORACLE})$, we define the worst-case expected regret of a data-driven policy $\pi \in \Pi$ for a degree of heterogeneity $\epsilon \geq 0$ as

$$\mathfrak{R}_{\mathcal{I},n}^\pi(\epsilon) := \sup_{\mu \in \mathcal{P}} \sup_{\nu_1, \dots, \nu_n \in \mathcal{U}_\epsilon(\mu)} \mathcal{R}_n(\pi, \mu, \nu_1, \dots, \nu_n). \quad (3)$$

We finally define the *worst-case asymptotic regret* of a data-driven policy $\pi \in \Pi$ in a ϵ -heterogeneous environment as the number of samples n grows large as

$$\mathfrak{R}_{\mathcal{I},\infty}^\pi(\epsilon) := \limsup_{n \rightarrow \infty} \mathfrak{R}_{\mathcal{I},n}^\pi(\epsilon). \quad (4)$$

We sometimes write $\mathfrak{R}_{K,\infty}^\pi(\epsilon)$ or $\mathfrak{R}_{W,\infty}^\pi(\epsilon)$ instead of $\mathfrak{R}_{\mathcal{I},\infty}^\pi(\epsilon)$ when the problem class is specified to highlight the dependence on the Kolmogorov or Wasserstein metrics.

3 Reduction to distributionally robust optimization in the asymptotic regime

A significant challenge in the analysis of problem (4) stems from the difficulty to characterize a potentially complex worst-case sequence of historical probability measures ν_1, \dots, ν_n and out-of-sample measure μ . We illustrate in the appendix that analyzing every element in the sequence of historical probability measures as a function of n is a priori challenging as nature will attempt to use different distributions for each of the past environments. We will show that, asymptotically, it suffices to consider cases in which nature selects a common distribution for all of the past environments. Formally, we define for a sample-size-agnostic policy $\pi \in \Pi$ the following optimization problem:

$$\sup_{\mu \in \mathcal{P}} \sup_{\nu \in \mathcal{U}_\epsilon(\mu)} |\text{opt}(\mu) - \mathcal{G}(\pi(\nu), \mu)|. \quad (5)$$

While Problem (5) resembles Problem (3), it differs on two crucial dimensions: *i.*) in Problem (5), nature selects the *same* distribution ν in the heterogeneity ball for all past environments; *ii.*) in Problem (5), the policy π is assumed to know the actual distribution of the past environments ν , whereas in Problem (3), the policy only observed the empirical distribution of past realizations.

Next, we will establish that Problems (5) and (3) are tightly connected under some mild assumptions. We introduce the following definition on the distance d defining the heterogeneity.

Definition 3 (Empirical triangular convergence). *We say that a distance d on the space of probability measures satisfies the empirical triangular convergence (ETC) property if and only if for every triangular array sequence of probability measures $(\mu_{i,n})_{1 \leq i \leq n, n \in \mathbb{N}^*}$ all belonging to \mathcal{P} , we have*

$$\lim_{n \rightarrow \infty} d(\hat{\mu}_n, \bar{\mu}_n) = 0 \quad \text{a.s.},$$

where $\hat{\mu}_n(\xi) := \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{\xi = \xi_{i,n}\}$ for samples $\xi_{i,n} \sim \mu_{i,n}$ and $\bar{\mu}_n = \frac{1}{n} \sum_{i=1}^n \mu_{i,n}$.

The ETC property requires that the sequence of empirical distributions converges to the average of ground truth distributions for arbitrary triangular arrays. In what follows we also impose convexity on the metric d . We note that both properties are satisfied by many common distances on probability spaces, including the Kolmogorov and Wasserstein distances on compact sets, as we discuss in the appendix. We now state our first main result which establishes a fundamental reduction.

Theorem 1 (Upper bound reduction). *Given an instance $\mathcal{I} = (\mathcal{X}, \Xi, \mathcal{P}, d, g, \text{ORACLE})$, let $\pi \in \Pi$ be a sample-size-agnostic policy. Assume that g is bounded on $\mathcal{X} \times \Xi$ and that d is a convex metric which satisfies the empirical triangular convergence property. Then*

$$\mathfrak{R}_{\mathcal{I},\infty}^\pi(\epsilon) \leq \lim_{\eta \rightarrow \epsilon^+} \sup_{\mu \in \mathcal{P}} \sup_{\nu \in \mathcal{U}_\eta(\mu)} |\text{opt}(\mu) - \mathcal{G}(\pi(\nu), \mu)|. \quad (6)$$

Theorem 1 establishes, that under some mild conditions, in the asymptotic regime, one may restrict attention to problem (5). This reduction significantly simplifies the analysis of sample-size-agnostic policies² as problem (5) does not involve any form of randomness in the input received by the policy, and only allows the adversary to select a single distribution.

