

---

# Conic Blackwell Algorithm: Parameter-Free Convex-Concave Saddle-Point Solving

---

**Julien Grand-Clément\***  
ISOM Department  
HEC Paris  
grand-clement@hec.fr

**Christian Kroer**  
IEOR Department  
Columbia University  
christian.kroer@columbia.edu

## Abstract

We develop new parameter-free and scale-free algorithms for solving convex-concave saddle-point problems. Our results are based on a new simple regret minimizer, the Conic Blackwell Algorithm<sup>+</sup> (CBA<sup>+</sup>), which attains  $O(1/\sqrt{T})$  average regret. Intuitively, our approach generalizes to other decision sets of interest ideas from the Counterfactual Regret minimization (CFR<sup>+</sup>) algorithm, which has very strong practical performance for solving sequential games on simplexes. We show how to implement CBA<sup>+</sup> for the simplex,  $\ell_p$  norm balls, and ellipsoidal confidence regions in the simplex, and we present numerical experiments for solving matrix games and distributionally robust optimization problems. Our empirical results show that CBA<sup>+</sup> is a simple algorithm that outperforms state-of-the-art methods on synthetic data and real data instances, without the need for any choice of step sizes or other algorithmic parameters.

## 1 Introduction

We are interested in solving *saddle-point problems* (SPPs) of the form

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{y} \in \mathcal{Y}} F(\mathbf{x}, \mathbf{y}), \quad (1)$$

where  $\mathcal{X} \subset \mathbb{R}^n$ ,  $\mathcal{Y} \subset \mathbb{R}^m$  are convex, compact sets, and  $F : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  is a subdifferentiable convex-concave function. Convex-concave SPPs arise in a number of practical problems. For example, the problem of computing a Nash equilibrium of a zero-sum game can be formulated as a convex-concave SPP, and this is the foundation of most methods for solving sequential zero-sum games [vS96, ZJBP07, TBJB15, KWKKS20]. They also arise in imaging [CP11],  $\ell_\infty$ -regression [ST18], Markov Decision Processes [Iye05, WKR13, ST18], and in distributionally robust optimization, where the max term represents the distributional uncertainty [ND16, BTHKM15]. In this paper we propose efficient, *parameter-free* algorithms for solving (1) in many settings, i.e., algorithms that do not require any tuning or choices of step sizes.

**Repeated game framework** One way to solve convex-concave SPPs is by viewing the SPP as a repeated game between two players, where each step  $t$  consists of one player choosing  $\mathbf{x}_t \in \mathcal{X}$ , the other player choosing  $\mathbf{y}_t \in \mathcal{Y}$ , and then the players observe the payoff  $F(\mathbf{x}_t, \mathbf{y}_t)$ . If each player employs a regret-minimization algorithm, then a well-known folk theorem says that the uniform average strategy generated by two regret minimizers repeatedly playing an SPP against each other converges to a solution to the SPP. We will call this the “repeated game framework” (see Section 2). There are already well-known algorithms for instantiating the above repeated game framework

---

\*Julien Grand-Clément acknowledges the financial support of Hi! Paris for this project.

for (1). For example, one can employ the *online mirror descent* (OMD) algorithm, which generates iterates as follows for the first player (and similarly for the second player):

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \langle \eta \mathbf{f}_t, \mathbf{x} \rangle + D(\mathbf{x} \| \mathbf{x}_t), \quad (2)$$

where  $\mathbf{f}_t \in \partial_{\mathbf{x}} F(\mathbf{x}_t, \mathbf{y}_t)$  ( $\partial_{\mathbf{x}}$  denotes here the set of subgradients as regards the variable  $\mathbf{x}$ ),  $D(\cdot \| \cdot)$  is a Bregman divergence which measures distance between pairs of points, and  $\eta > 0$  is an appropriate step size. By choosing  $D$  appropriately for  $\mathcal{X}$ , the update step (2) becomes efficient, and one can achieve an overall regret on the order of  $O(\sqrt{T})$  after  $T$  iterations. This regret can be achieved either by choosing a fixed step size  $\eta = \alpha/L\sqrt{T}$ , where  $L$  is an upper bound on the norms of the subgradients visited  $(\mathbf{f}_t)_{t \geq 0}$ , or by choosing adaptive step sizes  $\eta_t = \alpha/\sqrt{t}$ , for  $\alpha > 0$ . This is problematic, as 1) the upper bound  $L$  may be hard to obtain in many applications and may be too conservative in practice, and 2) adequately tuning the parameter  $\alpha$  can be time- and resource-consuming, and even practically infeasible for very large instances, since we won't know if the step size will cause a slow convergence or a divergence until late in the optimization process. This is not just a theoretical issue, as we highlight in our numerical experiments (Section 4) and in the appendices (Appendices F). Similar results and challenges hold for the popular *follow the regularized leader* (FTRL) algorithm (see Appendix F).

