Online First-Order Framework for Robust Convex Optimization

Nam Ho-Nguyen[1] and Fatma Kılınç-Karzan[1]

[1]Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, 15213, USA.

Abstract

Robust optimization (RO) has emerged as one of the leading paradigms to efficiently model parameter uncertainty. The recent connections between RO and problems in statistics and machine learning domains demand for solving RO problems in ever more larger scale. However, the traditional approaches for solving RO formulations based on building and solving robust counterparts or the iterative approaches utilizing nominal feasibility oracles can be prohibitively expensive and thus significantly hinder the scalability of RO paradigm. In this paper, we present a general and flexible iterative framework to approximately solve robust convex optimization problems that is built on a fully online first-order paradigm. In comparison to the existing literature, a key distinguishing feature of our approach is that it only requires access to first-order oracles that are remarkably cheaper than pessimization or nominal feasibility oracles, while maintaining the same convergence rates. This, in particular, makes our approach much more scalable and hence preferable in large-scale applications, specifically those from machine learning and statistics domains. We also provide new interpretations of existing iterative approaches in our framework and illustrate our framework on robust quadratic programming.

1 Introduction

Robust optimization (RO) is one of the leading modeling paradigms for optimization problems under uncertainty. As opposed to the other approaches, RO seeks a solution that is immunized against *all* possible realizations of uncertain model parameters (noises) from a given uncertainty set. It is widely adopted in practice mainly because of its computational tractability. We refer the reader to the paper by Ben-Tal and Nemirovski [6], the book by Ben-Tal et al. [4] and surveys [8, 9, 11, 16] for a detailed account of RO theory and numerous applications.

Recently, fascinating connections have been established between problems from the statistics and machine learning domains and robust optimization. More precisely, it is demonstrated that RO can be used to achieve desirable statistical properties such as stability, sparsity, and consistency. For example, for linear regression problems, El Ghaoui and Lebret [19] and Xu et al. [45] respectively establish the equivalence of the ridge regression and Lasso to specific RO formulations of unregularized regression problems. Moreover, Xu et al. [44] exhibit similar results in the context of regularizing support vector machines (SVMs), and [44, 45] validate the statistical consistency of methods such as SVM and Lasso via RO methodology. In addition to these RO interpretations of regularization techniques used in statistics and machine learning, robust versions of many problems from these domains are gaining traction. For example, [43] examines robust variants of SVMs and

other classification problems, and [2] explores a robust formulation for kernel classification problems. We refer the reader to [16, 5] and references therein for further examples and details on connections between robust optimization and statistics and machine learning.

These recent connections not only highlight the importance of RO methodology but also present algorithmic challenges where the scalability of RO algorithms with problem dimension becomes crucial. The primary method for solving a robust convex optimization problem is to transform it into an equivalent deterministic problem called the *robust counterpart*. Under mild assumptions, this yields a convex and tractable robust counterpart problem (see [4, 11, 3]), which can then be solved using existing convex optimization software and tools. This traditional approach has seen much success in decision making domain, nevertheless it has a major drawback that the reformulated robust counterpart is often not as scalable as the deterministic nominal program. In particular, the robust counterpart can easily belong to a different class of optimization problems as opposed to the underlying original deterministic problem. For example, a linear program (LP) with ellipsoidal uncertainty is equivalent to a convex quadratic program (QP), and similarly, a conic-quadratic program with ellipsoidal uncertainty is equivalent to a semidefinite program (SDP) (see e.g., [4, 11]). It is well-known that convex QPs as opposed to LPs, and SDPs as opposed to convex QPs are much less scalable in practice. This then presents a critical challenge in applying RO methodology in big data applications frequently encountered in machine learning and statistics, where even solving the original deterministic nominal problem to high accuracy is prohibitively time-consuming.

The iterative schemes that alternate between the generation/update of candidate solutions and the realizations of noises offer a convenient remedy to the scalability issues associated with the robust counterpart approach. Thus far, such approaches [31] and [5] have relied on two oracles: (*i*) *solution oracles* to solve instances of extended (or nominal) problems with constraint structures similar to (or the same as) the deterministic problem, and (*ii*) *noise oracles* to generate/update particular realizations of the uncertain parameters. At each iteration of these schemes, both solution and noise oracles are called, and their outputs are used to update the inputs of each other oracle in the next iteration. Because solution oracles rely on a solver of the same class capable of solving the deterministic problem, these iterative approaches circumvent the issue of the robust counterpart approach potentially relying on a different solver. Nevertheless, these iterative approaches still suffer from a serious drawback: the solution oracles in [31, 5] themselves can be expensive as they require solving extended or nominal optimization problems completely. While solving the nominal problem is not as computationally demanding as solving the robust counterpart, the overall procedure depending on repeated calls to such oracles can be prohibitive. In fact, each such call to a solution oracle may endure a significant computational cost, which is at least as much as the computational cost of solving an instance of the deterministic nominal problem. Note that, to ensure scalability, most applications in machine learning and statistics already need to rely on cheap first-order methods for solving deterministic nominal problems.

In this paper, we propose an efficient iterative framework for solving robust convex optimization problems which can rely on, in an *online* fashion, much cheaper *first-order oracles* in place of full solution and noise oracles. In particular, in each iteration, instead of solving a complete optimization problem within the solution and/or noise oracles, we show that simple simultaneous updates on the solution and noise in an online fashion using only first-order information from the deterministic constraint structure is sufficient to solve robust convex optimization problems. Moreover, we show that the number of calls to such online first-order (OFO) oracles is not only at most that of the

state-of-the-art iterative approaches utilizing full optimization based oracles for solution and/or noise, but also almost independent of the dimension of the problem. Therefore, this makes our approach especially attractive for applications in statistics and machine learning domains where it is critical to maintain that the overall approach has both gracious dependence on the dimension of the problem and cheap iterations. We outline our contribution more concretely after discussing the most relevant literature.

Related Work

Thus far, the iterative approaches, which bypass the restrictions of the robust counterparts, work with *extended* nominal problems that belong to the same class as the deterministic nominal one by carefully controlling the constraints included in the formulation corresponding to noise realizations.

For robust binary linear optimization problems with only objective function uncertainty and a polyhedral uncertainty set, Bertsimas and Sim [13] suggest an approach which relies on solving $n + 1$ number of instances of the nominal problem, where n is the dimension of the problem.

For robust convex optimization problems, Calafiore and Campi [15] study a 'constraint sampling' approach based on forming a single extended nominal problem of the same class as the deterministic one via i.i.d. sampling of noise realizations. They show that the optimal solution to this extended nominal problem is robust feasible with high probability where the probability depends on the sampling procedure, the number of samples drawn, and the dimension.

Mutapcic and Boyd [31] follow a 'cutting-plane' type approach where in each iteration, a solution oracle is called to solve an extended nominal problem of the same class as the deterministic problem and a noise oracle, referred to as *pessimization oracle*, is invoked to iteratively expand and refine the extended nominal problem. Given a candidate solution, a pessimization oracle either certifies its feasibility with respect to the robust constraints or returns a new noise realization from the uncertainty set for which the solution is infeasible; then the nominal constraint associated with that particular noise realization is included in the extended problem. This process is repeated until a robust feasible solution is found or the last extended problem is found to be infeasible. In the overall procedure, the number of iterations (or calls to the pessimization oracle) can be exponential in the dimension. Despite this, [31] reports impressive computational results. The cutting-plane approach is also further tested on mixed integer linear problems in [12] and it is demonstrated that the same computational phenomenon holds.

Both of the approaches from [15] and [31] pose issues for high-dimensional problems. In [15], as the dimension grows, an extended problem with linearly more nominal constraints is required to ensure the high probability guarantee on finding a good quality solution. In [31] at each iteration, a nominal constraint is added to the extended nominal problem. The theoretical bound on the number of constraints that need to be added is exponential, so the extended problem in [31] can grow to be exponentially large. Moreover, in both cases the extended nominal problem may no longer have certain favorable problem structure of the deterministic nominal problem, such as a network flow structure.

To address these issues, in particular, the issue of solving extended nominal problems that are not only larger-in-size than the deterministic problem but also may lack certain favorable problem structure of the deterministic problem, Ben-Tal et al. [5] introduce a new iterative approach to approximately solve robust feasibility problems via a *nominal feasibility oracle* and running an online learning algorithm to choose noise realizations. Given a particular noise realization, the nominal feasibility oracle solves an instance of the deterministic nominal feasibility problem obtained by simply fixing the noise to the given value. Hence, the problem solved by this oracle

has the *same* number of constraints and the *same structure* as the original nominal problem; in particular its size does not grow in each iteration. This is an important distinguishing feature of this approach. The other distinguishing feature is that Ben-Tal et al. [5] replace the pessimization oracle of [31] by employing an online learning algorithm, which simply requires first-order information of the noise from the constraint functions. Moreover, [5] provides a dimension independent bound on the number of iterations (nominal feasibility oracle calls).

Because the approaches of both [31] and [5] are closely related to our work, we give a detailed summary of these in Sections 4.2 and 4.3 respectively and highlight their connections to our work. In fact, we show that they both can be seen as special cases of our framework.

We close with a brief summary of the assumptions on the computational requirements of these methods. The constraint sampling approach of [15] requires access to a sampling procedure on the uncertainty sets as well as an oracle capable of solving the extended nominal problem. The cutting plane approach of [31] replaces the sampling procedure of [15] with a noise oracle, namely the pessimization oracle that works with the uncertainty sets but still requires the same type of optimization oracle as a solution oracle to solve the extended problems. Ben-Tal et al. [5] substitute the pessimization oracle with an online learning-based procedure, which requires merely first-order information from the constraint functions and simple projection type operations on the associated uncertainty sets, but it still relies on a solution oracle capable of solving the original nominal problem, which is essentially the same (up to log factors) as the optimization oracles in [15] and [31]. If the deterministic problem admits special structure such as network flows etc., a specific solver can be used in the framework of [5], but this is not possible for [15] and [31].

Summary of Our Contributions

It is possible to view all of these iterative approaches as two iterative processes that run simultaneously and in conjunction with each other to generate/update solutions and noise realizations. This naturally leads to a dynamic game environment where in each round Player 1 chooses a solution and Player 2 chooses a realization of uncertain parameters. In this framework, the policies employed by these players in their decision making determine the nature of the final approach. In the case of [31], Player 1 considers all of the previous noise realizations when making his decision, whereas Player 2 simply reacts to the current solution when choosing the noise. In [5], Player 1 reacts to only the current noise in generating/updating the solution while Player 2 minimizes the regret associated with past solutions in choosing noise.

In this paper, we further analyze this interaction between Player 1 and Player 2, with the aim of deriving a simpler and computationally much less demanding iterative approach to solving RO problems. Our contributions can be summarized as follows.

1. We build a *general and flexible framework* for iteratively solving robust feasibility problems, and demonstrate its flexibility by describing it as a meta-template. By customizing our framework appropriately, we modify the pessimization oracle-based approach of [31] by replacing the extended nominal solver used in [31] with efficient first-order updates. We call this the *FO-based pessimization approach*, and demonstrate that as opposed to [31] it has both a much better bound on the number of oracle calls and far superior practical performance. We also provide a new interpretation of the nominal feasibility oracle-based approach of [5] as a special case within our framework. Furthermore, we extend the analysis of the approach of [5] under the same assumptions, e.g., access to a nominal optimization oracle, and show that it can solve the robust optimization problem directly without relying on a binary search (see

Remark 4.3).

2. When the original deterministic problem admits first-order oracles capable of providing gradient/subgradient information on each constraint function, we demonstrate that *online first-order* (OFO) algorithms can be used to iteratively generate/update solutions and noise realizations simultaneously in an online manner leading to robust feasibility/infeasibility certificates within our framework. In contrast to the approaches of [31] and [5], which rely on full nominal feasibility oracles to generate points, our OFO-based approach only requires simple update rules in each iteration and thus has much lower per-iteration cost. Besides, our noise oracle generates a realization of the noise in an online learning fashion as was done in [5], and hence it is less expensive than the pessimization oracle of [31].

3. In our framework, the number of iterations (or oracle calls) needed to obtain approximate robust solution or a robust infeasibility certificate is a function of the approximation guarantee ϵ and the complexities of the domains for the solution and the uncertainty set; in particular, our convergence rate is (almost) independent of both the number of robust constraints and the dimension of the deterministic problem. We also demonstrate that the iteration complexity of our OFO-based approach is at least as good as that of the efficient approach of [5], and better than the exponential complexity of [31]. Overall, our OFO-based approach leads to computational savings over the approach of [5] by a factor as large as $O(1/(\epsilon^2 \log(1/\epsilon)))$ arithmetic operations when the number updates of the solution is smaller than or equal to the number of updates of the noise realization, which is the case in many applications. For further comparisons and discussion, see Section 4.4. In addition, our framework is amenable to exploiting favorable structural properties of the constraint functions such as strong concavity, smoothness, etc., through which better convergence rates can be achieved.

4. Our framework is based on formulating the robust feasibility problem as a convex-nonconcave saddle point (SP) problem, and explicitly analyzing its structure. While convex-concave SP problems are well-studied in the literature, and many efficient first-order algorithms exist for these (see for example [38, 27, 28]), the convex-*nonconcave* SP problem is not as well-studied. To our knowledge, an explicit study of convex-nonconcave SP problems and their relation to RO has not been conducted previously; in this respect, the most closely related work [5] neither provides an explicit connection between robust feasibility and SP problems, nor analyzes their structure explicitly.

To demonstrate the application and effectiveness of our proposed framework, we walk through a detailed example on robust QPs. In particular, for robust QPs, we are able to leverage a recent convex QP-based reformulation of the classical trust region subproblem [25, 23] in order to avoid working with a nonconvex reformulation in a lifted space as in [5, Section 4.2] and relying on a probabilistic follow-the-perturbed-leader type algorithm [5, Section 3.2]. While using such nonconvex techniques will work within our framework, our convex reformulation allows us to work directly in the original space of the variables with a deterministic subgradient-based algorithm while still achieving asymptotically similar iteration complexity guarantees as [5]. Moreover, each iteration of our approach requires only first-order updates where the most expensive operation is the computation of a maximum eigenvector; thus our per-iteration cost is significantly less.

We also conduct a preliminary numerical study on the comparison of our approach with other iterative approaches [31] and [5] on robust QPs arising in portfolio optimization. Our results show

that when the problem size is small, the nominal solver approaches of [31] and [5] are more efficient. However, when problem size increases, replacing the nominal solvers with first-order updates using our framework allows us to achieve faster solution times. This highlights the benefits and potential of investigating first-order based approaches such as ours in iterative RO methods.

Outline

The rest of the paper is organized as follows. We begin with some notation and preliminaries in Section 2. We introduce our robust feasibility problem and robust feasibility/infeasibility certificates in Section 2.1, convex-concave SP problems in Section 2.2, and briefly summarize important online convex optimization (OCO) tools as well as a useful OFO algorithm in Section 2.3. We formulate the robust feasibility problem as a convex-nonconcave SP problem in Section 3; this formulation and certain bounds associated with its SP gap function form the basis of our general framework for solving robust feasibility problems. In Section 4 we specify an assortment of approaches obtained in our general framework by using different oracles. We examine our OFO-based approach in Section 4.1 by interpreting various terms in our framework in the context of OCO. In Section 4.2, we modify the pessimization oracle-based approach of [31] to obtain an efficient bound on the number of iterations required. In Section 4.3 we show how the nominal feasibility oracle-based approach of [5] fits within our framework. Finally, we discuss the convergence rates and accelerations attainable in our framework and compare our work with the existing approaches in Section 4.4. In Section 5 we illustrate our OFO-based approach through an example application on robust QPs. We provide in Section 6 a preliminary numerical study comparing our framework with other iterative approaches [31] and [5]. We close with a summary of our results and a few compelling further research directions in Section 7. In Appendix A we give an alternative formulation of the robust feasibility problem as a convex-concave SP problem in an extended space, and discuss its advantages and disadvantages over the convex-nonconcave SP formulation.