Remark 1 (Relation to [46]). *We note that [46, eq. (7)] presents a finite-time bound on the regret of SAA in a similar setting to ours. Therefore by taking the limit in their bound, one may derive an asymptotic performance guarantee for SAA. In our notation, their result imply that,*

$$\mathfrak{R}_{\mathcal{I},\infty}^{\pi^{\text{SAA}}} \leq 2 \cdot \sup_{\mu \in \mathcal{P}} \sup_{\nu \in \mathcal{U}_\epsilon(\mu)} \sup_{x \in \mathcal{X}} |\mathcal{G}(x, \mu) - \mathcal{G}(x, \nu)|. \quad (7)$$

Theorem 1 differs from (7) because our reduction applies to any policy in the wide range of sample-size-agnostic policies. In particular, this enables us to characterize the minimax optimal asymptotic regret even when SAA is not (rate) optimal (see Sections 5.1 and 5.2).

We also show in the appendix that Theorem 1 holds whenever Ξ is finite without restricting the policy.

4 Lipschitz-continuous problems

In this section, we study a broad class of problems for which the underlying objective function $g(x, \cdot)$ is Lipschitz-continuous as a function of the environment.

Theorem 2 (Upper bounds for Lipschitz-continuous objectives). *Suppose that $\Xi = [0, M]$ for some $M > 0$. Let $\mathcal{I}_K = (\mathcal{X}, \Xi, \mathcal{P}, d_K, g, \text{ORACLE})$ and $\mathcal{I}_W = (\mathcal{X}, \Xi, \mathcal{P}, d_W, g, \text{ORACLE})$ be instances of the data-driven decision problem under the Kolmogorov and Wasserstein distances, respectively. Assume that for every $x \in \mathcal{X}$, the function $g(x, \cdot)$ is L -Lipschitz-continuous and that ORACLE is either a minimization or maximization oracle. Then for every $\epsilon \geq 0$,*

$$\mathfrak{R}_{\mathcal{I}_K,\infty}^{\pi^{\text{SAA}}}(\epsilon) \leq 2L \cdot M \cdot \epsilon \quad \text{and} \quad \mathfrak{R}_{\mathcal{I}_W,\infty}^{\pi^{\text{SAA}}}(\epsilon) \leq 2L \cdot \epsilon.$$

Theorem 2 implies that for this class of problems, the asymptotic regret of SAA vanishes as the radius of the heterogeneity ball goes to 0. Furthermore, this theorem shows that the dependence of the asymptotic worst-case regret in the heterogeneity parameter ϵ is at most linear. We show next that this dependence is the best possible for this class of problems and this notion of heterogeneity.

4.1 Newsvendor problem

Recall the newsvendor problem described in Section 1.3. Using $M > 0$ to denote an upper bound on demand, we let $\Xi = \mathcal{X} = [0, M]$, and our regret bounds may (or may not) depend on M . The objective function is the newsvendor cost which is parameterized by two quantities: c_u the underage cost and c_o the overage cost. The cost of an inventory level $x \in \mathcal{X}$ when observing demand $\xi \in \Xi$ is,

$$g(x, \xi) = c_u (\xi - x)^+ + c_o (x - \xi)^+,$$

where $(\cdot)^+ := \max\{\cdot, 0\}$ is the positive part operator. The measure space \mathcal{P} is the set of probability measures on Ξ without restrictions and for every $\mu \in \mathcal{P}$, we define the cost of the oracle as $\text{opt}(\mu) = \min_{x \in \mathcal{X}} \mathcal{G}(x, \mu)$. Therefore the oracle mapping is defined as $\text{ORACLE} : \mu \mapsto \arg \min_{x \in \mathcal{X}} \mathcal{G}(x, \mu)$.

We note that the objective function g satisfies the L -Lipschitz-continuity property on the environment realization variable ξ with $L = \max(c_u, c_o)$. Theorem 2 thus directly yields an upper bound on the achievable regret for this class of problems, for both Kolmogorov and Wasserstein, through the regret of SAA. We next show that these dependencies are the best possible, under any policy.

