
Distributionally Robust Optimization with Markovian Data

Mengmeng Li¹ Tobias Sutter¹ Daniel Kuhn¹

Abstract

We study a stochastic program where the probability distribution of the uncertain problem parameters is unknown and only indirectly observed via finitely many correlated samples generated by an unknown Markov chain with d states. We propose a data-driven distributionally robust optimization model to estimate the problem's objective function and optimal solution. By leveraging results from large deviations theory, we derive statistical guarantees on the quality of these estimators. The underlying worst-case expectation problem is nonconvex and involves $\mathcal{O}(d^2)$ decision variables. Thus, it cannot be solved efficiently for large d . By exploiting the structure of this problem, we devise a customized Frank-Wolfe algorithm with convex direction-finding subproblems of size $\mathcal{O}(d)$. We prove that this algorithm finds a stationary point efficiently under mild conditions. The efficiency of the method is predicated on a dimensionality reduction enabled by a dual reformulation. Numerical experiments indicate that our approach has better computational and statistical properties than the state-of-the-art methods.

1. Introduction

Decision problems under uncertainty are ubiquitous in machine learning, engineering and economics. Traditionally, such problems are modeled as stochastic programs that seek a decision $x \in X \subseteq \mathbb{R}^n$ minimizing an expected loss $\mathbb{E}_{\mathbb{P}}[L(x, \xi)]$, where the expectation is taken with respect to the distribution \mathbb{P} of the random problem parameter ξ . However, more often than not, \mathbb{P} is unknown to the decision maker and can only be observed indirectly via a finite number of training samples ξ_1, \dots, ξ_T . The classical sample average approximation (SAA) replaces the unknown probability distribution \mathbb{P} with the empirical distribution cor-

responding to the training samples and solves the resulting empirical risk minimization problem (Shapiro et al., 2014). Fuelled by modern applications in machine learning, however, there has recently been a surge of alternative methods for data-driven optimization complementing SAA. Ideally, any meaningful approach to data-driven optimization should display the following desirable properties.

(i) **Avoiding the optimizer's curse.** It is well understood that decisions achieving a low empirical cost on the training dataset may perform poorly on a test dataset generated by the same distribution. In decision analysis, this phenomenon is called the *optimizer's curse*, and it is closely related to overfitting in statistics (Smith & Winkler, 2006). Practically useful schemes should mitigate this detrimental effect.

(ii) **Statistical guarantees.** Most statistical guarantees for existing approaches to data-driven optimization critically rely on the assumption that the training samples are independent and identically distributed (i.i.d.). This assumption is often not justifiable or even wrong in practice. Practically useful methods should offer guarantees that remain valid when the training samples display serial dependencies.

(iii) **Computational tractability.** For a data-driven scheme to be practically useful it is indispensable that the underlying optimization problems can be solved efficiently.

While the SAA method is computationally tractable, it is susceptible to the optimizer's curse if training data is scarce; see, e.g., (Van Parys et al., 2021). *Distributionally robust optimization* (DRO) is an alternative approach that mitigates overfitting effects (Delage & Ye, 2010; Goh & Sim, 2010; Wiesemann et al., 2014). DRO seeks worst-case optimal decisions that minimize the expected loss under the most adverse distribution from within a given ambiguity set, that is, a distribution family characterized by certain known properties of the unknown data-generating distribution. DRO has been studied since Scarf's seminal treatise on the ambiguity-averse newsvendor problem (Scarf, 1958), but it has gained thrust only with the advent of modern robust optimization techniques (Bertsimas & Sim, 2004; Ben-Tal et al., 2009). In many cases of practical interest, DRO problems can be reformulated exactly as finite convex programs that are solvable in polynomial time. Indeed, such reformulations are available for many different ambiguity sets defined through generalized moment constraints (Delage & Ye, 2010; Goh & Sim, 2010; Wiesemann et al., 2014; Bertsimas et al., 2018),

¹Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne. Correspondence to: Mengmeng Li <mengmeng.li@epfl.ch>.

ϕ -divergences (Ben-Tal et al., 2013; Namkoong & Duchi, 2016), Wasserstein distances (Mohajerin Esfahani & Kuhn, 2018; Kuhn et al., 2019), or maximum mean discrepancy distances (Staib & Jegelka, 2019; Kirschner et al., 2020).

With the notable exceptions of (Dou & Anitescu, 2019; Derman & Mannor, 2020; Sutter et al., 2020; Duchi et al., 2021), we are not aware of any data-driven DRO models for non-i.i.d. data. In this paper we apply the general framework by Sutter et al. (2020) to data-driven DRO models with Markovian training samples and propose an efficient algorithm for their solution. Our DRO scheme is perhaps most similar to the one studied by Duchi et al. (2021), which can also handle Markovian data. However, this scheme differs from ours in two fundamental ways. First, while Duchi et al. (2021) work with ϕ -divergence ambiguity sets, we use ambiguity sets inspired by a statistical optimality principle recently established by Van Parys et al. (2021) and Sutter et al. (2020) using ideas from large deviations theory. Second, the statistical guarantees by Duchi et al. (2021) depend on unknown constants, whereas our confidence bounds are explicit and easy to evaluate. Dou & Anitescu (2019) assume that the training samples are generated by a first-order autoregressive process, which is neither a generalization nor a special case of a finite-state Markov chain. Indeed, while Dou & Anitescu (2019) can handle continuous state spaces, our model does not assume a linear dependence on the previous state. Derman & Mannor (2020) investigate a finite-state Markov decision process (MDP) with an unknown transition kernel, and they develop a DRO approach using a Wasserstein ambiguity set for estimating its value function. While MDPs are more general than the static optimization problems considered here, the out-of sample guarantees in (Derman & Mannor, 2020, Theorem 4.1) rely on the availability of several i.i.d. sample trajectories. In contrast, our statistical guarantees require only one single trajectory of (correlated) training samples. We also remark that MDP models can be addressed with online mirror descent methods (Jin & Sidford, 2020). Unfortunately, it is doubtful whether such methods could be applied to our DRO problems because the statistically optimal ambiguity sets considered in this paper are nonconvex.

In summary, while many existing data-driven DRO models are tractable and can mitigate the optimizer’s curse, there are hardly any statistical performance guarantees that apply when the training samples fail to be i.i.d. and when there is only one trajectory of correlated training samples. This paper addresses this gap. Specifically, we study data-driven decision problems where the training samples are generated by a time-homogeneous, ergodic Markov chain with d states. Sutter et al. (2020) show that statistically optimal data-driven decisions for such problems are obtained by solving a DRO model with a conditional relative entropy ambiguity set. The underlying worst-case expectation prob-

lem is nonconvex and involves d^2 decision variables. To our best knowledge, as of now there exist no efficient algorithms for solving this hard optimization problem. We highlight the following main contributions of this paper.