2 Notation and Preliminaries

Given $a \in \mathbb{R}$, $\text{sign}(a)$ denotes the sign of the number a. For a positive integer $n \in \mathbb{N}$, we let $[n] = \{1, \ldots, n\}$ and define $\Delta_n := \{x \in \mathbb{R}^n_+ : \sum_{i \in [n]} x_i = 1\}$ to be the standard simplex. Throughout the paper, the superscript, e.g., f^i, u^i, U^i, is used to attribute items to the i-th constraint, whereas the subscript, e.g., x_t, f_t, ϕ_t, is used to attribute items to the t-th iteration. Therefore, we sometimes use u^i, x_t, as well as u^i_t to denote vectors in $\mathbb{R}^n$. We use the notation $\{x_t\}_{t=1}^T$ to denote the collection of items $\{x_1, \ldots, x_T\}$. Given a vector $x \in \mathbb{R}^n$, we let $x^{(k)}$ denote its k-th coordinate for $k \in [n]$. One exception we make to this notation is that we always denote the convex combination weights $\theta \in \Delta_T$ with θ_t. For $x \in \mathbb{R}^n$ and $p \in [1, \infty]$, we use $\|x\|_p$ to denote the ℓ_p-norm of x defined as

$$\|x\|_p = \begin{cases} \left(\sum_{i \in [n]} |x^{(i)}|^p\right)^{1/p} & \text{if } p \in [1, \infty) \\ \max_{i \in [n]} |x^{(i)}| & \text{if } p = \infty \end{cases}.$$

Throughout this paper, we use Matlab notation to denote vectors and matrices, i.e., $[x; y]$ denotes the concatenation of two column vectors x, y. $\mathbb{S}^n$ denotes the space of $n \times n$ symmetric matrices, and we let $\mathbb{S}^n_+$ be the positive semidefinite cone in $\mathbb{S}^n$. We let I_n denote the identity matrix in $\mathbb{S}^n$. For a matrix $A \in \mathbb{S}^n$, $\lambda_{\max}(A)$, $\|A\|_{\text{Fro}}$, and $\|A\|_{\text{Spec}}$ correspond to its maximal eigenvalue, Frobenius norm, and spectral norm, respectively. Given a set V, we denote its closure by $\text{cl}(V)$. We abuse notation slightly by denoting $\nabla f(x)$ for both the gradient of function f at x if f is

differentiable and a subgradient of f at x, even if f is not differentiable. If f is of the form $f(x, u)$, then $\nabla_x f(x, u)$ denotes the subgradient of f at x while keeping the other variables fixed at u.

2.1 Robust Feasibility Problem

Consider a convex *deterministic* or *nominal* mathematical program

$$\min_x \left\{ f^0(x) : \ x \in X, \ f^i(x, u^i) \le 0, \ \forall i \in [m] \right\}, \tag{1}$$

where the domain $X \subset \mathbb{R}^n$ is closed and convex, the functions $f^0(x)$ and $f^i(x, u^i)$ for $i \in [m]$ are convex functions of x, and $u = (u^1, \ldots, u^m)$ is a fixed parameter vector. Without loss of generality we assume the objective function $f^0(x)$ does not have uncertainty. The *robust convex optimization problem* associated with (1) is

$$\mathrm{Opt} := \min_x \left\{ f^0(x) : \ x \in X, \ \sup_{u^i \in U^i} f^i(x, u^i) \le 0, \ \forall i \in [m] \right\}, \tag{2}$$

where $U^1, \ldots, U^m$ are the *uncertainty sets* given for the parameter u^i of constraint $i \in [m]$. Because we assume formulation (1) is convex, the overall optimization problem in (2) is convex.

In this paper, we work under the following mild regularity assumption:

Assumption 2.1. *The constraint functions $f^i(x, u^i)$ for all $i \in [m]$ are finite-valued on the domain $X \times U^i$, convex in x and concave in u^i. X, the domain for x, is closed and convex, and U^i, the domains for u^i, are closed and bounded.*

We take Assumption 2.1 as given for all our results and proofs. Without loss of generality, we assume that the uncertainty set has a Cartesian product form $U^1 \times \ldots \times U^m$, see e.g., [8]; we let $U = U^1 \times \ldots \times U^m$ and write $u = [u^1; \ldots ; u^m] \in U$. We do not further assume that the sets U^i are convex. However, for some algorithms we consider, convexity of U^i for $i \in [m]$ will be required.

A convex optimization problem can be solved by solving a polynomial number of associated feasibility problems in a standard way, via a binary search over its optimal value. In particular, let $[\underline{v}_0, \overline{v}_0]$ be an initial interval containing the optimal value of (2). At each iteration k of the binary search, we update the domain $X_k := X \cap \{x : \ f^0(x) \le v_k\}$ for some $v_k \in [\underline{v}_k, \overline{v}_k]$ and arrive at the following robust feasibility problem:

$$\text{find} \ \ x \in X_k \ \ \text{s.t.} \ \ \sup_{u^i \in U^i} f^i(x, u^i) \le 0 \ \ \forall i \in [m]. \tag{3}$$

Then based on the feasibility/infeasibility status of (3), we update our range $[\underline{v}_{k+1}, \overline{v}_{k+1}]$ and go to iteration $k + 1$. In this scheme, we are guaranteed to find a solution $x^* \in X$ whose objective value is within $\delta > 0$ of the optimum value of (2) in at most $\left\lfloor \log_2 \left(\frac{\overline{v}_0 - \underline{v}_0}{\delta} \right) \right\rfloor$ iterations. Therefore, one can equivalently study the complexity of solving robust feasibility problem (3) as opposed to (2). From now on, we focus on solving robust feasibility problem and assume that the constraint on the objective function $f^0(x)$ is already included in the domain X for simplicity in our notation.

Given functional constraints $f^i(x) \le 0$, $i \in [m]$, most convex optimization methods will declare infeasibility or return an approximate solution $x \in X$ such that $f^i(x) \le \epsilon$ for $i \in [m]$ for some tolerance level $\epsilon > 0$. Therefore, we consider the following *robust approximate feasibility problem*:

$$\begin{cases} \textit{Either}: \text{find} \ \ x \in X \ \ \text{s.t.} \ \ \sup_{u^i \in U^i} f^i(x, u^i) \le \epsilon \ \ \ \forall i \in [m]; \\ \textit{or}: \text{declare infeasibility}, \forall x \in X, \ \exists i \in [m] \ \ \text{s.t.} \ \ \sup_{u^i \in U^i} f^i(x, u^i) > 0. \end{cases} \tag{4}$$

7

We refer to any feasible solution x to (4), i.e., $x \in X$ such that $\sup_{u^i \in U^i} f^i(x, u^i) \leq \epsilon$ holds for all $i \in [m]$ as a *robust ϵ-feasibility certificate*. Similarly, any realization of the uncertain parameters $\bar{u} \in U$ such that there exists no $x \in X$ satisfying $f^i(x, \bar{u}^i) \leq 0$ for all $i \in [m]$ is referred to as a *robust infeasibility certificate*.

2.2 Saddle Point Problems

Saddle point (SP) problems play a vital role in our developments. In its most general form, a convex-concave SP problem is given by

$$\text{SV} = \inf_{x \in X} \sup_{y \in Y} \phi(x, y), \qquad (\mathcal{S})$$

where the function $\phi(x, y)$ is convex in x and concave in y and the domains X, Y are nonempty closed convex sets in Euclidean spaces $\mathbb{E}_x, \mathbb{E}_y$.

Any convex-concave SP problem $(\mathcal{S})$ gives rise to two convex optimization problems that are dual to each other:

$$
\begin{array}{rcll}
\text{Opt}(P) & = & \inf_{x \in X} [\overline{\phi}(x) := \sup_{y \in Y} \phi(x, y)] & (P) \\
\text{Opt}(D) & = & \sup_{y \in Y} [\underline{\phi}(y) := \inf_{x \in X} \phi(x, y)] & (D)
\end{array}
$$

with $\text{Opt}(P) = \text{Opt}(D) = \text{SV}$. It is well-known that the solutions to $(\mathcal{S})$ — the saddle points of ϕ on $X \times Y$ — are exactly the pairs $[x; y]$ formed by optimal solutions to the problems (P) and (D).

We quantify the accuracy of a candidate solution $[\bar{x}, \bar{y}]$ to SP problem $(\mathcal{S})$ with the *saddle point gap* given by

$$\epsilon_{\text{sad}}^{\phi}(\bar{x}, \bar{y}) := \overline{\phi}(\bar{x}) - \underline{\phi}(\bar{y}) = \underbrace{\left[\overline{\phi}(\bar{x}) - \text{Opt}(P)\right]}_{\geq 0} + \underbrace{\left[\text{Opt}(D) - \underline{\phi}(\bar{y})\right]}_{\geq 0}. \qquad (5)$$

Because convex-concave SP problems are simply convex optimization problems, they can in principle be solved by polynomial-time interior point methods (IPMs). However, the computational complexity of such methods depends heavily on the dimension of the problem. Thus, scalability of resulting algorithms becomes an issue in large-scale applications. As a result, for large-scale SP problems, one has to resort to first-order subgradient-type methods. On a positive note, there are many efficient first-order methods (FOMs) for convex-concave SP problems. These in particular include Nesterov's accelerated gradient descent algorithm [38] and Nemirovski's Mirror-Prox algorithm [34], both of which bound the saddle point gap at a rate of $\epsilon_{\text{sad}}^{\phi}(\bar{x}_T, \bar{u}_T) \leq O\left(\frac{1}{T}\right)$ where $\bar{x}_T, \bar{u}_T$ are solutions obtained after T iterations.

2.3 Online Convex Optimization Tools

Our efficient framework for RO employs tools from the online convex optimization domain. We now briefly outline these and refer to [17, 21, 42] for further details and applications of OCO.

OCO is used to capture decision making in dynamic environments. We are given a finite time horizon T, closed, bounded, and convex domain Z, and in each time period $t \in [T]$, a convex loss function $f_t : Z \to \mathbb{R}$ is revealed. At time periods $t \in [T]$ we must choose a decision $z_t \in Z$, and based on this we suffer a loss of $f_t(z_t)$ and receive some feedback typically in the form of first-order information on f_t. Our goal is to minimize the *weighted regret*

$$\sum_{t=1}^{T} \theta_t f_t(z_t) - \inf_{z \in Z} \sum_{t=1}^{T} \theta_t f_t(z), \qquad (6)$$

where $\theta \in \Delta_T$ is a vector of convex combination weights.[1]

Most OCO algorithms are closely related to offline iterative FOMs. In this paper, we will make use of the proximal setup of [27] to choose the sequence $\{z_t\}_{t=1}^T$ which ensures that the weighted regret (6) converges to 0 as $T \to \infty$. Thus, we make the following assumption on Z for the existence of a proximal setup.

Assumption 2.2. *Let $\mathbb{E}_z$ be the Euclidean space containing Z. There exists a norm $\|\cdot\|$ and its dual norm $\|\cdot\|_*$ on $\mathbb{E}_z$, a distance-generating function $\omega : Z \to \mathbb{R}$ which is 1-strongly convex with respect to $\|\cdot\|$ and leads to an easy-to-compute prox function $\mathrm{Prox}_z(\xi) := \arg\min_{w \in Z} \{\langle \xi, w \rangle + \omega(w) - \langle \omega'(z), w - z \rangle\}$ and set width $\Omega := \max_{z \in Z} \omega(z) - \min_{z \in Z} \omega(z)$ which is finite when Z is bounded.*

The proximal setup of Assumption 2.2 allows us to adjust to the geometry of domain Z. The standard basic domains satisfying Assumption 2.2 include simplex, Euclidean ball, and spectahedron; see [27, Section 1.7] for the standard proximal setups (i.e., Assumption 2.2) for these basic domains in terms of selection of $\|\cdot\|$ and resulting ω, Prox computation, and set width Ω.

Under Assumption 2.2 and various structural properties, the straightforward extension of the standard online mirror descent algorithm (see, e.g., [29]) from uniform weigths to weighted regret achieves the following convergence rate.

Theorem 2.1 ([29, Theorem 5]). *Suppose there exists $G \in (0, \infty)$ such that $\|\nabla f_t(z)\|_* \leq G$ for all $z \in Z$, $t \in [T]$. Define $\gamma = \sqrt{\frac{2\Omega}{G^2 \sum_{t=1}^T \theta_t^2}}$. Choose $z_1 = \arg\min_{z \in Z} \omega(z)$ and $z_{t+1} = \mathrm{Prox}_{z_t}(\gamma \theta_t \nabla f_t(z_t))$ for $t \in [T]$. Then*

$$\sum_{t=1}^T \theta_t f_t(z_t) - \inf_{z \in Z} \sum_{t=1}^T \theta_t f_t(z) \leq \sqrt{2\Omega G^2 \sum_{t=1}^T \theta_t^2}.$$

In particular, for uniform weights $\theta_t = 1/T$, the upper-bound becomes $O(1/\sqrt{T})$.

We refer to [29] for details of the proof. When $\omega(z) = z^\top z / 2$ and weights $\theta_t = 1/T$ for $t \in [T]$, the update rule $z_{t+1} = \mathrm{Prox}_{z_t}(\gamma \nabla f_t(z_t))$ becomes simply gradient descent, and Theorem 2.1 reduces to the standard bound of online gradient descent from [47].

3   General Framework for Robust Feasibility Problems

In this section, we build a general framework to solve the robust feasibility problem (4) by working with its natural saddle point formulation.

Given constraint functions $f^i(x, u^i)$, $i \in [m]$, let us define $\Phi(x, u) := \max_{i \in [m]} f^i(x, u^i)$. Then $\Phi(x, u)$ is a convex function of x, but not necessarily concave in u. In addition, with this definition of $\Phi(\cdot)$, the robust approximate feasibility problem (4) is equivalent to simply verifying

$$\text{either} \quad \inf_{x \in X} \sup_{u \in U} \Phi(x, u) = \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) \leq \epsilon \quad \text{or} \quad \inf_{x \in X} \sup_{u \in U} \Phi(x, u) > 0, \tag{7}$$

---

[1]Note that in the OCO literature, regret is usually defined with uniform weights $\theta_t = 1/T$. Nonuniform weights introduce flexibility to our framework by allowing selection of specific customization of OCO algorithms for exploiting structural properties of the constraint functions f^i to achieve better convergence rates. A prime example for this is when the functions are strongly convex.

which is nothing but solving a specific SP problem and checking its value. Analogous to the convex-concave SP gap (5), for a given solution $[\bar{x}, \bar{u}]$, we define the SP gap of problem (7) as

$$\epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u}) := \overline{\Phi}(\bar{x}) - \underline{\Phi}(\bar{u}) = \sup_{u \in U} \Phi(\bar{x}, u) - \inf_{x \in X} \Phi(x, \bar{u}).$$

In general, solving a convex-nonconcave SP problem of form (7), i.e., finding a solution $[\bar{x}, \bar{u}]$ such that $\epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u}) \leq \epsilon$, can be difficult. That said, a bound on the SP gap $\epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u})$ along with the value of $\Phi(\bar{x}, \bar{u})$ leads to robust feasibility certificates for (7) as follows.

Theorem 3.1. *Let $\Psi : X \times U \to \mathbb{R}$ be a given function associated with a SP (not necessarily admitting a convex-concave structure). Suppose we have $\bar{x} \in X$, $\bar{u} \in U$, and $\tau \in (0, 1)$ such that $\epsilon_{\text{sad}}^{\Psi}(\bar{x}, \bar{u}) \leq \tau\epsilon$. Then if $\Psi(\bar{x}, \bar{u}) \leq (1 - \tau)\epsilon$, we have $\sup_{u \in U} \Psi(\bar{x}, u) \leq \epsilon$. Moreover, if $\Psi(\bar{x}, \bar{u}) > (1 - \tau)\epsilon$ and $\tau \leq \frac{1}{2}$, we have $\inf_{x \in X} \Psi(x, \bar{u}) > 0$.*

Proof. Suppose $\Psi(\bar{x}, \bar{v}) \leq (1 - \tau)\epsilon$. Because $\epsilon_{\text{sad}}^{\Psi}(\bar{x}, \bar{v}) = \sup_{u \in U} \Psi(\bar{x}, u) - \inf_{x \in X} \Psi(x, \bar{u}) \leq \tau\epsilon$, we have $\sup_{u \in U} \Psi(\bar{x}, u) \leq \inf_{x \in X} \Psi(x, \bar{u}) + \tau\epsilon \leq \Psi(\bar{x}, \bar{u}) + \tau\epsilon \leq \epsilon$. On the other hand, when $\Psi(\bar{x}, \bar{u}) > (1 - \tau)\epsilon$, we have $(1 - \tau)\epsilon < \Psi(\bar{x}, \bar{u}) \leq \sup_{u \in U} \Psi(\bar{x}, u) \leq \inf_{x \in X} \Psi(x, \bar{u}) + \tau\epsilon$, which implies $\inf_{x \in X} \sup_{u \in U} \Psi(x, u) \geq \inf_{x \in X} \Psi(x, \bar{u}) > (1 - 2\tau)\epsilon \geq 0$ when $\tau \leq \frac{1}{2}$. $\qquad\square$

Remark 3.1. When $m = 1$, $\Phi(x, u) = f^1(x, u^1)$, and it is thus convex in x and concave in u due to Assumption 2.1. Therefore, in the case of a single robust constraint, i.e., $m = 1$, under Assumption 2.1 and assuming $U = U^1$ is a closed convex set, the optimization problem in (7) reduces to a standard convex-concave SP problem. ∎

While it is not very common, a few robust convex optimization problems come with a single robust constraint and convex uncertainty set U; see for example [2] for a robust version of a SVM problem with one constraint. In such cases, based on Remark 3.1, the resulting convex-concave SP problems can directly be solved via efficient FOMs. On the other hand, in the presence of multiple constraints, the function $\Phi(x, u)$ is not concave in $u = [u^1; \ldots; u^m]$ even under Assumption 2.1. Nevertheless, when $m > 1$, it is still possible to have a convex-concave SP reformulation of the optimization problem in (7) in an extended space via perspective transformations, which we present in Appendix A. While this reformulation has the benefit of reducing the robust feasibility problem to a well-known and well-studied problem, it destroys the simplicity of the original domains and constraint functions and hence comes with some challenges. Therefore, we develop a framework where we work directly with the convex-nonconcave SP formulation in (7) in the space of original variables. Moreover, because we work in the original space of variables, we simply utilize the first-order information on the original constraint functions f^i and original domains X and U^i. This direct approach in particular allows us to take greater advantage of the structure of the original formulation such as the availability of efficient projection (prox) computations over domains X, U^i, and/or better parameters for smoothness, Lipschitz continuity, etc., of the functions f^i.

