
A Unifying Framework for Online Optimization with Long-Term Constraints

Matteo Castiglioni*
Politecnico di Milano

Andrea Celli†
Bocconi University

Alberto Marchesi*
Politecnico di Milano

Giulia Romano*
Politecnico di Milano

Nicola Gatti*
Politecnico di Milano

Abstract

We study online learning problems in which a decision maker has to take a sequence of decisions subject to m *long-term constraints*. The goal of the decision maker is to maximize their total reward, while at the same time achieving small cumulative constraints violation across the T rounds. We present the first *best-of-both-world* type algorithm for this general class of problems, with no-regret guarantees both in the case in which rewards and constraints are selected according to an unknown stochastic model, and in the case in which they are selected at each round by an adversary. Our algorithm is the first to provide guarantees in the adversarial setting with respect to the optimal fixed strategy that satisfies the long-term constraints. In particular, it guarantees a $\rho/(1 + \rho)$ fraction of the optimal reward and sublinear regret, where ρ is a feasibility parameter related to the existence of strictly feasible solutions. Our framework employs traditional regret minimizers as black-box components. Therefore, by instantiating it with an appropriate choice of regret minimizers it can handle the *full-feedback* as well as the *bandit-feedback* setting. Moreover, it allows the decision maker to seamlessly handle scenarios with non-convex rewards and constraints. We show how our framework can be applied in the context of budget-management mechanisms for repeated auctions in order to guarantee long-term constraints that are not *packing* (e.g., ROI constraints).

1 Introduction

We study online learning problems where a decision maker takes decisions over T rounds. At each round t , the decision $\mathbf{x}_t \in \mathcal{X}$ is chosen before observing a reward function f_t together with a set of m *time-varying* constraint functions g_t . The decision maker is allowed to make decisions that are *not* feasible, provided that the overall sequence of decisions obeys the *long-term constraints* $\sum_{t=1}^T g_t(\mathbf{x}_t) \leq \mathbf{0}$, up to a small cumulative violation across the T rounds. The problem becomes that of finding a sequence of decisions \mathbf{x}_t which guarantees a reward close to that of the best fixed decision in hindsight while satisfying long-term constraints. This type of framework was first proposed by Mannor et al. [36], and it has numerous applications ranging from wireless communication [36] and multi-objective online classification [13], to *safe* online learning [4].

Mannor et al. [36] show that guaranteeing sublinear regret and sublinear cumulative constraints violation is impossible even when f_t and g_t are simple linear functions. Therefore, previous works either focus on the case in which constraints are generated i.i.d. according to some unknown stochastic

*DEIB, Politecnico di Milano, {matteo.castiglioni, alberto.marchesi, giulia.romano, nicola.gatti}@polimi.it.
† Computing Sciences Department, Bocconi University, andrea.celli2@unibocconi.it.

Algorithm	Constr.	Non-convex f_t and g_t	Bound — constant ρ Reward	Bound — constant ρ Violation	Bound — arbitrary ρ Reward	Bound — arbitrary ρ Violation
Yu et al. [44]	STOC	\times	$\text{OPT} - \tilde{O}(T^{1/2})$	$\tilde{O}(T^{1/2})$	—	—
Ours	STOC	\checkmark	$\text{OPT} - \tilde{O}(T^{1/2})$	$\tilde{O}(T^{1/2})$	$\text{OPT} - \tilde{O}(T^{3/4})$	$\tilde{O}(T^{3/4})$
	ADV	\checkmark	$\frac{\rho}{1+\rho}\text{OPT} - \tilde{O}(T^{1/2})$	$\tilde{O}(T^{1/2})$	—	—

Table 1: Comparison between the performance of our algorithm and previous work using the same baseline as ours. Bounds for settings that were *not* previously tractable are highlighted in gray. OPT is the reward of the baseline.

model, without providing any guarantees for the adversarial case, or provide results for adversarially-generated constraints under some strong assumptions on the structure of the problem or using a weaker baseline (a detailed discussion of related works can be found in Appendix A). A few examples in the latter case are [16, 21, 40, 42]. In the former setting (*i.e.*, stochastic constraints), Wei et al. [41] consider a weaker baseline that is feasible for each constraint g_t , going against the basic idea of long-term constraints. A notable exception is the work by Yu et al. [44], who employ the same baseline as ours, and provide an upper bound of $\tilde{O}(T^{1/2})$ for both regret and constraints violation (see Table 1). We also mention that there are some works studying the problem in which constraints are *static* (see, *e.g.*, [31, 35, 43, 45]), or focus on specific types of constraints, such as *knapsack constraints* [7, 30]. Our framework differs from those works as we deal with *arbitrary* and *time-varying* constraints. Moreover, it also extends the *online convex optimization* framework introduced by Zinkevich [46] by allowing for general non-convex loss functions f_t , arbitrary feasibility sets \mathcal{X} , and arbitrary time-varying long-term constraints.

1.1 Original contributions

Given the negative result by Mannor et al. [36], a natural question is what kind of guarantees we can reach in the adversarial setting, when adopting the standard baseline of the best fixed decision in hindsight satisfying (in expectation) the long-term constraints. We provide the first positive result going in this direction, by designing a no- α -regret algorithm that guarantees a sublinear cumulative constraints violation. Moreover, we make a step forward in the line of work initiated by Bubeck and Slivkins [14], by showing that our algorithm is also the first *best-of-both-worlds* algorithm for problems with arbitrary long-term constraints. This allows our algorithm to guarantee good worst-case performance (adversarial case), while being able to exploit well-behaved problem instances (stochastic case). The only assumption which we require is the existence of a decision that is strictly feasible with respect to the sequence of constraints. We denote by ρ the “margin” by which this decision is strictly feasible (see Section 2 for a definition). At the same time, we show that even without this assumption, we can recover sublinear regret and violation with stochastic constraints.

Previous work usually assumes that ρ is a given *constant*. In that case, our algorithm matches the guarantees by Yu et al. [44] when constraints are generated i.i.d. according to an unknown distribution, and has no- α -regret with $\alpha = \rho/(1 + \rho)$ in the adversarial case (see Table 1). Our algorithm only requires a lower bound on the real value of the feasibility parameter ρ . In the stochastic case, the lower bound may even be unknown, and the algorithm can efficiently estimate it from data. Moreover, we argue that if ρ is allowed to depend on T and take arbitrarily small values, then there are certain values ($\rho \leq T^{-1/4}$), for which any regret bound depending on $1/\rho$ would be useless (*i.e.*, *not* sublinear in T , see Section 3). This setting is usually overlooked by previous work, which assumes ρ to be a given constant. We show that, in the case of an arbitrary feasibility parameter ρ , in the stochastic setting our algorithm guarantees an upper bound of $\tilde{O}(T^{3/4})$ for regret and cumulative constraints violation.

Our framework employs traditional regret minimizers as black-box components. Therefore, by instantiating it with an appropriate choice of regret minimizers it can handle *full-feedback* as well as *bandit-feedback* settings. In the former case, after playing x_t , the decision maker gets to observe f_t and g_t , while in the latter case only the realized values $f_t(x_t)$ and $g_t(x_t)$ are observed. Moreover, this allows the decision maker to seamlessly handle scenarios with non-convex reward and constraints, by employing a suitable regret minimizer for non-convex losses (see, *e.g.*, [39]). Our algorithm is based on a two-stage approach in which *primal* and *dual* players interact through *Lagrangian games*.

In the first (*play*) phase, the primal player tries to balance out the maximization of their rewards with constraints violation. In the second (*recovery*) phase, the primal player only makes “safe decisions” to avoid violating constraints too much. It is possible to prove that, in the case of stochastic rewards and constraints, the algorithm never enters phase two. This property is particularly relevant for budget-pacing mechanisms in repeated auctions, since it is related to how budget is allocated. Our framework can also be instantiated to perform budget allocation subject to constraints that were previously *not* tractable by traditional mechanisms, such as ROI constraints [11, 22].

2 Preliminaries

The decision maker has a non-empty set of available strategies \mathcal{X} (this set may be non-convex, integral, and even non-compact). In each round $t \in [T]$,² the decision maker first chooses $\mathbf{x}_t \in \mathcal{X}$, and the environment selects a reward function $f_t : \mathcal{X} \rightarrow [0, 1]$ and a constraint function $g_t : \mathcal{X} \rightarrow [-1, 1]^m$ conditioned on the past history of play up to time $t - 1$ (*i.e.*, the environment chooses f_t and g_t without knowledge of \mathbf{x}_t). Notice that both f_t and g_t need *not* be convex. The latter specifies a set of m constraints of the form $g_t(\mathbf{x}) \leq \mathbf{0}$, with $g_{t,i}(\mathbf{x}) \leq 0$ denoting the i -th constraint.³ In the following, we denote as \mathcal{F} , respectively \mathcal{G} , the set of all the possible f_t , respectively g_t , functions (*e.g.*, \mathcal{F} and \mathcal{G} may contain all the Lipschitz-continuous functions defined over \mathcal{X}). At each round $t \in [T]$, the decision maker can condition their decision on prior feedbacks and on the sequence of prior decisions $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$, but no information about future rewards and constraint functions is available.

2.1 Strong duality through strategy mixtures

In the following, we define the optimization problem (Problem $\text{LP}_{f,g}$) which is used to define the baselines against which we compare the performances of the decision maker. Such a problem involves probabilistic mixtures of strategies in \mathcal{X} , which are crucial in order to recover strong duality.⁴

First, we introduce the set of probability measures on the Borel sets of \mathcal{X} . We refer to such a set as the set of *strategy mixtures*, denoted as Ξ . We endow \mathcal{X} with the Lebesgue σ -algebra, and we assume that all the functions in \mathcal{F} and \mathcal{G} are measurable with respect to every probability measure $\xi \in \Xi$. This ensures that the various expectations taken are well-defined, since the functions are assumed to be bounded above, and they are therefore integrable. In the following, for ease of presentation and with a slight abuse of notation, whenever we write a $\xi \in \Xi$ in place of an $\mathbf{x} \in \mathcal{X}$, we mean that we are taking the expectation with respect to the probability measure ξ . For instance, given $f \in \mathcal{F}$ and $g \in \mathcal{G}$, we have that $f(\xi) = \mathbb{E}_{\mathbf{x} \sim \xi} f(\mathbf{x})$ and $g(\xi) = \mathbb{E}_{\mathbf{x} \sim \xi} g(\mathbf{x})$.

Then, given two functions $f \in \mathcal{F}$ and $g \in \mathcal{G}$, we define the following optimization problem, which chooses the strategy mixture $\xi \in \Xi$ that maximizes the expected reward encoded by f , while guaranteeing that the constraints encoded by g are satisfied in expectation.

$$\text{OPT}_{f,g} := \begin{cases} \sup_{\xi \in \Xi} & f(\xi) & \text{s.t.} \\ & g(\xi) \leq 0. \end{cases} \quad (\text{LP}_{f,g})$$

We denote by $d_g \in [-1, 1]$ the largest possible value for which there exists a strategy mixture $\xi \in \Xi$ satisfying the constraints $g(\xi) \leq 0$ by a margin of at least d_g . Formally,

$$d_g := \sup_{\xi \in \Xi} \min_{i \in [m]} -g_i(\xi). \quad (1)$$

In order to ensure that $\text{OPT}_{f,g}$ is always well defined, we assume that it is always the case that $d_g \geq 0$. Notice that, if $d_g > 0$, then Problem $\text{LP}_{f,g}$ satisfies Slater’s condition.

In the following, we prove some auxiliary results relating to Problem $\text{LP}_{f,g}$ that will be useful in the rest of the paper. First, we introduce a Lagrangian relaxation of the problem.

²In this work, we denote by $[x]$ the set $\{1, \dots, x\}$ of the first x natural numbers.

³Focusing on the case $g_t(\mathbf{x}) \leq \mathbf{0}$ is w.l.o.g. since any set of constraints can be represented in such a form.

⁴The optimal fixed strategy mixture provides an arguably stronger baseline than the optimal fixed strategy. In stochastic settings, this baseline is related to the best dynamic policy. In particular, if we consider the case in which the observed functions are defined as the average of functions f_t and g_t across the T rounds, then the optimal mixture provides the same utility as the best dynamic policy (see [7] for a similar result).

Definition 2.1 (Lagrangian Function). Given two arbitrary functions $f \in \mathcal{F}$ and $g \in \mathcal{G}$, the *Lagrangian function* $\mathcal{L}_{f,g} : \Xi \times \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}$ of Problem $\text{LP}_{f,g}$ is defined as

$$\mathcal{L}_{f,g}(\xi, \lambda) := f(\xi) - \langle \lambda, g(\xi) \rangle.$$

If Problem $\text{LP}_{f,g}$ satisfies Slater's condition, then Theorem 1 of Chapter 8.3 in [34] readily gives us that strong duality holds even if f and g are arbitrary non-convex functions. Formally:

Corollary 2.2. *Given $f \in \mathcal{F}$ and $g \in \mathcal{G}$ such that $d_g > 0$, it holds*

$$\sup_{\xi \in \Xi} \inf_{\lambda \in \mathbb{R}_{\geq 0}^m} \mathcal{L}_{f,g}(\xi, \lambda) = \inf_{\lambda \in \mathbb{R}_{\geq 0}^m} \sup_{\xi \in \Xi} \mathcal{L}_{f,g}(\xi, \lambda) = \text{OPT}_{f,g}.$$

Next, we show that, if $d_g > 0$, then strong duality holds even when we restrict the admissible dual vectors $\lambda \in \mathbb{R}_{\geq 0}^m$ to the set \mathcal{D}_{d_g} , where, for any $q \in \mathbb{R}_{>0}$, we let $\mathcal{D}_q := \{\lambda \in \mathbb{R}_{\geq 0}^m : \|\lambda\|_1 \leq 1/q\}$ (omitted proofs can be found in Appendix B).

Theorem 2.3. *Given $f \in \mathcal{F}$ and $g \in \mathcal{G}$ such that $d_g > 0$, it holds*

$$\sup_{\xi \in \Xi} \inf_{\lambda \in \mathcal{D}_{d_g}} \mathcal{L}_{f,g}(\xi, \lambda) = \inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}_{f,g}(\xi, \lambda) = \text{OPT}_{f,g}.$$

2.2 Stochastic vs. adversarial: baselines and feasibility

We consider several settings, differing in how functions f_t and g_t are selected, either *stochastically* or *adversarially*. We say that functions f_t (respectively g_t) are selected stochastically, when they are independently drawn according to a given probability measure $\mu_{\mathcal{F}}$ over \mathcal{F} (respectively $\mu_{\mathcal{G}}$ over \mathcal{G}). Instead, we say that functions f_t (respectively g_t) are selected adversarially if each f_t (respectively g_t) is chosen by an adversary based on the sequence of prior decisions, namely $\mathbf{x}_1, \dots, \mathbf{x}_{t-1}$.

Consistently with previous work (see, e.g., [36]), we compare the performance of the decision maker (in terms of reward cumulated over the T rounds) against the baseline $T \text{OPT}_{\bar{f}, \bar{g}}$ (as defined by Problem $\text{LP}_{\bar{f}, \bar{g}}$), where \bar{f} and \bar{g} are suitably-defined functions. In particular:

- When functions f_t , respectively g_t , are selected stochastically, then we define function \bar{f} , respectively \bar{g} , so that $\bar{f}(\mathbf{x}) := \mathbb{E}_{f \sim \mu_{\mathcal{F}}}[f(\mathbf{x})]$, respectively $\bar{g}(\mathbf{x}) := \mathbb{E}_{g \sim \mu_{\mathcal{G}}}[g(\mathbf{x})]$.
- When functions f_t , respectively g_t , are selected adversarially, then we define function \bar{f} , respectively \bar{g} , so that $\bar{f}(\mathbf{x}) := \frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x})$, respectively $\bar{g}(\mathbf{x}) := \frac{1}{T} \sum_{t=1}^T g_t(\mathbf{x})$.