- We apply the general framework by Sutter et al. (2020), which uses ideas from large deviations theory to construct statistically optimal data-driven DRO models, to decision problems where the training data is generated by a time-homogeneous, ergodic finite-state Markov chain. We prove that the resulting DRO models are asymptotically consistent.
- We develop a customized Frank-Wolfe algorithm for solving the underlying worst-case expectation problems in an efficient manner. This is achieved via the following steps.
 - (i) We first reparametrize the problems to move the nonconvexities from the feasible set to the objective function.
 - (ii) We then develop a Frank-Wolfe algorithm for solving the reparametrized nonconvex problem of size $\mathcal{O}(d^2)$.
 - (iii) Using a duality argument, we show that the direction-finding subproblems with a linearized objective function are equivalent to convex programs of size $\mathcal{O}(d)$ with a rectangular feasible set that can be solved highly efficiently.
- We prove that the proposed Frank-Wolfe algorithm converges to a stationary point of the nonconvex worst-case expectation problem at a rate $\mathcal{O}(1/\sqrt{M})$, where M denotes the number of iterations. Each iteration involves the solution of a convex d -dimensional minimization problem with a smooth objective function and a rectangular feasible set. The solution of this minimization problem could be further accelerated by decomposing it into d one-dimensional optimization problems that can be processed in parallel.
- We propose a Neyman-Pearson-type hypothesis test for determining the Markov chain that generated a given sequence of training samples, and we construct examples of Markov chains that are arbitrarily difficult to distinguish on the basis of a finite training dataset alone.
- In the context of a revenue maximization problem where customers display a Markovian brand switching behavior, we show that the proposed DRO models for Markovian data outperform classical DRO models, which are either designed for i.i.d. data or based on different ambiguity sets.

Notation. The inner product of two vectors $a, b \in \mathbb{R}^m$ is denoted by $\langle a, b \rangle = a^\top b$, and the n -dimensional probability simplex is defined as $\Delta_n = \{x \in \mathbb{R}_+^n : \sum_{i=1}^n x_i = 1\}$. The relative entropy between two probability vectors $p, q \in \Delta_n$ is defined as $D(p||q) = \sum_{i=1}^n p_i \log(p_i/q_i)$, where we use the conventions $0 \log(0/q) = 0$ for $q \geq 0$ and $p \log(p/0) = \infty$ for $p > 0$. The closure and the interior of a subset \mathcal{D} of a topological space are denoted by $\text{cl } \mathcal{D}$ and $\text{int } \mathcal{D}$, respectively. For any $n \in \mathbb{N}$ we set $[n] = \{1, \dots, n\}$. We use A_i and $A_{\cdot i}$ to denote the i -th row and column of a matrix $A \in \mathbb{R}^{m \times n}$, respectively. All proofs are relegated to the Appendix.

2. Data-driven DRO with Markovian Data

In this paper we study the single-stage stochastic program

$$\min_{x \in X} \mathbb{E}_{\mathbb{P}}[L(x, \xi)], \quad (1)$$

where $X \subseteq \mathbb{R}^n$ is compact, ξ denotes a random variable valued in $\Xi = \{1, 2, \dots, d\}$, and $L : X \times \Xi \rightarrow \mathbb{R}$ represents a known loss function that is continuous in x . We assume that all random objects (e.g., ξ) are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ whose probability measure \mathbb{P} is unknown but belongs to a known ambiguity set \mathcal{A} to be described below. We further assume that \mathbb{P} can only be observed indirectly through samples $\{\xi_t\}_{t \in \mathbb{N}}$ from a time-homogeneous ergodic Markov chain with state space Ξ and deterministic initial state $\xi_0 \in \Xi$. We finally assume that the distribution of ξ coincides with the unique stationary distribution of the Markov chain. In the following we let $\theta^* \in \mathbb{R}_{++}^{d \times d}$ be the unknown stationary probability mass function of the doublet $(\xi_t, \xi_{t+1}) \in \Xi^2$, that is, we set

$$\theta_{ij}^* = \lim_{t \rightarrow \infty} \mathbb{P}(\xi_t = i, \xi_{t+1} = j) \quad \forall i, j \in \Xi. \quad (2)$$

By construction, the row sums of θ^* must coincide with the respective column sums, and thus θ^* must belong to

$$\Theta = \left\{ \theta \in \mathbb{R}_{++}^{d \times d} : \sum_{i,j=1}^d \theta_{ij} = 1, \sum_{j=1}^d \theta_{ij} = \sum_{j=1}^d \theta_{ji} \quad \forall i \in \Xi \right\},$$

that is, the set of all strictly positive doublet probability mass functions with balanced marginals. Below we call elements of Θ simply *models*. Note that any model $\theta \in \Theta$ encodes an ergodic time-homogeneous Markov chain with a unique row vector $\pi_\theta \in \mathbb{R}_{++}^{1 \times d}$ of stationary probabilities and a unique transition probability matrix $P_\theta \in \mathbb{R}_{++}^{d \times d}$ defined through $(\pi_\theta)_i = \sum_{j \in \Xi} \theta_{ij}$ and $(P_\theta)_{ij} = \theta_{ij} / (\pi_\theta)_i$, respectively. Hence, the elements of π_θ sum to 1, P_θ represents a row-stochastic matrix, and $\pi_\theta = \pi_\theta P_\theta$. Moreover, any $\theta \in \Theta$ induces a probability measure \mathbb{P}_θ on (Ω, \mathcal{F}) with

$$\mathbb{P}_\theta(\xi_t = i_t \quad \forall t = 1, \dots, T) = (P_\theta)_{\xi_0 i_1} \prod_{t=1}^{T-1} (P_\theta)_{i_t i_{t+1}}$$

for all $(i_1, \dots, i_T) \in \Xi^T$ and $T \in \mathbb{N}$. We are now ready to define the ambiguity set as $\mathcal{A} = \{\mathbb{P}_\theta : \theta \in \Theta\}$. Below we use \mathbb{E}_θ to denote the expectation operator with respect to \mathbb{P}_θ . By construction, there exists a model $\theta^* \in \Theta$ corresponding to the unknown true probability measure \mathbb{P} such that $\mathbb{P}_{\theta^*} = \mathbb{P}$. Estimating \mathbb{P} is thus equivalent to estimating the unknown true parameter θ^* . Given a finite training dataset $\{\xi_t\}_{t=1}^T$, a natural estimator for θ^* is the empirical doublet distribution $\hat{\theta}_T \in \mathbb{R}_{++}^{d \times d}$ defined through

$$(\hat{\theta}_T)_{ij} = \frac{1}{T} \sum_{t=1}^T \mathbf{1}_{(\xi_{t-1}, \xi_t) = (i, j)} \quad (3)$$

for all $i, j \in \Xi$. The ergodic theorem ensures that $\hat{\theta}_T$ converges \mathbb{P}_{θ^*} -almost surely to θ^* for any model $\theta \in \Theta$ (Ross, 2010, Theorem 4.1). In the remainder we define $\Theta' = \Delta_{d \times d}$ as the state space of $\hat{\theta}_T$, which is strictly larger than Θ . For ease of notation, we also define the *model-based predictor* $c(x, \theta) = \mathbb{E}_\theta[L(x, \xi)] = \sum_{i,j=1}^d L(x, i) \theta_{ij}$ as the expected loss of x under model θ . Note that $c(x, \theta)$ is jointly continuous in x and θ due to the continuity of $L(x, \xi)$ in x .

As we will show below, it is natural to measure the discrepancy between a model $\theta \in \Theta$ and an estimator realization $\theta' \in \Theta'$ by the conditional relative entropy.

Definition 1 (Conditional relative entropy). The conditional relative entropy of $\theta' \in \Theta'$ with respect to $\theta \in \Theta$ is

$$\begin{aligned} D_c(\theta' \parallel \theta) &= \sum_{i=1}^d (\pi_{\theta'})_i D((P_{\theta'})_i \parallel (P_\theta)_i) \\ &= \sum_{i,j=1}^d \theta'_{ij} \left(\log \left(\frac{\theta'_{ij}}{\sum_{k=1}^d \theta'_{ik}} \right) - \log \left(\frac{\theta_{ij}}{\sum_{k=1}^d \theta_{ik}} \right) \right), \end{aligned}$$

where the invariant distributions $\pi_\theta, \pi_{\theta'} \in \mathbb{R}_{++}^{1 \times d}$ and the transition probability matrices $P_\theta, P_{\theta'} \in \mathbb{R}_{++}^{d \times d}$ corresponding to θ and θ' , respectively, are defined in the usual way.