Because $\Phi(x, u)$ is not concave in u, we cannot bound the SP gap $\epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u})$ by using traditional FOMs designed for solving convex-concave SP problems. However, we next show that by just *partially* upper bounding $\epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u})$, we can derive a general iterative framework to obtain robust feasibility/infeasibility certificates. We describe the further specifics of this framework in Section 4.

Henceforth we will no longer use the shorthand notation $\Phi(x, u) = \max_{i \in [m]} f^i(x, u^i)$, but we will denote the SP gap $\epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u})$ as

$$\epsilon(\bar{x}, \bar{u}) := \epsilon_{\text{sad}}^{\Phi}(\bar{x}, \bar{u}) = \max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) - \inf_{x \in X} \max_{i \in [m]} f^i(x, \bar{u}^i). \tag{8}$$

The robust feasibility certificate result from Theorem 3.1 indicates the importance of bounding the SP gap $\epsilon(\bar{x}, \bar{u})$. Often, FOMs achieve this by iteratively generating points $x_t \in X$, $u_t \in U$ for $t \in [T]$ and tracking the points $\bar{x}$ and $\bar{u}$ obtained from a convex combination of $\{x_t, u_t\}_{t=1}^T$. In order to simplify our notation, given convex combination weights $\theta \in \Delta_T$ and points $\{x_t, u_t\}_{t=1}^T$, we let

$$\bar{x}_T := \sum_{t=1}^T \theta_t x_t \quad \text{and} \quad \bar{u}_T := \sum_{t=1}^T \theta_t u_t.$$

We now present an upper bound on $\epsilon(\bar{x}_T, \bar{u}_T)$ that follows naturally from the convex-concave structure of functions f^i. To this end, given a set of vectors $y_t \in \Delta_m$ for $t \in [T]$, we also define

$$\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) := \max_{i \in [m]} \left\{ \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \right\}, \quad \text{and}$$

$$\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) := \max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) - \inf_{x \in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i),$$

together with

$$\widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) := \inf_{x \in X} \sum_{t=1}^T \theta_t \sum_{i=1}^m y_t^{(i)} f^i(x, u_t^i) - \inf_{x \in X} \max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x, u_t^i).$$

Our next result relates these quantities to the value of the SP gap function $\epsilon(\bar{x}_T, \bar{u}_T)$.

Proposition 3.1. *Let $x_t \in X$ and $u_t \in U$ for $t \in [T]$ be given a set of vectors. Then for any set of vectors $y_t \in \Delta_m$ for $t \in [T]$ and any $\theta \in \Delta_T$, we have*

$$\epsilon\left(\sum_{t=1}^T \theta_t x_t, \sum_{t=1}^T \theta_t u_t\right) \leq \epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) + \epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) + \widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T). \quad (9)$$

Proof. Given $y_t \in \Delta_m$ for $t \in [T]$ and $\theta \in \Delta_T$, let us define $\bar{x} := \sum_{t=1}^T \theta_t x_t$ and $\bar{u} := \sum_{t=1}^T \theta_t u_t$. We first partition $\epsilon(\bar{x}, \bar{u})$ as $\epsilon(\bar{x}, \bar{u}) = \bar{\epsilon}(\bar{x}) + \underline{\epsilon}(\bar{u})$ where

$$\bar{\epsilon}(\bar{x}) := \max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) - \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i),$$

$$\underline{\epsilon}(\bar{u}) := \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) - \inf_{x \in X} \max_{i \in [m]} f^i(x, \bar{u}^i),$$

and then derive upper bounds on $\bar{\epsilon}(\bar{x})$ and $\underline{\epsilon}(\bar{u})$.

We start with bounding $\bar{\epsilon}(\bar{x})$. Because the functions $f^i(x, u^i)$ are convex in x for all i and $\theta \in \Delta_T$, we have $\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \max_{i \in [m]} \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i)$. Therefore,

$$\bar{\epsilon}(\bar{x}) = \max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) - \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i)$$

$$\leq \max_{i \in [m]} \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) + \max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) - \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i)$$

$$\leq \max_{i \in [m]} \left\{ \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \right\} + \max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) - \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i),$$

$$(10)$$

11

where the last inequality follows since $\max_{i\in[m]}\{\alpha_i - \beta_i\} \geq \max_{i\in[m]}\alpha_i - \max_{i\in[m]}\beta_i$ for any sequence of numbers α_i, β_i, $i \in [m]$.

Note that $\inf_{x\in X}\max_{i\in[m]} f^i(x, u^i) \geq \inf_{x\in X}\max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x, u^i)$ because under Assumption 2.1 the functions $f^i(x, u^i)$ are concave in u^i for all i. Thus, we arrive at

$$\underline{\epsilon}(\bar{u}) = \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x, u^i) - \inf_{x\in X}\max_{i\in[m]} f^i(x, \bar{u}^i)$$

$$\leq \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x, u^i) - \inf_{x\in X}\sum_{t=1}^{T}\theta_t\sum_{i=1}^{m} y_t^{(i)} f^i(x, u_t^i) + \inf_{x\in X}\sum_{t=1}^{T}\theta_t\sum_{i=1}^{m} y_t^{(i)} f^i(x, u_t^i)$$

$$- \inf_{x\in X}\max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x, u_t^i)$$

$$= \inf_{x\in X}\max_{i\in[m]}\sup_{u^i\in U^i} f^i(x, u^i) - \inf_{x\in X}\sum_{t=1}^{T}\theta_t\sum_{i=1}^{m} y_t^{(i)} f^i(x, u_t^i) + \widehat{\epsilon}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}). \tag{11}$$

Then by summing (10) and (11) and rearranging the terms, we deduce the result. $\qquad\square$

We are now ready to state our main result. This is analogous to Theorem 3.1 except that we do not need to bound all three terms in (9), but instead it suffices to guarantee that

$$\epsilon^{\circ}(\{x_t, u_t, \theta_t\}_{t=1}^{T}) + \epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}) \leq \epsilon.$$

We show that when the above condition holds, based on the value of $\max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i)$ we can then obtain either a robust ϵ-feasible solution, or an infeasibility certificate.

Theorem 3.2. *Suppose we have sequences $\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}$ with $x_t \in X$, $u_t \in U$, $y_t \in \Delta_m$ for all $t \in [T]$, $\theta \in \Delta_T$. Let $\tau \in (0,1)$. If $\epsilon^{\circ}(\{x_t, u_t, \theta_t\}_{t=1}^{T}) \leq \tau\epsilon$ and $\max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$, then the solution $\bar{x}_T := \sum_{t=1}^{T}\theta_t x_t$ is ϵ-feasible with respect to (4). If $\epsilon^{\bullet}(\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}) \leq (1-\tau)\epsilon$ and $\max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$, then (4) is infeasible.*

Proof. First suppose there exists a $\tau \in (0,1)$ and corresponding vectors $\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}$ such that $\epsilon^{\circ}(\{x_t, u_t, \theta_t\}_{t=1}^{T}) \leq \tau\epsilon$ and $\max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$ holds as well. Note that

$$\tau\epsilon \geq \epsilon^{\circ}(\{x_t, u_t, \theta_t\}_{t=1}^{T}) = \max_{i\in[m]}\left\{\sup_{u^i\in U^i}\sum_{t=1}^{T}\theta_t f^i(x_t, u^i) - \sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i)\right\}$$

$$\geq \max_{i\in[m]}\sup_{u^i\in U^i}\sum_{t=1}^{T}\theta_t f^i(x_t, u^i) - \max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i), \tag{12}$$

where the last inequality follows since $\max_{i\in[m]}\{\alpha_i - \beta_i\} \geq \max_{i\in[m]}\alpha_i - \max_{i\in[m]}\beta_i$ for any sequence of numbers α_i, β_i, $i \in [m]$. Then $\bar{x}_T$ is an ϵ-feasible solution for (4) because

$$\max_{i\in[m]}\sup_{u^i\in U^i} f^i(\bar{x}_T, u^i) = \max_{i\in[m]}\sup_{u^i\in U^i} f^i\left(\sum_{t=1}^{T}\theta_t x_t, u^i\right)$$

$$\leq \max_{i\in[m]}\sup_{u^i\in U^i}\sum_{t=1}^{T}\theta_t f^i(x_t, u^i) \leq \tau\epsilon + \max_{i\in[m]}\sum_{t=1}^{T}\theta_t f^i(x_t, u_t^i) \leq \epsilon,$$

where the first inequality follows from the convexity of the functions f^i and the fact that $\theta \in \Delta_T$, the second inequality from (12), and the last inequality holds since $\max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$.

On the other hand, suppose $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}) \leq (1-\tau)\epsilon$ and $\max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$. Note that

$$
\inf_{x \in X} \sum_{t=1}^{T} \theta_t \sum_{i=1}^{m} y_t^{(i)} f^i(x, u_t^i) \leq \inf_{x \in X} \sum_{t=1}^{T} \theta_t \max_{i \in [m]} f^i(x, u_t^i)
$$

$$
\leq \inf_{x \in X} \sum_{t=1}^{T} \theta_t \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) = \inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i), \qquad (13)
$$

where the first inequality follows since $y_t \in \Delta_m$ for all $t \in [T]$, the second inequality holds because $f^i(x, u_t^i) \leq \sup_{u^i \in U^i} f^i(x, u^i)$ for all $i \in [m]$ and $y_t^{(i)} \geq 0$ for $i \in [m]$, $t \in [T]$, and the last equation follows from $\theta \in \Delta_T$. Then using the bound

$$
(1-\tau)\epsilon \geq \epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T}) = \max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t^i) - \inf_{x \in X} \sum_{t=1}^{T} \theta_t \sum_{i=1}^{m} y_t^{(i)} f^i(x, u_t^i), \qquad (14)
$$

we arrive at

$$
\inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) \geq \inf_{x \in X} \sum_{t=1}^{T} \theta_t \sum_{i=1}^{m} y_t^{(i)} f^i(x, u_t^i) \geq \max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t^i) - (1-\tau)\epsilon > 0,
$$

where the first inequality follows from inequality (13), the second inequality from (14) and the last inequality holds because $\max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$. This implies (4) is infeasible. $\qquad \square$

In Section 4.1 we will show that $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^{T})$ can be interpreted as a weighted regret term (6). On the other hand, the term $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^{T})$ has no such direct interpretation. In order to upper-bound it by a weighted regret term, we need the following result.

Corollary 3.1. *Given sequences $\{x_t, u_t, \theta_t\}_{t=1}^{T}$ with $x_t \in X$, $u_t \in U$, for all $t \in [T]$, $\theta \in \Delta_T$, there is an appropriate choice of sequence $\{\bar{y}_t\}_{t=1}^{T}$ where $\bar{y}_t \in \Delta_m$ for all $t \in [T]$, such that $\epsilon^\bullet(\{x_t, u_t, \bar{y}_t, \theta_t\}_{t=1}^{T})$ is upper-bounded by*

$$
\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^{T}) := \sum_{t=1}^{T} \theta_t \max_{i \in [m]} f^i(x_t, u_t) - \inf_{x \in X} \sum_{t=1}^{T} \theta_t \max_{i \in [m]} f^i(x, u_t). \qquad (15)
$$

Thus, if $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^{T}) \leq (1-\tau)\epsilon$ and $\max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t^i) > (1-\tau)\epsilon$, then (4) is infeasible.

Proof. Given $\{u_t\}_{t=1}^{T}$, let $x^* \in \arg\min_{x \in X} \sum_{t=1}^{T} \theta_t \max_{i \in [m]} f^i(x, u_t^i)$ and for all $t \in [T]$ define $\bar{y}_t \in \mathbb{R}^m$ to be the i-th unit vector where i is the smallest index satisfying $i \in \arg\max_{i' \in [m]} f^i(x^*, u_t^{i'})$. Then $\inf_{x \in X} \sum_{t=1}^{T} \theta_t \sum_{i=1}^{m} \bar{y}_t^{(i)} f^i(x, u_t^i) = \inf_{x \in X} \sum_{t=1}^{T} \theta_t \max_{i \in [m]} f^i(x, u_t^i)$, and the bound follows from $\max_{i \in [m]} \sum_{t=1}^{T} \theta_t f^i(x_t, u_t) \leq \sum_{t=1}^{T} \theta_t \max_{i \in [m]} f^i(x_t, u_t)$. We deduce the last result from Theorem 3.2. $\qquad \square$

The following corollary demonstrates how we can choose τ in Theorem 3.2.

13

Corollary 3.2. *Suppose* $\{x_t, u_t, y_t, \theta_t\}_{t=1}^T$ *with* $x_t \in X$, $u_t \in U$, $y_t \in \Delta_m$ *for all* $t \in [T]$, *and* $\theta \in \Delta_T$ *is such that there exists* $\kappa^\circ, \kappa^\bullet \in (0,1)$ *satisfying* $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \le \epsilon \kappa^\circ$ *and* $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \le \epsilon \kappa^\bullet$ *with* $\kappa^\circ + \kappa^\bullet \le 1$. *Let* $\tau \in [\kappa^\circ, 1 - \kappa^\bullet]$. *Whenever* $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le (1-\tau)\epsilon$ *as well, the solution* $\bar{x}_T := \sum_{t=1}^T \theta_t x_t$ *is* ϵ-*feasible with respect to* (4). *Also, whenever* $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) > (1 - \tau)\epsilon$, *then* (4) *is infeasible.*

Proof. Note that $\tau \in (0,1)$ follows from its definition, $\kappa^\circ, \kappa^\bullet \ge 0$, and $\kappa^\circ + \kappa^\bullet \le 1$. Furthermore, the interval $[\kappa^\circ, 1 - \kappa^\bullet]$ is well-defined since $\kappa^\circ \le 1 - \kappa^\bullet$ always holds. Moreover, $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \le \epsilon \kappa^\circ \le \epsilon \tau$ and $\epsilon^\bullet(\{x_t, u_t, y_t, \theta_t\}_{t=1}^T) \le \epsilon \kappa^\bullet \le \epsilon(1-\tau)$ holds from the definition of τ. The result now follows from Theorem 3.2. $\qquad\square$

Theorem 3.2 and Corollary 3.1 points to our general iterative framework for finding robust feasibility/infeasibility certificates of (4): generate sequences $\{x_t, u_t\}_{t=1}^T$ iteratively to bound $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ and $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$, and then evaluate the term $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i)$. We provide a description of our framework in Algorithm 1. We assume that we have access to weights $\{\theta_t\}_{t=1}^T$ and update algorithms $\mathcal{A}_i$ and $\mathcal{A}_x$ for choosing $u_t^i \in U^i$ and $x_t \in X$ based on past observations $\{x_s, u_s\}_{s=1}^{t-1}$. We denote the updates by

$$u_t^i = \mathcal{A}_i(\{x_s, u_s\}_{s=1}^{t-1}) \in U^i \; \forall i \in [m], \quad x_t = \mathcal{A}_x(\{x_s, u_s\}_{s=1}^{t-1}) \in X,$$

and initializations $u_1^i = \mathcal{A}_i(\{\}) \; \forall i \in [m]$, $x_1 = \mathcal{A}_x(\{\})$. Moreover, we assume that these algorithms enjoy the following convergence guarantees: for any sequence $\{x_t\}_{t=1}^T$, let $u_t^i = \mathcal{A}_i(\{x_s, u_s\}_{s=1}^{t-1}) \; \forall i \in [m]$, then

$$\sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \le \mathcal{R}_i(T); \tag{16}$$

for any sequence $\{u_s\}_{s=1}^T$, let $x_t = \mathcal{A}_x(\{x_s, u_s\}_{s=1}^{t-1})$, then

$$\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) = \sum_{t=1}^T \theta_t \max_{i \in [m]} f^i(x_t, u_t^i) - \inf_{x \in X} \sum_{t=1}^T \theta_t \max_{i \in [m]} f^i(x, u_t^i) \le \mathcal{R}_x(T). \tag{17}$$

Explicit examples of $\mathcal{A}_i, \mathcal{A}_x$ and their bounds $\mathcal{R}_i, \mathcal{R}_x$ will be discussed in Section 4. Generally, we desire that the error bounds $\mathcal{R}_i(T), \mathcal{R}_x(T) \to 0$ as $T \to \infty$, which can be achieved by using online mirror descent as in Theorem 2.1. That said, our OFO-based approach in Algorithm 1 is quite flexible in terms of the selection of OFO algorithms $\mathcal{A}_i, \mathcal{A}_x$, and is certainly not restricted to only using online mirror descent.