Intuitively, in the stochastic case, the baseline is instantiated with an expectation of functions taken with respect to the probability measure $\mu_{\mathcal{F}}$ (respectively $\mu_{\mathcal{G}}$). Instead, in the adversarial case, the baseline uses the average of functions f_t (respectively g_t) observed over the T rounds.

Let us remark that, when the set \mathcal{X} is compact convex and functions f_t and g_t are convex, then Problem $\text{LP}_{\bar{f}, \bar{g}}$ defining our baselines can be equivalently re-written by using strategies $\mathbf{x} \in \mathcal{X}$ rather than strategy mixtures $\xi \in \Xi$, since there always exists an optimal solution to Problem $\text{LP}_{\bar{f}, \bar{g}}$ that places all the probability mass on a single strategy.

Our goal is to design online algorithms for the decision maker that output a sequence of decisions $\mathbf{x}_1, \dots, \mathbf{x}_T$ such that both the *cumulative regret* with respect to the performance of the baseline, defined as $R^T := T \text{OPT}_{\bar{f}, \bar{g}} - \sum_{t=1}^T f_t(\mathbf{x}_t)$, and the *cumulative constraints violation*, defined as $V^T := \max_{i \in [m]} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t)$, grow sublinearly in the number of rounds T .

In conclusion, we introduce a problem-specific parameter that is strictly related to the feasibility of Problem $\text{LP}_{\bar{f}, \bar{g}}$. We call it the *feasibility parameter* $\rho \in \mathbb{R}$, which is formally defined as follows:

- When functions g_t are selected stochastically, $\rho := \sup_{\xi \in \Xi} \min_{i \in [m]} -\bar{g}_i(\xi)$.
- When functions g_t are selected adversarially, $\rho := \sup_{\xi \in \Xi} \min_{t \in [T]} \min_{i \in [m]} -g_{t,i}(\xi)$.

Intuitively, in the stochastic case, ρ is equal to $d_{\bar{g}}$, while in the adversarial case it is computed similarly, but considering the worst case with respect to the functions g_t observed at each round t . Notice that, when $\rho > 0$, Slater's condition is satisfied for Problem $\text{LP}_{\bar{f}, \bar{g}}$.

In the following, we denote by $\xi^* \in \Xi$ a strategy mixture that is optimal for Problem $\text{LP}_{\bar{f}, \bar{g}}$. Moreover, we always assume that functions f_t and g_t are such that Problem $\text{LP}_{\bar{f}, \bar{g}}$ is feasible, and we let $\xi^\circ \in \Xi$ be the *feasible* strategy mixture that is optimal for the problem defining ρ .⁵

2.3 Regret minimizers

A *regret minimizer* (RM) for a set \mathcal{W} is an abstract model for a decision maker repeatedly interacting with an environment. At each t , a RM performs two operations: (i) `NEXTELEMENT()`, which outputs an element $w_t \in \mathcal{W}$; and (ii) `OBSERVEUTILITY(\cdot)`, which updates the internal state of the RM using the feedback received from the environment. This is either a utility function $u_t : \mathcal{W} \rightarrow [a, b]$ having range $[a, b] \subseteq \mathbb{R}$ (*full feedback*) or only the value $u_t(w_t)$ (*bandit feedback*), with u_t possibly depending adversarially on w_1, \dots, w_{t-1} . The objective of the RM is to output a sequence w_1, \dots, w_T of points in \mathcal{W} so that its *cumulative regret*, defined as $\sup_{w \in \mathcal{W}} \sum_{t=1}^T (u_t(w) - u_t(w_t))$, grows asymptotically sublinearly in T . See [19] for a review of the various RMs available in the literature.

For the ease of presentation, we introduce the concept of *regret minimizer constructor*, which is a procedure, say `INIT($\mathcal{W}, [a, b], \eta$)`, that builds a RM on the basis of the three parameters given as input. In particular, the procedure returns a RM instantiated for the set \mathcal{W} , working with utility functions having range $[a, b]$, and such that its cumulative regret is guaranteed to grow sublinearly in the time horizon T with probability at least $1 - \eta$.

3 A unifying meta-algorithm

In this section, we present our meta-algorithm. Its core idea is to instantiate suitable pairs of RMs, where one is working in the domain \mathcal{X} of primal variables and the other in a suitable subset of the domain \mathbb{R}_+^m of dual variables. At each round $t \in [T]$, the algorithm makes the RMs “play” against each other in a *Lagrangian game*, where the utility functions observed by them are related to the Lagrangian function $\mathcal{L}_{f_t, g_t}(x, \lambda)$ of Problem LP_{f_t, g_t} .⁶

Algorithm 1 provides the pseudo-code of the meta-algorithm, which takes as input: the total number of rounds T , a failure probability $\delta \in (0, 1)$ such that the guarantees provided by the algorithm hold with probability at least $1 - \delta$, and a lower bound $\hat{\rho} \geq 0$ on the value of the feasibility parameter ρ .

Algorithm description. The algorithm works in two phases. In the first one, called *play phase*, the algorithm builds a primal RM, called \mathcal{R}_1^P , working in the primal domain \mathcal{X} and a dual RM, called \mathcal{R}_1^D , operating on the subset $\mathcal{D}_{\hat{\rho}}$ of the dual domain \mathbb{R}_+^m , where $\hat{\rho}$ is set in Line 1. The algorithm makes the two RMs playing against each other (see the call `LAGRANGIANGAME($\mathcal{R}_1^P, \mathcal{R}_1^D, 1$)`) until either the cumulative violation V^t incurred by the algorithm exceeds a given threshold (see Line 4, where $M_{\hat{\rho}}$ is defined in Equation (2)) or round T is reached. Then, in the second phase, called *recovery phase*, the algorithm constructs a new pair of primal, dual RMs, with the latter working on the $(m - 1)$ -dimensional simplex Δ_m . The recovery phase uses the remaining rounds to make these new RMs play against each other, with the primal RM observing modified utility functions that do *not* account for functions f_t (see the call `LAGRANGIANGAME($\mathcal{R}_1^P, \mathcal{R}_1^D, 0$)`). Intuitively, this is needed in order to ensure that the algorithm plays strategies x_t that satisfy the constraints, thus balancing out the cumulative constraint violation accumulated in the first phase. The pseudo-code describing one “play” between two RMs, called \mathcal{R}^P and \mathcal{R}^D , is defined by the sub-procedure `LAGRANGIANGAME($\mathcal{R}^P, \mathcal{R}^D, v$)` in Algorithm 2. The additional parameter $v \in \{0, 1\}$ is used to control the feedback fed into the primal RM \mathcal{R}^P ; specifically, if $v = 1$, then \mathcal{R}^P observes a utility function that also accounts for f_t (play phase), otherwise, if $v = 0$, the observed utility function only accounts for the term depending on g_t (recovery phase).

Regret minimizer constructors. Algorithm 1 also needs access to two suitably-defined regret minimizer constructors, namely `INITP($\mathcal{W}, [a, b], \eta$)` and `INITD($\mathcal{W}, [a, b], \eta$)`, where the former is used

⁵Notice that ξ^* and ξ° may *not* be well defined in all the cases in which the problem that defines them does *not* admit a maximum. Nevertheless, in such cases, we assume that ξ^* (or ξ°) is a strategy mixture arbitrarily “close” to the supremum, so that all of our results continue to hold up to negligible additive approximations that are dominated by other approximation factors, and we can safely ignore them for ease of exposition.

⁶The idea of having pairs of primal, dual RMs playing a Lagrangian game was originally introduced by Immorlica et al. [30], restricted to the case of knapsack constraints.

Algorithm 1 META-ALGORITHM($T, \delta, \hat{\rho}$)

```

1:  $\tilde{\rho} \leftarrow \max\{\hat{\rho}/2, T^{-1/4}\}, \eta \leftarrow \delta/3, t \leftarrow 1$ 
    $\triangleright$  Phase I: Play
2:  $\mathcal{R}_I^P \leftarrow \text{INIT}^P(\mathcal{X}, [-1/\tilde{\rho}, 1+1/\tilde{\rho}], \eta)$ 
3:  $\mathcal{R}_I^D \leftarrow \text{INIT}^D(\mathcal{D}_{\tilde{\rho}}, [-1/\tilde{\rho}, 1/\tilde{\rho}], 0)$ 
4: while  $V^t \leq (T-t)\tilde{\rho} + M_{\tilde{\rho}} - 1 \wedge t \leq T$  do
5:    $\mathbf{x}_t \leftarrow \text{LAGRANGIANGAME}(\mathcal{R}_I^P, \mathcal{R}_I^D, 1)$ 
6:    $t \leftarrow t+1$ 
7:  $T_1 \leftarrow t-1$ 
    $\triangleright$  Phase II: Recovery
8:  $\mathcal{R}_{II}^P \leftarrow \text{INIT}^P(\mathcal{X}, [-1, 1], \eta)$ 
9:  $\mathcal{R}_{II}^D \leftarrow \text{INIT}^D(\Delta_m, [-1, 1], 0)$ 
10: while  $t \leq T$  do
11:    $\mathbf{x}_t \leftarrow \text{LAGRANGIANGAME}(\mathcal{R}_{II}^P, \mathcal{R}_{II}^D, 0)$ 
12:    $t \leftarrow t+1$ 

```

Algorithm 2 LAGRANGIANGAME($\mathcal{R}^P, \mathcal{R}^D, v$)

```

1:  $\mathbf{x}_t \leftarrow \mathcal{R}^P.\text{NEXTELEMENT}()$ 
2:  $\boldsymbol{\lambda}_t \leftarrow \mathcal{R}^D.\text{NEXTELEMENT}()$ 
3: Play  $\mathbf{x}_t$  and get  $f_t$  and  $g_t$   $\triangleright$  Full f.
   Play  $\mathbf{x}_t$  and get  $f_t(\mathbf{x}_t)$  and  $g_t(\mathbf{x}_t)$   $\triangleright$  Bandit f.
    $\triangleright$  Primal RM update
4: Let  $u_t^P : \mathbf{x} \mapsto v f_t(\mathbf{x}) - \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}) \rangle$   $\triangleright$  Full f.
    $u_t^P(\mathbf{x}_t) \leftarrow v f_t(\mathbf{x}_t) - \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle$   $\triangleright$  Bandit f.
5:  $\mathcal{R}^P.\text{OBSERVEUTILITY}(u_t^P)$   $\triangleright$  Full f.
    $\mathcal{R}^P.\text{OBSERVEUTILITY}(u_t^P(\mathbf{x}_t))$   $\triangleright$  Bandit f.
    $\triangleright$  Dual RM update
6: Let  $u_t^D : \boldsymbol{\lambda} \mapsto \langle \boldsymbol{\lambda}, g_t(\mathbf{x}) \rangle$ 
7:  $\mathcal{R}^D.\text{OBSERVEUTILITY}(u_t^D)$ 

```

to build RMs working in the primal domain and the latter for those operating on the dual one. Their actual implementation depends on the specific problem at hand. In the following, we let $\mathcal{E}_{t,\eta}^P$ be the regret upper bound (on $t \in [T]$ rounds) for primal RMs \mathcal{R}^P dealing with utility functions having range $[0, 1]$, as returned by the call $\text{INIT}^P(\mathcal{X}, [0, 1], \eta)$. Notice that, when the range is $[a, b]$, the same RM can be adopted by first normalizing utility values, so that the resulting regret upper bound is $(b-a)\mathcal{E}_{t,\eta}^P$. As for dual RMs \mathcal{R}^D , we let \mathcal{E}_t^D be the regret upper bound (on $t \in [T]$ rounds) provided by the RM defined for the set Δ_m , while $\mathcal{E}_t^D/\tilde{\rho}$ is the upper bound for the dual RM instantiated on the set $\mathcal{D}_{\tilde{\rho}}$. Notice that, since dual RMs always have full feedback, we can safely assume that the regret bounds \mathcal{E}_t^D hold deterministically. We also assume that RMs provide bounds that increase with the number of rounds, *i.e.*, such that $\mathcal{E}_{t,\eta}^P \leq \mathcal{E}_{t',\eta}^P$ and $\mathcal{E}_t^D \leq \mathcal{E}_{t'}^D$ for all $t \leq t'$.

How to construct RMs. INIT^D can be implemented by using *online mirror descent* (OMD) with domain Δ_m (or \mathcal{D}_1) and a negative entropy regularizer. Since the utility function u_t^D is linear in $\boldsymbol{\lambda}$, we get a regret bound for the primal RM of $\mathcal{E}_T^D = O(\sqrt{T \log(m)})$ (see, *e.g.*, [12, 37]). The design of INIT^P depends on the structure of \mathcal{X} and functions f_t and g_t . For instance, in convex settings with full feedback we can employ OMD [29], while with bandit feedback we can use [15]. Finally, for non-convex functions we can employ, *e.g.*, the RMs in [39]. All these RMs guarantee $\tilde{O}(\sqrt{T})$ regret.

How to get away with no knowledge of ρ . In Section 7, we show that a lower bound $\hat{\rho}$ is *not* necessary when functions g_t are selected stochastically. Indeed, it is sufficient to add a preliminary phase to Algorithm 1, which is used to infer a suitable lower bound on ρ from experience. In order to do this, only \sqrt{T} rounds are needed, so that the bounds on the cumulative regret and the cumulative constraint violation achieved by the algorithm are *not* compromised. When functions g_t are chosen adversarially, as we show in Section 7, it is impossible to compute a lower bound on the feasibility parameter ρ by only adding a preliminary phase to Algorithm 1 which uses \sqrt{T} rounds.

Remark 3.1 (Dependence on the lower bound $\hat{\rho}$). *Algorithm 1 can take as input any $\hat{\rho} \geq 0$. However, since our regret bounds include a factor $1/\tilde{\rho}$, by choosing the trivial lower bound $\hat{\rho} = 0$ we incur in a regret of $\tilde{O}(\sqrt{T}/\hat{\rho}) = \tilde{O}(T^{3/4})$. In order to obtain optimal bounds, we would like to have $\tilde{\rho} = \Omega(\rho)$.*

Remark 3.2 (Dependence on the feasibility parameter ρ). *In our analysis, we include the dependence on the feasibility parameter ρ in the regret and constraint violation bounds achieved by the algorithm. As customary, the goal is devising bounds of the form $\text{poly}(\text{instance}) \cdot h(T)$, where the first term is a polynomial function of the parameters defining the problem instance, and $h(T) = o(T)$. Therefore, we cannot include a factor $1/\rho$ in the regret bounds if ρ can be arbitrarily small. Even from a practical standpoint, when ρ is too small, a $1/\rho$ regret bound is too large to be significant. For those reasons, we set $\tilde{\rho}$ in Algorithm 1 to be the maximum between the feasibility parameter lower bound $\hat{\rho}$ and $T^{-1/4}$. Notice that the value $T^{-1/4}$ has been carefully chosen so as to minimize the maximum between the cumulative regret and the cumulative constraint violation when the lower bound on the feasibility parameter $\hat{\rho}$ is too small.*

4 Analysis with stochastic constraints and adversarial rewards

We start by analyzing the performance of our meta-algorithm (Algorithm 1) when the reward and constraint functions are selected stochastically and adversarially, respectively.

Given $t \in [T]$ and $\eta \in (0, 1)$, we let $\mathcal{E}_{t,\eta} := \sqrt{8t \log(18mt^2/\eta)}$ be the value bounding differences between expectations and empirical means of constraint functions, obtained by applying the Azuma-Hoeffding inequality, and holding with probability at least $1 - \eta$. Given $\gamma \in (0, 1)$, we also let

$$M_\gamma := \frac{2}{\gamma} \sqrt{T} + \left(2 + \frac{3}{\gamma}\right) \mathcal{E}_{t,\eta} + \left(1 + \frac{2}{\gamma}\right) \mathcal{E}_{t,\eta}^{\mathbb{P}} + \frac{1}{\gamma} \mathcal{E}_t^{\mathbb{D}}, \quad (2)$$

which is a recurring term related to the maximum violation that Algorithm 1 accepts in play phase.