Proposition 1 (Heterogeneous newsvendor). *For the data-driven newsvendor problem in a heterogeneous environment under the Kolmogorov and Wasserstein distances, given any $\epsilon \geq 0$,*

$$\begin{aligned} \frac{c_u + c_o}{2} \cdot \epsilon \cdot M &\leq \inf_{\pi \in \Pi} \mathfrak{R}_{K,\infty}^{\pi}(\epsilon) \leq \mathfrak{R}_{K,\infty}^{\pi^{\text{SAA}}}(\epsilon) \leq 2 \max(c_u, c_o) \cdot \epsilon \cdot M, \\ \frac{c_u + c_o}{2} \cdot \epsilon &\leq \inf_{\pi \in \Pi} \mathfrak{R}_{W,\infty}^{\pi}(\epsilon) \leq \mathfrak{R}_{W,\infty}^{\pi^{\text{SAA}}}(\epsilon) \leq 2 \max(c_u, c_o) \cdot \epsilon. \end{aligned}$$

²Note that considering sample-size-agnostic policies is not very restrictive. As we will see in Section 4 and 5 this class is sufficient to achieve optimal rates for the asymptotic worst-case regret as a function of ϵ .

Proposition 1 first shows that SAA is robust to deviation under multiple forms of heterogeneity and that the regret scales linearly with the radius of the heterogeneity ball for both distances. We show furthermore that the linear dependence is, for the newsvendor problem, the best rate achievable for any data-driven policy. We also note that the lower bound obtained for the newsvendor depends linearly on the Lipschitz constant L through the factor $(c_u + c_o)$.

Beyond newsvendor, the asymptotic worst-case regret of SAA over the general class of L -Lipschitz-continuous functions is of the order $\Theta(L \cdot \epsilon)$. This suggests that the performance of SAA may deteriorate for problems in which the objective function g is not smooth.

5 Beyond Lipschitz-continuous problems

5.1 Pricing problem

Recall the pricing problem presented in Section 1.3. The environment space $\Xi = [0, M]$ represents the set of possible wtp. Similarly, the set of possible prices \mathcal{X} is $[0, M]$. The set of considered probability measures \mathcal{P} is the whole set of probability measures on $[0, M]$ without restriction. The revenue generated by the DM with a price decision $x \in \mathcal{X}$ when facing a customer with wtp $\xi \in \Xi$ is,

$$g(x, \xi) = x \cdot \mathbb{1}\{x \leq \xi\}.$$

Therefore for any probability measure $\mu \in \mathcal{P}$ associated with the cumulative distribution F , we have that $\mathcal{G}(x, \mu) = x \cdot (1 - F(x))$. The goal of the DM is to set a price x in order to maximize its revenue. Equivalently, the goal is to minimize the regret against the oracle which posts the optimal price given a wtp distribution. Formally, the oracle operator is defined as, $\text{ORACLE} : \mu \mapsto \arg \max_{x \in \mathcal{X}} \mathcal{G}(x, \mu)$.

Kolmogorov pricing. Remark that the pricing objective is not Lipschitz-continuous and that one cannot use the argument presented in Section 4.1. We show in the appendix that despite the lack of smoothness of the objective, SAA is still near-optimal in the asymptotic regime under the Kolmogorov distance. We derive an upper bound on the asymptotic worst-case regret of SAA and a matching universal lower bound showing that the best achievable performance is order $\Theta(M\epsilon)$.

Failure of SAA in pricing with Wasserstein heterogeneity. In opposition to the strong performance achieved by SAA for pricing against Kolmogorov heterogeneity, our next proposition shows that this result does not hold if one considers instead the Wasserstein heterogeneity ball.

Proposition 2 (SAA and Wasserstein pricing). *For the data-driven pricing problem in a heterogenous environment under the Wasserstein distance, we have for $\epsilon > 0$ that*

$$\mathfrak{R}_{W, \infty}^{\pi^{\text{SAA}}}(\epsilon) = M.$$

Proposition 2 shows that, in stark contrast with Lipschitz problems, for the pricing problem, the asymptotic regret of SAA under Wasserstein heterogeneity does not shrink to zero as ϵ goes to zero. As a matter of fact, the regret of M is the worst possible, as the revenue of the oracle is bounded from above by M and the revenue of any data-driven policy is bounded from below by 0. Quite notably, this result holds for any level of heterogeneity ϵ , and an arbitrarily small deviation from the i.i.d. case leads to extremely poor performance, even with infinite data. We next propose a policy which "robustifies" SAA and we analyze it under the Wasserstein distance.