Remark 3.2. Notice that in Algorithm 1 we generate u_t *before* generating x_t. Thus, nothing stops us from choosing x_t based on the knowledge of u_t, or vice versa. Indeed, the pessimization oracle approach of [31] and the nominal feasibility oracle approach of [5] fit within our framework if we rewrite Algorithm 1 to reflect this, and it is a trivial matter to do so. However, a conflict may arise if we encounter a situation where generating x_t requires knowledge of u_t, and generating u_t also requires knowledge of x_t. Thus, when selecting the update algorithms $\mathcal{A}_i, \mathcal{A}_x$, care must be taken to avoid such situations. Our suggested OFO approach in Section 4.1 utilizes Theorem 2.1. Moreover, Theorem 2.1 generates the current decision in a *non-anticipatory manner* based on only the knowledge of f_{t-1}, and not of f_t. That is, it ensures that we will only use u_{t-1} to generate x_t, and similarly we only use x_{t-1} to generate u_t, thus no conflicts will arise. $\qquad\blacksquare$

Algorithm 1 OFO-based approximate robust feasibility solver.

---

input: update algorithms $\mathcal{A}_i$, $i \in [m]$, $\mathcal{A}_x$, tolerance level $\epsilon > 0$, sufficiently large $T = T(\epsilon)$ such that $\max_{i \in [m]} \mathcal{R}_i(T) + \mathcal{R}_x(T) \leq \epsilon$, and convex combination weights $\theta_1, \ldots, \theta_T > 0$.
output: either $\bar{x} \in X$ such that $\sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon$ for all $i \in [m]$, or an infeasibility certificate for (4).
initialize $u_1^i = \mathcal{A}_i(\{\})$ for $i \in [m]$ and $x_1 = \mathcal{A}_x(\{\})$.
for $t = 2, \ldots, T$ **do**
   for $i = 1, \ldots, m$ **do**
      compute $u_t^i = \mathcal{A}_i(\{x_s, u_s\}_{s=1}^{t-1}) \in U^i$.
   end for
   compute $x_t = \mathcal{A}_x(\{x_s, u_s\}_{s=1}^{t-1}) \in X$.
   obtain upper bounds $\max_{i \in [m]} \mathcal{R}_i(t) \geq \epsilon^\circ(\{x_s, u_s, \theta_s\}_{s=1}^t)$ and $\mathcal{R}_x(t) \geq \epsilon^\bullet(\{x_s, u_s, \theta_s\}_{s=1}^t)$.
   compute $\kappa_t^\circ = \max_{i \in [m]} \mathcal{R}_i(t)/\epsilon$, $\kappa_t^\bullet = \mathcal{R}_x(t)/\epsilon$.
   if $\kappa_t^\circ + \kappa_t^\bullet \leq 1$ **then**
      set $\vartheta_t := \max_{i \in [m]} \sum_{s=1}^t \theta_s f^i(x_s, u_s^i)$ and $\tau_t := 1 - \kappa_t^\bullet$.
      if $\vartheta_t > (1 - \tau_t)\epsilon$ **then return** 'infeasible'.
      if $\vartheta_t \leq (1 - \tau_t)\epsilon$ **then return** $\bar{x}_t = \frac{1}{t} \sum_{s=1}^t x_s$ as a robust ϵ-feasible solution to (4).
   end if
end for

Remark 3.3. Note that Algorithm 1 chooses $\tau_t = 1 - \kappa_t^\bullet$, whereas Corollary 3.2 allows us to choose from a range $\tau_t \in [\kappa_t^\circ, 1 - \kappa_t^\bullet]$. This is because it is theoretically possible for (4) to simultaneously be infeasible and robust ϵ-feasible, but in practice we would like to discover infeasibility of (4) rather than an approximately feasible solution. Then the best value for $\tau_t \in [\kappa_t^\circ, 1 - \kappa_t^\bullet]$ in detecting infeasibility of (4) is given by $\tau_t = 1 - \kappa_t^\bullet$. ∎

In the next section, we describe some approaches to implement Algorithm 1 in practice by providing explicit examples of $\mathcal{A}_i$ and $\mathcal{A}_x$.

4 Customizations of the General Framework

In this section, we examine how to generate the sequences $\{x_t, u_t\}_{t=1}^T$ in practice. In Section 4.1, we first interpret the terms $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ and $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ from Section 3 as weighted regret terms, which gives rise to our OFO-based approach. In Section 4.2, we modify the pessimization oracle-based approach of [31] to solving (4) within our framework. In Section 4.3, we examine the nominal feasibility oracle-based approach of [5] within the context of our general framework. Finally, in Section 4.4, we summarize and compare the convergence rates achievable via various customizations of our framework using these different approaches.

4.1 The OFO-based Approach

Let us first consider $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$. For any $i \in [m]$, given x_t, we define the function $f_t^i : U^i \to \mathbb{R}$ as $f_t^i(u^i) = -f^i(x_t, u^i)$. Then the function $f_t^i(u^i)$ is convex in u^i under Assumption 2.1, and the subterm of $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ given by

$$\sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \tag{18}$$

is the weighted regret (6) corresponding to the sequence of functions $\{f_t^i\}_{t=1}^T$. When the uncertainty sets U^i, $i \in [m]$ admit proximal setups as in Assumption 2.2, Theorem 2.1 from Section 2.3 gives

an efficient OFO algorithm for choosing $\{u_t^i\}_{i=1}^m$ to bound the regret subterms (18) with $O(1/\sqrt{T})$. Therefore, by using the online mirror descent algorithm of Theorem 2.1 as $\mathcal{A}_i$ in the computation of our u_t^i, we guarantee that

$$\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) = \max_{i \in [m]} \left\{ \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \right\} \leq \max_{i \in [m]} \mathcal{R}_i(T),$$

where $\mathcal{R}_i(T) = O(1/\sqrt{T})$ with uniform weights $\theta_t = 1/T$.

On the other hand, given $u_t^i \in U^i$ for $i \in [m]$, let us define $\varphi_t(x) := \max_{i \in [m]} f^i(x, u_t^i)$. Then $\varphi_t(x)$ is convex in x over X since the functions f^i are convex in x by Assumption 2.1. We can then rewrite $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ as

$$\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) = \sum_{t=1}^T \theta_t \max_{i \in [m]} f^i(x_t, u_t^i) - \inf_{x \in X} \sum_{t=1}^T \theta_t \max_{i \in [m]} f^i(x, u_t^i)$$

$$= \sum_{t=1}^T \theta_t \varphi_t(x_t) - \inf_{x \in X} \sum_{t=1}^T \theta_t \varphi_t(x). \tag{19}$$

Then $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ is also a weighted regret term (6) corresponding to the sequence of functions $\{\varphi_t\}_{t=1}^T$. When the domain X admits a proximal setup as in Assumption 2.2, Theorem 2.1 again gives an efficient OFO algorithm for choosing x_t to bound (19). Once again, we may choose $\mathcal{A}_x$ to be the online mirror descent, and get $\mathcal{R}_x(T) = O(1/\sqrt{T})$ with uniform weights $\theta_t = 1/T$.

Algorithm 1 can now be employed, provided we choose $T = \Omega(1/\epsilon^2)$, to solve the robust feasibility problem (4). Since the online mirror descent algorithm of Theorem 3.2 only uses first-order information, we can solve the robust feasibility problem (4) while avoiding reliance on a pessimization oracle for u as in [31] or a nominal feasibility oracle for x as in [5].

4.2 The Pessimization Oracle-Based Approach

Mutapcic and Boyd [31] generate solutions $x_t \in X$ at each iteration t by solving an extended nominal problem

$$\min_{x \in X} \left\{ f^0(x) : \ f^i(x, u^i) \leq 0, \ \forall u^i \in \hat{U}_{t-1}^i, \ i \in [m] \right\}, \tag{20}$$

where $\hat{U}_{t-1}^i \subset U^i$ are finite approximate uncertainty sets based on past noise realizations $\{u_{t'}^i\}_{i=1}^m$ for $t' \in [t-1]$. New noises u_t are then generated by calling the *pessimization oracles* on the current solution x_t. More precisely, given $x_t \in X$, the pessimization oracles solve $\sup_{u^i \in U^i} f^i(x_t, u^i)$ and return

$$u_t^i \in U^i \quad \text{s.t.} \quad f^i(x_t, u_t^i) \geq \sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon. \tag{21}$$

In terms of our framework of Algorithm 1, the update policy of generating new noises u_t in this approach of [31] corresponds to selecting the algorithms $\mathcal{A}_x$ to be a extended nominal solver for (20) and the algorithms $\mathcal{A}_i$ to be pessimization oracles that solve (21). Note that computing u_t^i requires knowledge of x_t (see Remark 3.2), and consequently the bound for the regret term (16) is $\mathcal{R}_i(T) \leq \tau\epsilon$ for any T. We show this in the proof of Theorem 4.1. If for all $i \in [m]$ we have $f^i(x_t, u_t^i) \leq (1 - \tau)\epsilon$, then we terminate and declare x_t is a robust ϵ-feasible and optimal solution; otherwise, we append $\hat{U}_t^i = \hat{U}_{t-1}^i \cup \{u_t^i\}$ and re-solve (20) with the new approximate sets $\hat{U}_t^i$. It is

16

shown in [31, Section 5.2] that the number of iterations T needed before termination with a robust ϵ-feasible solution x_T is upper bounded by $(1 + O(1/\epsilon))^n$ where n is the dimension of x.

Suppose now that we are interested in robust feasibility (4). [31, Section 5.3] discusses a number of variations for generating x_t by modifying (20). In contrast, we propose the following modification: instead of solving (20), generate $\{x_t\}_{t=1}^T$ via a non-anticipatory algorithm $\mathcal{A}_x$ (see Remark 3.2) to bound $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \mathcal{R}_x(T) = (1 - \tau)\epsilon$. We call our modification *FO-based pessimization*. Then the FO-based pessimization approach fits within our framework as a special case.

Theorem 4.1. *Let $\tau \in (0, 1)$. Suppose $\{x_t\}_{t=1}^T$ are generated iteratively to guarantee that $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq (1-\tau)\epsilon$ for any sequence $\{u_t\}_{t=1}^T$. Suppose u_t^i are generated by pessimization oracles (21) for $i \in [m]$. If there exists $t \in [T]$ such that for all $i \in [m]$ we have $f^i(x_t, u_t^i) \leq (1 - \tau)\epsilon$, then x_t is a robust ϵ-feasible solution to (4). If $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \leq (1 - \tau)\epsilon$, then $\bar{x}_T = \sum_{t=1}^T \theta_t x_t$ is a robust ϵ-feasible solution to (4). If $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) > (1 - \tau)\epsilon$, then we certify that (4) is robust infeasible.*

Proof. It is clear that if there exists $t \in [T]$ such that for all $i \in [m]$ we have $f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$, then x_t is a robust ϵ-feasible solution to (4). Furthermore, the fact that $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) > (1 - \tau)\epsilon$ implies robust infeasibility of (4) follows from our assumption that $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq (1-\tau)\epsilon$ and Corollary 3.1. To show that $\max_{i \in [m]} \sum_{t=1}^T f^i(x_t, u_t^i) \leq (1 - \tau)\epsilon$ implies that $\bar{x}_T$ is robust ϵ-feasible, we only need to show that $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \tau\epsilon$. Observe that by our definition of u_t^i in (21), we have $f^i(x_t, u_t^i) \geq \sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon$, hence the regret terms in $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ satisfy

$$\sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \leq \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t \left(\sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon \right)$$

$$= \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t \sup_{u^i \in U^i} f^i(x_t, u^i) + \tau\epsilon$$

$$\leq \tau\epsilon.$$

Then

$$\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) = \max_{i \in [m]} \left\{ \sup_{u^i \in U^i} \sum_{t=1}^T \theta_t f^i(x_t, u^i) - \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \right\} \leq \tau\epsilon,$$

and the result follows from Corollary 3.1. □

Theorem 4.1 can only be used to certify robust feasibility/infeasibility. Hence, to find a robust ϵ-optimal solution in FO-based pessimization approach, we must perform a binary search and solve at most $O(\log(1/\epsilon))$ instances of robust feasibility problems. Despite this, in Section 4.4, we discuss how FO-based pessimization approach which uses OFO algorithms to generate $\{x_t\}_{t=1}^T$ to bound $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ results in much better complexity guarantees than using an extended nominal feasibility solver (20) as proposed by [31], even when taking into account the additional $O(\log(1/\epsilon))$ factor.

Remark 4.1. In the pessimization oracle-based approach, the noises u_t need to be generated with knowledge of x_t, because it is not possible to guarantee $f^i(x_t, u_t^i) \geq \sup_{u^i \in U^i} f^i(x_t, u^i) - \tau\epsilon$ if the vectors u_t^i were chosen with only the knowledge of $x_1, \ldots, x_{t-1}$. ∎

4.3 The Nominal Feasibility Oracle-Based Approach

The nominal feasibility oracle-based approach of Ben-Tal et al. [5] suggest using OFO algorithms to choose a sequence $\{u_t\}_{t=1}^T$ that guarantees $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ is small, in a non-anticipatory fashion, for *any* sequence $\{x_t\}_{t=1}^T$. In this aspect, it essentially matches with our OFO-based approach outlined in Section 4.1 i.e., the choice of $\mathcal{A}_i$ is essentially the same. The key differentiating point between our OFO-based approach and that of [5] lies in which algorithm is chosen for $\mathcal{A}_x$. At step t, [5] utilizes a *nominal feasibility oracle*. That is, given parameters u_t, they call a powerful, and potentially expensive, nominal feasibility oracle that solves the following feasibility problem to ϵ-accuracy

$$\begin{cases} \textit{Either}: \text{ find } x \in X \quad \text{s.t.} \quad f^i(x, u_t^i) \leq (1-\tau)\epsilon \quad \forall i \in [m]; \\ \textit{or}: \text{ declare infeasibility, } \forall x \in X, \ \exists i \in [m] \quad \text{s.t.} \quad f^i(x, u_t^i) > 0. \end{cases} \tag{22}$$

We denote $x_t \in X$ to be the point returned by this oracle at step t, if it exists. For this approach, the outputs of a nominal feasibility oracle can be used to deduce a result similar to Corollary 3.1, except that we no longer need to evaluate $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$, we just need to bound $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$.

Theorem 4.2. *Given weights $\theta \in \Delta_T$, suppose that the sequence $\{u_t\}_{t=1}^T$ is generated in a non-anticipatory manner to guarantee $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \tau\epsilon$ for any sequence $\{x_t\}_{t=1}^T$. Also, suppose that at each step $t \in [T]$, x_t is generated by the nominal feasibility oracle which solves (22). If there exists $t \in [T]$ such that (22) declares infeasibility, then (4) is infeasible. Otherwise, if x_t satisfies $f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$ for all $t \in [T]$ and $i \in [m]$, we have a robust ϵ-feasibility certificate for (4).*

Proof. If (22) declares infeasibility, then it is obvious that the robust feasibility problem is infeasible. We focus on the latter case. By the premise of the theorem, we have $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \tau\epsilon$. Let us evaluate $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i)$. Because $\theta \in \Delta_T$ and from the definition of the nominal feasibility oracle we have $f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$ for all $t \in [T]$ and $i \in [m]$, we conclude $\max_{i \in [m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$. The conclusion now follows from Theorem 3.2. $\square$

Thus, the approach of [5], which works with nominal feasibility oracles, fits within our framework of Algorithm 1 right away. We next make three important remarks.