First, we introduce a useful event \mathbf{E} that encompasses all the cases in which Algorithm 1 successfully terminates. Then, Lemma 4.2 shows that such an event holds with probability at least $1 - \delta$. In particular, \mathbf{E} holds when the regret bounds of $\mathcal{R}_I^{\mathbb{P}}$ and $\mathcal{R}_{II}^{\mathbb{P}}$ hold, and, additionally, the differences between expectations and empirical means of constraint functions are bounded as desired.

Definition 4.1. We denote with \mathbf{E} the event in which Algorithm 1 satisfies the following conditions (recall that $\eta = \delta/3$): (i) the regret incurred by $\mathcal{R}_I^{\mathbb{P}}$ after T_1 rounds is upper bounded by $\mathcal{E}_{T_1,\eta}^{\mathbb{P}}$; (ii) the regret cumulated by $\mathcal{R}_{II}^{\mathbb{P}}$ after the remaining $T - T_1$ rounds is upper bounded by $\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}}$; and (iii) for every pair of rounds $t, t' \in [T] : t \leq t'$ and resource $i \in [m]$ it holds:

- $\left| \sum_{\tau=t}^{t'} g_{\tau,i}(\mathbf{x}_\tau) - \sum_{\tau=t}^{t'} \bar{g}_i(\mathbf{x}_\tau) \right| \leq \mathcal{E}_{t'-t,\eta}$,
- $\left| \sum_{\tau=t}^{t'} \lambda_\tau g_{\tau,i}(\mathbf{x}_\tau) - \sum_{\tau=t}^{t'} \lambda_\tau \bar{g}_i(\mathbf{x}_\tau) \right| \leq \mathcal{E}_{t'-t,\eta} \max_{\tau \in [T]: t \leq \tau \leq t'} \|\lambda_\tau\|_1$,
- $\left| \sum_{\tau=t}^{t'} g_{\tau,i}(\xi) - \sum_{\tau=t}^{t'} \bar{g}_i(\xi) \right| \leq \mathcal{E}_{t'-t,\eta}$ for $\xi \in \{\xi^*, \xi^\circ\}$,
- $\left| \sum_{\tau=t}^{t'} \lambda_\tau g_{\tau,i}(\xi) - \sum_{\tau=t}^{t'} \lambda_\tau \bar{g}_i(\xi) \right| \leq \mathcal{E}_{t'-t,\eta} \max_{\tau \in [T]: t \leq \tau \leq t'} \|\lambda_\tau\|_1$ for $\xi \in \{\xi^*, \xi^\circ\}$.

Lemma 4.2. *After running Algorithm 1, the event \mathbf{E} holds with probability at least $1 - \delta$.*

Next, we lower bound the cumulative reward obtained by Algorithm 1 during the play phase. Intuitively, we show that, if the cumulative constraints violation is large, then the decisions \mathbf{x}_t in the first T_1 rounds provide a per-round reward much higher than that achievable by ξ^* . This allows us to employ the following recovery phase to decrease constraints violation cumulated in the play phase, while also ensuring that the cumulative regret stays low at the end of the algorithm. Formally:

Lemma 4.3. *If event \mathbf{E} holds, then after round T_1 of Algorithm 1 the following inequality holds:*

$$\sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \geq \sum_{t=1}^{T_1} f_t(\xi^*) + (T - T_1) - \frac{1}{\rho} \mathcal{E}_{T_1,\eta} - \left(1 + \frac{2}{\rho}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\rho} \mathcal{E}_{T_1}^{\mathbb{D}}.$$

In the recovery phase, the only goal of Algorithm 1 is to decrease constraints violation. In the following Lemma 4.4, we show that, at each round of the recovery phase, the algorithm is “close” to satisfying (in expectation) all the constraints by at least ρ . Formally:

Lemma 4.4. *If event \mathbf{E} holds, then after Algorithm 1 halts, the following holds for every $i \in [m]$:*

$$\sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \leq -(T - T_1)\rho + 2\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}} + \mathcal{E}_{T-T_1}^{\mathbb{D}} + \mathcal{E}_{T-T_1,\eta}.$$

Now, we are ready to present the two main results of this section. First, we provide a bound on the cumulative regret and constraints violation when the lower bound $\hat{\rho}$ is sufficiently large.

Condition 4.5. It holds that $\hat{\rho} \geq 2T^{-1/4}$.

Notice that, under Condition 4.5, $\bar{\rho} = \hat{\rho}/2$. This gives us the following result:

Theorem 4.6. *Suppose that functions f_t and g_t are selected adversarially and stochastically, respectively. If Condition 4.5 is satisfied, then, with probability at least $1 - \delta$, Algorithm 1 provides*

$$R^T \leq \frac{1}{\bar{\rho}} \mathcal{E}_{T,\eta} + \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T,\eta}^{\mathbb{P}} + \frac{1}{\bar{\rho}} \mathcal{E}_T^{\mathbb{D}} \text{ and } V^T \leq M_{\bar{\rho}} + 2\mathcal{E}_{T,\eta}^{\mathbb{P}} + \mathcal{E}_T^{\mathbb{D}} + \mathcal{E}_{T,\eta}.$$

Finally, we also prove that even if Condition 4.5 is *not* satisfied, *i.e.*, the lower bound $\hat{\rho}$ is *not* sufficiently large, the following holds:

Theorem 4.7. *Suppose that functions f_t and g_t are selected adversarially and stochastically, respectively. Algorithm 1 guarantees that the following bounds hold with probability at least $1 - \delta$: $R_T \leq T^{1/4} \mathcal{E}_{T,\eta} + (1 + 2T^{1/4}) \mathcal{E}_{T,\eta}^p + T^{1/4} \mathcal{E}_T^D$ and $V_T \leq T^{3/4} + M_{T^{-1/4}} + 2\mathcal{E}_{T,\eta}^p + \mathcal{E}_T^D + \mathcal{E}_{T,\eta}$.*

Remark 4.8. *Notice that, by using primal and dual RMs whose regret bounds are of the order of $\tilde{O}(\sqrt{T})$, Theorem 4.6 allows us to recover $\tilde{O}(\sqrt{T}/\hat{\rho})$ regret and $\tilde{O}(\sqrt{T}/\hat{\rho})$ constraints violation for the case in which Condition 4.5 holds. Theorem 4.7 still provides $\tilde{O}(T^{3/4})$ regret and constraints violation when the condition is not met, which is necessary the case when $\rho = 0$.*

5 Analysis with stochastic constraints and stochastic rewards

In this section, we focus on the case in which both reward and constraint functions are selected stochastically. In this setting, we are able to show that Algorithm 1 never enters the recovery phase. As we argue in Section 8, this is an important property for budget-management applications, since it is related to the round in which the budget is fully depleted.

In order to prove our result, we extend the event \mathbf{E} to capture also the Azuma-Hoeffding bounds for the reward functions, which are stochastic in this setting.⁷ The core idea that we exploit to prove our result is that we can think of the two RMs as if they are playing a stochastic repeated zero-sum game, which is the repeated Lagrangian game whose functions are sampled according to the probability measures $\mu_{\mathcal{F}}$ and $\mu_{\mathcal{G}}$. By Theorem 2.3, strong duality holds, and the game has an equilibrium. Hence, it is possible to show that the per-round utility of the primal RM is close to the value of the game, which is $\text{OPT}_{\bar{f},\bar{g}}$. At the same time, it is possible to show that, if the cumulative constraints violation becomes large during the play phase (and, thus, $T_1 < T$), then the per-round utility of the primal RM is below $\text{OPT}_{\bar{f},\bar{g}}$, reaching a contradiction that proves the following theorem.

Theorem 5.1. *Suppose that functions f_t and g_t are selected stochastically. With probability at least $1 - \delta$, Algorithm 1 never enters the recovery phase, namely $T_1 = T$.*

Notice that regret bounds analogous to the ones in Theorems 4.6 and 4.7 also hold in the case in which both reward and constraint functions are selected stochastically.

6 Analysis with adversarial constraints

In this section, we study settings in which the constraint functions g_t are selected adversarially. As shown by Mannor et al. [36], it is impossible to obtain sublinear cumulative regret and constraints violation when using our baseline, *i.e.*, the best fixed strategy mixture ξ^* satisfying (in expectation) the long-term constraints. However, we show that it is possible to achieve a $\rho/(1 + \rho)$ fraction of the cumulative reward obtained by always playing ξ^* , while guaranteeing sublinear constraints violation. The dependence of the approximation factor on the feasibility parameter ρ is similar to the dependence on the per-round budget in problems with budget constraints (see the related works in Appendix A for more details). Moreover, as we discuss later in Section 8, when restricted to the case of budget constraints and adversarial reward/cost functions, our approximation factor matches the state-of-the-art bounds provided by Castiglioni et al. [17].

As a first step to prove our result, we provide a lower bound on the cumulative reward of the primal RM during the play phase. We show that it achieves at least a $\rho/(1 + \rho)$ fraction of the value obtained by the optimal solution in the first T_1 rounds. Moreover, the algorithm provides an additional utility compensating for the last rounds in which the algorithm only focuses in satisfying the constraints. Finally, we show that, in the recovery phase, the constraints are satisfied by at least ρ at each round, up to a term related to the regret of \mathcal{R}_{Π}^p and \mathcal{R}_{Π}^D , proving the following theorem.

Theorem 6.1. *Suppose that functions f_t and g_t are selected adversarially. If Condition 4.5 is satisfied, then, with probability at least $1 - \frac{2}{3}\delta$, Algorithm 1 guarantees that the following holds:*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) \geq \frac{\rho}{1+\rho} \sum_{t=1}^T \text{OPT}_{\bar{f},\bar{g}} - \left(1 + \frac{2}{\rho}\right) \mathcal{E}_{T,\eta}^p - \frac{1}{\rho} \mathcal{E}_T^D \text{ and } V^T \leq M_{\bar{\rho}} + 2\mathcal{E}_{T,\eta}^p + \mathcal{E}_T^D.$$

A similar result can be also derived for the case of stochastic rewards and adversarial constraints.

⁷Accounting for the martingale difference sequences $f_t(\mathbf{x}_t) - \bar{f}(\mathbf{x}_t)$ and $f_t(\xi^*) - \bar{f}(\xi^*)$.

Corollary 6.2. *Suppose functions f_t and g_t are selected stochastically and adversarially, respectively. If Condition 4.5 is satisfied, then, with probability at least $1 - \delta$, Algorithm 1 provides $\sum_{t=1}^T f_t(\mathbf{x}_t) \geq \frac{\rho}{1+\rho} \sum_{t=1}^T \text{OPT}_{\bar{f}, \bar{g}} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T, \eta}^p - \frac{1}{\bar{\rho}} \mathcal{E}_T^d - 2\mathcal{E}_{T, \eta}$ and $V^T \leq M_{\bar{\rho}} + 2\mathcal{E}_{T, \eta}^p + \mathcal{E}_T^d + \mathcal{E}_{T, \eta}$.*

Remark 6.3. *By using primal and dual RMs whose regret bounds are of the order of $\tilde{O}(\sqrt{T})$, Theorem 6.1 and Corollary 6.2 allows us to recover $\sum_{t=1}^T f_t(\mathbf{x}_t) \geq \frac{\rho}{1+\rho} \sum_{t=1}^T \text{OPT}_{\bar{f}, \bar{g}} - \tilde{O}(\sqrt{T}/\hat{\rho})$, and $\tilde{O}(\sqrt{T}/\hat{\rho})$ constraints violation for the case in which Condition 4.5 holds.*

7 How to get away with no knowledge about the feasibility parameter

We show how to extend Algorithm 1 in order to deal with settings in which a lower bound on the feasibility parameter ρ is *not* known, when functions g_t are selected stochastically. We propose Algorithm 3, which runs Algorithm 1 by first devoting a given number $T_0 < T$ of rounds to inferring a suitable lower bound $\hat{\rho}$ on the feasibility parameter ρ . Ideally, we would like to have $\hat{\rho} = \Omega(\rho)$, so as to recover bounds of the order $\tilde{O}(\sqrt{T}/\rho)$. In particular, we show that we can run Algorithm 3 with $T_0 = T^{1/2}$ to obtain an approximation of ρ that has an additive approximation error of the order $T^{1/4}$. This is sufficient to get $\hat{\rho} = \Omega(\rho)$, since a good approximation of ρ is only needed when $\rho \geq T^{1/4}$.

Let us remark that our approach only works when constraints functions g_t are selected stochastically. When these are chosen adversarially, it is easy to see that it is impossible to compute a lower bound on the feasibility parameter ρ by only using the first rounds. For instance, think of a setting in which ρ is very large in the first rounds, while it becomes small during the last ones.

To exploit the guarantees of Algorithm 1 presented in the previous sections, it is enough to show that, after the first T_0 rounds of Algorithm 3, $\hat{\rho} \leq \rho$ holds with high probability.

Algorithm 3 META-ALGORITHM(T, T_0, δ)

- 1: $\mathcal{R}^p \leftarrow \text{INIT}^p(\mathcal{X}, [-1, 1], \delta)$
 - 2: $\mathcal{R}^d \leftarrow \text{INIT}^d(\Delta_m, [-1, 1], 0)$
 - 3: $t \leftarrow 1$
 - 4: **while** $t \leq T_0$: **do**
 - 5: $\mathbf{x}_t \leftarrow \text{LAGRANGIANGAME}(\mathcal{R}^p, \mathcal{R}^d, 0)$
 - 6: $t \leftarrow t + 1$
 - 7: $\hat{\rho} \leftarrow -\frac{1}{T_0} \left(\max_{i \in [m]} \sum_{t=1}^{T_0} g_{t,i}(\mathbf{x}_t) + \mathcal{E}_{T_0, \delta} \right)$
 - 8: Run Algorithm 1 with $T - T_0, \delta$, and $\hat{\rho}$ as inputs
-

Lemma 7.1. *If $T_0 = \sqrt{T}$, then Algorithm 3 guarantees that $\hat{\rho} \leq \rho$ with probability at least $1 - \delta$.*

In order to recover a good estimate of ρ , we need the value of ρ to be sufficiently large.

Condition 7.2. It holds that $\rho \geq \frac{2}{T_0} (2\mathcal{E}_{T_0, \delta} + 2\mathcal{E}_{T_0, \delta}^p + \mathcal{E}_{T_0}^d)$.

Remark 7.3. *Notice that, by using primal and dual RMs whose regret bounds are of the order $\tilde{O}(\sqrt{T})$ and by setting $T_0 = \sqrt{T}$, Condition 7.2 is satisfied when $\rho = \omega(T^{-1/4})$.*

Next, we show that $\hat{\rho} = \Omega(\rho)$, which allows us to exploit the guarantees proved for Algorithm 1 in order to provide analogous ones for Algorithm 3. Formally:

Lemma 7.4. *By setting $T_0 = \sqrt{T}$, and assuming that Condition 7.2 is satisfied, after T_0 rounds of Algorithm 3 we have that $\hat{\rho} \geq \rho/2$ with probability at least $1 - 2\delta$.*

By applying the results of the previous sections on the guarantees of Algorithm 1, and by using primal and dual RMs whose regret bounds are of the order $\tilde{O}(\sqrt{T})$, we get $\tilde{O}(\sqrt{T}/\rho)$ and $\tilde{O}(\sqrt{T}/\rho)$ regret and violation bounds, respectively, when the functions g_t are selected stochastically.

8 Applications to repeated auctions settings

Internet advertising platforms usually operationalize large auction markets by using *proxy bidders* that place bids in repeated auctions on the advertisers' behalf. A proxy-bidder selects bids according to a *budget-pacing mechanism*, which manages the usage of the advertisers' budget over time [1, 22, 9]. In this section, we discuss the application of our framework to budget-management in auctions, arguing that it can deal with more general constraints on ad slots allocation with respect to what is currently achievable with multiplicative pacing algorithms, which manage only *knapsack constraints*.