The above issues can be circumvented by employing *adaptive* variants of OMD or FTRL, which lead to parameter- and scale-free algorithms that estimate the parameters through the observed subgradients, e.g., AdaHedge for the simplex setting [DRVEGK14] or AdaFTRL for general compact convex decisions sets [OP15]. Yet these adaptive variants have not seen practical adoption in large-scale game-solving, where regret-matching variants are preferred (see the next paragraph). As we show in our experiments, adaptive variants of OMD and FTRL perform much worse than our proposed algorithms. While these adaptive algorithms are referred to as parameter-free, this is only true in the sense that they are able to learn the necessary parameters. Our algorithm is parameter-free in the stronger sense that there are no parameters that even require learning. Formalizing this difference may be one interesting avenue for explaining the performance discrepancy on saddle-point problems.

**Regret Matching** In this paper, we introduce alternative regret-minimization schemes for instantiating the above framework. Our work is motivated by recent advances on solving large-scale zero-sum sequential games. In the zero-sum sequential game setting,  $\mathcal{X}$  and  $\mathcal{Y}$  are simplexes, the objective function becomes  $F(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{A}\mathbf{y} \rangle$ , and thus (1) reduces to a *bilinear* SPP. Based on this bilinear SPP formulation, the best practical methods for solving large-scale sequential games use the repeated game framework, where each player minimizes regret via some variant of *counterfactual regret minimization* (CFR, [ZJBP07]). Variants of CFR were used in every recent poker AI challenge, where poker AIs beat human poker players [BBJT15, MSB<sup>+</sup>17, BS18, BS19]. The CFR framework itself is a decomposition of the overall regret of the bilinear SPP into local regrets at each decision point in a sequential game [FKS19a]. The key to the practical performance of CFR-based algorithms seems to be three ingredients (beyond the CFR decomposition itself): (1) a particular regret minimizer called *regret matching<sup>+</sup>* (RM<sup>+</sup>) [TBJB15] which is employed at each decision point, (2) aggressive iterate averaging schemes that put greater weight on recent iterates (e.g. *linear averaging*, which weights iterate at period  $t$  by  $2t/T(T+1)$ ), and (3) an alternation scheme where the updates of the repeated game framework are performed in an asymmetric fashion. The CFR framework itself is specific to sequential bilinear games on simplexes, but these last three ingredients could potentially be generalized to other problems of the form (1). That is the starting point of the present paper.

The most challenging aspect of generalizing the above ingredients is that RM<sup>+</sup> is specifically designed for minimizing regret over a simplex. However, many problems of the form (1) have convex sets  $\mathcal{X}, \mathcal{Y}$  that are not simplexes, e.g., box constraints or norm-balls for distributionally robust optimization [BTHKM15]. In principle, regret matching arises from a general theory called *Blackwell approachability* [Bla56, HMC00], and similar constructions can be envisioned for other convex sets. However, in practice the literature has only focused on developing concrete implementable instantiations of Blackwell approachability for simplexes. A notable deviation from this is the work of [ABH11], who showed a general reduction between regret minimization over general convex compact sets and Blackwell approachability. However, their general reduction still does not yield a practically implementable algorithm: among other things, their reduction relies on certain black-box projections that are not always efficient. We show how to implement these necessary projections for the setting where  $\mathcal{X}$  and  $\mathcal{Y}$  are simplexes,  $\ell_p$  balls, and intersections of the  $\ell_2$  ball with a hyperplane

(with a focus on the case where an  $\ell_2$  ball is intersected with a simplex, which arises naturally as a confidence region). This yields an algorithm which we will refer to as the *conic Blackwell algorithm* (CBA), which is similar in spirit to the regret matching algorithm, but crucially generalizes to other decision sets. Motivated by the practical performance of  $\text{RM}^+$ , we construct a variant of CBA which uses a thresholding operation similar to the one employed by  $\text{RM}^+$ . We call this algorithm  $\text{CBA}^+$ .