Note that if the training samples $\{\xi_t\}_{t \in \mathbb{N}}$ are i.i.d., then the conditional relative entropy of θ' with respect to θ collapses to the relative entropy of $\pi_{\theta'}$ with respect to π_θ . Using the conditional relative entropy, we can now define the *distributionally robust predictor* $\hat{c}_r : X \times \Theta' \rightarrow \mathbb{R}$ through

$$\hat{c}_r(x, \theta') = \max_{\theta \in \text{cl } \Theta} \{c(x, \theta) : D_c(\theta' \parallel \theta) \leq r\} \quad (4a)$$

if (4a) is feasible and $\hat{c}_r(x, \theta') = \max_{\theta \in \text{cl } \Theta} c(x, \theta)$ otherwise. Note that $\hat{c}_r(x, \theta')$ represents the worst-case expected cost of x with respect to all probability measures in the ambiguity set $\{\mathbb{P}_\theta : D_c(\theta' \parallel \theta) \leq r\}$, which can be viewed as an image (in \mathcal{A}) of a conditional relative entropy ball (in Θ) of radius $r \geq 0$ around θ' . We know from (Sutter et al., 2020, Proposition 5.1) that the sublevel sets of $D_c(\theta' \parallel \theta)$ are compact for all fixed θ' , and thus the maximum in (4a) is attained whenever the problem is feasible. The distributionally robust predictor (4a) also induces a *distributionally robust prescriptor* $\hat{x}_r : \Theta' \rightarrow X$ defined through

$$\hat{x}_r(\theta') \in \arg \min_{x \in X} \hat{c}_r(x, \theta') \quad (4b)$$

Embracing the distributionally robust approach outlined above, if we only have access to T training samples, we will implement the data-driven decision $\hat{x}_r(\hat{\theta}_T)$ and predict the corresponding expected cost as $\hat{c}_r(\hat{x}_r(\hat{\theta}_T), \hat{\theta}_T)$, which we will refer to, by slight abuse of terminology, as the *in-sample risk*. To assess the quality of the data-driven decision $\hat{x}_r(\hat{\theta}_T)$

we use two performance measures that stand in direct competition with each other: the *out-of-sample risk* $c(\hat{x}_r(\hat{\theta}_T), \theta)$ and the *out-of-sample disappointment*

$$\mathbb{P}_\theta \left(c(\hat{x}_r(\hat{\theta}_T), \theta) > \hat{c}_r(\hat{x}_r(\hat{\theta}_T), \hat{\theta}_T) \right)$$

with respect to any fixed model θ . Intuitively, the out-of-sample risk represents the true expected loss of $\hat{x}_r(\hat{\theta}_T)$, and the out-of-sample disappointment quantifies the probability that the in-sample risk (the predicted loss) strictly underestimates the out-of-sample risk (the actual loss) under \mathbb{P}_θ . Following Sutter et al. (2020), we will consider a data-driven decision as desirable if it has a *low* out-of-sample risk and a *low* out-of-sample disappointment under the unknown true model θ^* . This is reasonable because the objective of the original problem (2) is to *minimize* expected loss. Optimistically underestimating the loss could incentivize naïve decisions that overfit to the training data and therefore lead to disappointment in out-of-sample tests. Thus, it makes sense to keep the out-of-sample disappointment small.

3. Statistical Guarantees

The distributionally robust optimization model (4) with a conditional relative entropy ambiguity set has not received much attention in the extant literature and may thus seem exotic at first sight. As we will show below, however, this model offers—in a precise sense—*optimal* statistical guarantees if the training samples are generated by an (unknown) time-homogeneous ergodic Markov chain. The reason for this is that the empirical doublet distribution (3) is a sufficient statistic for θ and satisfies a large deviation principle with rate function $D_c(\theta' \parallel \theta)$ under \mathbb{P}_θ for any $\theta \in \Theta$.

Before formalizing these statistical guarantees, we first need to recall some basic concepts from large deviation theory. We refer to the excellent textbooks by Dembo & Zeitouni (2010) and den Hollander (2008) for more details.

Definition 2 (Rate function (Dembo & Zeitouni, 2010, Section 2.1)). A function $I : \Theta' \times \Theta \rightarrow [0, \infty]$ is called a rate function if $I(\theta', \theta)$ is lower semi-continuous in θ' .

Definition 3 (Large deviation principle). An estimator $\hat{\theta}_T \in \Theta'$ satisfies a large deviation principle (LDP) with rate function I if for all $\theta \in \Theta$ and Borel sets $\mathcal{D} \subseteq \Theta'$ we have

$$- \inf_{\theta' \in \text{int } \mathcal{D}} I(\theta', \theta) \leq \liminf_{T \rightarrow \infty} \frac{1}{T} \log \mathbb{P}_\theta \left(\hat{\theta}_T \in \mathcal{D} \right) \quad (5a)$$

$$\leq \limsup_{T \rightarrow \infty} \frac{1}{T} \log \mathbb{P}_\theta \left(\hat{\theta}_T \in \mathcal{D} \right) \quad (5b)$$

$$\leq - \inf_{\theta' \in \text{cl } \mathcal{D}} I(\theta', \theta). \quad (5c)$$

Conveniently, the empirical doublet distribution (3) obeys an LDP with the conditional relative entropy as rate function.

Lemma 1 (LDP for Markov chains (Dembo & Zeitouni, 2010, Theorem 3.1.13)). *The empirical doublet distribution (3) satisfies an LDP with rate function $D_c(\theta' \parallel \theta)$.*

Lemma 1 is the key ingredient to establish the following out-of-sample guarantees for the distributionally robust predictor and the corresponding prescriptor.

Theorem 2 (Out-of-sample guarantees (Sutter et al., 2020, Theorems 3.1 & 3.2)). *The distributionally robust predictor and prescriptor defined in (4) offer the following guarantees.*

(i) *For all $\theta \in \Theta$, $x \in X$, we have*

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \log \mathbb{P}_\theta \left(c(x, \theta) > \hat{c}_r(x, \hat{\theta}_T) \right) \leq -r.$$

(ii) *For all $\theta \in \Theta$, we have*

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \log \mathbb{P}_\theta \left(c(\hat{x}_r(\hat{\theta}_T), \theta) > \hat{c}_r(\hat{x}_r(\hat{\theta}_T), \hat{\theta}_T) \right) \leq -r.$$

Theorem 2 asserts that the out-of-sample disappointment of the distributionally robust predictor and the corresponding prescriptor decay at least as fast as $e^{-rT+o(T)}$ asymptotically as the sample size T grows, where the decay rate r coincides with the radius of the conditional relative entropy ball in (4). Note also that $\hat{x}_r(\hat{\theta}_T)$ can be viewed as an instance of a data-driven prescriptor, that is, any Borel measurable function that maps the training samples $\{\xi_t\}_{t=1}^T$ to a decision in X . One can then prove that $\hat{x}_r(\hat{\theta}_T)$ offers the best (least possible) out-of-sample risk among the vast class of data-driven prescriptors whose out-of-sample disappointment decays at a prescribed rate of at least r . Maybe surprisingly, this optimality property holds uniformly across all models $\theta \in \Theta$ and thus in particular for the unknown true model θ^* ; see (Sutter et al., 2020, Theorem 3.2). It is primarily this Pareto dominance property that makes the distributionally robust optimization model (4) interesting.