Robustification of SAA for Wasserstein pricing. The key challenge in pricing is that for two distributions μ and ν such that $d_W(\mu, \nu)$ is arbitrarily small, a price x could have a high revenue for μ but a low revenue for ν , and the difference in revenue between distributions may be arbitrarily large. This enabled us to construct an instance for which the optimal price for the in-sample distribution ν generates a low revenue for the out-of-sample distribution μ .

Remark 2 (Proof Sketch, Comparison with Standard Techniques). *To the best of our understanding, our analysis of policies beyond SAA requires a new technique. Assume $M = 1$ for brevity, and consider the problem of pricing. We need to upper-bound the loss $\mathcal{G}(\text{ORACLE}(\mu), \mu) - \mathcal{G}(\pi(\nu), \mu)$. A standard analysis of SAA, which satisfies $\pi^{\text{SAA}}(\nu) = \text{ORACLE}(\nu)$, would write,*

$$\mathcal{G}(\text{ORACLE}(\mu), \mu) - \mathcal{G}(\text{ORACLE}(\nu), \mu) \leq 2 \sup_{x \in \mathcal{X}} |\mathcal{G}(x, \mu) - \mathcal{G}(x, \nu)|.$$

However, for certain decisions x , the error $|\mathcal{G}(x, \mu) - \mathcal{G}(x, \nu)|$ can be non-vanishing even if the distance $d_W(\mu, \nu)$ is vanishing. Our fix is the following. For a generic policy π , we write

$$\begin{aligned} \mathcal{G}(\text{ORACLE}(\mu), \mu) - \mathcal{G}(\pi(\nu), \mu) &\leq (\mathcal{G}(\text{ORACLE}(\mu), \mu) - \mathcal{G}(\pi(\mu), \nu)) \\ &\quad + (\mathcal{G}(\text{ORACLE}(\nu), \nu) - \mathcal{G}(\pi(\nu), \mu)) \end{aligned} \quad (8)$$

The key observation is that if we set $\pi(\nu)$ not to $\text{ORACLE}(\nu)$, but to $\text{ORACLE}(\nu) - \delta$, then both large parentheses in (8) can be (identically) upper-bounded using the relation we introduce in Proposition 3.

Proposition 3. *Let μ and $\nu \in \mathcal{P}$ and let x_1 and x_2 be two prices such that $x_1 < x_2$. We have that,*

$$\mathcal{G}(x_2, \nu) - \mathcal{G}(x_1, \mu) \leq (x_2 - x_1) + M \cdot \frac{d_W(\mu, \nu)}{x_2 - x_1}.$$

Proposition 3 unveils an interesting tradeoff faced by a DM who prices in a heterogeneous environment. By selecting appropriately the difference between two prices x_1 and x_2 , one can ensure that the gap in revenues for two different distributions can be controlled by the Wasserstein distance.

We next restrict attention to the class of policies which additively *deviate* from the action selected by SAA and leverage Proposition 3 to define the correct deviation parameter.

Definition 4 (δ -SAA policies). *For any one-dimensional problem we say that a policy is a δ -SAA policy if for every empirical distribution $\hat{\mu}$, the policy selects the closest action in \mathcal{X} to $\text{ORACLE}(\hat{\mu}) + \delta$, for some $\delta \in \mathbb{R}$ that could be positive or negative. We denote this policy by $\pi^{\text{SAA}(\delta)}$.*

We note that when $\delta = 0$, we recover the usual SAA policy. Our next result establishes the best possible regret dependence on ϵ and M for the class of pricing problems, using a δ -SAA policy.