**Our contributions** We introduce  $\text{CBA}^+$ , a parameter-free algorithm which achieves  $O(\sqrt{T})$  regret in the worst-case and generalizes the strong performances of  $\text{RM}^+$  for bilinear, simplex saddle-points solving to other more general settings. A major selling point for  $\text{CBA}^+$  is that it does not require any step size choices. Instead, the algorithm implicitly adjusts to the structure of the domains and losses by being instantiations of Blackwell’s approachability algorithm. After developing the  $\text{CBA}^+$  algorithm, we then develop analogues of another crucial components for large-scale game solving. In particular, we prove a generalization of the folk theorem for the repeated game framework for solving (1), which allows us to incorporate polynomial averaging schemes such as linear averaging. We then show that  $\text{CBA}^+$  is compatible with linear averaging on the iterates. This mirrors the case of  $\text{RM}$  and  $\text{RM}^+$ , where only  $\text{RM}^+$  is compatible with linear averaging on the iterates. We also show that both CBA and  $\text{CBA}^+$  are compatible with polynomial averaging when simultaneously performed on the regrets and the iterates. Combining all these ingredients, we arrive at a new class of algorithms for solving convex-concave SPPs. As long as efficient projection operations can be performed (which we show for several practical domains, including the simplex,  $\ell_p$  balls and confidence regions in the simplex), one can apply the repeated game framework on (1), where one can use either CBA or  $\text{CBA}^+$  as a regret minimizer for  $\mathcal{X}$  and  $\mathcal{Y}$  along with polynomial averaging on the generated iterates to solve (1) at a rate of  $O\left(1/\sqrt{T}\right)$ .

We highlight the practical efficacy of our algorithmic framework on several domains. First, we solve two-player zero-sum matrix games and extensive-form games, where  $\text{RM}^+$  regret minimizer combined with linear averaging and alternation, and  $\text{CFR}^+$ , lead to very strong practical algorithms [TBJB15]. We find that  $\text{CBA}^+$  combined with linear averaging and alternation leads to a comparable performance in terms of the iteration complexity, and may even slightly outperform  $\text{RM}^+$  and  $\text{CFR}^+$ . On this simplex setting, we also find that  $\text{CBA}^+$  outperforms both AdaHedge and AdaFTRL. Second, we apply our approach to a setting where  $\text{RM}^+$  and  $\text{CFR}^+$  do not apply: distributionally robust empirical risk minimization (DR-ERM) problems. Across two classes of synthetic problems and four real data sets, we find that our algorithm based on  $\text{CBA}^+$  performs orders of magnitude better than online mirror descent and FTRL, as well as their optimistic variants, when using their theoretically-correct fixed step sizes. Even when considering adaptive step sizes, or fixed step sizes that are up to 10,000 larger than those predicted by theory, our  $\text{CBA}^+$  algorithm performs better, with only a few cases of comparable performance (at step sizes that lead to divergence for some of the other non parameter-free methods). The fast practical performance of our algorithm, combined with its simplicity and the total lack of step sizes or parameters tuning, suggests that it should be seriously considered as a practical approach for solving convex-concave SPPs in various settings.

Finally, we make a brief note on accelerated methods. Our algorithms have a rate of convergence towards a saddle point of  $O(1/\sqrt{T})$ , similar to OMD and FTRL. In theory, it is possible to obtain a faster  $O(1/T)$  rate of convergence when  $F$  is differentiable with Lipschitz gradients, for example via mirror prox [Nem04] or other primal-dual algorithms [CP16]. However, our experimental results show that  $\text{CBA}^+$  is faster than optimistic variants of FTRL and OMD [SALS15], the latter being almost identical to the mirror prox algorithm, and both achieving  $O(1/T)$  rate of convergence. A similar conclusion has been drawn in the context of sequential game solving, where the fastest  $O(1/\sqrt{T})$  CFR-based algorithms have better practical performance than the theoretically-superior  $O(1/T)$ -rate methods [KWKKS20, KFS18]. In a similar vein, using *error-bound conditions*, it is possible to achieve a linear rate, e.g., when solving bilinear saddle-point problems over polyhedral decision sets using the extragradient method [Tse95] or optimistic gradient descent-ascent [WLZL20]. However, these linear rates rely on unknown constants, and may not be indicative of practical performance.