Next, we prove that if the radius r of the conditional relative entropy ball tends to 0 at a sufficiently slow speed as T increases, then the distributionally robust predictor converges \mathbb{P}_θ -almost surely to the true expected loss of the decision x .

Theorem 3 (Strong asymptotic consistency). *If $r_T \geq \frac{d}{T}$ for all $T \in \mathbb{N}$ and $\lim_{T \rightarrow \infty} r_T = 0$, then*

$$\lim_{T \rightarrow \infty} \hat{c}_{r_T}(x, \hat{\theta}_T) = c(x, \theta) \quad \mathbb{P}_\theta\text{-a.s.} \quad \forall x \in X, \theta \in \Theta.$$

4. Numerical Solution

Even though the distributionally robust optimization model (4) offers powerful statistical guarantees, it is practically useless unless the underlying optimization problems can be solved efficiently. In this section we develop a numerical procedure to compute the predictor (4a) for a fixed x and θ' . This is challenging for the following reasons. First, the conditional relative entropy $D_c(\theta' \parallel \theta)$ is nonconvex in θ

(see Remark 13 in Appendix A), and thus problem (4a) has a nonconvex feasible set. In addition, problem (4a) involves d^2 decision variables (the entries of the matrix θ) and thus its dimension is high already for moderate values of d .

In the following we reformulate (4a) as an equivalent optimization problem with a convex feasible set and a nonconvex objective function, which is amenable to efficient numerical solution via a customized Frank-Wolfe algorithm (Frank & Wolfe, 1956; Jaggi, 2013). To this end, recall from Definition 1 that the conditional relative entropy $D_c(\theta' \parallel \theta)$ can be expressed as a weighted sum of relative entropies between corresponding row vectors of the transition probability matrices $P_{\theta'}$ and P_θ . By construction, P_θ is a row-stochastic matrix, and the stationary distribution π_θ corresponding to P_θ is a normalized non-negative row vector with $\pi_\theta P_\theta = \pi_\theta$. Routine calculations show that

$$(\pi_\theta)^\top = (A_d(P_\theta))^{-1} [0 \ 0 \ \dots \ 1]^\top, \quad (6)$$

where the matrix $A_d(P_\theta) \in \mathbb{R}^{d \times d}$ is defined through

$$(A_d(P_\theta))_{ij} = \begin{cases} (P_\theta)_{ii} - 1 & \text{for } i = j, i < d \\ (P_\theta)_{ji} & \text{for } i \neq j, i < d \\ 1 & \text{for } i = d. \end{cases}$$

From now on we denote by $\mathcal{P} \subseteq \mathbb{R}^{d \times d}$ the set of all transition probability matrices of ergodic Markov chains on Ξ with strictly positive entries. Applying the variable substitution $P_\theta \leftarrow \theta$ and using (6), we can now reformulate the optimization problem (4a) in terms of P_θ .

Lemma 4 (Reformulation). *If $\theta' > 0$, then the worst-case expectation problem (4a) is equivalent to*

$$\hat{c}_r(x, \theta') = \max_{P \in \mathcal{D}_r(\theta')} \Psi(x, P), \quad (7)$$

where $\Psi(x, P) = \sum_{i=1}^d L(x, i)(A_d(P)^{-1})_{id}$ and

$$\mathcal{D}_r(\theta') = \left\{ P \in \mathcal{P} : \sum_{i=1}^d (\pi_{\theta'})_i D((P_{\theta'})_i \parallel (P)_i) \leq r \right\}.$$

In the following we will sometimes abbreviate $\mathcal{D}_r(\theta')$ by \mathcal{D} when the dependence on θ' and r is irrelevant. Problem (7) is more attractive than (4a) from a computational point of view because its feasible set is closely related to the popular relative entropy uncertainty sets, which have received considerable attention in robust optimization (Ben-Tal et al., 2013). On the other hand, while the objective function $c(x, \theta)$ of the original problem (4a) was linear in θ , the objective function $\Psi(x, P)$ of (7) is nonconvex in P . These properties suggest that (7) can be solved efficiently and exactly if its objective function is replaced with a linear approximation. We thus develop a customized Frank-Wolfe algorithm that generates approximate solutions for

Algorithm 1 Frank-Wolfe algorithm for solving (7)

- 1: **Input:** $r > 0$, $\theta' > 0$, $x \in X$, $\varepsilon > 0$, $\pi_{\theta'}$, $P_{\theta'}$
- 2: **Output:** P
- 3: Initialize $P^{(0)}$, $g^{(0)} = 2\varepsilon$, $m = 0$
- 4: **while** $g^{(m)} > \varepsilon$ **do**
- 5: Compute $S^{(m)} \in \operatorname{argmax}_{S \in \mathcal{D}} \langle S, \nabla_P \Psi(x, P^{(m)}) \rangle$
- 6: Compute $g^{(m)} = \langle S^{(m)} - P^{(m)}, \nabla_P \Psi(x, P^{(m)}) \rangle$
- 7: **if** $g^{(m)} \leq \varepsilon$ **then**
- 8: Return $P = P^{(m)}$
- 9: **end if**
- 10: Compute $\gamma_m \in \operatorname{argmax}_{\gamma \in [0,1]} \Psi(x, (P^{(m)} + \gamma(S^{(m)} - P^{(m)})))$
- 11: $P^{(m+1)} = P^{(m)} + \gamma_m(S^{(m)} - P^{(m)})$
- 12: $m \rightarrow m + 1$
- 13: **end while**




Figure 1. Flow diagram of Algorithm 2.

problem (7) by solving a sequence of linearized oracle subproblems. While our main Frank-Wolfe routine is presented in Algorithm 1, its key subroutine for solving the oracle subproblems is described in Algorithm 2.

Algorithm 1 generates a sequence of approximate solutions for problem (7) that enjoy rigorous optimality guarantees.

Theorem 5 (Convergence of Algorithm 1). *For any fixed $\varepsilon > 0$ and $\theta' > 0$, Algorithm 1 finds an approximate stationary point for problem (7) with a Frank-Wolfe gap of at most ε in $\mathcal{O}(1/\varepsilon^2)$ iterations.*

The proof of Theorem 5 leverages convergence results for generic Frank-Wolfe algorithms by Lacoste-Julien (2016). A detailed proof is relegated to Appendix C.

We now explain the implementation of Algorithm 1. Note that the only nontrivial step is the solution of the oracle subproblem in line 5. To construct this subproblem we need the gradient $\nabla_P \Psi(x, P)$, which can be computed in closed form. Full details are provided in Appendix C. Even though the oracle subproblem is convex by construction (see Proposition 14), it is still computationally expensive because it involves d^2 decision variables $P \in \mathbb{R}^{d \times d}$ subject

to $\mathcal{O}(d)$ linear constraints as well as a single conditional relative entropy constraint. Algorithm 2 thus solves the dual oracle subproblem, which involves only d decision variables $\eta \in \mathbb{R}^d$ subject to box constraints. Strong duality holds because $P_{\theta'}$ represents a Slater point for the primal problem (see Proposition 18). The dual oracle subproblem is amenable to efficient numerical solution via Stochastic Gradient Descent (SGD), and a primal maximizer can be recovered from any dual minimizer by using the problem's first order optimality conditions (see Proposition 18). The high-level structure of Algorithm 2 is visualized in Figure 1.