Remark 4.2. Similar to Remark 4.1, a critical property required in the approach of [5] of the vectors x_t is that $f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$. This is possible only if x_t were chosen with the knowledge of $\{u_1^i\}_{i=1}^m, \ldots, \{u_t^i\}_{i=1}^m$. ∎

Remark 4.3. Theorem 4.2 states that the nominal feasibility oracle-based approach can solve robust feasibility problems (4). This then recovers [5, Theorems 1,2]. In addition, we next make a nice and practical observation that was overlooked in [5]. We show that slightly adjusting this oracle will let us directly solve the robust *optimization* problem (2), i.e., optimize a convex objective function $f^0(x)$ instead of relying on a binary search over the optimal objective value. Recall that Opt is the optimal value of the RO problem (see (2)). Naively, to solve for Opt, we would embed f^0 into the constraint set, and then perform a binary search over the robust feasible set by repeatedly applying the oracle-based approach and Theorem 4.2 to check for robust feasibility. Suppose that now, instead of using a nominal feasibility oracle to solve (22), we work with a *nominal optimization oracle*. That is, given fixed parameters $\{u_t^i\}_{i=1}^m$, we have access to an oracle that solves

$$\text{Opt}_t = \inf_x \left\{ f^0(x) : f^i(x, u_t^i) \leq 0, \ i \in [m], \ x \in X \right\}.$$

18

When solving for Opt_t, most convex optimization solvers will either declare that the constraints are infeasible, or return a point $x_t \in X$ such that $f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$ and $f^0(x_t) \leq \text{Opt}_t + \epsilon$. It is clear that $f^0(x_t) \leq \text{Opt}_t + \epsilon \leq \text{Opt} + \epsilon$. Given such a sequence of points $\{x_t\}_{t=1}^T$, from Theorem 4.2 we deduced that $\bar{x}_T = \sum_{t=1}^T \theta_t x_t$ is a robust ϵ-feasible solution. Moreover, convexity of f^0 implies

$$ f^0(\bar{x}_T) \leq \sum_{t=1}^T \theta_t f^0(x_t) \leq \sum_{t=1}^T \theta_t(\text{Opt} + \epsilon) = \text{Opt} + \epsilon. $$

Hence, not only do we claim that $\bar{x}_T$ is robust ϵ-feasible, but that it is also ϵ-optimal. Thus, when our oracle can return ϵ-optimal solutions, which most solvers can, we eliminate the need to perform a binary search. ∎

Below we elaborate on the differences between Theorem 4.2 and Corollary 3.1.

Remark 4.4. In contrast to Corollary 3.1, Theorem 4.2 does not need to control the term $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$. There are two reasons for this: (*i*) due to (22), each point x_t satisfies $f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$, hence $\max_{i\in[m]} \sum_{t=1}^T \theta_t f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$ always holds. Therefore, the infeasibility part of Corollary 3.1 never becomes relevant, and (*ii*) due to the oracle solving (22), infeasibility may be declared at any step $t \in [T]$ in Theorem 4.2. This offers the possibility of stopping early rather than having to wait until all T steps are completed. Thus, the nominal feasibility oracle-based approach trades off using more effort at each iteration t to solve (22) for the ability to terminate early. In contrast, our OFO-based approach opts to keep the per-iteration cost cheap while giving up the ability to terminate early. More formally, let us examine a particular way of solving (22) within a nominal feasibility oracle. Note that (22) is equivalent to checking $F_t \leq (1-\tau)\epsilon$ or $F_t > 0$, where

$$ F_t := \inf_{x\in X} \left\{ \max_{i\in[m]} f^i(x, u_t^i) \right\}. \tag{23} $$

Since each $f^i(x, u_t^i)$ is convex in x for fixed u_t^i, $\max_{i\in[m]} f^i(x, u_t^i)$ is convex in x also, hence standard convex optimization methods may be employed to find $x_t \in X$ such that

$$ F_t \leq \max_{i\in[m]} f^i(x_t, u_t^i) \leq F_t + (1-\tau)\epsilon. $$

Then, by checking whether $\max_{i\in[m]} f^i(x_t, u_t^i) \leq (1-\tau)\epsilon$ or $\max_{i\in[m]} f^i(x_t, u_t^i) > (1-\tau)\epsilon$, we can determine whether $F_t \leq (1-\tau)\epsilon$ or $F_t > 0$ respectively. In particular, if we find that $F_t \leq (1-\tau)\epsilon$, our point x_t is feasible for (22).

Also, when all the vectors x_t satisfy (23), we have the bound

$$
\begin{aligned}
\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) &= \sum_{t=1}^T \theta_t \max_{i\in[m]} f^i(x_t, u_t^i) - \inf_{x\in X} \sum_{t=1}^T \theta_t \max_{i\in[m]} f^i(x, u_t^i) \\
&\leq \sum_{t=1}^T \theta_t \left[\max_{i\in[m]} f^i(x_t, u_t^i) - \inf_{x\in X} \max_{i\in[m]} f^i(x, u_t^i) \right] \leq (1-\tau)\epsilon.
\end{aligned}
$$

Consequently, we deduce that the nominal feasibility oracle, implemented as a convex optimization problem, also naturally bounds $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ although this bound is not utilized in Theorem 4.2. In terms of our framework of Algorithm 1, the update policy of generating new solutions x_t in this approach corresponds to selecting the algorithm $\mathcal{A}_x$ to be a convex optimization solver that

solves Opt_t. Then whenever the solver returns a feasible solution, the regret bound (17) satisfies $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \mathcal{R}_x(T) = (1 - \tau)\epsilon$ for any T. Note that the term $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ inherently includes the objective functions $\max_{i \in [m]} f^i(x, u_t^i)$ of each problem F_t. At each iteration t, instead of evaluating F_t to $(1 - \tau)\epsilon$ accuracy, our OFO-based approach performs only a simple update based on the first-order information, and it yields a bound on $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ from the overall collection of these simple updates. ∎

4.4 Convergence Rates and Discussion

We summarize the convergence rates achievable in our general RO framework for various cases. We first examine the number of iterations required for each approach discussed, then proceed to analyze the per-iteration cost of each approach. A summary of our discussion is given in Table 1. We use the notation $r_u(\epsilon)$ to denote the number of iterations T required for algorithms $\mathcal{A}_i$ to guarantee $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \max_{i \in [m]} \mathcal{R}_i(T) \leq \epsilon/2$. Similarly, we let $r_x(\epsilon)$ be the number of iterations T required for algorithm $\mathcal{A}_x$ to guarantee that $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \mathcal{R}_x(T) \leq \epsilon/2$. Then the resulting worst-case number of iterations needed in Algorithm 1 to obtain robust ϵ-feasibility/infeasibility certificates is $\max\{r_u(\epsilon), r_x(\epsilon)\}$.

As outlined in Section 4, employing standard OFO-based algorithms, i.e., Theorem 2.1, on the terms $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T)$ and $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ requires $r_u(\epsilon) = O(1/\epsilon^2)$ and $r_x(\epsilon) = O(1/\epsilon^2)$ iterations to ensure they are no larger than $\epsilon/2$. Thus, our OFO-based approach from Section 4.1 requires $O(1/\epsilon^2)$ iterations to solve (3). Since our OFO-based approach returns only robust ϵ-feasible solutions, we need to perform a binary search and repeatedly invoke our method $O(\log(1/\epsilon))$ times to obtain ϵ-optimal solutions, so the total number of iterations is $O(\log(1/\epsilon)/\epsilon^2)$.

Our FO-based pessimization approach, i.e., our modification of the pessimization oracle-based approach of [31] outlined in Section 4.2, requires $r_x(\epsilon)$ iterations to solve (3) because by Theorem 4.1 we only need to guarantee $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \mathcal{R}_x(T) \leq \epsilon/2$. Taking into account the binary search factor $O(\log(1/\epsilon))$ to find a robust ϵ-optimal solution, the total number of iterations required is $O(\log(1/\epsilon)/\epsilon^2)$, which is much better than the exponential $(1 + O(1/\epsilon))^n$ bound of [31, Section 5.2] that uses a full nominal solution oracle (20). Similarly, the nominal feasibility/optimization oracle-based approach of [5] outlined in Section 4.3 requires $r_u(\epsilon) = O(1/\epsilon^2)$ iterations (or $r_u(\epsilon) \log(1/\epsilon) = O(\log(1/\epsilon)/\epsilon^2)$ iterations if only a feasibility oracle is used) to obtain robust ϵ-optimal solutions because by Theorem 4.2 we only need to bound $\epsilon^\circ(\{x_t, u_t, \theta_t\}_{t=1}^T) \leq \max_{i \in [m]} \mathcal{R}_i(T) \leq \epsilon/2$.

Remark 4.5. The flexibility of our general framework in terms of the selection of algorithms $\mathcal{A}_i, \mathcal{A}_x$ extends beyond just using Theorem 2.1. Depending on the structure of functions f^i and uncertainty domains U^i, the algorithms $\mathcal{A}_i$ and $\mathcal{A}_x$ may be replaced by more appropriate OCO algorithms. For example, when f^i are strongly convex, certain OCO algorithms achieve faster convergence rates. Moreover, unless explicitly required by the algorithms $\mathcal{A}_i$, we do not need to assume convexity of the sets U^i. As a result, the follow-the-leader or follow-the-perturbed-leader type algorithms from [30] can be utilized as $\mathcal{A}_i$ in our framework even when U^i are nonconvex but certain assumptions ensuring applicability of these algorithms are satisfied. Such assumptions are satisfied for example when $f^i(x, u^i)$ are linear in u^i and the nonconvex sets U^i admit a certain linear optimization oracle. This is for example the case in a certain lifted representation of the robust convex quadratic constraint discussed in [5, Section 4.2]. Similarly, when the functions $f^i(x, u^i)$ are exp-concave in u^i, applying the online Newton step algorithm of [22] for $\mathcal{A}_i$ results in a weighted regret bound of at most $O(\log(T)/T)$ in T iterations. Such f^i that are exp-concave in u^i satisfying Assumption 2.1 arise in optimization under uncertainty problems where variance is used as a risk measure, e.g.,

mean-variance portfolio optimization problems, see for example [3, Example 25]. Essentially, the same flexibility for acceleration and/or working with nonconvex sets U^i is present in [5] as well.

In the presence of favorable problem structure, based on Table 1, if an accelerated algorithm to exploit problem structure is employed in the place of $\mathcal{A}_i$, the overall number of iterations of the nominal feasibility approach is immediately reduced accordingly. Analogous result holds for $\mathcal{A}_x$ and the FO-based pessimization approach. However, in the case of our OFO-based approach, we need to have favorable structure in *both* x and u and utilize the corresponding accelerated algorithms $\mathcal{A}_x, \mathcal{A}_i$ to attain the acceleration of the overall approach. ∎

We now discuss the per-iteration cost for each approach. In order to discuss the total *arithmetic complexity* of each approach, we let k be the maximum dimension of the uncertain parameters u^i for $i \in [m]$ and recall that n denotes the dimension of the decision variables x. In the case where our domains $X, \{U^i\}_{i=1}^m$ have favorable geometry, such as Euclidean ball or simplex, the vectors $x_t, \{u_t^i\}_{i=1}^m$ are updated via simple closed-form prox operations, which cost $O(n)$ and $O(km)$ per iteration respectively. The cost of computing the subgradients $\nabla_x f^i(x, u^i), \nabla_u f^i(x, u^i)$ is at least $O(km + mn)$ each iteration. This cost is incurred in each iteration of all of the approaches we discuss. From this, we deduce that the per-iteration cost of our OFO-based approach is at most $O(km + mn)$.

The per-iteration cost of the pessimization oracle based approaches involve calling m pessimization oracles (21) and the costs related to updating x_t. We denote by $\text{Pess}(\epsilon, k)$ the complexity of a pessimization oracle with tolerance ϵ and k variables. A summary of different possible implementations is given in Table 2. If $\sup_{u^i \in U^i} f^i(x, u^i)$ has a simple closed form solution, then the resulting arithmetic cost for $\text{Pess}(\epsilon, k)$ is $O(k)$ for each pessimization oracle. If we can use polynomial-time IPMs, this cost becomes $O(k^3 \log(1/\epsilon))$ (see [7, Section 6.6]), and using FOMs has cost $O(k \log(1/\epsilon))$ in the best case when the functions f^i are smooth *and* strongly convex in u^i. In the case of our FO-based pessimization approach, the update involving x_t will be given by simple closed form formulas for prox operations when X has favorable geometry, resulting in a cost of $O(mn)$. The full pessimization approach of [31] incurs the cost of solving an extended nominal feasibility problem for the update of x_t.

The per-iteration cost of the nominal feasibility/optimization oracle-based approach of [5], as well as that of of [31], depends on the type of solver used to solve the nominal optimization/feasibility problem (22). We denote by $\text{Nom}(\epsilon, m, n)$ the complexity of a nominal oracle with tolerance ϵ, m constraints and n variables. Note that nominal solvers can be either optimization or feasibility solvers. If it is the latter, an extra $\log(1/\epsilon)$ factor is incurred to perform binary search. A summary of different possible implementations for $\text{Nom}(\epsilon, m, n)$ is given in Table 3. When applicable for $\text{Nom}(\epsilon, m, n)$ implementation, polynomial-time IPMs are guaranteed to terminate in $O(\sqrt{m} \log(1/\epsilon))$ iterations with a solution to (22) and thus offer the best rates in terms of their dependence on ϵ. They also have the advantage that they can act as a nominal optimization oracle, and hence by Remark 4.3 there will be no need to perform an additional binary search to find an ϵ-optimal solution. On the other hand, they demand significantly more memory, and their per-iteration cost is quite high in terms of the dimension, usually around the order of $O(n^3 + mn)$, see [7, Chapter 6.6]. In order to keep both the memory requirements and the per-iteration cost associated with implementing the nominal feasibility oracle $\text{Nom}(\epsilon, m, n)$ low, one may opt for a FOM called the CoMirror algorithm that can work with functional constraints, see [1] and [27, Section 1.3]. CoMirror algorithm is guaranteed to find a solution to the nominal ϵ-feasibility problem within $O(1/\epsilon^2)$ iterations, with a much cheaper per-iteration cost of $O(mn)$. Because CoMirror method

Table 1: Summary of different approaches to generate $\{x_t, u_t\}_{t=1}^T$.

Approach	Binary search	No. iterations	Per-iteration cost
OFO-based	$\log(1/\epsilon)$	$\max\{r_u(\epsilon), r_x(\epsilon)\}$	$O(km + mn)$
FO-based pessimization	$\log(1/\epsilon)$	$r_x(\epsilon)$	$m\operatorname{Pess}(\epsilon, k) + O(mn)$
Nominal oracle	see Table 3	$r_u(\epsilon)$	$O(km) + \operatorname{Nom}(\epsilon, m, n)$
Full pessimization	see Table 3	$O(1/\epsilon^n)$	$m\operatorname{Pess}(\epsilon, k) + \operatorname{Nom}(\epsilon, m+t, n)^*$
			*(number of constraints is $m + t$ as it grows by at least 1 each iteration t)
Direct FOM via CoMirror	1	$O(1/\epsilon^2)$	$m\operatorname{Pess}(\epsilon, k) + O(mn)$

Table 2: Arithmetic complexity for different implementations of pessimization oracles.

Implementation	$\operatorname{Pess}(\epsilon, k)$
Closed form	$O(k)$
IPM	$O(k^3 \log(1/\epsilon))$
FOM*	$O(k \log(1/\epsilon))$

*(when f^i are smooth, strongly convex in u^i)

can optimize as well, it does not need binary search. However, to the best of our knowledge, its possibility to exploit further structural properties of the functions f^i, such as smoothness in x, to improve the dependence on ϵ are not known. In order to exploit such properties in the implementation of $\operatorname{Nom}(\epsilon, m, n)$, it is possible to cast (22) as a convex-concave SP problem, and then apply efficient FOMs such as Nesterov's algorithm [38] or Nemirovski's Mirror Prox algorithm [34] to achieve a convergence rate of $O(\log(m)/\epsilon)$ and per-iteration cost of $O(mn)$. This convex-concave SP approach can only be used as a nominal feasibility oracle, so we must repeat the process $\log(1/\epsilon)$ times to obtain an ϵ-optimal solution.

Table 3: Arithmetic complexity for different implementations of nominal oracles.