We consider the problem faced by a bidder who takes part in a sequence of repeated auctions. We focus on the case of *second-price* and *first-price* auctions, since they are the *de facto* standard in large Internet advertising platforms. At each round $t \in [T]$, the bidder observes their valuation v_t from a finite set of n_v possible valuations $\mathcal{V} \subset [0, 1]$. Such valuation models targeting preferences of the advertiser. Then, the bidder chooses a bid $b_t \in \mathcal{B}$, where $\mathcal{B} \subset [0, 1]$ is a finite set of n_b possible bids such that $0 \in \mathcal{B}$ (i.e., the bidder is allowed to skip items without incurring in any cost). The utility of the bidder depends on the largest among competing bids, denoted by β_t . In particular, the utility is computed as $f_t(b_t) = (v_t - c_t(b_t))\mathbb{1}\{b_t \geq \beta_t\}$, where the cost c_t is such that $c_t(b_t) = \beta_t\mathbb{1}\{b_t \geq \beta_t\}$ in second-price auctions, and $c_t(b_t) = b_t\mathbb{1}\{b_t \geq \beta_t\}$ for first-price ones. Finally, the bidder has a target *per-round* budget of $\rho > 0$, which yields an overall budget $B := \rho T$ that limits the total spending over the T rounds. In the case of budget-constrained bidding, a strictly feasible solution can be easily achieved by always bidding 0. Using the target per-round budget $\rho = B/T$ we can write the budget constraint as $\sum_{t \in [T]} g_t(b_t) \leq 0$, with $g_t(b) = c_t(b) - \rho$ for any $b \in \mathcal{B}$. Notice that, in this setting, we have the same feasibility parameter ρ for both the stochastic and the adversarial case.

As a benchmark to evaluate the algorithm, we consider the best feasible static policy $\pi : \mathcal{V} \rightarrow \mathcal{B}$. The set of static policies can be represented by $\mathcal{X} := \mathcal{B}^{n_v}$, where a vector $\mathbf{b} \in \mathcal{B}^{n_v}$ encodes the policy's bids for each possible valuation. To apply our framework to this problem, it is sufficient to design a primal regret minimizer constructor (recall that, in order to design dual RMs, we can employ OMD). This can be implemented by instantiating a regret minimizer EXP3.P [5] for each possible valuation in \mathcal{V} . Given a failure probability $\nu \in (0, 1)$, each RM guarantees a regret bound $O(\sqrt{T n_b \log(n_b/\nu)})$ with probability at least $1 - \nu$. Thus, given a desired failure probability $\eta \in (0, 1)$, by setting $\nu = \eta/n_v$ we get that, with probability at least $1 - \eta$, the bounds of all the RMs hold. Hence, by a union bound, we get that the regret of a primal RM is $\mathcal{E}_{T,\eta}^p = O(n_v \sqrt{T n_b \log(n_b n_v/\eta)})$.

Guaranteed budget completion in the stochastic case. The crux of budget-pacing mechanisms is ensuring that the advertisers' budget is not depleted too early (thereby missing potentially valuable future advertising opportunities), while being fully depleted within the planned duration of the campaign. Theorem 5.1 shows that, when inputs are generated according to some stochastic model, Algorithm 1 never enters the recovery phase. This is crucial in the context of budget-pacing mechanisms, because whenever the algorithm enters the recovery phase it will converge to always bid 0 in order to mitigate constraints violation. Therefore, the bidder could miss out on potentially valuable items. Moreover, if the platform wanted to guarantee that the bidder does not spend more than the budget B , it would be enough to set a *virtual budget* of $B - \tilde{O}(T^{1/2})$ to compensate for the potential constraints violation. Finally, we argue that, in large-scale markets, an individual bidder has almost no impact on prices, and, thus, stochastic behavior of costs is a reasonable assumption.

Adversarial case. Theorem 6.1 of [17] shows how to construct an algorithm that provides a ρ fraction of the optimal utility for problems with budget constraints and adversarial inputs. The ratio $\rho/(1 + \rho)$ obtained in Theorem 6.1 matches such result. The latter assumes that rewards and costs are in $[0, 1]$, and, thus, $g_t \in [-\rho, 1 - \rho]$ (as they only model budget constraints). However, in our case we have $g_t \in [-1, 1]$. By normalizing the former range to match with ours, we get $g_t \in [-\rho/(1 - \rho), 1]$. Therefore, the feasibility parameter would be $\rho' = \rho/(1 - \rho)$. By rewriting our guarantees as a function of ρ , we get $\rho'/(1 + \rho') = \rho$, which is the same guarantee of [17].

Handling ROI constraints. Traditional budget-pacing mechanisms (see, e.g., [11, 8]) are based on primal-dual algorithms that are near optimal in settings with knapsack constraints only, and they cannot be generalized to deal with other types of long-term constraints. However, there are many real-world situations in which guaranteeing other types of constraints is crucial for practical applications (see, e.g., [26, 25]). One example is the case of *return on investment* (ROI) constraints [6, 26, 32]. The recent work by Golrezaei et al. [25] presents a threshold-based algorithm for repeated second-price auctions under budget and ROI constraints. Our framework allows advertisers to reach a target ROI while keeping expenses under control also in the setting of repeated first-price auctions. In particular, given a target ROI ω , we define the ROI constraints as $g_t(b_t) = (\omega - v_t/b_t)\mathbb{1}\{b_t \geq \beta_t\} \leq 0$. Then, it is enough to instantiate our framework as described before to immediately get that the cumulative violation of the budget and ROI constraints are upper bounded by $\tilde{O}(T^{1/2})$. This holds both in the fully stochastic and in the fully adversarial setting.

See Appendix C for a detailed discussion on the types of constraints that our framework can handle.

References

- [1] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. Budget pacing for targeted online advertisements at linkedin. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1613–1619, 2014.
- [2] Shipra Agrawal and Nikhil R Devanur. Bandits with concave rewards and convex knapsacks. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 989–1006. ACM, 2014.
- [3] Shipra Agrawal and Nikhil R Devanur. Bandits with global convex constraints and objective. *Operations Research*, 67(5):1486–1502, 2019.
- [4] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [5] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [6] Jason Auerbach, Joel Galenson, and Mukund Sundararajan. An empirical analysis of return on investment maximization in sponsored search auctions. In *Proceedings of the 2nd International Workshop on Data Mining and Audience Intelligence for Advertising*, pages 1–9, 2008.
- [7] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with knapsacks. *J. ACM*, 65(3), 2018.
- [8] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. The best of many worlds: Dual mirror descent for online allocation problems. *arXiv preprint arXiv:2011.10124*, 2020.
- [9] Santiago Balseiro, Anthony Kim, Mohammad Mahdian, and Vahab Mirrokni. Budget-management strategies in repeated auctions. *Operations Research*, 2021.
- [10] Santiago Balseiro, Haihao Lu, and Vahab Mirrokni. Regularized online allocation problems: Fairness and beyond. In *International Conference on Machine Learning*, pages 630–639. PMLR, 2021.
- [11] Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.
- [12] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [13] Andrey Bernstein, Shie Mannor, and Nahum Shimkin. Online classification with specificity constraints. *Advances in Neural Information Processing Systems*, 23, 2010.
- [14] Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Conference on Learning Theory*, pages 42–1. JMLR Workshop and Conference Proceedings, 2012.
- [15] Sébastien Bubeck, Yin Tat Lee, and Ronen Eldan. Kernel-based methods for bandit convex optimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, page 72–85, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450345286.
- [16] Xuanyu Cao and KJ Ray Liu. Online convex optimization with time-varying constraints and bandit feedback. *IEEE Transactions on automatic control*, 64(7):2665–2680, 2018.
- [17] Matteo Castiglioni, Andrea Celli, and Christian Kroer. Online learning with knapsacks: the best of both worlds. In *International Conference on Machine Learning, ICML 2022*, pages 2767–2783, 2022.
- [18] Andrea Celli, Riccardo Colini-Baldeschi, Christian Kroer, and Eric Sodomka. The parity ray regularizer for pacing in auction markets. In *Proceedings of the ACM Web Conference 2022*, pages 162–172, 2022.
- [19] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [20] Tianyi Chen and Georgios B Giannakis. Bandit convex optimization for scalable and dynamic iot management. *IEEE Internet of Things Journal*, 6(1):1276–1286, 2018.

- [21] Tianyi Chen, Qing Ling, and Georgios B Giannakis. An online convex optimization approach to proactive network resource allocation. *IEEE Transactions on Signal Processing*, 65(24): 6350–6364, 2017.
- [22] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 2021.
- [23] Jérémie Gallien and Shobhit Gupta. Temporary and permanent buyout prices in online auctions. *Management Science*, 53(5):814–833, 2007.
- [24] Lodewijk Gelauff, Ashish Goel, Kamesh Munagala, and Sravya Yandamuri. Advertising for demographically fair outcomes. *arXiv preprint arXiv:2006.03983*, 2020.
- [25] Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni. Bidding and pricing in budget and roi constrained markets. *arXiv preprint arXiv:2107.07725*, 2021.
- [26] Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. Auction design for roi-constrained buyers. In *Proceedings of the Web Conference 2021*, pages 3941–3952, 2021.
- [27] Elad Hazan. *Efficient algorithms for online convex optimization and their applications*. Princeton University, 2006.
- [28] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2):169–192, 2007.
- [29] Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [30] Nicole Immorlica, Karthik Abinav Sankararaman, Robert Schapire, and Aleksandrs Slivkins. Adversarial bandits with knapsacks. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019*, pages 202–219. IEEE Computer Society, 2019.
- [31] Rodolphe Jenatton, Jim Huang, and Cédric Archambeau. Adaptive algorithms for online convex optimization with long-term constraints. In *International Conference on Machine Learning*, pages 402–411. PMLR, 2016.
- [32] Bin Li, Xiao Yang, Daren Sun, Zhi Ji, Zhen Jiang, Cong Han, and Dong Hao. Incentive mechanism design for roi-constrained auto-bidding. *arXiv preprint arXiv:2012.02652*, 2020.
- [33] Nikolaos Liakopoulos, Apostolos Destounis, Georgios Paschos, Thrasyvoulos Spyropoulos, and Panayotis Mertikopoulos. Cautious regret minimization: Online optimization with long-term budget constraints. In *International Conference on Machine Learning*, pages 3944–3952. PMLR, 2019.
- [34] David G Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.
- [35] Mehrdad Mahdavi, Rong Jin, and Tianbao Yang. Trading regret for efficiency: online convex optimization with long term constraints. *The Journal of Machine Learning Research*, 13(1): 2503–2528, 2012.
- [36] Shie Mannor, John N Tsitsiklis, and Jia Yuan Yu. Online learning with sample path constraints. *Journal of Machine Learning Research*, 10(3), 2009.
- [37] Arkadij Semenovič Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- [38] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- [39] Arun Sai Suggala and Praneeth Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. In Aryeh Kontorovich and Gergely Neu, editors, *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, volume 117 of *Proceedings of Machine Learning Research*, pages 845–861. PMLR, 08 Feb–11 Feb 2020.
- [40] Wen Sun, Debadeepta Dey, and Ashish Kapoor. Safety-aware algorithms for adversarial contextual bandit. In *International Conference on Machine Learning*, pages 3280–3288. PMLR, 2017.
- [41] Xiaohan Wei, Hao Yu, and Michael J Neely. Online primal-dual mirror descent under stochastic constraints. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(2):1–36, 2020.

- [42] Xinlei Yi, Xiuxian Li, Lihua Xie, and Karl H Johansson. Distributed online convex optimization with time-varying coupled inequality constraints. *IEEE Transactions on Signal Processing*, 68: 731–746, 2020.
- [43] Hao Yu and Michael J. Neely. A low complexity algorithm with $o(\sqrt{T})$ regret and $o(1)$ constraint violations for online convex optimization with long term constraints. *Journal of Machine Learning Research*, 21(1):1–24, 2020.
- [44] Hao Yu, Michael Neely, and Xiaohan Wei. Online convex optimization with stochastic constraints. *Advances in Neural Information Processing Systems*, 30, 2017.
- [45] Jianjun Yuan and Andrew Lamperski. Online convex optimization for cumulative constraints. *Advances in Neural Information Processing Systems*, 31, 2018.
- [46] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936, 2003.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** We present the goal of our paper and summarize the theoretical results. In the Original Contribution section, we underline the novelty of our work and in the Related Work section we compare our results with those in the literature (see also Table 1).
 - (b) Did you describe the limitations of your work? **[Yes]** The assumptions characterizing the settings in which our results hold are clearly stated throughout the paper.
 - (c) Did you discuss any potential negative societal impacts of your work? **[N/A]** The results presented in this work does not have any foreseeable negative societal consequence.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** In each theorem statement we provide all the assumptions needed for the result to hold.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** They are provided in Appendix B and Appendix C.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[N/A]**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[N/A]**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[N/A]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[N/A]**
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...

- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Related works

The *online convex optimization* (OCO) framework was first proposed in the machine learning literature by Zinkevich [46], and since then it has significantly expanded becoming widely influential in the learning community (see, *e.g.*, [27, 29, 38]). In what follows, we highlight the most relevant works with respect to ours from the literature related to online convex optimization problems with constraints. The analysis and the results are quite different depending on the nature of the constraints, which may be static, *i.e.*, time-invariant, or stochastic/adversarial, *i.e.*, time-variant.

Static constraints. Zinkevich [46] first addressed online convex optimization problems with static constraints by developing a projection-based *online gradient descent* (OGD) algorithm. This method guarantees a regret upper bound of $\mathcal{O}(\sqrt{T})$ for an arbitrary sequence of convex objective functions with bounded subgradients. Hazan et al. [28] showed that this is a tight bound up to constant factors. When the set defined by the static constraints is complex, the conventional projection-based online algorithms can be difficult to implement due to the potentially high computational cost of carrying out the projection operation. To overcome this difficulty, Mahdavi et al. [35] propose an efficient algorithm which is an adaptation of OGD achieving a cumulative regret of order $\mathcal{O}(\sqrt{T})$ and a cumulative constraints violation of $\mathcal{O}(T^{3/4})$. These bounds are generalized by Jenatton et al. [31] who propose an algorithm that achieves a cumulative regret of $\mathcal{O}(T^{\max\{\beta, 1-\beta\}})$ and a cumulative violation of $\mathcal{O}(T^{1-\beta/2})$, where $\beta \in (0, 1)$ is a user-defined parameter. Other works, such as, *e.g.*, [45, 43], propose primal-dual algorithms and achieve better bounds by making further assumptions. In particular, Yu and Neely [43] achieve bounds on the cumulative regret of $\mathcal{O}(\sqrt{T})$ and on the cumulative violation of $\mathcal{O}(1)$ by assuming that the Slater’s condition holds (*i.e.*, the existence of a strictly feasible solution). Then, Yuan and Lamperski [45] achieve a cumulative regret of $\mathcal{O}(\log T)$ and a constraint violation of $\mathcal{O}(\sqrt{T})$ under the assumption that the objective functions are strongly convex. In all the the papers cited above, the regret is computed with respect to the best fixed action in hindsight, that does *not* violate the constraints at each round t . This metric is called *static regret*.

Stochastic constraints. Yu et al. [44] consider an online convex optimization framework with stochastic constraints, where the objective functions are chosen by an adversary, and the constraint functions are independent and identically distributed (i.i.d.) over time. Yu et al. [44] provide a primal-dual proximal gradient algorithm achieving $\mathcal{O}(\sqrt{T})$ cumulative regret and constraint violation by assuming Slater’s condition. Moreover, Wei et al. [41] provide bounds of the same order by assuming a less stringent version of the Slater’s condition. As a performance metric, the latter work use *static regret*.

Adversarial constraints. Various works in the literature address the online learning setting with adversarial reward and constraint functions. This problem was first studied by Mannor et al. [36] in a two-player game setting. The regret is computed with respect to the best strategy from the set of fixed strategies that satisfy the constraints on average. Mannor et al. [36] show that in general it is impossible to compete against the best decision in such a set. In particular, they construct a two-player game where there exists a policy for the adversary such that, among the policies of the player that violate sublinearly the constraints, there is no policy that can achieve the no-regret property in terms of maximizing the player’s reward. Sun et al. [40] study a similar problem related to contextual bandits and show that also in their setting the decision maker is unable to compete again this baseline by adapting the result from Mannor et al. [36] to their setting. To circumvent this issue and provide some guarantees, they rely on a weaker baseline to compute the regret. In particular, they assume that the decision set is rich enough that, in hindsight, there exist a fixed action that satisfies the constraints at each round: they are using the *static regret* as a performance metric. Then, by employing static regret as a baseline, Sun et al. [40] show that the approaches of Mahdavi et al. [35] and Jenatton et al. [31] can be extended to the online learning framework with adversarial sequential constraints. Therefore, they provide an algorithm which is a generalization of that from Mahdavi et al. [35] achieving sublinear cumulative regret and constraint violations.