To construct the dual oracle subproblem, fix a decision $x \in X$ and an anchor point $P^{(m)} \in \mathcal{P}$, and set $C = \nabla_P \Psi(x, P^{(m)})$. In addition, define $\underline{\eta}_i = \max_j \{C_{ij}\}$ and

$$\bar{\eta}_i = \frac{1}{1 - e^{-r}} \left(d \max_{i,j} \{C_{ij}\} - e^{-r} \text{tr}(C^\top P_{\theta'}) \right) - \sum_{k \neq i} \underline{\eta}_k.$$

The dual subproblem minimizes $Q(\eta) = \sum_{i=1}^d Q_i(\eta)$ over the box $[\underline{\eta}, \bar{\eta}]$, where $Q_i(\eta) = \eta_i - \lambda^*(\eta)/d$ and

$$\lambda^*(\eta) = \exp \left(\sum_{i,j=1}^d (\pi_{\theta'})_i (P_{\theta'})_{ij} \log \left(\frac{\eta_i - C_{ij}}{(\pi_{\theta'})_i} \right) - r \right)$$

is a convex function. As $Q(\eta)$ is reminiscent of an empirical average, the dual oracle subproblem lends itself to efficient numerical solution via SGD, which obviates costly high-dimensional gradient calculations. In addition, the partial derivatives of $Q_i(\eta)$ are readily available in closed form as

$$\frac{\partial}{\partial \eta_i} Q_i(\eta^{(n)}) = d - (\pi_{\theta'})_i \sum_{j=1}^d \frac{(P_{\theta'})_{ij}}{\eta_i^{(n)} - C_{ij}}.$$

We set the SGD learning rate to $\ell = KN^{-1/2}$ for some constant $K > 0$, which ensures that the suboptimality of the iterates generated by the SGD algorithm decays as $\mathcal{O}(1/\sqrt{N})$ after N iterations; see (Nemirovski et al., 2009, Section 2.2). A summary of Algorithm 2 in pseudocode is given below.

Algorithm 2 Solution of the oracle subproblem

- 1: **Input:** $\underline{\eta}, \bar{\eta}, \ell, P_{\theta'}, N$
 - 2: **Output:** $S^{(N)} \in \arg \min_{S \in \mathcal{D}} \langle S, -\nabla_P \Psi(x, P^{(m)}) \rangle$
 - 3: Initialize $\eta^{(0)}$ with $\eta_i \leq \eta_i^{(0)} \leq \bar{\eta}_i, i = 1, \dots, d$
 - 4: **for** $n = 0, \dots, N$ **do**
 - 5: **for** $i = 1, \dots, d$ **do**
 - 6: $y_i^{(n)} = \eta_i^{(n)} - \ell \frac{\partial}{\partial \eta_i} Q_i(\eta^{(n)})$
 - 7: **end for**
 - 8: $\eta^{(n+1)} \in \arg \min_{\underline{\eta}_i \leq s_i \leq \bar{\eta}_i, i=1, \dots, d} \|s - y^{(n)}\|_2^2$
 - 9: **end for**
 - 10: $\eta^* = \eta^{N+1}, \lambda^* = \lambda^*(\eta^*)$
 - 11: $(S^{(N)})_{d(i-1)+j} = (\pi_{\theta'})_i (P_{\theta'})_{ij} \lambda^* / (\eta_i^* - \tilde{c}_j^{(i)})$ for $i, j = 1, \dots, d$
-

Theorem 6 (Convergence of Algorithm 2). *For any fixed $\varepsilon > 0$, Algorithm 2 outputs an ε -suboptimal solution of the primal oracle subproblem in $\mathcal{O}(1/\varepsilon^2)$ iterations.*

Theorem 6 follows from standard convergence results in (Nemirovski et al., 2009, Section 2.2) applied to the dual oracle subproblem. We finally remark that the stepsize γ_m to be computed in line 1 of Algorithm 1 can be obtained via a direct line search method.

Remark 7 (Overall complexity of Algorithm 1). *Denote by ε_1 the error of the Frank-Wolfe algorithm (Algorithm 1) provided by Theorem 5, and by ε_2 the error of the subroutine (Algorithm 2) provided by Theorem 6. That is, suppose we solve the Frank-Wolfe subproblem approximately in the sense that $\langle S^{(m)}, \nabla_P \Psi(x, P^{(m)}) \rangle \geq \max_{S \in \mathcal{D}} \langle S, \nabla_P \Psi(x, P^{(m)}) \rangle - \kappa \varepsilon_2$, for some constant $\kappa > 0$. We denote by C_f the curvature constant of the function Ψ . By (Lacoste-Julien, 2016), after running $\mathcal{O}((1 + 2\kappa \varepsilon_2 / C_f) / \varepsilon_1^2)$ many iterations of Algorithm 1 we achieve an accumulated error below ε_1 . As Algorithm 1 calls Algorithm 2 in total $\mathcal{O}(1/\varepsilon_1^2)$ times, the overall complexity coincides with the number of operations needed to solve $\mathcal{O}(1/\varepsilon_1^2)$ times a d -dimensional minimization problems of a smooth convex function over a compact box. This significantly reduces the computational cost that would be needed to solve in each step the original d^2 -dimensional nonconvex problem (4a) directly. A further reduction of complexity is possible by applying a randomized Frank-Wolfe algorithm, as suggested by Reddi et al. (2016).*

Given an efficient computational method to solve the predictor problem (4a), we now address the solution of the prescriptor problem (4b). Since our Frank-Wolfe algorithm only enables us to access the values of the objective function $\hat{c}_r(x, \theta')$ for fixed values of $x \in X$ and $\theta' \in \Theta'$, we cannot use gradient-based algorithms to optimize over x . Nevertheless, we can resort to derivative-free optimization (DFO) approaches, which can find local or global minima under appropriate regularity conditions on $\hat{c}_r(x, \theta')$.

Besides popular heuristic algorithms such as simulated annealing (Metropolis et al., 1953) or genetic algorithms (Holland, 1992), there are two major classes of DFO methods: direct search and model-based methods. The direct search methods use only comparisons between function evaluations to generate new candidate points, while the model-based methods construct a smooth surrogate objective function. We focus here on direct search methods due to their solid worst-case convergence guarantees. Descriptions of these algorithms and convergence proofs can be found in the recent textbook by Audet & Hare (2017). For a survey of model-based DFO see (Conn et al., 2009). Off-the-shelf software to solve constrained DFO problems includes `fminsearch` in Matlab or `NOMAD` in C++; see (Rios & Sahinidis, 2013) for an extensive list of solvers.

When applied to the extreme barrier function formulation of the prescriptor problem (4b), which is obtained by ignoring the constraints and adding a high penalty to each infeasible decision (Audet & Dennis, 2006), the directional direct-search method by Vicente (2013) offers an attractive worst-case convergence guarantee thanks to (Vicente, 2013, Theorem 2 & Corollary 1). The exact algorithm that we use to solve (4b) is outlined in Algorithm 3 in Appendix E.