Implementation	$\operatorname{Nom}(\epsilon, m, n)$	Type	Binary search
IPM	$O\left(km + \sqrt{m}(n^3 + mn)\log(1/\epsilon)\right)$	optimization	1
CoMirror	$O\left(mn/\epsilon^2\right)$	optimization	1
Convex-concave SP*	$O\left(\log(m)mn/\sqrt{\epsilon}\right)$	feasibility	$\log(1/\epsilon)$

*(when f^i are smooth, strongly convex in x)

Recall that Table 1 summarizes the rates for the various approaches, together with rates for the full pessimization approach of [31] and using the CoMirror with pessimization (discussed in Section 4.5). Note that the total *overall arithmetic complexity* of each approach is obtained by multiplying the quantities in each row in Table 1. The quantities $r_u(\epsilon), r_x(\epsilon)$ will generally be $O(1/\epsilon^2)$, with potential for application-specific acceleration when the functions f^i exhibit favorable structure. Table 1 indicates that our FO-based pessimization approach when it admits a closed form solution for the implementation of $\operatorname{Pess}(\epsilon, k)$ and the nominal feasibility oracle-based approach which uses a polynomial-time IPM solver to implement the nominal feasibility oracle $\operatorname{Nom}(\epsilon, m, n)$ give the best dependence on ϵ among all of the methods. These are better than our OFO-based

approach by factors of $\max\{1, r_u(\epsilon)/r_x(\epsilon)\}$ and $\max\{1, r_x(\epsilon)/r_u(\epsilon)\}$ respectively. However, in many applications, we can expect that $r_u(\epsilon) \approx r_x(\epsilon)$, so these factors will be constant. In this case, our OFO-based approach becomes competitive with having a closed form pessimization oracle in our FO-based pessimization approach or using a nominal IPM solver in [5]. That said, compared to IPMs, our OFO-based approach demands much less memory, and it is able to maintain a much lower dependence on the dimensions m, n and thus is much more scalable, whereas the cost per iteration of such IPMs has a rather high dependence on the dimension. In addition, the memory requirements of IPMs are far more than OFO algorithms, posing a critical disadvantage to their use in large-scale applications. Similar comparisons of our OFO-based approach against pessimization or nominal feasibility oracle-based approaches utilizing other methods point out its advantage, which is at least an order of magnitude better in terms of its dependence on ϵ. In fact, when $r_x(\epsilon) \approx r_u(\epsilon)$, our method can lead to savings over the approach of [5] with CoMirror algorithm used in its oracle by a factor as large as $O(1/(\epsilon^2 \log(1/\epsilon)))$.

4.5   Connections with Existing First-order Methods

Finally, we would like to discuss and contrast directly solving robust convex optimization problems (2) via general first-order methods. Many FOMs require domains that are simple so that the prox operations can be easily done. In that respect, domains defined by multiple functional constraints $g^i(x) \leq 0$ creates a challenge for directly applying many of these algorithms. We now discuss two existing classes of FOMs that are designed to handle such domains: primal-dual methods and the CoMirror approach. Applying these FOMs to the RO problem (2) can be viewed as another alternative solution methodology to solve RO problems without using the robust counterpart.

A general technique to address the functional constraints in the domain is to embed these constraints into the objective through Lagrange multipliers, and then solve the associated dual problem via FOMs (see e.g., [32]). Such methods are known as primal-dual methods. For the RO problem (2), this corresponds to solving

$$\max_{\lambda} \left\{ \mathcal{L}^*(\lambda) := \min_{x \in X} \left[f^0(x) + \sum_{i=1}^{m} \lambda^{(i)} g^i(x) \right] : \lambda \geq 0 \right\},$$

where we define $g^i(x) := \sup_{u^i \in U^i} f^i(x, u^i)$. Primal-dual methods (e.g., [32]) commonly require us to solve the inner minimization problem over $x \in X$ at each iteration. For RO, this means we must solve an expensive SP problem at each iteration. Our OFO-based approach aims to improve on this by reducing the per-iteration cost of each step to simple first-order updates. Two exceptions within the primal-dual methods are the work of Nedić and Ozdaglar [33] and Yu and Neely [46], which have cheap per-iteration cost based on only gradient computations and projection operations in the Euclidean setup. Nedić and Ozdaglar [33] provide a convergence rate of $O(1/\sqrt{T})$ in the non-smooth case. While using such a primal-dual method has the advantage that no binary search is needed, we note that this requires two assumptions to guarantee convergence: we have access to exact first-order information for the robust constraint functions $g^i(x) := \sup_{u^i \in U^i} f^i(x, u^i)$, and the standard Slater constraint qualification condition (i.e., strict feasibility) is satisfied. The first assumption is often not satisfied, since we may only be able to compute $g^i(x)$ up to accuracy ϵ. While there exists some FOMs that work with inexact objective gradients over simple domains, see e.g., [18], such methods have only been applied to specific max-type objectives, e.g., objectives obtained from smoothing. It is unclear how such methods can be extended for more general max-type functions which can arise in RO. Secondly, enforcing the Slater condition implicitly enforces

feasibility of (2). In contrast, our framework directly uses the functions $f^i(x, u^i)$, so it does not need to take into account the inexact gradient information, and can certify infeasibility of (2). Yu and Neely [46] present a method that can guarantee $O(1/T)$ convergence when all functions are smooth. However, for RO problems, the constraint functions $g^i(x)$ are non-smooth due to the supremum operation, thus their results do not apply to RO.

The only FOM that we are aware of that can solve convex problems with functional constraints without assuming feasibility is the CoMirror algorithm [1] and its earlier variations in the Euclidean setup [35, 37, 40]. The CoMirror[1] algorithm finds an ϵ-optimal ϵ-feasible solution in $O(1/\epsilon^2)$ iterations to a convex program $\min_{x \in X} \{f^0(x) : g^i(x) \leq 0, \ i \in [m]\}$ or certifies its infeasibility by using (sub)gradient information of the objective f^0 as well as the constraint functions g^i. In the RO problem (2) we defined $g^i(x) := \sup_{u^i \in U^i} f^i(x, u^i)$. As mentioned above, in many cases, we may only be able to compute $g^i(x)$ approximately, thus only have access to approximate/inexact gradient information. It is unknown to us whether or not techniques such as the ones from [18] can be applied to the CoMirror algorithm in the presence of this type of gradient information. While the CoMirror algorithm's complexity is $O(1/\epsilon^2)$ (see also Nesterov [37, Chapter 3.2.4] for a similar result in the Euclidean case), our iterative framework can exploit favorable structure on the functions f^i that can improve on the iteration complexity $r_u(\epsilon), r_x(\epsilon)$. For the Euclidean case, Nesterov [37, Chapters 2.3.4-2.3.5] shows also that convergence can be obtained in $O(\log(1/\epsilon))$ iterations when the objective and all constraint functions are both smooth and strongly convex in x. However, such an improvement does not apply to the RO problem, since we cannot in general guarantee that $g^i(x) = \sup_{u^i \in U^i} f^i(x, u^i)$ is smooth in x. It is unknown whether the iteration complexity of CoMirror algorithm can be improved when only the underlying function $f^i(x, u^i)$ is strongly convex or smooth, or when $g^i(x)$ is strongly convex but non-smooth.

Finally, let us get back to the case when we have a robust feasibility problem with a *single* constraint $m = 1$ and a convex uncertainty set $U = U^1$. In such a case, as discussed in Remark 3.1, we have a direct convex-concave SP problem ($\mathcal{S}$) under Assumption 2.1. The OFO-based approach then corresponds to bounding the whole SP gap (5), the FO-based pessimization corresponds to bounding the primal gap i.e., the first term in (5), and the nominal feasibility oracle approach corresponds to bounding the dual gap, i.e., the second term in (5). Without any further structural assumptions on f^1, convex-concave SP problems can be solved in $O(1/\epsilon^2)$ iterations. Our approaches also achieve this rate immediately, see Table 1. Moreover, when f^1 is smooth in x and strongly concave in u^1, our general framework can achieve a rate of $O(1/\epsilon)$. However, for specific applications involving a single robust constraint, directly working with a specialized convex-concave SP formulation can improve this rate further. For example, such an improved rate of $O(1/\sqrt{\epsilon})$ is achieved in [2] for a robust support vector machine problem.

5 Application Example: Robust Quadratic Programming

Our framework is general and can be applied to many robust convex optimization problems. In this section we walk through the setup and resulting convergence rates of our framework for a robust feasibility problem of a quadratically constrained quadratic program (QP) with ellipsoidal uncertainty. To be precise, our deterministic feasibility problem is

$$\text{find} \ \ x \in X \ \ \text{s.t.} \ \ \|A_i x\|_2^2 \leq b_i^\top x + c_i, \quad \forall i \in [m],$$

---

[1] Recall that the CoMirror algorithm is also discussed in Section 4.4 as a method to implement the nominal feasibility solver; in that case we are given the noises $\bar{u}^i$ resulting in $g^i(x) := f^i(x, \bar{u}^i)$, and thus the subgradient of $g^i(x)$ is simply the subgradient of $f^i(x, \bar{u}^i)$.

where $X \subseteq \mathbb{R}^n$ is the unit Euclidean ball, $A_i \in \mathbb{R}^{n \times n}$, $b_i \in \mathbb{R}^n$, and $c_i \in \mathbb{R}$ for all $i \in [m]$. We consider the robust quadratic feasibility problem given by

$$\text{find} \;\; x \in X \;\; \text{s.t.} \;\; \sup_{u \in \widehat{U}} \left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i \leq 0, \quad \forall i \in [m], \tag{24}$$

where $P_1^i, \ldots, P_K^i$ are uncertainty matrices for each constraint $i \in [m]$, for simplicity we assume uncertainty sets $U^i = \widehat{U} = \{u \in \mathbb{R}^K : \|u\|_2 \leq 1\}$ for all $i \in [m]$, and $u^{(k)}$ denotes the k-th entry of u.

It is well known that the robust counterpart of this feasibility problem is a semidefinite program [4, 11]. Because current state-of-the-art QP solvers can handle two to three orders of magnitude larger QPs than semidefinite programs (SDPs), Ben-Tal et al. [5, Section 4.2] suggest an approach that avoids solving SDPs associated with robust QPs. Their approach relies on running a probabilistic OCO algorithm in which a trust region subproblem (TRS)—a class of well-studied nonconvex QPs—is solved in each iteration. Our results here further enhance this approach. In particular, we show that we can achieve the same rate of convergence in our framework while working with a deterministic OCO algorithm and only carrying out first-order updates in each iteration. In fact, the most expensive operation involved with each iteration of our approach is a maximum eigenvalue computation. Because maximum eigenvalue computation is much cheaper than solving a TRS, we not only present a deterministic approach but also strikingly reduce the cost of each iteration.

To simplify our exposition, let us introduce some notation. For each $i \in [m]$, we define the matrix $\mathcal{P}_x^i \in \mathbb{R}^{n \times K}$ whose columns are given by the vectors $P_k^i x$ for $k \in [K]$ together with

$$Q_x^i := (\mathcal{P}_x^i)^\top \mathcal{P}_x^i \in \mathbb{S}_+^K, \quad r_x^i := (\mathcal{P}_x^i)^\top A_i x \in \mathbb{R}^K, \quad \text{and} \quad s_x^i := \|A_i x\|_2^2 - b_i^\top x - c_i \in \mathbb{R};$$

then it is easy to check that for all $i \in [m]$ and $u \in \mathbb{R}^K$ we have

$$\left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i = u^\top Q_x^i u + 2(r_x^i)^\top u + s_x^i.$$

For each $i \in [m]$, we define $f^i : X \times \widehat{U} \to \mathbb{R}$ as

$$f^i(x, u) := \left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i + \lambda_{\max}(Q_x^i) \left(1 - \|u\|_2^2 \right)$$

$$= u^\top Q_x^i u + 2(r_x^i)^\top u + s_x^i + \lambda_{\max}(Q_x^i) \left(1 - \|u\|_2^2 \right). \tag{25}$$

Lemma 5.1. *For each $i \in [m]$, the function $f^i(x, u)$ defined in (25) is convex in x for any fixed $u \in \widehat{U}$ and concave in u for any given x. Moreover, for all $i \in [m]$ and for any $x \in X$,*

$$\sup_{u \in \widehat{U}} \left\| \left(A_i + \sum_{k=1}^{K} P_k^i \, u^{(k)} \right) x \right\|_2^2 - b_i^\top x - c_i = \sup_{u \in \widehat{U}} f^i(x, u).$$

Proof. Fix $i \in [m]$. By rearranging terms in (25), we obtain $f^i(x, u) = u^\top (Q_x^i - \lambda_{\max}(Q_x^i) I_K) u + 2(r_x^i)^\top u + s_x^i$. Since $Q_x^i - \lambda_{\max}(Q_x^i) I_K \in \mathbb{S}_+^K$ for any given x, $f^i(x, u)$ is concave in u for any given x.

25

Now consider a fixed $u \in \widehat{U}$. Note that

$$\lambda_{\max}(Q_x^i) = \max_{\|v\|_2 \leq 1} v^\top (Q_x^i) v = \max_{\|v\|_2 \leq 1} \sum_{1 \leq j,k \leq K} v^{(j)} v^{(k)} x^\top (P_j^i)^\top P_k^i x = \max_{\|v\|_2 \leq 1} x^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right)^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right) x.$$

Because $\left(\sum_{k=1}^K P_k^i \, v^{(k)} \right)^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right) \in \mathbb{S}_+^n$, then $\lambda_{\max}(Q_x^i)$ is a maximum of convex quadratic functions of x and hence is convex in x. Thus, for fixed $u \in \widehat{U}$, $f^i(x, u)$ is convex in x.

Reformulation of the nonconvex QP over an ellipsoid into a convex QP over the ellipsoid via the relation between $u^\top Q_x^i u + 2(r_x^i)^\top u + s_x^i$ and $f^i(x, u)$ in (25) follows from [25, Theorem 2.1]. □

Lemma 5.1 implies that $\sup_{u \in \widehat{U}} f^i(x, u) \leq 0$ is an alternate representation of our robust quadratic constraint. We next state the convergence rate in our framework for the associated feasibility problem. For this, we define the quantities

$$\sigma^2 := \max_{i \in [m]} \sum_{k=1}^K \|P_k^i\|_{\mathrm{Fro}}^2, \quad \chi := \max_{i \in [m]} \max_{k \in [K]} \|P_k^i\|_{\mathrm{Spec}}, \quad \text{and}$$
$$\rho := \max_{i \in [m]} \|A_i\|_{\mathrm{Spec}}, \qquad \beta := \max_{i \in [m]} \|b_i\|_2. \tag{26}$$

Note that $\chi \leq \sigma$. Furthermore, [5, Lemma 7] proves that $\|Q_x^i\|_{\mathrm{Fro}} \leq \sigma^2$ and $\|r_x^i\|_2 \leq \sigma \rho$ holds for all x such that $\|x\|_2 \leq 1$.

Corollary 5.1. *Let our domain be given by $X = \{x \in \mathbb{R}^n : \|x\|_2 \leq 1\}$. The customization of our OFO-based approach to the problem* (24) *ensures that within $O\left(((\rho + \sqrt{K}\sigma)^2 + \beta)^2) \epsilon^{-2} \right)$ iterations, we obtain a robust feasibility/infeasibility certificate. Moreover, each iteration in our framework relies on a first-order update where the most expensive operation in the case of* (24) *is computing $\lambda_{\max}(Q_x^i)$, which can be done efficiently.*

Proof. In order to apply OFO-based approach, we need to customize our proximal setup. Given that the sets X and $\widehat{U}$ are Euclidean balls, we set the proximal setup for generating the iterates $\{x_t, u_t^i\}_{t=1}^T$ to be the standard Euclidean d.g.f. with $\|\cdot\|_2$-norm, and thus $\Omega_X = \Omega_{\widehat{U}} = \frac{1}{2}$. We must bound the magnitude of the gradients measured by the $\|\cdot\|_2$-norm. Note that for any $i \in [m]$, the gradients of f^i are given by

$$\nabla_u f^i(x, u) = 2 \left(Q_x^i - \lambda_{\max}(Q_x^i) I_K \right) u + 2r_x^i$$
$$\nabla_x f^i(x, u) = 2 \left(A_i + \sum_{k=1}^K P_k^i \, u^{(k)} \right)^\top \left(A_i + \sum_{k=1}^K P_k^i \, u^{(k)} \right) x + 2 \left(1 - \|u\|_2^2 \right) \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right)^\top \left(\sum_{k=1}^K P_k^i \, v^{(k)} \right) x - b_i,$$

where $v \in \widehat{U}$ is an eigenvector of Q_x^i corresponding to $\lambda_{\max}(Q_x^i)$.