Liakopoulos et al. [33] define a new notion of regret, to overcome the impossibility result from Mannor et al. [36]. They introduce a refined regret metric which compares the agent’s incurred losses to those of a *K-benchmark*, which is the best strategy in the hindsight such that, for each time window of length K , the long-term constraints over that window are satisfied. They provide parametric results

that are useful to balance the trade-off between regret minimization and long-term residual constraint violation. Moreover, instead of the Slater’s condition they consider a less stringent assumption related to the definition of their regret metric.

A recent line of works such as [21, 20] and [16] provide some results related to the regret against *dynamic policies*. As expected, comparing against a dynamic baseline require very strong assumptions. Chen et al. [21] compute a bound on the cumulative dynamic regret which is sublinear in the time horizon T only if the drift of the baseline sequence (i.e., $\sum_{t=1}^T \|\mathbf{x}_{t+1}^* - \mathbf{x}_t^*\|$) and that of the constraints (i.e., $\sum_{t=1}^T \max_{\mathbf{x}} \|\mathbf{g}_{t+1}(\mathbf{x}) - \mathbf{g}_t(\mathbf{x})\|^+$) are $o(T^{2/3})$. Cao and Liu [16] consider a bandit feedback setting and, in order to provide sublinear regret and constraint violations, they assume that all the loss functions have uniformly bounded difference (i.e., for each t and $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, $|f_t(\mathbf{x}) - f_t(\mathbf{x}')| \leq M$ for some positive constant M), and that the drift of the baseline sequence is sublinear. In other words, the underlying dynamic optimization problems vary *slowly* over time. Both Chen et al. [21] and Cao and Liu [16] need to assume the Slater’s condition. Yi et al. [42] provide similar results in a distributed online convex optimization setting with adversarial constraints. They analyze both the case in which the Slater’s condition holds, and the case without this assumption.

Others relevant related works, are those studying online learning problems in which the decision maker has to satisfy supply/budget constraints. In this setting, the decision maker wants to maximize their expected reward without violating a set of m resource constraints. The process stops at time horizon T , or when the total consumption of some resource exceeds its budget. Badanidiyuru et al. [7] first introduce and solve the *Bandits with Knapsacks* (BwK) framework, in which they consider bandit feedback, stochastic objective and constraint functions. Other optimal algorithms for Stochastic BwK were proposed by Agrawal and Devanur [2, 3] and by Immorlica et al. [30]. The *Adversarial Bandits with Knapsacks* setting was first studied by Immorlica et al. [30]. The authors shows that an appropriate baseline is the best fixed distribution over arms. Achieving no-regret is no longer possible under this baseline and, therefore, they provide no- α -regret guarantees for their algorithm.

We remark that in this paper we are able to handle more general constraints than Immorlica et al. [30], which can deal only with budget constraints. Moreover, we can compete with a baseline stronger than the static regret used by Sun et al. [40], without needing the strong assumptions required, for instance, by Cao and Liu [16].

B Omitted proofs

B.1 Proof omitted from Section 2.1

Theorem 2.3. *Given $f \in \mathcal{F}$ and $g \in \mathcal{G}$ such that $d_g > 0$, it holds*

$$\sup_{\xi \in \Xi} \inf_{\lambda \in \mathcal{D}_{d_g}} \mathcal{L}_{f,g}(\xi, \lambda) = \inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}_{f,g}(\xi, \lambda) = \text{OPT}_{f,g}.$$

Proof. As a first step, we prove that $\inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) = \inf_{\lambda \in \mathbb{R}_+^m} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda)$. Notice that for each $\lambda \in \mathbb{R}_+^m$ such that $\|\lambda\|_1 > 1/d_g$, it holds

$$\sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) \geq \mathcal{L}(\xi^\circ, \lambda) \geq -\langle \lambda^*, g(\xi^\circ) \rangle \geq d_g \|\lambda^*\|_1 > 1,$$

where, with an abuse of notation, $\xi^\circ \in \Xi$ denotes a strictly feasible strategy mixture for Problem $\text{LP}_{f,g}$. That is a strategy mixture $\xi \in \Xi$ which is optimal for the problem defining d_g in Equation (1), and, thus, it satisfies all the constraints by at least d_g (i.e., it holds $g_i(\xi^\circ) \leq -d_g$ for all $i \in [m]$).⁸ Thus, it holds that

$$\inf_{\lambda \in \mathbb{R}_+^m \setminus \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) > 1.$$

Moreover, since

$$\inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) \leq \sup_{\xi \in \Xi} \mathcal{L}(\xi, \mathbf{0}) \leq 1,$$

⁸Notice that ξ° may *not* be well defined when the problem in Equation 1 does *not* admit a maximum. In such cases, we can take a ξ° that is arbitrarily “close” to a supremum, so that the result still holds.

we can conclude that

$$\begin{aligned} \inf_{\lambda \in \mathbb{R}_+^m} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) &= \min \left\{ \inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda); \inf_{\lambda \in \mathbb{R}_+^m \setminus \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) \right\} \\ &= \inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda). \end{aligned} \quad (3)$$

Then,

$$\begin{aligned} \text{OPT}_{f,g} &= \sup_{\xi \in \Xi} \inf_{\lambda \in \mathbb{R}_+^m} \mathcal{L}(\xi, \lambda) \\ &\leq \sup_{\xi \in \Xi} \inf_{\lambda \in \mathcal{D}_{d_g}} \mathcal{L}(\xi, \lambda) \\ &\leq \inf_{\lambda \in \mathcal{D}_{d_g}} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) \\ &= \inf_{\lambda \in \mathbb{R}_+^m} \sup_{\xi \in \Xi} \mathcal{L}(\xi, \lambda) \\ &= \text{OPT}_{f,g}, \end{aligned}$$

where the first inequality holds since in the right-hand side the inf is taken over the more restrictive set \mathcal{D}_{d_g} , the second one by the *max-min inequality*, while the second-to-last equality holds by Equation (3). This concludes the proof. \square

B.2 Proofs omitted from Section 4

Lemma 4.2. *After running Algorithm 1, the event \mathbf{E} holds with probability at least $1 - \delta$.*

Proof. Given a desired failure probability $\delta \in (0, 1)$, recall that $\eta = \delta/3$ and set $\varepsilon = \eta/18mT^2$. Consider the following inequalities in which the differences between expectations and empirical means of constraint functions are bounded:

$$\left| \sum_{\tau=t}^{t'} g_{\tau,i}(\mathbf{x}_\tau) - \sum_{\tau=t}^{t'} \bar{g}_i(\mathbf{x}_\tau) \right| > 2\sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}, \quad (4)$$

$$\left| \sum_{\tau=t}^{t'} g_{\tau,i}(\xi^\circ) - \sum_{\tau=t}^{t'} \bar{g}_i(\xi^\circ) \right| > 2\sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}, \quad (5)$$

$$\left| \sum_{\tau=t}^{t'} g_{\tau,i}(\xi^*) - \sum_{\tau=t}^{t'} \bar{g}_i(\xi^*) \right| > 2\sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}, \quad (6)$$

$$\left| \sum_{\tau=t}^{t'} \lambda_\tau g_{\tau,i}(\mathbf{x}_\tau) - \sum_{\tau=t}^{t'} \lambda_\tau \bar{g}_i(\mathbf{x}_\tau) \right| > 2 \max_{\tau \in [T]: t \leq \tau \leq t'} \|\lambda_\tau\|_1 \sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}, \quad (7)$$

$$\left| \sum_{\tau=t}^{t'} \lambda_\tau g_{\tau,i}(\xi^*) - \sum_{\tau=t}^{t'} \lambda_\tau \bar{g}_i(\xi^*) \right| > 2 \max_{\tau \in [T]: t \leq \tau \leq t'} \|\lambda_\tau\|_1 \sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}, \quad (8)$$

$$\left| \sum_{\tau=t}^{t'} \lambda_\tau g_{\tau,i}(\xi^\circ) - \sum_{\tau=t}^{t'} \lambda_\tau \bar{g}_i(\xi^\circ) \right| \leq 2 \max_{\tau \in [T]: t \leq \tau \leq t'} \|\lambda_\tau\|_1 \sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}. \quad (9)$$

By applying Azuma-Hoeffding inequality to each martingale difference sequence, we get that each inequality holds with probability at most 2ε . We denote by \mathbf{E}_η the event in which Equations (4), (5), (6), and (7) are satisfied for all $t, t' \in [T]$ with $t < t'$, and for all $i \in [m]$. By a union bound that takes into account the six events above, the m constraints, and all the possible time intervals from t to t' , which are at most T^2 , we get:

$$\mathbb{P}(\mathbf{E}_\eta) \geq 1 - 6mT^2(2\varepsilon) = 1 - 12mT^2\varepsilon = 1 - \frac{2}{3}\eta \geq 1 - \eta.$$

Therefore, event \mathbf{E}_η holds with probability at least $1 - \eta$. Moreover, let us recall that:

$$\mathcal{E}_{t'-t,\eta} = \sqrt{8(t' - t) \ln \frac{18mT^2}{\eta}} = 2\sqrt{2(t' - t) \ln \frac{1}{\varepsilon}}.$$

Now, consider event \mathbf{E} in which Algorithms 1 satisfies the following conditions: (i) the regret incurred by $\mathcal{R}_I^{\mathbb{P}}$ after T_1 rounds is upper bounded by $\mathcal{E}_{T_1,\eta}^{\mathbb{P}}$; (ii) the regret cumulated by $\mathcal{R}_{II}^{\mathbb{P}}$ after the remaining $T - T_1$ rounds is upper bounded by $\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}}$; and (iii) event \mathbf{E}_η holds. Recall that each one of the conditions (i), (ii) and (iii) holds with probability at least $1 - \eta$; hence, by a union bound we get:

$$\mathbb{P}(\mathbf{E}) \geq 1 - 3\eta = 1 - \delta.$$

This concludes the proof. \square

Lemma 4.3. *If event \mathbf{E} holds, then after round T_1 of Algorithm 1 the following inequality holds:*

$$\sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \geq \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1,\eta} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}.$$

Proof. By the no-regret property of the primal regret minimizer, we have that:

$$\sum_{t=1}^{T_1} \left(f_t(\mathbf{x}_t) - \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \right) \geq \sum_{t=1}^{T_1} \left(f_t(\boldsymbol{\xi}^*) - \langle \boldsymbol{\lambda}_t, g_t(\boldsymbol{\xi}^*) \rangle \right) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}}. \quad (10)$$

Let $i^* \in \arg \max_{i \in [m]} \sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t)$ be one of the ‘‘most violated’’ constraints. We prove that:

$$\sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \geq (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}. \quad (11)$$

To do that, we consider the following two cases.

Case $T_1 = T$. We get:

$$\sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \geq \sum_{t=1}^{T_1} \langle \mathbf{0}, g_t(\mathbf{x}_t) \rangle - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}} = (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}.$$

Case $T_1 < T$. By the condition in Line 4 of Algorithm 1, we have that $\sum_{t=1}^{T_1} g_{t,i^*}(\mathbf{x}_t) \geq (T - T_1)\bar{\rho}$. Thus, we have that:

$$\sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \geq \sum_{t=1}^{T_1} \frac{1}{\bar{\rho}} g_{t,i^*}(\mathbf{x}_t) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}} \geq (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}},$$

where the first inequality follows from the no-regret property of the dual regret minimizer and the second one from the fact that $\sum_{t=1}^{T_1} g_{t,i^*}(\mathbf{x}_t) \geq (T - T_1)\bar{\rho}$ when $T_1 < T$.

Now, by using Equation (11), we can provide a lower bound on the cumulative reward obtained by Algorithm 1 during the play phase. We have that:

$$\begin{aligned} \sum_{t=1}^{T_1} f_t(\mathbf{x}_t) &\geq \sum_{t=1}^{T_1} \left(f_t(\boldsymbol{\xi}^*) - \langle \boldsymbol{\lambda}_t, g_t(\boldsymbol{\xi}^*) \rangle + \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \right) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}}, \\ &\geq \sum_{t=1}^{T_1} \left(f_t(\boldsymbol{\xi}^*) - \langle \boldsymbol{\lambda}_t, g_t(\boldsymbol{\xi}^*) \rangle \right) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}, \\ &\geq \sum_{t=1}^{T_1} \left(f_t(\boldsymbol{\xi}^*) - \langle \boldsymbol{\lambda}_t, \bar{g}(\boldsymbol{\xi}^*) \rangle \right) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1,\eta}, \\ &\geq \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1,\eta}, \end{aligned}$$

where the first inequality holds by Equation (10), the second one by Equation (11), the third one follows from the fact that the event \mathbf{E} holds, while the last one from the fact that $\bar{g}(\boldsymbol{\xi}^*) \leq 0$ by definition. \square

Lemma 4.4. *If event \mathbf{E} holds, then after Algorithm 1 halts, the following holds for every $i \in [m]$:*
 $\sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \leq -(T - T_1)\rho + 2\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}} + \mathcal{E}_{T-T_1}^{\mathbb{D}} + \mathcal{E}_{T-T_1,\eta}$.

Proof. Let $i^* \in \arg \max_{i \in [m]} \sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t)$ be one of the ‘‘most violated’’ constraints. Then,

$$\begin{aligned} (T - T_1)\rho &\leq - \sum_{t=T_1+1}^T \langle \boldsymbol{\lambda}_t, \bar{g}(\boldsymbol{\xi}^\circ) \rangle \\ &\leq - \sum_{t=T_1+1}^T \langle \boldsymbol{\lambda}_t, g_t(\boldsymbol{\xi}^\circ) \rangle + \mathcal{E}_{T-T_1,\eta} \\ &\leq - \sum_{t=T_1+1}^T \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle + 2\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}} + \mathcal{E}_{T-T_1,\eta} \\ &\leq - \sum_{t=T_1+1}^T g_{t,i^*}(\mathbf{x}_t) + \mathcal{E}_{T-T_1}^{\mathbb{D}} + 2\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}} + \mathcal{E}_{T-T_1,\eta}, \end{aligned}$$

where the first inequality comes from the definition of ρ , the second one from the fact that event \mathbf{E} holds, the third one from the no-regret property of the primal regret minimizer, and the last one from the no-regret property of the dual regret minimizer. Hence,

$$\sum_{t=T_1+1}^T g_{t,i^*}(\mathbf{x}_t) \leq -(T - T_1)\rho - \mathcal{E}_{T-T_1}^{\mathbb{D}} + 2\mathcal{E}_{T-T_1,\eta}^{\mathbb{P}} + \mathcal{E}_{T-T_1,\eta}. \quad (12)$$

It follows from the definition of i^* that, if Equation (12) holds for i^* , then, it holds for every $i \in [m]$. This concludes the proof. \square

Theorem 4.6. *Suppose that functions f_t and g_t are selected adversarially and stochastically, respectively. If Condition 4.5 is satisfied, then, with probability at least $1 - \delta$, Algorithm 1 provides $R^T \leq \frac{1}{\rho}\mathcal{E}_{T,\eta} + \left(1 + \frac{2}{\rho}\right)\mathcal{E}_{T,\eta}^{\mathbb{P}} + \frac{1}{\rho}\mathcal{E}_T^{\mathbb{D}}$ and $V^T \leq M_{\bar{\rho}} + 2\mathcal{E}_{T,\eta}^{\mathbb{P}} + \mathcal{E}_T^{\mathbb{D}} + \mathcal{E}_{T,\eta}$.*

Proof. By Lemma 4.2, event \mathbf{E} holds with probability at least $1 - \delta$. In the rest of the proof, we assume the event \mathbf{E} holds, providing a bound that holds with probability at least $1 - \delta$.