Theorem 8 (Convergence of Algorithm 3 (Vicente, 2013)). *Fix any $\theta' \in \Theta'$ and $\varepsilon > 0$. If $\nabla_x \hat{c}_r(x, \theta')$ is Lipschitz continuous in x , then Algorithm 3 outputs a solution x_N with $\|\nabla_x \hat{c}_r(x_N, \theta')\| \leq \varepsilon$ after $N \sim \mathcal{O}(\varepsilon^{-2})$ iterations. If $\hat{c}_r(x, \theta')$ is also convex in x , then we further have*

$$\|\hat{c}_r(x_N, \theta') - \min_{x \in X} \hat{c}_r(x, \theta')\| \leq \varepsilon.$$

5. An Optimal Hypothesis Test

We now develop a Neyman-Pearson-type hypothesis test for determining which one of two prescribed Markov chains has generated a given sequence ξ_1, \dots, ξ_T of training samples. Specifically, we assume that these training samples were either generated by the Markov chain corresponding to the doublet probability mass function $\theta^{(1)} \in \Theta$ or by the one corresponding to $\theta^{(2)} \in \Theta$, that is, we have $\theta \in \{\theta^{(1)}, \theta^{(2)}\}$. Next, we construct a decision rule encoded by the open set

$$\mathcal{B} = \{\theta' \in \Theta' : D_c(\theta' \parallel \theta^{(1)}) < D_c(\theta' \parallel \theta^{(2)})\}, \quad (8)$$

which predicts $\theta = \theta^{(1)}$ if $\hat{\theta}_T \in \mathcal{B}$ and predicts $\theta = \theta^{(2)}$ otherwise. The quality of any such decision rule is conveniently measured by the type I and type II error probabilities

$$\alpha_T = \mathbb{P}_{\theta^{(1)}}(\hat{\theta}_T \notin \mathcal{B}) \quad \text{and} \quad \beta_T = \mathbb{P}_{\theta^{(2)}}(\hat{\theta}_T \in \mathcal{B}) \quad (9)$$

for $T \in \mathbb{N}$, respectively. Specifically, α_T represents the probability that the proposed decision rule wrongly predicts $\theta = \theta^{(2)}$ if the training samples were generated by $\theta^{(1)}$, and β_T represents the probability that the decision rule wrongly predicts $\theta = \theta^{(1)}$ if the data was generated by $\theta^{(2)}$.

Proposition 9 (Decay of the error probabilities). *If $\theta^* \in \arg \min_{\theta' \notin \mathcal{B}} D_c(\theta' \parallel \theta^{(1)})$ and $r = D_c(\theta^* \parallel \theta^{(1)}) > 0$, then*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log \alpha_T = -r \quad \text{and} \quad \lim_{T \rightarrow \infty} \frac{1}{T} \log \beta_T = -r.$$

Proposition 9 ensures that the proposed decision rule is powerful in the sense that both error probabilities decay exponentially with the sample size T at some rate r that can at least principally be computed. The next theorem is inspired by the celebrated Neyman-Pearson lemma (Cover & Thomas, 2006, Theorem 11.7) and shows that this decision rule is actually optimal in a precise statistical sense.

Theorem 10 (Optimality of the hypothesis test). *Let $\bar{\alpha}_T$ and $\bar{\beta}_T$ represent the type I and type II error probabilities of the decision rule obtained by replacing \mathcal{B} with any other open set $\bar{\mathcal{B}} \subseteq \Theta'$. If $\lim_{T \rightarrow \infty} \frac{1}{T} \log \bar{\alpha}_T \leq \lim_{T \rightarrow \infty} \frac{1}{T} \log \alpha_T$, then $\lim_{T \rightarrow \infty} \frac{1}{T} \log \bar{\beta}_T \geq \lim_{T \rightarrow \infty} \frac{1}{T} \log \beta_T$.*

Theorem 10 implies that the proposed hypothesis test is Pareto optimal in the following sense. If any test with the same test statistic $\hat{\theta}_T$ but a different set $\bar{\mathcal{B}}$ has a faster decaying type I error probability, then it must necessarily have a slower decaying type II error probability. In Appendix D.1 we construct two different Markov chains that are almost indistinguishable on the basis of a finite training dataset. Indeed, we will show that the (optimal) decay rates of the type I and type II error probabilities can be arbitrarily small.

6. Revenue Maximization under a Markovian Brand Switching Model

We now test the performance of our approach in the context of a revenue maximization model, which aims to recognize and exploit repeat-buying and brand-switching behavior. This model assumes that the probability of a customer buying a particular brand depends on the brand purchased last (Herniter & Magee, 1961; Chintagunta et al., 1991; Gensler et al., 2007; Leeflang et al., 2015). Adopting the perspective of a medium-sized retailer, we then formulate an optimization problem that seeks to maximize profit from sales by anticipating the long-term average demand of each brand based on the brand transition probabilities between different customer segments characterized by distinct demographic attributes. For instance, younger people might prefer relatively new brands, while senior people might prefer more established brands. Once such insights are available, the retailer selects the brands to put on offer with the goal to maximize long-run average revenue. We assume that the brand purchasing behavior of the customers is exogenous, which is realistic unless the retailer is a monopolist or has at least significant market power. Therefore, the retailer's ordering decisions have no impact on the behavior of the customers. The ergodicity of the brand choice dynamics of each customer segment can be justified by the ergodicity of demographic characteristics (Ezzati, 1974).

6.1. Problem Formulation

Assume that there are d different brands of a particular good, and denote by $a \in \mathbb{R}^d$ the vector of retail prices per unit of the good for each brand. The seller needs to decide which brands to put on offer. This decision is encoded by a binary vector $x \in \{0, 1\}^d$ with $x_j = 1$ if brand j is offered and $x_j = 0$ otherwise. To quantify the seller's revenue, we need to specify the demand for each brand. To this end, assume that the customers are clustered into n groups with

similar demographic characteristics (Kuo et al., 2002). The percentage of customers in group k is denoted by $w_k \in \mathbb{R}_+$, and the aggregate brand preferences ξ_t^k of the customers in group k and period t represent an ergodic Markov chain on $\Xi = \{1, \dots, d\}$ with stationary distribution π_{θ^k} . Thus, the long-run average cost per customer and time period is

$$\begin{aligned} c(x, \theta) &= \lim_{T \rightarrow \infty} -\frac{1}{T} \left(\sum_{t=0}^{T-1} \sum_{j=1}^d a_j x_j \sum_{k=1}^n w_k \mathbf{1}_{\xi_t^k=j} \right) \\ &= \sum_{k=1}^n w_k \sum_{j=1}^d -a_j x_j (\pi_{\theta^k})_j = \sum_{k=1}^n w_k c^k(x, \theta^k), \end{aligned}$$

where $c^k(x, \theta^k) = \sum_{j=1}^d -a_j x_j (\pi_{\theta^k})_j$ and $\theta = (\theta^k)_{k=1}^n$. Here, the second equality exploits the ergodic theorem (Ross, 2010, Theorem 4.1). It is easy to show that $c^k(x, \theta^k)$ can be expressed as an expected loss with respect to a probability measure encoded by θ^k , much like the objective function of problem (1). Finally, we define $X = \{x \in \{0, 1\}^d : Cx \leq b\}$ for some $C \in \mathbb{R}^{m \times d}$ and $b \in \mathbb{R}^m$. The linear inequalities $Cx \leq b$ may capture budget constraints or brand exclusivity restrictions etc. The cost minimization problem $\min_{x \in X} c(x, \theta)$ can thus be viewed as an instance of (1). If θ is unknown and only T training samples are available, we construct the empirical doublet distributions $\hat{\theta}_T^k$ for all $k = 1, \dots, n$ and solve the DRO problem

$$\min_{x \in X} \sum_{k=1}^n w_k \hat{c}_r^k(x, \hat{\theta}_T^k), \quad (10)$$

where

$$\hat{c}_r^k(x, \hat{\theta}_T^k) = \max_{\theta^k \in \mathcal{C}_1} \left\{ c^k(x, \theta^k) : D_c(\hat{\theta}_T^k \| \theta^k) \leq r \right\}. \quad (11)$$

Below we denote by $\hat{c}_T(x, \hat{\theta}_T)$ the objective function and by $\hat{x}_r(\hat{\theta}_T)$ a minimizer of (10), where $\hat{\theta}_T = (\hat{\theta}_T^k)_{k=1}^n$.