Theorem 3 (Deviation for Wasserstein pricing). *For the data-driven pricing problem in a heterogeneous environment under the Wasserstein distance, for any $\epsilon > 0$, let $\tilde{\delta} = -\sqrt{M \cdot \epsilon}$. Then,*

$$\frac{1}{4} \cdot \sqrt{M \cdot \epsilon} \leq \inf_{\pi \in \Pi} \mathfrak{R}_{W, \infty}^{\pi}(\epsilon) \leq \mathfrak{R}_{W, \infty}^{\pi^{\text{SAA}(\tilde{\delta})}}(\epsilon) \leq 4 \cdot \sqrt{M \cdot \epsilon}.$$

Theorem 3 shows that by appropriately deflating the price posted by SAA, one is able to be robust against Wasserstein heterogeneity and achieve a worst-case asymptotic regret that vanishes to 0 as the degree of heterogeneity ϵ goes to 0. This proves that one may operate efficiently for pricing under Wasserstein heterogeneity as ϵ goes to 0. Our result establishes that in contrast with Lipschitz problems, it is now impossible to achieve a linear dependence in ϵ (since $\sqrt{\epsilon} > \epsilon$), with the regret growing at a rate of $\Omega(\sqrt{M \cdot \epsilon})$ under any policy. A small radius ϵ has a significantly higher impact in pricing problems than in, e.g., newsvendor problems.

We note that the deflated policy proposed in Theorem 3 improves over SAA when one can make use of the knowledge of ϵ . It is worth noting that SAA does not need such knowledge. Furthermore, one can show that this deflated policy incurs a $\Omega(\sqrt{\epsilon})$ regret for pricing under Kolmogorov distance and thus performs poorly compared to SAA in that setting. A natural question would be to understand the best achievable performance without knowledge of ϵ and/or d .

5.2 Ski-rental problem

We now consider the ski-rental problem in which renting skis costs \$1 per unit of time while buying them costs \$ b up-front, for some real value b . The environment space Ξ represents the set of possible lengths of the ski trip and a decision x represents the duration after which skis should be bought (if the ski trip has not ended by that time). We call x the *rental duration*. Let $\mathcal{X} = \Xi = [0, M]$, where we note that setting $x = M$ indicates that skis should never be bought. The set of probability measures \mathcal{P} is the whole space of probability measures supported on $[0, M]$. For every rental duration $x \in \mathcal{X}$ and any trip length $\xi \in \Xi$, the cost incurred is

$$g(x, \xi) = \xi \cdot \mathbb{1}\{\xi \leq x\} + (b + x) \cdot \mathbb{1}\{\xi > x\}.$$

Finally, as the goal is to minimize cost we have that, $\text{ORACLE} : \mu \mapsto \arg \min_{x \in \mathcal{X}} \mathcal{G}(x, \mu)$.

Wasserstein ski-rental. We proved in Section 5.1 that SAA performs arbitrarily poorly in pricing under Wasserstein heterogeneity. We now show that SAA falters similarly for the Wasserstein ski-rental problem. Furthermore, we design and analyze a δ -SAA policy which inflates the action selected by SAA (i.e., $\delta > 0$) and achieves the best rate possible as a function of the heterogeneity radius. Formally, we show the following.

Theorem 4. *For the data-driven ski-rental problem in a heterogenous environment under the Wasserstein distance, we have for $\epsilon > 0$ that,*

$$\mathfrak{R}_{W,\infty}^{\pi^{\text{SAA}}}(\epsilon) = \Theta(1) \quad \text{and} \quad \inf_{\pi \in \Pi} \mathfrak{R}_{W,\infty}^{\pi}(\epsilon) = \Theta(\sqrt{\epsilon}).$$

Theorem 4 formalizes the failure of SAA for ski-rental under Wasserstein distance. A notable fact about the ski-rental problem under Wasserstein heterogeneity is that both the regret of SAA and of the optimal data-driven decision do not scale with the size of the support. For pricing, SAA scaled linearly with M and the optimal policy scales in \sqrt{M} . The fact that $\mathfrak{R}_{W,\infty}^{\pi^{\text{SAA}}}(\epsilon) = O(1)$ in ski-rental requires a separate non-trivial proof in Theorem 4.³

Kolmogorov ski-rental. Under the Kolmogorov distance, we had seen that SAA was near-optimal for both the newsvendor and pricing problems. We establish next that the performance of SAA for the ski-rental problem under Kolmogorov distance is, surprisingly, highly suboptimal.

Theorem 5. *For the data-driven ski-rental problem in a heterogenous environment under the Kolmogorov distance, we have for $\epsilon > 0$ that,*

$$\mathfrak{R}_{K,\infty}^{\pi^{\text{SAA}}}(\epsilon) = \Theta(M \cdot \epsilon) \quad \text{and} \quad \inf_{\pi \in \Pi} \mathfrak{R}_{K,\infty}^{\pi}(\epsilon) = \tilde{\Theta}(\epsilon),$$

where $\tilde{\Theta}$ provides rates order up to logarithmic factors.