We first provide an upper bound on the cumulative regret. By Lemma 4.3, we have:

$$\sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \geq \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \frac{1}{\rho}\mathcal{E}_{T_1,\eta} - \left(1 + \frac{2}{\rho}\right)\mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\rho}\mathcal{E}_{T_1}^{\mathbb{D}}. \quad (13)$$

Hence, it holds:

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t) &\geq \sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \\ &\geq \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \frac{1}{\rho}\mathcal{E}_{T_1,\eta} - \left(1 + \frac{2}{\rho}\right)\mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\rho}\mathcal{E}_{T_1}^{\mathbb{D}} \\ &\geq \sum_{t=1}^T f_t(\boldsymbol{\xi}^*) - \frac{1}{\rho}\mathcal{E}_{T_1,\eta} - \left(1 + \frac{2}{\rho}\right)\mathcal{E}_{T_1,\eta}^{\mathbb{P}} - \frac{1}{\rho}\mathcal{E}_{T_1}^{\mathbb{D}} \\ &\geq \sum_{t=1}^T f_t(\boldsymbol{\xi}^*) - \frac{1}{\rho}\mathcal{E}_{T,\eta} - \left(1 + \frac{2}{\rho}\right)\mathcal{E}_{T,\eta}^{\mathbb{P}} - \frac{1}{\rho}\mathcal{E}_T^{\mathbb{D}}, \end{aligned}$$

where the second inequality holds by Equation (13) and the third one by $\sum_{t=T_1+1}^T f_t(\boldsymbol{\xi}^*) \leq T - T_1$, which follows from the fact that the range of f_t is $[0, 1]$.

By recalling that $\xi^* \in \Xi$ is defined as an optimal solution to Problem $\text{LP}_{\bar{f}, \bar{g}}$ and $R^T = T \circ \text{PT}_{\bar{f}, \bar{g}} - \sum_{t=1}^T f_t(\mathbf{x}_t)$, the following bound on the cumulative regret holds:

$$R^T = \sum_{t=1}^T f_t(\xi^*) - \sum_{t=1}^T f_t(\mathbf{x}_t) \leq \frac{1}{\bar{\rho}} \mathcal{E}_{T, \eta} + \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T, \eta}^{\text{P}} + \frac{1}{\bar{\rho}} \mathcal{E}_{T, \eta}^{\text{D}}.$$

Next, we provide an upper bound on the cumulative constraints violation.

By Lemma 4.4, for every $i \in [m]$, we have that:

$$\sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \leq -(T - T_1)\rho + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta}. \quad (14)$$

Hence, for every $i \in [m]$, it holds

$$\begin{aligned} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t) &= \sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t) + \sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \\ &\leq (T - T_1)\tilde{\rho} + M_{\tilde{\rho}} - (T - T_1)\rho + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta} \\ &\leq M_{\tilde{\rho}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta} \\ &\leq M_{\tilde{\rho}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}} + \mathcal{E}_{T, \eta}. \end{aligned}$$

The first inequality follows from Equation (14) and by the condition in Line 4 of Algorithm 1, which ensures $\sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t) \leq (T - T_1)\tilde{\rho} + M_{\tilde{\rho}}$ for every $i \in [m]$. Moreover, the second inequality follows from $\tilde{\rho} \leq \rho$, since Condition 4.5 holds. Let $i^* \in \arg \max_{i \in [m]} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t)$ be one of the most violated constraints. By recalling that $V^T = \max_{i \in [m]} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t)$, the following bound on the cumulative constraints violation holds:

$$V^T = \sum_{t=1}^T g_{t,i^*}(\mathbf{x}_t) \leq M_{\tilde{\rho}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}} + \mathcal{E}_{T, \eta}.$$

This concludes the proof. \square

Theorem 4.7. *Suppose that functions f_t and g_t are selected adversarially and stochastically, respectively. Algorithm 1 guarantees that the following bounds hold with probability at least $1 - \delta$: $R_T \leq T^{1/4} \mathcal{E}_{T, \eta} + (1 + 2T^{1/4}) \mathcal{E}_{T, \eta}^{\text{P}} + T^{1/4} \mathcal{E}_T^{\text{D}}$ and $V_T \leq T^{3/4} + M_{T^{-1/4}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}} + \mathcal{E}_{T, \eta}$.*

Proof. If $\hat{\rho} \geq 2T^{-1/4}$, the claim follows by Theorem 4.6. Thus, we prove the statement for the case $\tilde{\rho} = T^{-1/4}$. First, we provide an upper bound on the cumulative regret. By Lemma 4.2, we have that event \mathbf{E} holds with probability at least $1 - \delta$. In the rest of the proof, we assume that the event \mathbf{E} holds, and provide a bound that holds with probability at least $1 - \delta$. We have:

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t) &\geq \sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \\ &\geq \sum_{t=1}^{T_1} f_t(\xi^*) + (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1, \eta} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\text{D}} \\ &\geq \sum_{t=1}^T f_t(\xi^*) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1, \eta} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\text{D}} \\ &\geq \sum_{t=1}^T f_t(\xi^*) - \frac{1}{\bar{\rho}} \mathcal{E}_{T, \eta} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_T^{\text{D}} \\ &\geq \sum_{t=1}^T f_t(\xi^*) - T^{1/4} \mathcal{E}_{T, \eta} - \left(1 + 2T^{1/4}\right) \mathcal{E}_{T, \eta}^{\text{P}} - T^{1/4} \mathcal{E}_T^{\text{D}}. \end{aligned}$$

These steps are similar to those used to prove the regret bound in Theorem 4.6 (see the proof of Theorem 4.6 for further details). By recalling that ξ^* is an optimal solution to Problem $\text{LP}_{\bar{f}, \bar{g}}$ and $R^T = T \text{OPT}_{\bar{f}, \bar{g}} - \sum_{t=1}^T f_t(\mathbf{x}_t)$, the following bound on the cumulative regret holds:

$$R^T = \sum_{t=1}^T f_t(\xi^*) - \sum_{t=1}^T f_t(\mathbf{x}_t) \leq T^{1/4} \mathcal{E}_{T, \eta} + \left(1 + 2T^{1/4}\right) \mathcal{E}_{T, \eta}^{\text{P}} + T^{1/4} \mathcal{E}_T^{\text{D}}.$$

Next, we provide an upper bound on the cumulative constraints violation.

For every $i \in [m]$, the following holds

$$\begin{aligned} \sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) &\leq -(T - T_1)\rho + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta} \\ &\leq 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta}, \end{aligned} \quad (15)$$

where the first inequality follows from Lemma 4.4, while the second one from $\rho \geq 0$. Hence, for every $i \in [m]$, it holds

$$\begin{aligned} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t) &= \sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t) + \sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \\ &\leq (T - T_1)\tilde{\rho} + M_{\tilde{\rho}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta} \\ &\leq (T - T_1)T^{-1/4} + M_{T^{-1/4}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta} \\ &\leq T^{3/4} + M_{T^{-1/4}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} + \mathcal{E}_{T-T_1}^{\text{D}} + \mathcal{E}_{T-T_1, \eta} \\ &\leq T^{3/4} + M_{T^{-1/4}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}} + \mathcal{E}_{T, \eta}. \end{aligned}$$

The first inequality follows from Equation (15) and from the condition in Line 4 of Algorithm 1, which ensures that $\sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t) \leq (T - T_1)\tilde{\rho} + M_{\tilde{\rho}}$ for every $i \in [m]$. Moreover, the second inequality follows from $\tilde{\rho} = T^{-1/4}$. Thus, by letting $i^* \in \arg \max_{i \in [m]} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t)$, and by recalling that $V^T = \max_{i \in [m]} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t)$, the following bound on the cumulative constraints violation holds:

$$V^T = \sum_{t=1}^T g_{t,i^*}(\mathbf{x}_t) \leq T^{3/4} + M_{T^{-1/4}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}} + \mathcal{E}_{T, \eta}.$$

This concludes the proof. □

B.3 Proof omitted from Section 5

First, we provide a preliminary result on the value of the Lagrangian game when primal and dual players are constrained to specific sets of strategies.

Lemma B.1. *Let $f \in \mathcal{F}$ and $g \in \mathcal{G}$ be such that $d_g > 0$. Moreover, given any $\epsilon > 0$, let $\Xi_{\epsilon, g} := \{\xi \in \Xi : \max_{i \in [m]} g_i(\xi) \geq \epsilon\}$. The following holds:*

$$\sup_{\xi \in \Xi_{\epsilon}} \inf_{\lambda \in \mathcal{D}_{d_g/2}} \mathcal{L}_{f, g}(\xi, \lambda) \leq \text{OPT}_{f, g} - \frac{\epsilon}{d_g}.$$

Proof. Let $\xi \in \Xi_{\epsilon, g}$ and $i^* \in \arg \max_{i \in [m]} g_i(\xi)$. Then,

$$\begin{aligned}
\inf_{\lambda \in \mathcal{D}_{d_g/2}} \left\{ f(\xi) - \langle \lambda, g(\xi) \rangle \right\} &= f(\xi) - \frac{2}{d_g} g_{i^*}(\xi) \\
&= \inf_{\lambda \in \mathcal{D}_{d_g}} \left\{ f(\xi) - \langle \lambda, g(\xi) \rangle \right\} - \frac{1}{d_g} g_{i^*}(\xi) \\
&\leq \sup_{\xi \in \Xi} \inf_{\lambda \in \mathcal{D}_{d_g}} \mathcal{L}_{f, g}(\xi, \lambda) - \frac{1}{d_g} g_{i^*}(\xi) \\
&\leq \text{OPT}_{f, g} - \frac{1}{d_g} g_{i^*}(\xi) \\
&\leq \text{OPT}_{f, g} - \frac{\epsilon}{d_g},
\end{aligned}$$

where the second inequality follows from Theorem 2.3, while the last one holds by the definition of $\Xi_{\epsilon, g}$ and i^* . \square

Next, we introduce a new event that extends \mathbf{E} by considering also the (stochastic) sequence of reward functions f_t . Formally, the event is defined as follows.

Definition B.2. We denote with $\bar{\mathbf{E}}$ the event in which Algorithm 1 satisfies the following conditions (recall that $\eta = \delta/3$): (i) event \mathbf{E} holds; (ii) for every pair of rounds $t, t' \in [T] : t \leq t'$ it holds:

- $\left| \sum_{\tau=t}^{t'} f_\tau(\mathbf{x}_\tau) - \sum_{\tau=t}^{t'} \bar{f}(\mathbf{x}_\tau) \right| \leq \mathcal{E}_{t'-t, \eta}$,
- $\left| \sum_{\tau=t}^{t'} f_\tau(\xi^*) - \sum_{\tau=t}^{t'} \bar{f}(\xi^*) \right| \leq \mathcal{E}_{t'-t, \eta}$.

Lemma B.3. After running Algorithm 1, the event $\bar{\mathbf{E}}$ holds with probability at least $1 - \delta$.

Proof. Given a desired failure probability $\delta \in (0, 1)$, recall that $\eta = \delta/3$ and set $\epsilon = \eta/12mT^2$. Consider the following inequalities in which the differences between expectations and empirical means of reward functions are bounded:

$$\left| \sum_{\tau=t}^{t'} f_\tau(\mathbf{x}_\tau) - \sum_{\tau=t}^{t'} \bar{f}(\mathbf{x}_\tau) \right| > 2\sqrt{2(t'-t) \ln \frac{1}{\epsilon}}, \quad (16)$$

$$\left| \sum_{\tau=t}^{t'} f_\tau(\xi^*) - \sum_{\tau=t}^{t'} \bar{f}(\xi^*) \right| > 2\sqrt{2(t'-t) \ln \frac{1}{\epsilon}}. \quad (17)$$

By applying the Azuma-Hoeffding inequality to each martingale difference sequence, we get that each inequality holds with probability at most 2ϵ . We denote by $\bar{\mathbf{E}}_\eta$ the event in which Equations (16) and (17) hold for every $t \leq t' \in [T] : t < t'$ and event \mathbf{E}_η holds (see the proof of Lemma 4.2 for the definition of event \mathbf{E}_η). By a union bound, we have that:

$$\mathbb{P}(\bar{\mathbf{E}}_\eta) \geq 1 - 2\epsilon(4mT^2 + 2T^2) \geq 1 - \eta.$$

Therefore, event $\bar{\mathbf{E}}_\eta$ holds with probability at least $1 - \eta$. Moreover, let us recall that:

$$\mathcal{E}_{t'-t, \eta} = \sqrt{8(t'-t) \ln \frac{12mT^2}{\eta}} = 2\sqrt{2(t'-t) \ln \frac{1}{\epsilon}}.$$

Now, consider the event $\bar{\mathbf{E}}$ in which Algorithm 1 satisfies the following conditions: (i) the regret incurred by $\mathcal{R}_I^{\mathbb{P}}$ after T_1 rounds is upper bounded by $\mathcal{E}_{T_1, \eta}^{\mathbb{P}}$; (ii) the regret cumulated by $\mathcal{R}_{II}^{\mathbb{P}}$ after the remaining $T - T_1$ rounds is upper bounded by $\mathcal{E}_{T-T_1, \eta}^{\mathbb{P}}$; and (iii) event $\bar{\mathbf{E}}_\eta$ holds. Recall that each one of the conditions (i), (ii) and (iii) holds with probability at least $1 - \eta$; hence, by a union bound we get:

$$\mathbb{P}(\bar{\mathbf{E}}) \geq 1 - 3\eta = 1 - \delta.$$

This concludes the proof. \square

As a first step, we prove that the primal regret minimizer gets a per-round utility that is close to the value $\text{OPT}_{\bar{f}, \bar{g}}$. Formally:

Lemma B.4. *If the event $\bar{\mathbf{E}}$ holds, then, for every round $\tau \in [T_1]$ the following inequality holds:*

$$\sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}_t) \geq \tau \text{OPT}_{\bar{f}, \bar{g}}^{\text{LP}} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} - \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}.$$

Proof. Let $\boldsymbol{\xi}^*$ be an optimal solution to Problem $\text{LP}_{\bar{f}, \bar{g}}$, and let $\bar{\boldsymbol{\lambda}} = \frac{1}{\tau} \sum_{t=1}^{\tau} \boldsymbol{\lambda}_t$. Then, it holds

$$\begin{aligned} \sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}_t) &\geq \sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\boldsymbol{\xi}^*, \boldsymbol{\lambda}_t) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} \\ &\geq \sum_{t=1}^{\tau} \mathcal{L}_{\bar{f}, \bar{g}}(\boldsymbol{\xi}^*, \boldsymbol{\lambda}_t) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} - \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta} \\ &= \sum_{t=1}^{\tau} \mathcal{L}_{\bar{f}, \bar{g}}(\boldsymbol{\xi}^*, \bar{\boldsymbol{\lambda}}) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} - \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta} \\ &\geq \tau \inf_{\boldsymbol{\lambda} \in \mathcal{D}_{\bar{\rho}}} \mathcal{L}_{\bar{f}, \bar{g}}(\boldsymbol{\xi}^*, \boldsymbol{\lambda}) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} - \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta} \\ &= \tau \sup_{\boldsymbol{\xi} \in \Xi} \inf_{\boldsymbol{\lambda} \in \mathcal{D}_{\bar{\rho}}} \mathcal{L}_{\bar{f}, \bar{g}}(\boldsymbol{\xi}, \boldsymbol{\lambda}) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} - \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta} \\ &= \tau \text{OPT}_{\bar{f}, \bar{g}}^{\text{LP}} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} - \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta}, \end{aligned}$$

where the first inequality follows from the no-regret property of the primal regret minimizer, the second one from the definition of the event $\bar{\mathbf{E}}$, and the third one from the definition of $\boldsymbol{\xi}^*$. Moreover, the first equality follows from the fact that \bar{g} is independent from t . This concludes the proof. \square

Now, we show that the dual regret minimizer gets a per-round utility that is close to the value $\text{OPT}_{\bar{f}, \bar{g}}$. Moreover, the attained utility increases by an additive factor proportional to the primal violation. This can be proved only in the setting with stochastic reward functions. Indeed, in this setting the primal and dual regret minimizers are playing a stochastic repeated zero-sum game that converges to an equilibrium. Notice that this is *not* true when the reward functions are adversarial.