Let us fix an $i \in [m]$. We first bound $\|\nabla_u f^i(x, u)\|_2$ for any $u \in \widehat{U}$ as follows:

$$\|\nabla_u f^i(x, u)\|_2 = 2 \left\| \left(Q_x^i - \lambda_{\max}(Q_x^i) I_K \right) u + r_x^i \right\|_2$$
$$\leq 2 \left(\left\| \left(Q_x^i - \lambda_{\max}(Q_x^i) I_K \right) u \right\|_2 + \left\| r_x^i \right\|_2 \right) \leq 2\lambda_{\max}(Q_x^i) \|u\|_2 + 2\sigma\rho \leq 2(\sigma^2 + \sigma\rho),$$

where the second inequality follows from $\left\|Q_x^i - \lambda_{\max}(Q_x^i) I_K\right\|_{\mathrm{Spec}} \le \lambda_{\max}(Q_x^i)$ and $\|r_x^i\|_2 \le \sigma\rho$ which is implied by [5, Lemma 7], and the last inequality follows from the facts that $u \in \widehat{U}$, the definitions given in (26), and $\lambda_{\max}(Q_x^i) = \|\mathcal{P}_x^i\|_{\mathrm{Spec}}^2 \le \|\mathcal{P}_x^i\|_{\mathrm{Fro}}^2 \le \sum_{k=1}^K \|P_k^i\|_{\mathrm{Fro}}^2 \le \sigma^2$ for any $x \in X$. Therefore, we deduce from Theorem 2.1 with uniform weights $\theta_t = 1/T$ that the rate of convergence for bounding the weighted regret associated with constraint $i \in [m]$ using the online mirror descent algorithm is

$$\sup_{u \in U} \frac{1}{T}\sum_{t=1}^T f^i(x_t, u) - \frac{1}{T}\sum_{t=1}^T f^i(x_t, u_t) \le \frac{2(\sigma^2 + \sigma\rho)}{\sqrt{T}}.$$

This implies that $r_u(\epsilon) = O((\sigma^2 + \sigma\rho)^2 \epsilon^{-2})$.

We next bound the weighted regret of the functions $\varphi_t(x) = \max_{i \in [m]} f^i(x, u_t^i)$, i.e., the term $\epsilon^\bullet(\{x_t, u_t, \theta_t\}_{t=1}^T)$ by bounding the $\|\cdot\|_2$-norm of $\nabla_x \varphi_t(x)$. Notice that

$$\|\nabla_x \varphi_t(x)\|_2 \le \max_{i \in [m]} \|\nabla_x f^i(x, u_t^i)\|_2.$$

Thus, we must bound $\|\nabla_x f^i(x, u)\|_2$ for all $x \in X$, $u \in \widehat{U}$. To this end, note that for any $u \in \widehat{U}$

$$\left\|\sum_{k=1}^K P_k^i u^{(k)}\right\|_{\mathrm{Spec}} \le \sum_{k=1}^K \|P_k^i\|_{\mathrm{Spec}} |u^{(k)}| \le \sqrt{K} \max_{k \in [K]} \|P_k^i\|_{\mathrm{Spec}} \le \sqrt{K}\chi,$$

where the second inequality holds because $\|u\|_1 \le \sqrt{K}\|u\|_2 \le \sqrt{K}$ holds for all $u \in \widehat{U}$. Then for any $x \in X$, $u \in \widehat{U}$, and eigenvector $v \in \widehat{U}$, we have

$$\begin{aligned}
\|\nabla_x f^i(x, u)\|_2 &\le 2\left\|A_i + \sum_{k=1}^K P_k^i u^{(k)}\right\|_{\mathrm{Spec}}^2 \|x\|_2 + 2\left(1 - \|u\|_2^2\right)\left\|\sum_{k=1}^K P_k^i v^{(k)}\right\|_{\mathrm{Spec}}^2 \|x\|_2 + \|b_i\|_2 \\
&\le 2(\rho + \sqrt{K}\chi)^2 + 2K\chi^2 + \beta \\
&\le 4(\rho + \sqrt{K}\sigma)^2 + \beta.
\end{aligned}$$

Hence, $\|\nabla_x \varphi_t(x)\|_2 \le 4(\rho + \sqrt{K}\sigma)^2 + \beta$. Then Theorem 2.1 with weights $\theta_t = 1/T$ implies

$$\sum_{t=1}^T \theta_t \max_{i \in [m]} f^i(x_t, u_t^i) - \inf_{x \in X} \sum_{t=1}^T \max_{i \in [m]} f^i(x, u_t^i) \le \frac{\left(4(\rho + \sqrt{K}\sigma)^2 + \beta\right)}{\sqrt{T}}.$$

Thus, $r_x(\epsilon) = O\left(((\rho + \sqrt{K}\sigma)^2 + \beta)^2\right)\epsilon^{-2}$. Therefore, the number of iterations required for our OFO-based approach to obtain a robust feasibility/infeasibility certificate is $T = \max\{r_x(\epsilon), r_u(\epsilon)\}$.

Note that each iteration of our approach requires a first-order update that is composed of computing the gradients $\nabla_x f^i(x, u)$ and $\nabla_u f^i(x, u)$ and prox computations. Because our domains involve only direct products of Euclidean balls and simplices, they admit efficient prox computations which take $O(Km + mn)$ time. In order to evaluate the gradients $\nabla_x f^i(x, u)$ and $\nabla_u f^i(x, u)$, in addition to the elementary matrix vector operations, we need to compute $\lambda_{\max}(Q_x^i)$ which is the most expensive operation in our first-order update. Fortunately, computing the maximum eigenvalue of a matrix is a well-studied problem and can be computed very efficiently. □

In the case of robust QP feasibility problem (24), [5, Corollary 3] states that with probability $1 - \delta$, their framework returns robust feasibility/infeasibility certificates in at most $O\left(K^2\sigma^2(\rho^2 + \sigma^2)\log(m/\delta)\epsilon^{-2}\right)$ calls (iterations) to their oracle. In each call to their oracle, a nominal feasibility problem is solved to the accuracy $\epsilon/2$. In comparison we deduce from Corollary 5.1 that our framework requires comparable number of iterations as the approach of Ben-Tal et al. [5]. Even so, there are a number of reasons that considerably favor our approach. First, our approach is deterministic as opposed to the high $1 - \delta$ probability guarantee of [5] which requires using an adaptation of the follow-the-perturbed-leader type OCO. Second, each iteration of their approach requires solving a nominal feasibility problem for solution oracle as well as solving TRSs for the computation of noises u_t. In contrast to this, in each iteration we carry out mainly elementary operations such as matrix vector multiplications and our most computationally expensive operation is the maximum eigenvalue computations $\lambda_{\max}(Q_x^i)$. While there are established algorithms to solve the TRS, it is inherently more complicated than finding the maximum eigenvalue of a positive semidefinite matrix. Moreover, [5] suffers from the additional computational cost of their solution oracle which solves the nominal feasibility problem. Hence, our approach, while requiring a comparable number of iterations, reduces the cost per iteration remarkably.

6   Numerical Study

In this section, we conduct a numerical study comparing the approaches discussed so far. We consider the following quadratic program inspired by mean-variance portfolio optimization problems with a factor model for the return vector (see, e.g., [20]):

$$\min_x \left\{ \|Vx\|_2^2 + x^\top D x - \lambda \mu^\top x : \ x \in \Delta_n \right\}, \tag{27}$$

where $\mu \in \mathbb{R}^n$ is the expected return vector, the term $x^\top (V^\top V + D)x$ captures the risk associated with the portfolio via a factor model, and $\lambda \geq 0$ represents the trade-off between the expected return of the portfolio and the risk associated with the portfolio.

In the robust formulation of (27), we consider the case where the true parameters $\mu \in \mathbb{R}^n$ and $V \in \mathbb{R}^{m \times n}$ belong to uncertainty sets $\mathcal{M}$ and $\mathcal{V}$ of form

$$\mathcal{M} := \{\mu : \ \mu_0 - \gamma \leq \mu, \ \mu \leq \mu_0 + \gamma\}, \quad \mathcal{V} := \left\{V = V_0 + \sum_{k=1}^K P_k u_k : \ \|u\|_2 \leq 1\right\},$$

where the nominal data $\mu_0 \in \mathbb{R}^n$, $\gamma \in \mathbb{R}^n$, and $V_0 \in \mathbb{R}^{m \times n}$, $\{P_k \in \mathbb{R}^{m \times n}\}_{k=1}^K$ are given to us. Then the robust problem is given by

$$\min_x \left\{ \max_{V \in \mathcal{V}} \|Vx\|_2^2 + x^\top D x - \lambda \min_{\mu \in \mathcal{M}} \mu^\top x : \ x \in \Delta_n \right\}. \tag{28}$$

Our test instances are synthetically generated, largely following the random instance generation model from [20]. We begin by specifying three parameters: n, the number of variables; m, the number of factors (which controls the rank of V); and $\alpha \in (0, 1)$, a parameter controlling the size of the uncertainty sets. For each instance, we randomly generate matrices $V \in \mathbb{R}^{m \times n}$ and $F \in \mathbb{R}^{m \times m}$, where we ensure F is positive semidefinite, and define $D = 0.1 \, \mathrm{Diag}(V^\top F V)$. We then generate $p > m$ factor samples $f_{(l)} \in \mathbb{R}^m$, $l \in [p]$, where each $f_{(l)} \sim N(0, F)$, and we also generate $\mu \in \mathbb{R}^n$ where each entry $\mu_i \sim U(1, 5)$. We then set $\mu_{(l)} = \mu + V^\top f_{(l)} + \epsilon_l$, where $\epsilon_{(l)} \sim N(0, D)$ are

28

independent of the factor sample $f_{(l)}$. The matrices μ and V are estimated via linear regression on $\mu_{(l)}$ and $f_{(l)}$, to obtain $\bar{\mu}, \bar{V}$. The nominal data for (27) are set to be $\mu_0 = \bar{\mu}$, $V_0 = F^{1/2}\bar{V}$. To define the uncertainty sets, we first compute the scaled sum of squared errors for each $i \in [n]$, $s_i^2 = \frac{1}{p-m-1}\sum_{l=1}^p (\mu_{(l),i} - \mu_{0,i} - V_{0,i}^\top f_{(l)})^2$. Let $c_J(\alpha)$ be the α-critical value of an F-distribution with J degrees of freedom, and let ν be the top-left entry of A^{-1}, where $A \in \mathbb{R}^{(p+1)\times(p+1)}$ is the Gram matrix of the vectors $\mathbf{1}_m, \{f_{(l)}\}_{l=1}^p$. Then we set $\gamma_i = \sqrt{\nu c_1(\alpha) s_i^2}$ for $i \in [n]$, which defines the uncertainty set for μ. The uncertainty set for V is chosen by randomly generating matrices P_k, and then scaling them appropriately so that the norm of each column i of $V - V_0$ is at most $\sqrt{m c_m(\alpha) s_i^2}$ for every $V \in \mathcal{V}$.

We set $p = 90$ and $\alpha = 0.95$, while varying $m \in \{3, 5, 7, 10, 15, 20, 25\}$ and $n \in \{100, 200, 300, 400, 500, 600, 700\}$. We fix the underlying dimension of the uncertainty set $\mathcal{V}$ to be $K = \min\{2m, 15\}$. We generate five instances for each combination of m and n.

The four approaches we test are our OFO-based approach from Section 4.1, our FO-based pessimization approach from Section 4.2 (see Theorem 4.1), the nominal oracle-based approach of [5] from Section 4.3, and the full pessimization approach of [31], which requires both a pessimization and an extended nominal feasiblity oracle. Since (28) is an instance of a robust quadratic program, the form for nominal and pessimization oracles can be derived from Section 5. One-dimensional line search using Brent's algorithm [14] was used to choose step sizes for each iteration of FO-based methods. An error tolerance of $\epsilon = 0.002$ is used in all instances.

Experiments are performed on a Linux machine with 2.8GHz processor and 64GB memory using Python v3.5.2. Whenever the nominal (extended nominal) oracles and pessimization oracles do not have closed form solutions, they are implemented in Gurobi v7.0.2. We use standard Gurobi tolerances and parameter choices. We employ the implementation of Brent's algorithm in Python's `scipy.optimize` package.

Figure 1 plots the average solve times in seconds against different n for each of the approaches, averaging across all m. As we expect, for low dimensions n, the oracle-based approaches solve the instances very quickly compared to our first-order based approaches. However, as n increases to $400, 500, 600, 700$, we see that the solution times of our first-order based approaches beat the nominal oracle approach, and become comparable to the full pessimization approach for $n = 600, 700$. In particular, we observe that the FO-based pessimization approach (see Theorem 4.1) solves faster when $n = 700$.

The dimension m influences the rank of the nominal matrix $V_0^\top V_0$ and controls the difficulty of the problems. Examining the results for different m further highlights the benefits of utilizing the first-order based approaches. Figure 2 plots average solve times for different m while fixing $n = 400, 500, 600, 700$. For the oracle-based methods, the solution times increase with m, while the solution times for first-order based methods remains relatively constant with m. For $m \geq 20$, we observe that our first-order based approaches significantly outperforms the oracle-based methods which require a nominal solver. Notice that, while we expect our OFO-based approach to outperform the FO-based pessimization approach due to the burden of solving an eigenvalue problem in each iteration for computing the pessimization oracle, our results indicate the opposite. This is because for small values of K, calling a pessimization oracle is faster than the line search performed in the FO-based noise update. However, we believe that as K increases, one-dimensional line search will become more efficient.
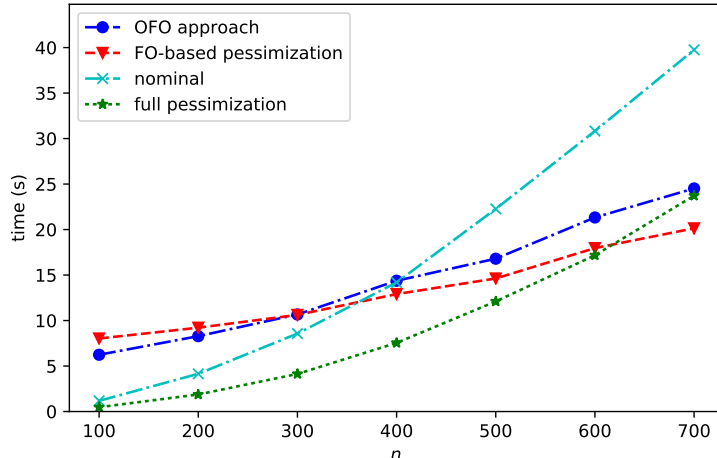
Figure 1: Average solve times (seconds) for different n.

Finally, we examine the number of iterations and cost per iteration of different approaches averaged across all instances in Table 4. We observe that, contrary to their theoretical iteration guarantees, the oracle-based approaches of [31, 5] need very few iterations to find a solution. However, as expected, the average time per iteration is significantly higher for these methods due to their reliance on full nominal optimization solvers. This further highlights the benefit of utilizing first-order methods for robust optimization when the deterministic version of the problem is already very expensive, and hence nominal oracles become expensive.

Table 4: Average number of iterations and average time per iteration for each approach.

	# iterations	seconds per iteration
first-order	961.487	0.015
FO-based pessimization	1009.054	0.013
nominal	3.708	4.841
full pessimization	1.975	4.875

7 Conclusion

In this paper, we advance the line of research in [13, 31, 5] that aims to solve robust optimization problems via iterative techniques, i.e., without transforming them into their equivalent robust counterparts. Thus far, the literature on iterative methods for RO has relied on more expensive nominal feasibility or pessimization oracles. However, in many applications of robust convex optimization, the original deterministic problem comes equipped with first-order oracles that provide gradient/subgradient information on the constraint functions f^i. In this paper, we present an efficient framework that can both work with cheap online first-order oracles and also capture the prior oracle-based approaches of [31] and [5]. We further show that working with these OFO oracles essentially does not increase the worst case theoretical bound on number of overall oracle calls, i.e., the worst case bound on number of main iterations of our approach is better than or comparable

(a) $n = 400$

(b) $n = 500$

(c) $n = 600$

(d) $n = 700$

Figure 2: Average solve times (seconds) for different n and m.

to the prior approaches. Moreover, when OFO oracles are utilized in our framework, the resulting overall arithmetic complexity including all of the basic operations in each iteration is remarkably cheaper than the prior approaches. The resulting framework is simple, easy-to-implement, flexible, and it can easily be customized to many applications. We demonstrate our framework via an illustrative robust QP example, where the most expensive operation in each iteration of our framework is a maximum eigenvalue computation. We further illustrate this with a preliminary numerical study on robust portfolio optimization problem.

Our framework is amenable to exploiting favorable structural properties of the functions f^i such as strong concavity, smoothness, etc., through which better convergence rates can be achieved. For example, when f^i are strongly concave in u^i, by exploiting this structural information and using a customization of the weighted regret online mirror descent for strongly convex functions, it is possible to achieve a better convergence rate of $O(1/\epsilon)$ in both our online first-order oracle setup and the nominal feasibility oracle framework of [5]. This then partially resolves/refines an open question stated in [5] for the lower bound on the number of iterations/calls needed in their nominal feasibility oracle based framework. However, it remains open whether $O(1/\epsilon^2)$ bound is tight when no further favorable structure is present in f^i or the tightness of $O(1/\epsilon)$ in the favorable case.