## 2 Game setup and Blackwell Approachability

As stated in Section 1, we will solve (1) using a repeated game framework. The first player chooses strategies from  $\mathcal{X}$  in order to minimize the sequence of payoffs in the repeated game, while the second player chooses strategies from  $\mathcal{Y}$  in order to maximize payoffs. There are  $T$  iterations with indices  $t = 1, \dots, T$ . In this framework, each iteration  $t$  consists of the following steps:

1. Each player chooses strategies  $\mathbf{x}_t \in \mathcal{X}, \mathbf{y}_t \in \mathcal{Y}$
2. First player observes  $\mathbf{f}_t \in \partial_{\mathbf{x}} F(\mathbf{x}_t, \mathbf{y}_t)$  and uses  $\mathbf{f}_t$  when computing the next strategy
3. Second player observes  $\mathbf{g}_t \in \partial_{\mathbf{y}} F(\mathbf{x}_t, \mathbf{y}_t)$  and uses  $\mathbf{g}_t$  when computing the next strategy

The goal of each player is to minimize their regret  $R_{T,\mathbf{x}}, R_{T,\mathbf{y}}$  across the  $T$  iterations:

$$R_{T,\mathbf{x}} = \sum_{t=1}^T \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \langle \mathbf{f}_t, \mathbf{x} \rangle, \quad R_{T,\mathbf{y}} = \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{y} \rangle - \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{y}_t \rangle.$$

The reason this repeated game framework leads to a solution to the SPP problem (1) is the following folk theorem. Relying on  $F$  being convex-concave and subdifferentiable, it connects the regret incurred by each player to the duality gap in (1).

**Theorem 2.1** (Theorem 1, [Kro20]). *Let  $(\bar{\mathbf{x}}_T, \bar{\mathbf{y}}_T) = \frac{1}{T} \sum_{t=1}^T (\mathbf{x}_t, \mathbf{y}_t)$  for any  $(\mathbf{x}_t)_{t \geq 1}, (\mathbf{y}_t)_{t \geq 1}$ . Then*

$$\max_{\mathbf{y} \in \mathcal{Y}} F(\bar{\mathbf{x}}_T, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}, \bar{\mathbf{y}}_T) \leq (R_{T,\mathbf{x}} + R_{T,\mathbf{y}})/T.$$

Therefore, when each player runs a regret minimizer that guarantees regret on the order of  $O(\sqrt{T})$ ,  $(\bar{\mathbf{x}}_T, \bar{\mathbf{y}}_T)_{T \geq 0}$  converges to a solution to (1) at a rate of  $O(1/\sqrt{T})$ . Later we will show a generalization of Theorem 2.1 that will allow us to incorporate more aggressive averaging schemes that put additional weight on the later iterates. Given the repeated game framework, the next question becomes which algorithms to employ in order to minimize regret for each player. As mentioned in Section 1, for zero-sum games, variants of regret matching are used in practice.