Theorem 5 shows that the asymptotic worst-case regret of SAA scales linearly with the radius of heterogeneity ϵ and with the size of the support M . We also characterize the rate of the optimal asymptotic worst-case regret as a function of ϵ . The upper bound on the asymptotic regret is obtained in [19] through a variant of the SAA policy, which caps the maximum number of days to rent. We complement this result by providing a matching lower bound. Note that, in contrast to the pricing and newsvendor problems, the scaling of SAA is not optimal in M . As the scale of the support grows, the asymptotic worst-case regret of SAA is considerably worse than the optimal achievable rate.

6 Conclusion

All in all, the present results offer a systematic way of understanding and quantifying the implications of operating in heterogeneous environments. Our framework enables us to develop a common language to compare achievable performance across central classes of problems and to unveil novel insights about the performance of a central policy, Sample Average Approximation, when slightly deviating from the widely studied i.i.d. regime. In settings where SAA fails, we also design robustification procedures achieving rate-optimal asymptotic guarantees. A key takeaway of this analysis across a broad class of problems and for different heterogeneity structures is that it is necessary to understand the structure of the problem we are facing but also the nature of the heterogeneity when designing data-driven policies that are robust to these environments.

To derive our performance guarantees, we established a crucial connection between data-driven decision making in heterogeneous environments and distributionally robust optimization. While this work leverages this connection essentially to derive bounds on achievable performances, we believe that this connection may be of interest to develop new policies from first principles.

This work also opens many additional questions. First, our analysis characterizes the performance in the asymptotic regime where the sample size grows large, and could be complemented by quantifying the performance of policies with finite samples in heterogeneous environments. Furthermore, our results for Lipschitz problems pin down a key driver of performance, whereas isolating the properties of the policies and the framework elements that drive different levels of achievable performance beyond the Lipschitz case remains open. Another key question would be to understand whether there exists a “best of both worlds” policy that does not use the knowledge of the type of heterogeneity and performs well across heterogeneity types (as opposed to the robustified policy designed for pricing and ski-rental). Finally, we believe that additional exciting and practical complement to this work include incorporating contexts and deriving statistical tests to characterize the type of heterogeneity along with its radius and provide an empirical validation of the procedures developed here.

³The fact that $\mathfrak{R}_{W,\infty}^{\pi^{\text{SAA}}}(\epsilon) \leq M$ in pricing was vacuous.