6.2. Numerical Experiments

We now validate the out-of-sample guarantees of the proposed Markov DRO approach experimentally when the observable data is indeed generated by ergodic Markov chains and compare our method against three state-of-the-art approaches to data-driven decision making from the literature. The first baseline is the Wasserstein DRO approach by Derman & Mannor (2020), which replaces the worst-case expectation problem (11) for customer group k with

$$\hat{c}_r^k(x, \hat{\theta}_T^k) = \max_{P \in \mathcal{D}_r(\hat{\theta}_T^k)} \Psi(x, P),$$

where the reparametrized objective function $\Psi(x, P)$ is defined as in Lemma 4, and the ambiguity set is defined as

$$\mathcal{D}_r(\hat{\theta}_T^k) = \left\{ P \in \mathcal{P} : d_W((P_{\hat{\theta}_T^k})_i, P_i) \leq r \forall i \right\}.$$

Here, d_W denotes the 1-Wasserstein distance, where the transportation cost between two states $i, j \in \Xi$ is set to $|i - j|$. We also compare our method against the classical SAA approach (Shapiro et al., 2014) and the DRO approach by Van Parys et al. (2021), which replaces the conditional relative entropy in the ambiguity set (11) with the ordinary relative entropy (*i.e.*, the Kullback-Leibler divergence). We stress that both of these approaches were originally designed for serially independent training samples. For small T it is likely that $\hat{\theta}_T \not\asymp 0$, in which case Algorithm 1 may not converge; see Theorem 5. In this case, we slightly perturb and renormalize $\hat{\theta}_T$ to ensure that $\hat{\theta}_T > 0$ and $\hat{\theta}_T \in \Theta'$.

Synthetic data. In the first experiment we solve a random instance of the revenue maximization problem with $n = 5$ customer groups and $d = 10$ brands, where the weight vector w and the price vector a are sampled from the uniform distributions on Δ_n and $\{1, \dots, 10\}^d$, respectively. To construct the transition probability matrix of the Markov chain reflecting the brand switching behavior of any group k , we first sample a random matrix from the uniform distribution on $[0, 1]^{d \times d}$, increase two random elements of this matrix to 4 and 5, respectively, and normalize all rows to render the matrix stochastic. As the data-generating Markov chains in this synthetic experiment are known, the true out-of-sample risk of any data-driven decision can be computed exactly.¹




Figure 2. Out-of-sample risk of different data-driven decisions based on synthetic data. Solid lines and shaded areas represent empirical averages and empirical 90% confidence intervals, respectively, computed using 100 independent training datasets.




Figure 3. Empirical out-of-sample disappointment of different data-driven decisions based on synthetic data, computed using 100 independent training datasets.

¹The Matlab code for reproducing all results is available from https://github.com/mkvdro/DRO_Markov.

Real-world data. The second experiment is based on a marketing dataset from Kaggle,² which tracks the purchasing behavior of 2,000 customers with respect to $d = 5$ brands of chocolates. The customers are clustered into $n = 5$ groups based on their age, education level and income by using the K -means++ clustering algorithm by Vassilvitskii & Arthur (2006). The resulting clusters readily induce a weight vector w . The price vector a is sampled randomly from the uniform distribution on $\{1, \dots, 10\}^d$. In addition, we concatenate the purchase histories of all customers in any group k and interpret the resulting time series as a trajectory of the unknown Markov chain corresponding to group k .




Figure 4. Out-of-sample risk of different data-driven decisions based on 200 test samples from the Kaggle dataset.

Results. Figures 2 and 4 show that the proposed Markov DRO method outperforms the SAA scheme and the DRO approach by Van Parys et al. (2021) tailored to i.i.d. data (denoted as ‘i.i.d. DRO’) in that its out-of-sample risk displays a smaller mean as well as a smaller variance. For small training sample sizes T , our method is also superior to the Wasserstein DRO approach by Derman & Mannor (2020), but for $T \gtrsim 300$ the two methods display a similar behavior. Note also that the mean as well as the variance of the out-of-sample risk increase with r for all three DRO approaches. This can be explained by the increasing inaccuracy of a model that becomes more and more pessimistic. In fact, if no transitions are observed from a certain brand A to another brand B given that the data follows an ergodic Markov process, then brand B’s true long-term value is overestimated when the data is falsely treated as i.i.d. This phenomenon explains the discrepancy between the out-of-sample risk incurred by methods that treat the data as Markovian or as i.i.d., respectively; see Figure 4. Note also that if the prescribed decay rate r drops to 0, then all methods should collapse to the SAA approach if $\theta' > 0$. If θ' has zero entries, however, the worst-case expected risk with respect to a conditional relative entropy ambiguity set does not necessarily reduce to the empirical risk if r tends to 0. We shed more light on this phenomenon in Remark 12. Figure 3 further shows that the Markov DRO approach results in the smallest out-of-sample disappointment among all tested methods for

²<https://www.kaggle.com/khalidnasereddin/retail-dataset-analysis>

a wide range of rate parameters r . Comparing the charts for $T = 10$, $T = 300$ and $T = 500$, we also see that the out-of-sample disappointment of the Markov DRO method decays with T , which is consistent with Theorem 2. In practice, the optimal choice of the decay rate r remains an open question. To tune r with the aim to trade off in-sample performance against out-of-sample disappointment, one could use the rolling window heuristic for model selection in time series models by Bergmeir & Benítez (2012).

Scalability. We also compare the scalability of the proposed Frank-Wolfe algorithm against that of an interior point method by Waltz et al. (2006) (with or without exact gradient information), which represents a state-of-the-art method in nonlinear optimization. To this end, we solve 10 instances of the worst-case expectation problem (11) with rate parameter $r = 1$ for a fixed decision x sampled from the uniform distribution on $\{0, 1\}^d$. The 10 instances involve independent training sets of size $T = 5,000$, each of which is sampled from the same fixed Markov chain constructed as in the experiments with synthetic data. The problem instances are modelled in MATLAB, and all experiments are run on an Intel i5-5257U CPU (2.7GHz) computer with 16GB RAM. Figure 5 shows that our Frank-Wolfe algorithm is significantly faster than both baseline methods whenever the Markov chain accommodates at least $d = 100$ states. In particular, note that the interior point methods run out of memory as soon as d exceeds 200.




Figure 5. Runtimes (in seconds) of different methods for solving problem (4a). Lines represent means and shaded areas represent the ranges between the smallest and the largest runtimes.

Acknowledgements. We thank Bart Van Parys for valuable comments on the paper and for suggesting the Markov coin example in Appendix D.1. This research was supported by the Swiss National Science Foundation under the NCCR Automation, grant agreement 51NF40_180545.