**Blackwell Approachability** Regret matching arises from the *Blackwell approachability* framework [Bla56]. In Blackwell approachability, a decision maker repeatedly takes decisions  $\mathbf{x}_t$  from some convex decision set  $\mathcal{X}$  (this set plays the same role as  $\mathcal{X}$  or  $\mathcal{Y}$  in (1)). After taking decision  $\mathbf{x}_t$  the player observes a vector-valued payoff function  $\mathbf{u}_t(\mathbf{x}) \in \mathbb{R}^n$ . The goal for the decision maker is to force the average payoff  $\frac{1}{t} \sum_{\tau=1}^t \mathbf{u}_\tau(\mathbf{x}_\tau)$  to approach some convex target  $\mathcal{S}$ . Blackwell proved that a convex target set  $\mathcal{S}$  can be approached if and only if for every halfspace  $\mathcal{H} \supseteq \mathcal{S}$ , there exists  $\mathbf{x} \in \mathcal{X}$  such that for every possible payoff function  $\mathbf{u}(\cdot)$ ,  $\mathbf{u}(\mathbf{x})$  is guaranteed to lie in  $\mathcal{H}$ . The action  $\mathbf{x}$  is said to *force*  $\mathcal{H}$ . Blackwell's proof is via an algorithm: at iteration  $t$ , his algorithm projects the average payoff  $\bar{\mathbf{u}} = \frac{1}{t-1} \sum_{\tau=1}^{t-1} \mathbf{u}_\tau(\mathbf{x}_\tau)$  onto  $\mathcal{S}$ , and then the decision maker chooses an action  $\mathbf{x}_t$  that forces the tangent halfspace to  $\mathcal{S}$  generated by the normal  $\bar{\mathbf{u}} - \pi_{\mathcal{S}}(\bar{\mathbf{u}})$ , where  $\pi_{\mathcal{S}}(\bar{\mathbf{u}})$  is the orthogonal projection of  $\bar{\mathbf{u}}$  onto  $\mathcal{S}$ . We call this algorithm *Blackwell's algorithm*; it approaches  $\mathcal{S}$  at a rate of  $O(1/\sqrt{T})$ . It is important to note here that Blackwell's algorithm is rather a meta-algorithm than a concrete algorithm. Even within the context of Blackwell's approachability problem, one needs to devise a way to compute the forcing actions needed at each iteration, i.e., to compute  $\pi_{\mathcal{S}}(\bar{\mathbf{u}})$ .

**Details on Regret Matching** Regret matching arises by instantiating Blackwell approachability with the decision space  $\mathcal{X}$  equal to the simplex  $\Delta(n)$ , the target set  $\mathcal{S}$  equal to the nonpositive orthant  $\mathbb{R}_-^n$ , and the vector-valued payoff function  $\mathbf{u}_t(\mathbf{x}_t) = \mathbf{f}_t - \langle \mathbf{f}_t, \mathbf{x}_t \rangle \mathbf{e}$  equal to the regret associated to each of the  $n$  actions (which correspond to the corners of  $\Delta(n)$ ). Here  $\mathbf{e} \in \mathbb{R}^n$  has one on every component. [HMC00] showed that with this setup, playing each action with probability proportional to its positive regret up to time  $t$  satisfies the forcing condition needed in Blackwell's algorithm. Formally, regret matching (RM) keeps a running sum  $\mathbf{r}_t = \sum_{\tau=1}^t (\mathbf{f}_\tau - \langle \mathbf{f}_\tau, \mathbf{x}_\tau \rangle \mathbf{e})$ , and then action  $i$  is played with probability  $x_{t+1,i} = [\mathbf{r}_{t,i}]^+ / \sum_{i=1}^n [\mathbf{r}_{t,i}]^+$ , where  $[\cdot]^+$  denotes thresholding at zero. By Blackwell's approachability theorem, this algorithm converges to zero average regret at a rate of  $O(1/\sqrt{T})$ . In zero-sum game-solving, it was discovered that a variant of

regret matching leads to extremely strong practical performance (but the same theoretical rate of convergence). In regret matching<sup>+</sup> (RM<sup>+</sup>), the running sum is thresholded at zero at every iteration:  $\mathbf{r}_t = [\mathbf{r}_{t-1} + \mathbf{f}_t - \langle \mathbf{f}_t, \mathbf{x}_t \rangle \mathbf{e}]^+$ , and then actions are again played proportional to  $\mathbf{r}_t$ . In the next section, we describe a more general class of regret-minimization algorithms based on Blackwell’s algorithm for general sets  $\mathcal{X}$ , introduced in [ABH11]. Note that a similar construction of a general class of algorithms can be achieved through the *Lagrangian Hedging* framework of [Gor07]. It would be interesting to construct a CBA<sup>+</sup>-like algorithm and efficient projection approaches for this framework as well.