References

- [1] Allouah, A., Bahamou, A. and Besbes, O. [2022], ‘Pricing with samples’, *Operations Research* **70**(2), 1088–1104.
- [2] Babaioff, M., Gonczarowski, Y. A., Mansour, Y. and Moran, S. [2018], Are two (samples) really better than one?, in ‘Proceedings of the 2018 ACM Conference on Economics and Computation’, pp. 175–175.
- [3] Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F. and Vaughan, J. W. [2010], ‘A theory of learning from different domains’, *Machine learning* **79**(1), 151–175.
- [4] Bertsimas, D., Gupta, V. and Kallus, N. [2018], ‘Robust sample average approximation’, *Mathematical Programming* **171**(1), 217–282.
- [5] Besbes, O., Gur, Y. and Zeevi, A. [2014], ‘Stochastic multi-armed-bandit problem with non-stationary rewards’, *Advances in neural information processing systems* **27**.
- [6] Besbes, O. and Mouchtaki, O. [2021], ‘How big should your data really be? data-driven newsvendor: learning one sample at a time’, *working paper, Columbia University*.
- [7] Bilodeau, B., Negrea, J. and Roy, D. M. [2020], ‘Relaxing the iid assumption: Adaptively minimax optimal regret via root-entropic regularization’, *arXiv preprint arXiv:2007.06552*.
- [8] Blanc, G., Lange, J., Malik, A. and Tan, L.-Y. [2021], ‘On the power of adaptivity in statistical adversaries’, *arXiv preprint arXiv:2111.10352*.
- [9] Blanchard, M. [2022], ‘Universal online learning: An optimistically universal learning rule’, *arXiv preprint arXiv:2201.05947*.
- [10] Block, A., Dagan, Y., Golowich, N. and Rakhlin, A. [2022], ‘Smoothed online learning is as easy as statistical learning’, *arXiv preprint arXiv:2202.04690*.
- [11] Borodin, A. and El-Yaniv, R. [2005], *Online computation and competitive analysis*, cambridge university press.
- [12] Brustle, J., Cai, Y. and Daskalakis, C. [2020], Multi-item mechanisms without item-independence: Learnability via robustness, in ‘Proceedings of the 21st ACM Conference on Economics and Computation’, pp. 715–761.
- [13] Cai, Y. and Daskalakis, C. [2017], Learning multi-item auctions with (or without) samples, in ‘2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)’, IEEE, pp. 516–527.
- [14] Chen, X. and Wang, Y. [2020], ‘Dynamic pricing with demand learning in the presence of outlier customers’, <https://ssrn.com/abstract=3650656>.
- [15] Cheung, W. C. and Simchi-Levi, D. [2019], ‘Sampling-based approximation schemes for capacitated stochastic inventory control models’, *Mathematics of Operations Research* **44**(2), 668–692.
- [16] Cheung, W. C., Simchi-Levi, D. and Zhu, R. [2022], ‘Hedging the drift: Learning to optimize under nonstationarity’, *Management Science* **68**(3), 1696–1713.
- [17] Daskalakis, C. and Zampetakis, M. [2020], More revenue from two samples via factor revealing sdps, in ‘Proceedings of the 21st ACM Conference on Economics and Computation’, pp. 257–272.
- [18] Diakonikolas, I., Kamath, G., Kane, D., Li, J., Moitra, A. and Stewart, A. [2019], ‘Robust estimators in high-dimensions without the computational intractability’, *SIAM Journal on Computing* **48**(2), 742–864.
- [19] Diakonikolas, I., Kontonis, V., Tzamos, C., Vakilian, A. and Zarifis, N. [2021], Learning online algorithms with distributional advice, in ‘International Conference on Machine Learning’, PMLR, pp. 2687–2696.

- [20] Duchi, J. C. and Namkoong, H. [2021], ‘Learning models with uniform performance via distributionally robust optimization’, *The Annals of Statistics* **49**(3), 1378 – 1406.
- [21] Fu, H., Immorlica, N., Lucier, B. and Strack, P. [2015], Randomization beats second price as a prior-independent auction, in ‘Proceedings of the Sixteenth ACM Conference on Economics and Computation’, pp. 323–323.
- [22] Gallego, G. and Moon, I. [1993], ‘The distribution free newsboy problem: review and extensions’, *Journal of the Operational Research Society* **44**(8), 825–834.
- [23] Guo, W., Jordan, M. and Zampetakis, E. [2021], ‘Robust learning of optimal auctions’, *Advances in Neural Information Processing Systems* **34**.
- [24] Haghtalab, N., Han, Y., Shetty, A. and Yang, K. [2022], ‘Oracle-efficient online learning for beyond worst-case adversaries’, *arXiv preprint arXiv:2202.08549* .
- [25] Haghtalab, N., Roughgarden, T. and Shetty, A. [2022], Smoothed analysis with adaptive adversaries, in ‘2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)’, IEEE, pp. 942–953.
- [26] Hampel, F. R. [1971], ‘A general qualitative definition of robustness’, *The annals of mathematical statistics* **42**(6), 1887–1896.
- [27] Hanneke, S. [2021], ‘Learning whenever learning is possible: Universal learning under general stochastic processes.’, *J. Mach. Learn. Res.* **22**, 130–1.
- [28] Haussler, D. [1992], ‘Decision theoretic generalizations of the pac model for neural net and other learning applications’, *Information and computation* **100**(1), 78–150.
- [29] Huang, Z., Mansour, Y. and Roughgarden, T. [2018], ‘Making the most of your samples’, *SIAM Journal on Computing* **47**(3), 651–674.
- [30] Huber, P. J. [1992], Robust estimation of a location parameter, in ‘Breakthroughs in statistics’, Springer, pp. 492–518.
- [31] Kantorovich, L. V. and Akilov, G. P. [2016], *Functional analysis*, Elsevier.
- [32] Kantorovich, L. V. and Rubinshtein, S. [1958], ‘On a space of totally additive functions’, *Vestnik of the St. Petersburg University: Mathematics* **13**(7), 52–59.
- [33] Karnin, Z. S. and Anava, O. [2016], ‘Multi-armed bandits: Competing with optimal sequences’, *Advances in Neural Information Processing Systems* **29**.
- [34] Kearns, M. J., Schapire, R. E. and Sellie, L. M. [1994], ‘Toward efficient agnostic learning’, *Machine Learning* **17**(2), 115–141.
- [35] Kearns, M. and Li, M. [1993], ‘Learning in the presence of malicious errors’, *SIAM Journal on Computing* **22**(4), 807–837.
- [36] Kim, S., Pasupathy, R. and Henderson, S. G. [2015], ‘A guide to sample average approximation’, *Handbook of simulation optimization* pp. 207–243.
- [37] Kleywegt, A. J., Shapiro, A. and Homem-de Mello, T. [2002], ‘The sample average approximation method for stochastic discrete optimization’, *SIAM Journal on Optimization* **12**(2), 479–502.
- [38] Klivans, A. R., Long, P. M. and Servedio, R. A. [2009], ‘Learning halfspaces with malicious noise.’, *Journal of Machine Learning Research* **10**(12).
- [39] Lam, H. [2021], ‘On the impossibility of statistically improving empirical optimization: A second-order stochastic dominance perspective’, *arXiv preprint arXiv:2105.13419* .
- [40] Levi, R., Perakis, G. and Uichanco, J. [2015], ‘The data-driven newsvendor problem: new bounds and insights’, *Operations Research* **63**(6), 1294–1306.
- [41] Luo, H., Wei, C.-Y., Agarwal, A. and Langford, J. [2018], Efficient contextual bandits in non-stationary worlds, in ‘Conference On Learning Theory’, PMLR, pp. 1739–1776.