Lemma B.5. *If event $\bar{\mathbf{E}}$ holds and Condition 4.5 is satisfied, then for each $\tau \in [T_1]$ and each $i \in [m]$*

$$\sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\boldsymbol{\xi}_t, \boldsymbol{\lambda}_t) \leq \tau \text{OPT}_{\bar{f}, \bar{g}}^{\text{LP}} + \frac{1}{\bar{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta} - \sum_{t=1}^{\tau} g_{t, i}(\mathbf{x}_t).$$

Proof. In the following, let $\boldsymbol{\lambda}^* \in \arg \min_{\boldsymbol{\lambda} \in \mathcal{D}_{\bar{\rho}}} \sum_{t=1}^{\tau} \mathcal{L}_{\bar{f}, \bar{g}}(\boldsymbol{\xi}_t, \boldsymbol{\lambda})$, $\epsilon := \frac{\max_{i \in [m]} \sum_{t=1}^{\tau} g_{t, i}(\mathbf{x}_t) - \mathcal{E}_{\tau, \eta}}{\tau}$, and $\bar{\boldsymbol{\xi}} := \frac{1}{\tau} \sum_{t=1}^{\tau} \boldsymbol{\xi}_t$, where $\boldsymbol{\xi}_t \in \Xi$ denotes the strategy mixture that plays deterministically \mathbf{x}_t . Moreover, let us define the set $\Xi_{\epsilon, \bar{g}} := \{\boldsymbol{\xi} \in \Xi : \max_{i \in [m]} \bar{g}_i(\boldsymbol{\xi}) \geq \epsilon\}$.

As a first step, we prove that $\bar{\boldsymbol{\xi}} \in \Xi_{\epsilon, \bar{g}}$. In particular, since the event $\bar{\mathbf{E}}$ holds, we have that

$$\max_{i \in [m]} \bar{g}_i(\bar{\boldsymbol{\xi}}) \geq \frac{\sum_{t=1}^{\tau} \max_{i \in [m]} g_i(\bar{\boldsymbol{\xi}}) - \mathcal{E}_{\tau, \eta}}{\tau} = \epsilon$$

For every $\tau \in [T_1]$, we have:

$$\sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}_t) \leq \sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}^*) + \frac{1}{\bar{\rho}} \mathcal{E}_{\tau}^{\text{D}} \quad (18a)$$

$$\leq \sum_{t=1}^{\tau} \mathcal{L}_{\bar{f}, \bar{g}}(\mathbf{x}_t, \boldsymbol{\lambda}^*) + \frac{1}{\bar{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\bar{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18b)$$

$$\leq \inf_{\lambda \in \mathcal{D}_{\tilde{\rho}}} \sum_{t=1}^{\tau} \mathcal{L}_{\tilde{f}, \tilde{g}}(\mathbf{x}_t, \lambda) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18c)$$

$$= \tau \inf_{\lambda \in \mathcal{D}_{\tilde{\rho}}} \mathcal{L}_{\tilde{f}, \tilde{g}}(\bar{\xi}, \lambda) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18d)$$

$$= \tau \inf_{\lambda \in \mathcal{D}_{\rho/2}} \mathcal{L}_{\tilde{f}, \tilde{g}}(\bar{\xi}, \lambda) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18e)$$

$$\leq \tau \sup_{\xi \in \Xi_{\epsilon, \tilde{g}}} \inf_{\lambda \in \mathcal{D}_{\rho/2}} \mathcal{L}_{\tilde{f}, \tilde{g}}(\xi, \lambda) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18f)$$

$$\leq \tau \sup_{\xi \in \Xi} \inf_{\lambda \in \mathcal{D}_{\rho}} \left(\mathcal{L}_{\tilde{f}, \tilde{g}}(\xi, \lambda) - \frac{\epsilon}{\rho} \right) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18g)$$

$$= \tau \left(\text{OPT}_{\tilde{f}, \tilde{g}} - \frac{\epsilon}{\rho} \right) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18h)$$

$$= \tau \text{OPT}_{\tilde{f}, \tilde{g}} - \tau \frac{\max_{i' \in [m]} \sum_{t=1}^{\tau} g_{t, i'}(\mathbf{x}_t) - \mathcal{E}_{\tau, \eta}}{\tau \rho} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \quad (18i)$$

$$\leq \tau \text{OPT}_{\tilde{f}, \tilde{g}} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} - \max_{i' \in [m]} \frac{\sum_{t=1}^{\tau} g_{t, i'}(2\mathbf{x}_t)}{\rho} \quad (18j)$$

$$\leq \tau \text{OPT}_{\tilde{f}, \tilde{g}} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} - \max_{i' \in [m]} \sum_{t=1}^{\tau} g_{t, i'}(\mathbf{x}_t) \quad (18k)$$

$$\leq \tau \text{OPT}_{\tilde{f}, \tilde{g}} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} - \sum_{t=1}^{\tau} g_{t, i}(\mathbf{x}_t) \quad \forall i \in [m], \quad (18l)$$

where Equation (18a) is given by the no-regret property of the dual regret minimizer, and Equation (18b) by the definition of the event $\bar{\mathbf{E}}$, which holds by assumption. Moreover, Equation (18d) follows from the fact that \tilde{f} and \tilde{g} are independent from t , Equation (18e) follows from $\tilde{\rho} = \hat{\rho}/2 \leq \rho/2$, and Equation (18f) from $\bar{\xi} \in \Xi_{\epsilon, \tilde{g}}$. Finally, Equation (18g) follows from Lemma B.1, Equation (18i) by definition of ϵ , and Equation (18j) by $\tilde{\rho} \leq \rho$. \square

Theorem 5.1. *Suppose that functions f_t and g_t are selected stochastically. With probability at least $1 - \delta$, Algorithm 1 never enters the recovery phase, namely $T_1 = T$.*

Proof. We prove the statement of the theorem by considering two cases.

Case ‘‘Condition 4.5 holds’’. By Lemma 4.2, event \mathbf{E} holds with probability at least $1 - \delta$. In the rest of the proof, we assume that the event \mathbf{E} holds, and we provide a bound that holds with probability at least $1 - \delta$. For every $\tau \in [T_1]$, we have:

$$\begin{aligned} \sum_{t=1}^{\tau} g_t(\mathbf{x}_t) &\leq \tau \text{OPT}_{\tilde{f}, \tilde{g}} - \sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \lambda_t) + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \\ &\leq \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} + \left(1 + \frac{1}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} \\ &= \left(2 + \frac{3}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} \\ &\leq \frac{2}{\tilde{\rho}} \sqrt{T} - 1 + \left(2 + \frac{3}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta} + \left(1 + \frac{2}{\tilde{\rho}}\right) \mathcal{E}_{\tau, \eta}^{\text{P}} + \frac{1}{\tilde{\rho}} \mathcal{E}_{\tau}^{\text{D}} \\ &= M_{\tilde{\rho}} - 1, \end{aligned}$$

where the first inequality follows from Lemma B.4, the second one from Lemma B.5, the third one from the fact that $\frac{2}{\tilde{\rho}} \sqrt{T} - 1 \geq 0$, being $\tilde{\rho} \leq 1$, and the last equation follows from the definition of $M_{\tilde{\rho}}$. This implies that the algorithm never enters the recovery phase when Condition 4.5 holds.

Case “Condition 4.5 does *not* hold”. By Lemma 4.2, event **E** holds with probability at least $1 - \delta$. In the rest of the proof, we assume that the event **E** holds, and we provide a bound that holds with probability at least $1 - \delta$. Suppose by contradiction that $T_1 < T$. This implies that a constraint $i \in [m]$ is violated by at least $M_{T^{-1/4}} - 1$. Let $i^* \in \arg \max_{i \in [m]} \sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t)$ be one of the most violated constraints during the play phase. Then, we have:

$$\begin{aligned} \sum_{t=1}^{T_1} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}_t) &= \sum_{t=1}^{T_1} \left(f(\mathbf{x}_t) - \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \right) \\ &\leq T_1 - \sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \\ &\leq T_1 - \sum_{t=1}^{T_1} \frac{1}{T^{-1/4}} g_{t, i^*}(\mathbf{x}_t) + T^{1/4} \mathcal{E}_{T_1}^D \\ &\leq T_1 - T^{1/4} (M_{T^{-1/4}} - 1) + T^{1/4} \mathcal{E}_{\tau_1}^D \\ &< - \left(1 + \frac{2}{T^{-1/4}} \right) \mathcal{E}_{\tau, \eta}^P - \frac{1}{T^{-1/4}} \mathcal{E}_{\tau, \eta}, \end{aligned}$$

where the second inequality follows from the no-regret property of the dual regret minimizer and the fact that, when Condition 4.5 does *not* hold, $\tilde{\rho} = T^{-1/4}$. The last inequality follows from the definition of $M_{T^{-1/4}}$. Then, the result above allows us to reach the desired contradiction when compared with the following one. In particular, for every $\tau \in [T_1]$, we have:

$$\begin{aligned} \sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}_t) &\geq \sum_{t=1}^{\tau} \mathcal{L}_{f_t, g_t}(\boldsymbol{\xi}^\circ, \boldsymbol{\lambda}_t) - \left(1 + \frac{2}{T^{-1/4}} \right) \mathcal{E}_{\tau, \eta}^P \\ &\geq \sum_{t=1}^{\tau} \mathcal{L}_{f_t, \bar{g}}(\boldsymbol{\xi}^\circ, \boldsymbol{\lambda}_t) - \frac{1}{T^{-1/4}} \mathcal{E}_{\tau, \eta} - \left(1 + \frac{2}{T^{-1/4}} \right) \mathcal{E}_{\tau, \eta}^P \\ &\geq - \frac{1}{T^{-1/4}} \mathcal{E}_{\tau, \eta} - \left(1 + \frac{2}{T^{-1/4}} \right) \mathcal{E}_{\tau, \eta}^P, \end{aligned}$$

where the first inequality follows from the no-regret property of the primal regret minimizer, the second one follows from the fact that event **E** holds, and the third one from the feasibility of $\boldsymbol{\xi}^\circ$. \square

B.4 Proofs omitted from Section 6

As a first step, we provide a lower bound for the cumulative reward achieved during the play phase. In particular, we show that it achieves at least a $\rho/(1 + \rho)$ fraction of the value obtained by an optimal solution in the first T_1 rounds.

Lemma B.6. *If Condition 4.5 is satisfied, then, with probability at least $1 - \eta$, at round T_1 of Algorithm 1 it holds that:*

$$\sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \geq \frac{\rho}{1 + \rho} \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \left(1 + \frac{2}{\tilde{\rho}} \right) \mathcal{E}_{T_1, \eta}^P - \frac{1}{\tilde{\rho}} \mathcal{E}_{T_1}^D.$$

Proof. Let $\bar{\boldsymbol{\xi}} \in \Xi$ be a strategy mixture obtained by playing with probability $1/(1 + \rho)$ the mixture $\boldsymbol{\xi}^\circ$ and with the remaining probability $\rho/(1 + \rho)$ an optimal mixture $\boldsymbol{\xi}^*$. Notice that the probabilities are well defined, since $\rho \geq 0$. Then, for every $t \in [T]$ and $i \in [m]$, it holds:

$$\frac{1}{1 + \rho} g_{t,i}(\boldsymbol{\xi}^\circ) + \frac{\rho}{1 + \rho} g_{t,i}(\boldsymbol{\xi}^*) \leq -\frac{\rho}{1 + \rho} + \frac{\rho}{1 + \rho} = 0$$

where the inequality follows from the fact that $g_{t,i}(\boldsymbol{\xi}^\circ) \leq -\rho$ and $g_{t,i}(\boldsymbol{\xi}^*) \leq 1$. Therefore, for every $t \in [T]$ and $i \in [m]$, it holds that $g_{t,i}(\bar{\boldsymbol{\xi}}) \leq 0$. Assume that the regret bounds of the regret minimizers hold. Notice that this happens with probability at least $1 - \eta$. Then, by the no-regret property of the

primal regret minimizer, we have that

$$\sum_{t=1}^{T_1} \mathcal{L}_{f_t, g_t}(\mathbf{x}_t, \boldsymbol{\lambda}_t) \geq \sum_{t=1}^{T_1} \mathcal{L}_{f_t, g_t}(\bar{\boldsymbol{\xi}}, \boldsymbol{\lambda}_t) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\mathbb{P}}. \quad (19)$$

Let $i^* \in \arg \max_{i \in [m]} \sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t)$ be one of the most violated constraints during the play phase. Next, we prove that

$$\sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \geq (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}.$$

We consider two cases. If $T_1 = T$, then

$$\sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \geq \sum_{t=1}^{T_1} \langle \mathbf{0}, g_t(\bar{\boldsymbol{\xi}}) \rangle - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}} = -\frac{1}{\bar{\rho}} \mathcal{E}_{\tau_2}^{\mathbb{D}} = (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}.$$

Otherwise, we have that $\sum_{t=1}^{T_1} g_{t,i^*}(\mathbf{x}_t) \geq \tilde{\rho}(T - T_1)$ and

$$\sum_{t=1}^{T_1} \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \geq \left(\sum_{t=1}^{T_1} \frac{1}{\tilde{\rho}} g_{t,i^*}(\mathbf{x}_t) \right) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}} \geq (T - T_1) - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}. \quad (20)$$

Thus,

$$\begin{aligned} \sum_{t=1}^{T_1} f_t(\mathbf{x}_t) &\geq \sum_{t=1}^{T_1} \left(f_t(\bar{\boldsymbol{\xi}}) - \langle \boldsymbol{\lambda}_t, g_t(\bar{\boldsymbol{\xi}}) \rangle + \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle \right) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\mathbb{P}} \\ &\geq \sum_{t=1}^{T_1} \left(f_t(\bar{\boldsymbol{\xi}}) - \langle \boldsymbol{\lambda}_t, g_t(\bar{\boldsymbol{\xi}}) \rangle \right) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{\tau_2}^{\mathbb{D}} \\ &\geq \sum_{t=1}^{T_1} f_t(\bar{\boldsymbol{\xi}}) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{\tau_2}^{\mathbb{D}} \\ &\geq \sum_{t=1}^{T_1} \left(\frac{1}{1 + \rho} f_t(\boldsymbol{\xi}^\circ) + \frac{\rho}{1 + \rho} f_t(\boldsymbol{\xi}^*) \right) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{\tau_2}^{\mathbb{D}} \\ &\geq \frac{\rho}{1 + \rho} \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\mathbb{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\mathbb{D}}, \end{aligned}$$

where the first inequality follows from Equation (19), the second one from Equation 20, the third one from the fact that for each $t \in [T]$ it holds $g_t(\bar{\boldsymbol{\xi}}) \leq 0$, while the fourth inequality follows from the definition of $\bar{\boldsymbol{\xi}}$. This concludes the proof. \square

Notice that, for a small T_1 , we have a large lower bound on the cumulative reward. Intuitively, this means that when the play phase is short, the primal regret minimizer accumulated so much regret in the play phase that the recovery phase can be addressed without worrying about the reward.

As a second step, we provide an upper bound on the cumulative constraints violation during the recovery phase. In particular, we show that the constraints are satisfied by at least ρ at each round up to a term related to the regret of $\mathcal{R}^{\mathbb{P}}$ and $\mathcal{R}^{\mathbb{D}}$.