### 3 Conic Blackwell Algorithm

We present the Conic Blackwell Algorithm Plus (CBA<sup>+</sup>), a no-regret algorithm which uses a variation of Blackwell’s approachability procedure [Bla56] to perform regret minimization on general convex compact decision sets  $\mathcal{X}$ . We will assume that losses are coming from a bounded set; this occurs, for example, if there exists  $L_x, L_y$  (that we do not need to know), such that

$$\|\mathbf{f}\| \leq L_x, \|\mathbf{g}\| \leq L_y, \forall \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}, \forall \mathbf{f} \in \partial_x F(\mathbf{x}, \mathbf{y}), \forall \mathbf{g} \in \partial_y F(\mathbf{x}, \mathbf{y}). \quad (3)$$

CBA<sup>+</sup> is best understood as a combination of two steps. The first is the basic CBA algorithm, derived from Blackwell’s algorithm, which we describe next. To convert Blackwell’s algorithm to a regret minimizer on  $\mathcal{X}$ , we use the reduction from [ABH11], which considers the conic hull  $\mathcal{C} = \text{cone}(\{\kappa\} \times \mathcal{X})$  where  $\kappa = \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$ . The Blackwell approachability problem is then instantiated with  $\mathcal{X}$  as the decision set, target set equal to the polar  $\mathcal{C}^\circ = \{\mathbf{z} : \langle \mathbf{z}, \hat{\mathbf{z}} \rangle \leq 0, \forall \hat{\mathbf{z}} \in \mathcal{C}\}$  of  $\mathcal{C}$ , and payoff vectors  $(\langle \mathbf{f}_t, \mathbf{x}_t \rangle, -\mathbf{f}_t)$ . The conic Blackwell algorithm (CBA) is implemented by projecting the average payoff vector onto  $\mathcal{C}$ , calling this projection  $\alpha(\kappa, \mathbf{x})$  with  $\alpha \geq 0$  and  $\mathbf{x} \in \mathcal{X}$ , and playing the action  $\mathbf{x}$ .

The second step in CBA<sup>+</sup> is to modify CBA to make it analogous to RM<sup>+</sup> rather than to RM. To do this, the algorithm does not keep track of the average payoff vector. Instead, we keep a running aggregation of the payoffs, where we always add the newest payoff to the aggregate, and then project the aggregate onto  $\mathcal{C}$ . More concretely, pseudocode for CBA<sup>+</sup> is given in Algorithm 1. This pseudocode relies on two functions:  $\text{CHOOSEDECISION}_{\text{CBA}^+} : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$ , which maps the aggregate payoff vector  $\mathbf{u}_t$  to a decision in  $\mathcal{X}$ , and  $\text{UPDATEPAYOFF}_{\text{CBA}^+}$  which controls how we aggregate payoffs. Given an aggregate payoff vector  $\mathbf{u} = (\tilde{u}, \hat{\mathbf{u}}) \in \mathbb{R} \times \mathbb{R}^n$ , we have

$$\text{CHOOSEDECISION}_{\text{CBA}^+}(\mathbf{u}) = (\kappa/\tilde{u})\hat{\mathbf{u}}.$$

If  $\tilde{u} = 0$ , we just let  $\text{CHOOSEDECISION}_{\text{CBA}^+}(\mathbf{u}) = \mathbf{x}_0$  for some chosen  $\mathbf{x}_0 \in \mathcal{X}$ .

The function  $\text{UPDATEPAYOFF}_{\text{CBA}^+}$  is implemented by adding the most recent payoff to the aggregate payoffs, and then projecting onto  $\mathcal{C}$ . More formally, it is defined as

$$\text{UPDATEPAYOFF}_{\text{CBA}^+}(\mathbf{u}, \mathbf{x}, \mathbf{f}, \omega, S) = \pi_{\mathcal{C}} \left( \frac{S}{S+\omega} \mathbf{u} + \frac{\omega}{S+\omega} (\langle \mathbf{f}, \mathbf{x} \rangle / \kappa, -\mathbf{f}) \right),$$

where  $\omega$  is the weight assigned to the most recent payoff and  $S$  the weight assigned to the previous aggregate payoff  $\mathbf{u}$ . Because of the projection step in  $\text{UPDATEPAYOFF}_{\text{CBA}^+}$ , we always have  $\mathbf{u} \in \mathcal{C}$ , which in turn guarantees that  $\text{CHOOSEDECISION}_{\text{CBA}^+}(\mathbf{u}) \in \mathcal{X}$ , since  $\mathcal{C} = \text{cone}(\{\kappa\} \times \mathcal{X})$ .