- [42] Lykouris, T., Mirrokni, V. and Paes Leme, R. [2018], Stochastic bandits robust to adversarial corruptions, in ‘Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing’, pp. 114–122.
- [43] Lykouris, T., Simchowitz, M., Slivkins, A. and Sun, W. [2021], Corruption-robust exploration in episodic reinforcement learning, in ‘Conference on Learning Theory’, PMLR, pp. 3242–3245.
- [44] Mansour, Y., Mohri, M. and Rostamizadeh, A. [2009], ‘Domain adaptation: Learning bounds and algorithms’, *arXiv preprint arXiv:0902.3430* .
- [45] Mohajerin Esfahani, P. and Kuhn, D. [2018], ‘Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations’, *Mathematical Programming* **171**(1), 115–166.
- [46] Mohri, M. and Muñoz Medina, A. [2012], New analysis and algorithm for learning with drifting distributions, in ‘International Conference on Algorithmic Learning Theory’, Springer, pp. 124–138.
- [47] Morse, P. M. and Kimball, G. E. [1946], Methods of operations research, Technical report, Center for Naval Analyses, Alexandria VA, Operations Evaluation Group.
- [48] Myerson, R. B. [1981], ‘Optimal auction design’, *Mathematics of operations research* **6**(1), 58–73.
- [49] Perakis, G. and Roels, G. [2008], ‘Regret in the newsvendor model with partial information’, *Operations Research* **56**(1), 188–203.
- [50] Purohit, M., Svitkina, Z. and Kumar, R. [2018], ‘Improving online algorithms via ml predictions’, *Advances in Neural Information Processing Systems* **31**.
- [51] Rakhlin, A., Sridharan, K. and Tewari, A. [2011], ‘Online learning: Stochastic and constrained adversaries’, *arXiv preprint arXiv:1104.5070* .
- [52] Riley, J. and Zeckhauser, R. [1983], ‘Optimal selling strategies: When to haggle, when to hold firm’, *The Quarterly Journal of Economics* **98**(2), 267–289.
- [53] Roughgarden, T. [2021], *Beyond the worst-case analysis of algorithms*, Cambridge University Press.
- [54] Scarf, H. [1958], ‘A min-max solution of an inventory problem’, *Studies in the mathematical theory of inventory and production* .
- [55] Servedio, R. A. [2003], ‘Smooth boosting and learning with malicious noise’, *The Journal of Machine Learning Research* **4**, 633–648.
- [56] Shorack, G. R. [1979], ‘The weighted empirical process of row independent random variables with arbitrary distribution functions’, *Statistica Neerlandica* **33**(4), 169–189.
- [57] Williams, D. [1991], *Probability with martingales*, Cambridge university press.