Lemma B.7. *With probability at least $1 - \eta$, when Algorithm 1 halts it holds that for each $i \in [m]$:*

$$\sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \leq -(T - T_1)\rho + \mathcal{E}_{T-T_1}^{\mathbb{D}} + 2\mathcal{E}_{T-T_1, \eta}^{\mathbb{P}}.$$

Proof. Let i^* be one of the most violated constraints, i.e., $i^* \in \arg \max_{i \in [m]} \sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t)$. Then, we have that:

$$\begin{aligned} (T - \tau)\rho &\leq - \sum_{t=T_1+1}^T \langle \boldsymbol{\lambda}_t, g_t(\boldsymbol{\xi}^\circ) \rangle \\ &\leq - \sum_{t=T_1+1}^T \langle \boldsymbol{\lambda}_t, g_t(\mathbf{x}_t) \rangle + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} \\ &\leq - \sum_{t=T_1+1}^T g_{t,i^*}(\mathbf{x}_t) + \mathcal{E}_{T-T_1}^{\text{D}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}}, \end{aligned}$$

where the first inequality follows from the definition of $\boldsymbol{\xi}^\circ$ and the fact that it is always feasible of at least ρ , the second one follows from the assumption that the primal regret minimizer satisfies the regret bound, and the last inequality from the guarantee on the regret of the dual regret minimizer. We conclude the proof by noticing that the regret bound of the primal regret minimizer holds with probability at least $1 - \eta$. \square

Now, we can provide our bounds for adversarial constraints.

Theorem 6.1. *Suppose that functions f_t and g_t are selected adversarially. If Condition 4.5 is satisfied, then, with probability at least $1 - \frac{2}{3}\delta$, Algorithm 1 guarantees that the following holds: $\sum_{t=1}^T f_t(\mathbf{x}_t) \geq \frac{\rho}{1+\rho} \sum_{t=1}^T \text{OPT}_{f, \bar{g}} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_T^{\text{D}}$ and $V^T \leq M_{\bar{\rho}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}}$.*

Proof. In the following, we assume that both Lemma B.6 and Lemma B.7. By an union bound, this holds with probability $1 - 2\eta = 1 - \frac{2}{3}\delta$. Then, it holds

$$\begin{aligned} \sum_{t=1}^T f_t(\mathbf{x}_t) &\geq \sum_{t=1}^{T_1} f_t(\mathbf{x}_t) \\ &\geq \sum_{t=1}^{T_1} \frac{\rho}{1+\rho} f_t(\boldsymbol{\xi}^*) + (T - T_1) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\text{D}} \\ &\geq \frac{\rho}{1+\rho} \sum_{t=1}^{T_1} f_t(\boldsymbol{\xi}^*) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T_1, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_{T_1}^{\text{D}} \\ &\geq \frac{\rho}{1+\rho} \sum_{t=1}^T f_t(\boldsymbol{\xi}^*) - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_T^{\text{D}}, \end{aligned}$$

where the second inequality comes from Lemma B.6. This proves the bound on the regret.

By Lemma B.7, for each $i \in [m]$,

$$\begin{aligned} \sum_{t=1}^T g_{t,i}(\mathbf{x}_t) &= \sum_{t=1}^{T_1} g_{t,i}(\mathbf{x}_t) + \sum_{t=T_1+1}^T g_{t,i}(\mathbf{x}_t) \\ &\leq (T - T_1)\bar{\rho} + M_{\bar{\rho}} - (T - T_1)\rho + \mathcal{E}_{T-T_1}^{\text{D}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} \\ &\leq M_{\bar{\rho}} + \mathcal{E}_{T-T_1}^{\text{D}} + 2\mathcal{E}_{T-T_1, \eta}^{\text{P}} \\ &\leq M_{\bar{\rho}} + \mathcal{E}_T^{\text{D}} + 2\mathcal{E}_{T, \eta}^{\text{P}}, \end{aligned}$$

where the second inequality comes from $\bar{\rho} \leq \rho$. \square

Corollary 6.2. *Suppose functions f_t and g_t are selected stochastically and adversarially, respectively. If Condition 4.5 is satisfied, then, with probability at least $1 - \delta$, Algorithm 1 provides $\sum_{t=1}^T f_t(\mathbf{x}_t) \geq \frac{\rho}{1+\rho} \sum_{t=1}^T \text{OPT}_{f, \bar{g}} - \left(1 + \frac{2}{\bar{\rho}}\right) \mathcal{E}_{T, \eta}^{\text{P}} - \frac{1}{\bar{\rho}} \mathcal{E}_T^{\text{D}} - 2\mathcal{E}_{T, \eta}$ and $V^T \leq M_{\bar{\rho}} + 2\mathcal{E}_{T, \eta}^{\text{P}} + \mathcal{E}_T^{\text{D}} + \mathcal{E}_{T, \eta}$.*

Proof. It is easy to see that Theorem 6.1 can be extended to consider the definition of ξ^* for stochastic rewards. formally, it holds $\sum_{t=1}^T f_t(x_t) \geq \frac{\rho}{1+\rho} \sum_{t=1}^T f_t(\xi^*) - \left(1 + \frac{2}{\rho}\right) \mathcal{E}_{T,\eta}^{\mathbb{P}} - \frac{1}{\rho} \mathcal{E}_T^{\mathbb{D}}$. Consider the two martingale difference sequences $\sum_{t=1}^T f_t(\mathbf{x}_t) - \bar{f}(\mathbf{x}_t)$ and $\sum_t f_t(\xi^*) - \bar{f}(\xi^*)$. We can apply Azuma-Hoeffding inequality to prove that, with probability at least $1 - \eta$, it holds $\sum_t |f_t(\mathbf{x}_t) - \bar{f}(\mathbf{x}_t)| \leq \mathcal{E}_{T,\eta}$ and $\sum_t |f_t(\xi^*) - \bar{f}(\xi^*)| \leq \mathcal{E}_{T,\eta}$. Then,

$$\begin{aligned} \sum_{t=1}^T \bar{f}(\mathbf{x}_t) &\geq \sum_{t=1}^T f_t(\mathbf{x}_t) - \mathcal{E}_{T,\eta} \\ &\geq \frac{\rho}{1+\rho} \sum_{t=1}^T f_t(\xi^*) - \left(1 + \frac{2}{\rho}\right) \mathcal{E}_{T,\eta} - \frac{1}{\rho} \mathcal{E}_T^{\mathbb{D}} - \mathcal{E}_{T,\eta} \\ &\geq \frac{\rho}{1+\rho} \sum_{t=1}^T \bar{f}(\xi^*) - \left(1 + \frac{2}{\rho}\right) \mathcal{E}_{T,\eta} - \frac{1}{\rho} \mathcal{E}_T^{\mathbb{D}} - 2\mathcal{E}_{T,\eta}, \end{aligned}$$

proving the statement. \square

B.5 Proofs omitted from Section 7

Lemma 7.1. *If $T_0 = \sqrt{T}$, then Algorithm 3 guarantees that $\hat{\rho} \leq \rho$ with probability at least $1 - \delta$.*

Proof. By Azuma-Hoeffding inequality, we have that with probability at least $1 - \delta$, for each $i \in [m]$ it holds $\left| \sum_{t=1}^{T_0} g_{t,i}(\mathbf{x}_t) - \bar{g}_i(\mathbf{x}_t) \right|$. Hence,

$$-\max_{i \in [m]} \sum_{t=1}^{T_0} g_{t,i}(\mathbf{x}_t) \leq -\max_{i \in [m]} \sum_{t=1}^{T_0} \bar{g}_i(\mathbf{x}_t) + \mathcal{E}_{T_0,\delta} \leq T_0 \bar{g}(\xi^\circ) + \mathcal{E}_{T_0,\delta} = T_0 \rho + \mathcal{E}_{T_0,\delta},$$

where the second and third inequality follow from the definition of ξ° . Then,

$$\begin{aligned} \hat{\rho} &= -\frac{1}{T_0} \left(\max_{i \in [m]} \sum_{t=1}^{T_0} g_{t,i}(\mathbf{x}_t) + \mathcal{E}_{T_0,\delta} \right) \\ &\leq \frac{1}{T_0} (T_0 \rho + \mathcal{E}_{T_0,\delta} - \mathcal{E}_{T_0,\delta}) \\ &= \rho. \end{aligned}$$

This concludes the proof. \square

Lemma 7.4. *By setting $T_0 = \sqrt{T}$, and assuming that Condition 7.2 is satisfied, after T_0 rounds of Algorithm 3 we have that $\hat{\rho} \geq \rho/2$ with probability at least $1 - 2\delta$.*

Proof. First, notice that with probability $1 - \delta$, the primal regret minimizer has regret bounded by $\mathcal{E}_{T_0,\delta}^{\mathbb{P}}$. Moreover, by the Azuma-Hoeffding inequality, it holds $\left| \sum_{t=1}^{T_0} \lambda_t g_t(\xi^\circ) - \lambda_t \bar{g}(\xi^\circ) \right| \leq \mathcal{E}_{T_0,\delta}$ with probability $1 - \delta$. Consider the case in which both the conditions hold. This happens with probability at least $1 - 2\delta$ by a union bound.

Then,

$$\begin{aligned} -\max_{i \in [m]} \sum_{t=1}^{T_0} g_{t,i}(\mathbf{x}_t) &\geq -\sum_{t=1}^{T_0} \langle \lambda_t, g_t(\mathbf{x}_t) \rangle - \mathcal{E}_{T_0}^{\mathbb{D}} \\ &\geq -\sum_{t=1}^{T_0} \langle \lambda_t, g_t(\xi^\circ) \rangle - \mathcal{E}_{T_0}^{\mathbb{D}} - 2\mathcal{E}_{T_0,\delta}^{\mathbb{P}} \\ &\geq -\sum_{t=1}^{T_0} \langle \lambda_t, \bar{g}(\xi^\circ) \rangle - \mathcal{E}_{T_0}^{\mathbb{D}} - 2\mathcal{E}_{T_0,\delta}^{\mathbb{P}} - \mathcal{E}_{T_0,\delta} \\ &\geq T_0 \rho - \mathcal{E}_{T_0}^{\mathbb{D}} - 2\mathcal{E}_{T_0,\delta}^{\mathbb{P}} - \mathcal{E}_{T_0,\delta}. \end{aligned}$$

Hence,

$$\begin{aligned}
\hat{\rho} &= -\frac{1}{T_0} \left(\max_{i \in [m]} \sum_{t=1}^{T_0} g_{t,i}(\mathbf{x}_t) + \mathcal{E}_{T_0,\delta} \right) \\
&\geq \frac{1}{T_0} (T_0 \rho - \mathcal{E}_{T_0}^D - 2\mathcal{E}_{T_0,\delta}^P - \mathcal{E}_{T_0,\delta} - \mathcal{E}_{T_0,\delta}) \\
&\geq \rho/2 + \frac{1}{T_0} (T_0 \rho/2 - \mathcal{E}_{T_0}^D - 2\mathcal{E}_{T_0,\delta}^P - 2\mathcal{E}_{T_0,\delta}) \\
&\geq \rho/2,
\end{aligned}$$

where the last inequality comes from Condition 7.2. This concludes the proof. \square

C Applications

In this section, we provide more details on possible applications of our techniques to settings related to Internet advertising platforms.

C.1 Budget-management with ROI constraints

Traditional budget-pacing mechanisms (see, *e.g.*, [8, 11]) are based on primal-dual algorithms which are shown to be near optimal in settings that involve only knapsack constraints, but cannot be generalized to deal with other types of long-term constraints. However, there are many real-world scenarios in which guaranteeing other types of long-term constraints is crucial for practical applications. One example is the case of *return on investment* (ROI) constraints [6, 26, 32].⁹ The recent work by Golrezaei et al. [25] presents a threshold-based algorithms for repeated *second-price* auctions under budget and ROI constraints. Our framework allows advertisers to reach a target ROI while keeping budget expenditures under control also in the setting of repeated *first-price* auctions, which is a frequent setting in practice.¹⁰ In particular, given a target ROI $\omega \geq 0$ and the largest among competing bids β_t , we define the ROI constraints as

$$g_t(b_t) = \left(\omega - \frac{v_t}{b_t} \right) \mathbb{1}\{b_t \geq \beta_t\} \leq 0.$$

Then, it is enough to instantiate the framework with the same setup of Section 8, that is, EXP3.P [5] for each of the possible valuations in \mathcal{V} , and OMD equipped with negative-entropy regularizer for the dual RM. Therefore, we immediately get that the cumulative violations of the budget and ROI constraints are upper bounded by $\tilde{O}(T^{1/2})$. This holds both in the fully stochastic and in the fully adversarial setting under the assumption of having a strictly feasible solution, which is reasonable since it is enough to have a sufficiently *small* bid in the set of available bids \mathcal{B} . We observe that always bidding such a *small* bid is sufficient to satisfy the ROI constraints but will penalize the cumulative rewards obtained by the advertiser.

C.2 Future research direction: fairness constraints

Consider the setting in which each item appearing at time t is characterized by one or more of n_c categories according to the vector $e_t \in [0, 1]^{n_c}$. A bidder may have distributional preferences over such categories, such as ensuring that at least a certain fraction of impressions is allocated to each category. This is the case, for example, of advertisers who need to perform online outreach to a population of users while achieving a distribution over different demographics *close* to that of the real underlying population. For example, Gelauff et al. [24] provide an interesting field study about running advertising campaigns for Participatory Budgeting elections. In Participatory Budgeting elections, community members are asked to vote between various public projects in order to allocate a total budget. The election organizer may use online advertising to try to promote the initiative,

⁹This is a frequent advertising objective in large Internet advertising platform. See, *e.g.*, <https://tinyurl.com/c86rezhd> and <https://tinyurl.com/mr49vz8a>.

¹⁰For example, in 2019 Google announced a shift to first-price auctions for its AdManager exchange. See <https://tinyurl.com/chv5nxys>.

and in doing so the goal is to reach a “demographic mix” comparable to that of the local population. Surprisingly, Gelauff et al. [24] show that advertisers currently have to resort to complex segmentation strategies through subcampaigns in order to achieve that goal.

Two recent works propose to achieve such distributional preferences within budget-pacing mechanisms by embedding them into a concave regularization term in the advertiser’s objective [10, 18]. Such frameworks specifically consider the case of repeated second-price auctions, and can directly handle only packing constraints. Encoding distributional preferences via a regularization term in the objective implies that they cannot provide any formal guarantee w.r.t. how *close* the realized distribution of impressions is to the target, despite showing promising performance in practice.

Differently from previous work, our framework can *explicitly* handle distributional constraints within second- and first-price auction frameworks. Let vector $\hat{e} \in [0, 1]^{n_c}$ be such that \hat{e}_j is the fraction of impressions that we want to be allocated to users of category j . Then, for each category $j \in [n_c]$, we could enforce the following type of constraints

$$g_{t,j}(b_t) := \hat{e}_j - e_{t,j} \mathbb{1}\{b_t \geq \beta_t\} \leq 0.$$

Assuming the existence of a strictly feasible bidding strategy, our framework guarantees that, for each category j ,

$$\hat{e}_j - \frac{1}{T} \sum_{t=1}^T e_{t,j} \mathbb{1}\{b_t \geq \beta_t\} \leq \tilde{O}(T^{-1/2}),$$

which guarantees that, in the limit, the difference between the average distribution of impressions and the target thresholds is vanishing.

The main question which still needs to be answered in order to apply our framework in the case of fairness constraints is whether we can motivate the existence of a strictly feasible solution. One reasonable requirement is to constrain the target vector \hat{e} to be a point in the full-dimensional simplex with dimension n_c . On top of that, the advertiser would need a way to “buy what’s necessary” in order to match the distributional constraints. This desideratum could be achieved, for example, by introducing *buyout options* for advertisers, in the spirit of Gallien and Gupta [23] (*i.e.*, when the advertiser needs impression from a certain category, they always have the option of bidding the buyout value to be sure to win the relevant items). Therefore, assuming the population of users is large enough, an advertiser could achieve a strictly feasible solution by bidding according to the fixed strategy mixture recommending to bid the buyout option for each category j with a probability greater than or equal to \hat{e}_j .

The model we described is clearly a simplification of real budget-pacing systems. Moreover, the practical implications of introducing buyout options should be further investigated, in order to understand if they constitute a viable solution both for the platform and advertisers. Finally, we leave as interesting future research directions the problem of studying the general setting (with arbitrary sets \mathcal{V} and \mathcal{B}), and that of providing an empirical evaluation of the above techniques on real-world data.