Let us give some intuition on the effect of projection onto  $\mathcal{C}$ . In a geometric sense, it is easier to visualize things in  $\mathbb{R}^2$  with  $\mathcal{C} = \mathbb{R}_+^2$  and  $\mathcal{C}^\circ = \mathbb{R}_-^2$ . The projection on  $\mathcal{C}$  moves the vector along the edges of  $\mathcal{C}^\circ$ , maintaining the distance to  $\mathcal{C}^\circ$  and moving toward the vector  $\mathbf{0}$ . This is illustrated in Figure 5 in Appendix B.1. From a game-theoretic standpoint, the projection on  $\mathcal{C} = \mathbb{R}_+^2$  eliminates the components of the payoffs that are negative. It enables CBA<sup>+</sup> to be less pessimistic than CBA, which may accumulate negative payoffs on actions for a long time and never resets the components of the aggregated payoff to 0, leading to some actions being chosen less frequently.

We will see in the next section that RM<sup>+</sup> is related to CBA<sup>+</sup> but replaces the exact projection step  $\pi_{\mathcal{C}}(\mathbf{u})$  in  $\text{UPDATEPAYOFF}_{\text{CBA}^+}$  by a suboptimal solution to the projection problem. Let us also note the difference between CBA<sup>+</sup> and the algorithm introduced in [ABH11], which we have called

---

**Algorithm 1** Conic Blackwell Algorithm Plus (CBA<sup>+</sup>)
 

---

- 1: **Input** A convex, compact set  $\mathcal{X} \subset \mathbb{R}^n$ ,  $\kappa = \max\{\|\mathbf{x}\|_2 \mid \mathbf{x} \in \mathcal{X}\}$ .
  - 2: **Algorithm parameters** Weights  $(\omega_\tau)_{\tau \geq 1} \in \mathbb{R}^{\mathbb{N}}$ .
  - 3: **Initialization**  $t = 1$ ,  $\mathbf{x}_1 \in \mathcal{X}$ .
  - 4: Observe  $\mathbf{f}_1$  then set  $\mathbf{u}_1 = (\langle \mathbf{f}_1, \mathbf{x}_1 \rangle / \kappa, -\mathbf{f}_1) \in \mathbb{R} \times \mathbb{R}^n$ .
  - 5: **for**  $t \geq 1$  **do**
  - 6:   Choose  $\mathbf{x}_{t+1} = \text{CHOOSEDECISION}_{\text{CBA}^+}(\mathbf{u}_t)$ .
  - 7:   Observe the loss  $\mathbf{f}_{t+1} \in \mathbb{R}^n$ .
  - 8:   Update  $\mathbf{u}_{t+1} = \text{UPDATEPAYOFF}_{\text{CBA}^+}(\mathbf{u}_t, \mathbf{x}_{t+1}, \mathbf{f}_{t+1}, \omega_{t+1}, \sum_{\tau=1}^t \omega_\tau)$ .
  - 9:   Increment  $t \leftarrow t + 1$ .
- 

CBA. CBA uses different UPDATEPAYOFF and CHOOSEDECISION functions. In CBA the payoff update is defined as

$$\text{UPDATEPAYOFF}_{\text{CBA}}(\mathbf{u}, \mathbf{x}, \mathbf{f}, \omega, S) = \frac{S}{S + \omega} \mathbf{u} + \frac{\omega}{S + \omega} (\langle \mathbf{f}, \mathbf{x} \rangle / \kappa, -\mathbf{f}).$$

Note in particular the lack of projection as compared to CBA<sup>+</sup>, analogous to the difference between RM and RM<sup>+</sup>. The CHOOSEDECISION<sub>CBA</sub> function then requires a projection onto  $\mathcal{C}$ :

$$\text{CHOOSEDECISION}_{\text{CBA}}(\mathbf{u}) = \text{CHOOSEDECISION}_{\text{CBA}^+}(\pi_{\mathcal{C}}(\mathbf{u})).$$

Based upon the analysis in [Bla56], [ABH11] show that CBA with uniform weights (both on payoffs and decisions) guarantees  $O(1/\sqrt{T})$  average regret. The difference between CBA<sup>+</sup> and CBA is similar to the difference between the RM and RM<sup>+</sup> algorithms. In practice, RM<sup>+</sup> performs significantly better than RM for solving matrix games, when combined with *linear averaging* on the decisions (as opposed to the uniform averaging used in Theorem 2.1). In the next theorem, we show that CBA<sup>+</sup> is compatible with linear averaging on decisions only. We present a detailed proof in Appendix B.