There are several other compelling avenues for future research. From a practical perspective, it is well-known, and also confirmed by our preliminary proof-of-concept computational experiments, that the computation of gradients/subgradients constitute a major bottleneck in the practical performance of FOMs. Thus, as a step to reduce the efforts involved in such computations, possible incorporation of stochastic [41, 36] and/or randomized FOMs [26, 10] working with stochastic subgradients into our framework is of great practical and theoretical interest. A critical assumption in our approach as well as others, e.g., see [5] and references therein, is that the domain X is convex. Removing the convexity requirement on the domain X will be an important theoretical development on its own. Besides, this will open up possibilities for more principled approaches to solving robust combinatorial optimization problems (see [11, 13]) where such a convexity assumption on X is not satisfied. Finally, another attractive research direction is develop analogous frameworks for multi-stage RO problems such as robust Markov decision processes (see [39, 24]).

Acknowledgments

References

[1] A. Beck, A. Ben-Tal, N. Guttmann-Beck, and L. Tetruashvili. The CoMirror algorithm for solving nonsmooth constrained convex problems. *Operations Research Letters*, 38(6):493 – 498, 2010.

[2] A. Ben-Tal, S. Bhadra, C. Bhattacharyya, and A. Nemirovski. Efficient methods for robust classification under uncertainty in kernel matrices. *Journal of Machine Learning Research*, 13(1):2923–2954, Oct. 2012.

[3] A. Ben-Tal, D. den Hertog, and J.-P. Vial. Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming*, 149(1-2):265–299, 2015.

[4] A. Ben-Tal, L. Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, 2009.

[5] A. Ben-Tal, E. Hazan, T. Koren, and S. Mannor. Oracle-based robust optimization via online learning. *Operations Research*, 63(3):628–638, 2015.

[6] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[7] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization*. Society for Industrial and Applied Mathematics, 2001.

[8] A. Ben-Tal and A. Nemirovski. Robust optimization – methodology and applications. *Mathematical Programming*, 92(3):453–480, 2002.

[9] A. Ben-Tal and A. Nemirovski. Selected topics in robust convex optimization. *Mathematical Programming*, 112(1):125–158, 2008.

[10] A. Ben-Tal and A. Nemirovski. On solving large-scale polynomial convex problems by randomized first-order algorithms. *Mathematics of Operations Research*, 40(2):474–494, 2015.

[11] D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[12] D. Bertsimas, I. Dunning, and M. Lubin. Reformulation versus cutting-planes for robust optimization. *Computational Management Science*, 13(2):195–217, Apr 2016.

[13] D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98(1-3):49–71, 2003.

[14] R. Brent. *Algorithms for Minimization Without Derivatives*. Dover Books on Mathematics. Dover Publications, 1973.

[15] G. Calafiore and M. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.

[16] C. Caramanis, S. Mannor, and H. Xu. Robust optimization in machine learning. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[17] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[18] O. Devolder, F. Glineur, and Y. Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1):37–75, Aug 2014.

[19] L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.

[20] D. Goldfarb and G. Iyengar. Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1):1–38, 2003.

[21] E. Hazan. The convex optimization approach to regret minimization. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[22] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

[23] N. Ho-Nguyen and F. Kılınç-Karzan. A second-order cone based approach for solving the trust-region subproblem and its variants. *SIAM Journal on Optimization*, 27(3):1485–1512, 2017.

[24] G. N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

[25] V. Jeyakumar and G. Y. Li. Trust-region problems with linear inequality constraints: Exact SDP relaxation, global optimality and robust optimization. *Mathematical Programming*, 147(1):171–206, 2013.

[26] A. Juditsky, F. Kılınç-Karzan, and A. Nemirovski. Randomized first order algorithms with applications to ℓ_1 minimization. *Mathematical Programming*, 142(1-2):269–310, 2013.

[27] A. Juditsky and A. Nemirovski. First-order methods for nonsmooth convex large-scale optimization, I: General purpose methods. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[28] A. Juditsky and A. Nemirovski. First-order methods for nonsmooth convex large-scale optimization, II: Utilizing problem's structure. In S. Sra, S. Nowozin, and S. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2012.

[29] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari. Regularization techniques for learning with matrices. *Journal of Machine Learning Research*, 13(Jun):1865–1890, 2012.

[30] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291 – 307, October 2005.

[31] A. Mutapcic and S. Boyd. Cutting-set methods for robust convex optimization with pessimizing oracles. *Optimization Methods and Software*, 24(3):381–406, June 2009.

[32] A. Nedić and A. Özdağlar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, 2009.

[33] A. Nedić and A. Özdağlar. Subgradient methods for saddle-point problems. *Journal of Optimization Theory and Applications*, 142(1):205–228, 2009.

[34] A. Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

[35] A. Nemirovski and D. B. Yudin. *Problem complexity and method efficiency in optimization.* Wiley-Interscience series in discrete mathematics. Wiley, Chichester, New York, 1983. A Wiley-Interscience publication.

[36] A. S. Nemirovski, A. B. Juditsky, G. Lan, and A. Shapiro. Stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

[37] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course.* Kluwer Academic Publishers, 2004.

[38] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.

[39] A. Nilim and L. El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[40] B. T. Polyak. A general method of solving extremum problems. *Soviet Mathematics Doklady*, 8(3):593–597, 1967.

[41] H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407, 1951.

[42] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.

[43] P. K. Shivaswamy, C. Bhattacharyya, and A. J. Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7(Jul):1283–1314, 2006.

[44] H. Xu, C. Caramanis, and S. Mannor. Robustness and regularization of support vector machines. *Journal of Machine Learning Research*, 10(Jul):1485–1510, 2009.

[45] H. Xu, C. Caramanis, and S. Mannor. Robust regression and Lasso. *IEEE Transactions on Information Theory*, 56(7):3561–3574, July 2010.

[46] H. Yu and M. J. Neely. A primal-dual type algorithm with the o(1/t) convergence rate for large scale constrained convex programs. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 1900–1905, Dec 2016.

[47] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pages 928–936, 2003.

A Convex-Concave Saddle Point Reformulation

The SP problem (7) based on the function $\Phi(x, u)$ which is not necessarily concave in u admits a convex-concave SP representation in a lifted space via perspective transformations. To present this reformulation, we start by defining the following sets with additional variables $y \in \mathbb{R}_+^m$ and new variables v^i for $i \in [m]$:

$$V^i = \left\{ [v^i; y^{(i)}] : \ 0 < y^{(i)} \le 1, \ \frac{v^i}{y^{(i)}} \in U^i \right\} \quad \forall i \in [m],$$

$$W = \left\{ w = [v^1; \ldots; v^m; y] : \ [v^i; y^{(i)}] \in \mathrm{cl}(V^i), \ i \in [m], \ \sum_{i=1}^m y^{(i)} = 1 \right\}.$$

Note that for all $i \in [m]$, $\mathrm{cl}(V^i) = V^i \cup \{[0; 0]\}$ because we assumed U^i to be closed sets. For the point $[v^i; y^{(i)}] = [0; 0]$, we set $y^{(i)} f^i \left(x, \frac{v^i}{y^{(i)}} \right) = 0$ for any $x \in X$. Note that setting $y^{(i)} f^i \left(x, \frac{v^i}{y^{(i)}} \right) = 0$ for $[v^i; y^{(i)}] = [0; 0]$ is well-defined as the continuation since from Assumption 2.1, $f^i(x, u^i)$ is continuous and finite-valued on U^i, and U^i is compact, so we deduce that $f^i(x, u^i)$ will be bounded on U^i.

We also define the function $\psi : X \times W \to \mathbb{R}$ as

$$\psi(x, w) = \psi(x, v, y) := \sum_{i=1}^m y^{(i)} f^i \left(x, \frac{v^i}{y^{(i)}} \right).$$

Lemma A.1. *For fixed $w \in W$, the function $\psi(x, w)$ is convex in x over X, and $\psi(x, w)$ is a concave function of w over W for any fixed x. Moreover, W is closed, and when U^i for $i \in [m]$ are convex, the sets V^i for $i \in [m]$ and W are all convex.*

Proof. For any $w = [v^1; \ldots; v^m; y]$, the function ψ is convex in x since in all of the nonzero terms in the summation over all $i \in [m]$ defining ψ, we have $y^{(i)} > 0$ and in each such nonzero term each function $f^i \left(x, \frac{v^i}{y^{(i)}} \right)$ is convex in x for the given $\frac{v^i}{y^{(i)}} \in U^i$ (see Assumption 2.1). In addition, for any given $x \in X$, the function ψ is jointly concave in v and y because it is written as a sum of the perspective functions of functions f^i which are concave in u^i (see Assumption 2.1).

The closedness of W is immediate, and the convexity of the sets V^i and W follows immediately from their definition and the convexity assumption on U^i. □

With these definitions and Lemma A.1, we observe that (7) is equivalent to evaluating the convex-concave SP problem defined by the function ψ over the convex domains X and W:

$$\inf_{x \in X} \sup_{w \in W} \psi(x, w) \le \epsilon \quad \text{or} \quad \inf_{x \in X} \sup_{w \in W} \psi(x, w) > 0. \tag{29}$$

We state this formally in the following lemma.

Lemma A.2. *For any $\epsilon > 0$ and $\bar{x} \in X$,*

$$\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \le \epsilon \quad \text{if and only if} \quad \sup_{w \in W} \psi(\bar{x}, w) \le \epsilon.$$

As a result,

$$\inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) \le \epsilon \quad \text{if and only if} \quad \inf_{x \in X} \sup_{w \in W} \psi(x, w) \le \epsilon.$$

Proof. Fix $\bar{x} \in X$ and $\epsilon > 0$. Suppose $\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon$; then for all $u^i \in U^i$, $i \in [m]$, we have $f^i(\bar{x}, u^i) \leq \epsilon$. Now consider any $w = [v^1; \ldots; v^m; y] \in W$. Then $0 \leq y^{(i)} \leq 1$ for all $i \in [m]$ and $\sum_{i=1}^m y^{(i)} = 1$. For all $i \in [m]$, define $u^i = \frac{v^i}{y^{(i)}} \in U^i$ whenever $y^{(i)} > 0$. Then $y^{(i)} f^i(\bar{x}, \frac{v^i}{y^{(i)}}) = y^{(i)} f^i(\bar{x}, u^i) \leq y^{(i)} \epsilon$ for $0 < y^{(i)} \leq 1$. In addition, when $y^{(i)} = 0$, because $w \in W$ we must have $v^i = 0$ and then by definition we have $y^{(i)} f^i(\bar{x}, \frac{v^i}{y^{(i)}}) = 0$. Therefore, from $\sum_{i=1}^m y^{(i)} = 1$, we deduce $\psi(\bar{x}, w) = \sum_{i=1}^m y^{(i)} f^i\left(\bar{x}, \frac{v^i}{y^{(i)}}\right) \leq \epsilon$ holds for any $w \in W$.

Now suppose that $\sup_{w \in W} \psi(\bar{x}, w) \leq \epsilon$ holds. Given $i \in [m]$ and $u^i \in U^i$, set w to have components $y^{(i)} = 1$, $v^i = u^i$, and $[v^j; y^{(j)}] = [0; 0]$ for $j \neq i$. Then $f^i(\bar{x}, u^i) = \psi(\bar{x}, w) \leq \epsilon$. Hence, $\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon$ follows. $\qquad \square$

Remark A.1. When $m = 1$, i.e., we have only one function $f^1(x, u^1)$ and only one uncertainty set U^1, hence $W = U^1$ and $\inf_{x \in X} \sup_{v \in W} \psi(x, w) = \inf_{x \in X} \sup_{u^1 \in U^1} f^1(x, u^1)$. Also, under Assumption 2.1, $\psi(x, w)$ is convex in x and concave in u^1. Thus, the preceding perspective transformation resulting in (29) directly generalizes this case of a convex-concave SP formulation for $m = 1$ discussed in Remark 3.1. $\qquad \blacksquare$

As a result, Lemma A.2 and Theorem 3.1 combined with any FOM that provides bounds on the saddle point gap $\epsilon_{\text{sad}}^\psi(\bar{x}, \bar{w})$ lead to an efficient way of verifying robust feasibility of (7) as follows:

Theorem A.1. *Suppose $\bar{x} \in X$, $\bar{w} \in W$, and $\tau \in (0,1)$ are such that $\epsilon_{\text{sad}}^\psi(\bar{x}, \bar{w}) \leq \tau\epsilon$. If $\psi(\bar{x}, \bar{w}) \leq (1 - \tau)\epsilon$, then $\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon$. If $\psi(\bar{x}, \bar{w}) > (1 - \tau)\epsilon$ and $\tau \leq \frac{1}{2}$, then $\inf_{x \in X} \max_{i \in [m]} \sup_{u^i \in U^i} f^i(x, u^i) > 0$.*

Proof. Suppose $\psi(\bar{x}, \bar{w}) \leq \tau\epsilon$. By Theorem 3.1, we have $\inf_{x \in X} \sup_{w \in W} \psi(x, w) \leq \sup_{w \in W} \psi(\bar{x}, w) \leq \epsilon$. By Lemma A.2, $\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) \leq \epsilon$ as well.

On the other hand, when $\psi(\bar{x}, \bar{w}) > (1 - \tau)\epsilon$ and $\tau \leq \frac{1}{2}$, Theorem 3.1 implies $\inf_{x \in X} \sup_{w \in W} \psi(x, w) \geq \inf_{x \in X} \psi(x, \bar{w}) > 0$. Then by Lemma A.2, $\max_{i \in [m]} \sup_{u^i \in U^i} f^i(\bar{x}, u^i) > 0$ follows. $\qquad \square$

Because of the existence of efficient FOMs to solve convex-concave SP problems, Theorem A.1 suggests a possible advantage of using the convex-concave SP problem given in (29). Nevertheless, working with the SP reformulation given by (29) in the extended space $X \times W$ presents a number of critical challenges. First, efficient FOMs associated with convex-concave SP problems often require computing prox operations or projections onto the domains X and W. Unfortunately, even if projection (or prox-mappings) onto U^i admits a closed form solution or an efficient procedure, it is unclear how to extend such projections onto W. Furthermore, while the perspective transformations involved in constructing the function ψ preserves certain desirable properties of the functions f^i, such as Lipschitz continuity and smoothness, the parameters associated with ψ are in general larger than those associated with the original functions f^i. Such parameters are critical for FOM convergence rates, and thus the FOMs when applied to solve (29) will have slower convergence rates.

To address the issues outlined above, in the main paper we discuss how to obtain robust feasibility/infeasibility certificates for the convex-nonconcave SP problem (7) directly, i.e., we work with the functions f^i and the sets U^i directly. This direct approach in particular allows us to take greater advantage of the structure of the original formulation such as the availability of efficient

projection (prox) computations over domains, and/or better parameters for smoothness, Lipschitz continuity, etc., of the functions.

B Supplementary Numerical Results

Below, we provide the exact numerical values corresponding to the data used to generate the figures in our numerical study in Section 6.

Table 5: Average solve time (seconds) of each approach for different n (Figure 1).

		Approach			
		OFO-based	FO-based pessimization	nominal	full pessimization
	100	6.24*	8.03†	1.18	0.47
	200	8.28‡	9.21‡	4.14	1.88
	300	10.67	10.62	8.57	4.13
n	400	14.38‡	12.91‡	14.19	7.55
	500	16.80	14.63	22.27	12.10
	600	21.33	17.96	30.83	17.17
	700	24.52	20.14	39.76	23.72

Table 6: Average solve time (seconds) of each approach for different m and n (Figure 2).

			m						
			3	5	7	10	15	20	25
		OFO-based	13.17§	12.21	14.78	14.03	14.73	16.23	15.24
	400	FO-based pessimization	12.47§	11.04	11.92	12.97	14.27	14.81	12.77
		nominal	2.16	4.23	6.77	11.31	15.82	24.86	31.77
		full pessimization	1.61	2.69	3.86	6.03	8.87	12.53	16.05
		OFO-based	15.96	15.55	15.88	17.22	17.41	17.52	18.04
	500	FO-based pessimization	13.04	12.26	14.08	17.09	16.28	14.55	15.07
		nominal	4.02	5.92	10.22	17.86	22.06	45.16	50.63
n		full pessimization	2.57	4.29	6.06	9.95	13.92	22.44	25.46
		OFO-based	18.62	19.52	20.97	21.25	20.86	23.46	24.62
	600	FO-based pessimization	15.75	16.48	18.11	16.55	17.16	20.89	20.80
		nominal	3.59	8.68	13.01	23.71	38.05	55.65	73.11
		full pessimization	2.91	6.16	8.68	13.25	21.25	30.91	37.05
		first-order	19.49	21.86	24.46	25.04	25.12	28.48	27.18
	700	pessimization	16.28	17.26	21.51	22.07	20.86	22.10	20.87
		nominal	6.38	12.53	14.24	29.81	48.79	77.03	89.51
		full pessimization	4.98	8.39	12.02	18.66	29.02	42.89	50.10

---

*Three instances out of 35 did not solve due to numerical issues.

†Two instances out of 35 did not solve due to numerical issues.

‡One instance out of 35 did not solve due to numerical issues.

§One instance out of five did not solve due to numerical issues.