References

- Audet, C. and Dennis, J. E. Mesh adaptive direct search algorithms for constrained optimization. **SIAM Journal on Optimization**, 17(1):188–217, 2006.
- Audet, C. and Hare, W. **Derivative-Free and Blackbox Optimization**. Springer, 2017.
- Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. **Robust Optimization**. Princeton University Press, 2009.
- Ben-Tal, A., Den Hertog, D., De Waegenare, A., Melenberg, B., and Rennen, G. Robust solutions of optimization problems affected by uncertain probabilities. **Management Science**, 59(2):341–357, 2013.
- Bergmeir, C. and Benítez, J. M. On the use of cross-validation for time series predictor evaluation. **Information Sciences**, 191:192–213, 2012.
- Berman, A. and Plemmons, R. J. **Nonnegative Matrices in the Mathematical Sciences**. SIAM, 1994.
- Bertsimas, D. and Sim, M. The price of robustness. **Operations Research**, 52(1):35–53, 2004.
- Bertsimas, D., Gupta, V., and Kallus, N. Data-driven robust optimization. **Mathematical Programming**, 167(2):235–292, 2018.
- Chintagunta, P. K., Jain, D. C., and Vilcassim, N. J. Investigating heterogeneity in brand preferences in logit models for panel data. **Journal of Marketing Research**, 28(4):417–428, 1991.
- Conn, A. R., Scheinberg, K., and Vicente, L. N. **Introduction to Derivative-Free Optimization**. SIAM, 2009.
- Cover, T. M. and Thomas, J. A. **Elements of Information Theory**. John Wiley & Sons, 2006.
- Delage, E. and Ye, Y. Distributionally robust optimization under moment uncertainty with application to data-driven problems. **Operations Research**, 58(3):595–612, 2010.
- Dembo, A. and Zeitouni, O. **Large Deviations Techniques and Applications**. Springer, 2010.
- den Hollander, F. **Large Deviations**. American Mathematical Society, 2008.
- Derman, E. and Mannor, S. Distributional robustness and regularization in reinforcement learning. **arXiv e-prints**, 2020.
- Dou, X. and Anitescu, M. Distributionally robust optimization with correlated data from vector autoregressive processes. **Operations Research Letters**, 47(4):294 – 299, 2019.
- Duchi, J., Glynn, P., and Namkoong, H. Statistics of robust optimization: A generalized empirical likelihood approach. **Mathematics of Operations Research**, 2021. Forthcoming.
- Eriksson, K., Estep, D., and Johnson, C. **Applied Mathematics: Body and Soul. Volume 1: Derivatives and Geometry in \mathbb{R}^3** . Springer, 2004.
- Ezzati, A. Forecasting market shares of alternative home-heating units by Markov process using transition probabilities estimated from aggregate time series data. **Management Science**, 21(4):462–473, 1974.
- Frank, M. and Wolfe, P. An algorithm for quadratic programming. **Naval Research Logistics Quarterly**, 3:95–110, 1956.
- Gensler, S., Dekimpe, M. G., and Skiera, B. Evaluating channel performance in multi-channel environments. **Journal of Retailing and Consumer Services**, 14(1):17–23, 2007.
- Goh, J. and Sim, M. Distributionally robust optimization and its tractable approximations. **Operations Research**, 58(4):902–917, 2010.
- Herniter, J. D. and Magee, J. F. Customer behavior as a Markov process. **Operations Research**, 9(1):105–122, 1961.
- Holland, J. H. **Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence**. MIT Press, 1992.
- Jaggi, M. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In **International Conference on Machine Learning**, pp. 427–435, 2013.
- Jin, Y. and Sidford, A. Efficiently solving MDPs with stochastic mirror descent. In **International Conference on Machine Learning**, pp. 4890–4900, 2020.
- Kirschner, J., Bogunovic, I., Jegelka, S., and Krause, A. Distributionally robust Bayesian optimization. In **Artificial Intelligence and Statistics**, pp. 2174–2184, 2020.
- Kuhn, D., Esfahani, P. M., Nguyen, V. A., and Shafieezadeh-Abadeh, S. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In **Operations Research & Management Science in the Age of Analytics**, pp. 130–166. 2019.
- Kuo, R., Ho, L., and Hu, C. M. Integration of self-organizing feature map and k-means algorithm for market segmentation. **Computers & Operations Research**, 29(11):1475–1493, 2002.

- Lacoste-Julien, S. Convergence rate of Frank-Wolfe for non-convex objectives. **arXiv e-prints**, 2016.
- Leeftang, P. S., Wieringa, J. E., Bijmolt, T. H., and Pauwels, K. H. **Modeling Markets**. Springer, 2015.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. Equation of state calculations by fast computing machines. **The Journal of Chemical Physics**, 21(6):1087–1092, 1953.
- Mohajerin Esfahani, P. and Kuhn, D. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. **Mathematical Programming**, 171(1-2):115–166, 2018.
- Namkoong, H. and Duchi, J. C. Stochastic gradient methods for distributionally robust optimization with f-divergences. In **Advances in Neural Information Processing Systems**, pp. 2208–2216, 2016.
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust stochastic approximation approach to stochastic programming. **SIAM Journal on Optimization**, 19(4):1574–1609, 2009.
- Perron, O. Zur Theorie der Matrices. **Mathematische Annalen**, 64(2):248–263, 1907.
- Reddi, S. J., Sra, S., Póczos, B., and Smola, A. Stochastic Frank-Wolfe methods for nonconvex optimization. In **Annual Conference on Communication, Control, and Computing**, pp. 1244–1251, 2016.
- Rios, L. M. and Sahinidis, N. V. Derivative-free optimization: A review of algorithms and comparison of software implementations. **Journal of Global Optimization**, 56(3):1247–1293, 2013.
- Ross, M. **Introduction to Probability Models**. Elsevier, 2010.
- Scarf, H. A min-max solution of an inventory problem. In Arrow, K. J., Karlin, S., and Scarf, H. (eds.), **Studies in the Mathematical Theory of Inventory and Production**, pp. 201–209. Stanford University Press, 1958.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. **Lectures on Stochastic Programming: Modeling and Theory**. SIAM, 2014.
- Smith, J. and Winkler, R. The optimizer’s curse: Skepticism and postdecision surprise in decision analysis. **Management Science**, 52(3):311–322, 2006.
- Staib, M. and Jegelka, S. Distributionally robust optimization and generalization in kernel methods. In **Advances in Neural Information Processing Systems**, pp. 9134–9144, 2019.
- Sutter, T., Van Parys, B. P. G., and Kuhn, D. A general framework for optimal data-driven optimization. **arXiv e-prints**, 2020.
- Van Parys, B. P. G., Mohajerin Esfahani, P., and Kuhn, D. From data to decisions: Distributionally robust optimization is optimal. **Management Science**, 2021. Articles in Advance.
- Vassilvitskii, S. and Arthur, D. k-means++: The advantages of careful seeding. In **ACM-SIAM Symposium on Discrete Algorithms**, pp. 1027–1035, 2006.
- Vicente, L. N. Worst case complexity of direct search. **EURO Journal on Computational Optimization**, 1(1-2):143–153, 2013.
- Waltz, R. A., Morales, J. L., Nocedal, J., and Orban, D. An interior algorithm for nonlinear optimization that combines line search and trust region steps. **Mathematical Programming**, 107(3):391–408, 2006.
- Wiesemann, W., Kuhn, D., and Sim, M. Distributionally robust convex optimization. **Operations Research**, 62(6):1358–1376, 2014.