**Theorem 3.1.** Consider  $(\mathbf{x}_t)_{t \geq 0}$  generated by CBA<sup>+</sup> with uniform weights:  $\omega_\tau = 1, \forall \tau \geq 1$ . Let  $L = \max\{\|\mathbf{f}_t\|_2 \mid t \geq 1\}$  and  $\kappa = \max\{\|\mathbf{x}\|_2 \mid \mathbf{x} \in \mathcal{X}\}$ . Then

$$\frac{\sum_{t=1}^T t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T t \langle \mathbf{f}_t, \mathbf{x} \rangle}{T(T+1)} = O\left(\kappa L / \sqrt{T}\right).$$

Note that in Theorem 3.1, we have *uniform* weights on the sequence of payoffs  $(\mathbf{u}_t)_{t \geq 0}$ , but *linearly increasing* weights on the sequence of decisions. The proof relies on properties specific to CBA<sup>+</sup>, and it does not extend to CBA. Numerically it also helps CBA<sup>+</sup> but not CBA. In Appendix B, we show that both CBA and CBA<sup>+</sup> achieve  $O\left(\kappa L / \sqrt{T}\right)$  convergence rates when using a weighted average on *both* the decisions and the payoffs (Theorems B.2-B.3). In practice, using linear averaging only on the decisions, as in Theorem 3.1, performs vastly better than linear averaging on both decisions and payoffs. We present empirical evidence of this in Appendix B.

We can compare the  $O\left(\kappa L / \sqrt{T}\right)$  average regret for CBA<sup>+</sup> with the  $O\left(\Omega L / \sqrt{T}\right)$  average regret for OMD [NY83, BTN01] and FTRL [AHR09, McM11], where  $\Omega = \max\{\|\mathbf{x} - \mathbf{x}'\|_2 \mid \mathbf{x}, \mathbf{x}' \in \mathcal{X}\}$ . We can always recenter  $\mathcal{X}$  to contain  $\mathbf{0}$ , in which case the bounds for OMD/FTRL and CBA<sup>+</sup> are equivalent since  $\kappa \leq \Omega \leq 2\kappa$ . Note that the bound on the average regret for Optimistic OMD (O-OMD, [CYL<sup>+</sup>12]) and Optimistic FTRL (O-FTRL, [RS13]) is  $O\left(\Omega^2 L / T\right)$  in the game setup, a priori better than the bound for CBA<sup>+</sup> as regards the number of iterations  $T$ . Nonetheless, we will see in Section 4 that the empirical performance of CBA<sup>+</sup> is better than that of  $O(1/T)$  methods. A similar situation occurs for RM<sup>+</sup> compared to O-OMD and O-FTRL for solving poker games [FKS19b, KWKKS20].

The following theorem gives the convergence rate of CBA<sup>+</sup> for solving saddle-points (1), based on our convergence rate on the regret of each player (Theorem 3.1). The proof is in Appendix C.

**Theorem 3.2.** Let  $(\bar{\mathbf{x}}_T, \bar{\mathbf{y}}_T) = 2 \sum_{t=1}^T t (\mathbf{x}_t, \mathbf{y}_t) / (T(T+1))$ , where  $(\mathbf{x}_t)_{t \geq 0}, (\mathbf{y}_t)_{t \geq 0}$  are generated by the repeated game framework with CBA<sup>+</sup> with uniform weights:  $\omega_\tau = 1, \forall \tau \geq 1$ . Let  $L = \max\{L_x, L_y\}$  defined in (3) and  $\kappa = \max\{\max\{\|\mathbf{x}\|_2, \|\mathbf{y}\|_2\} \mid \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}\}$ . Then

$$\max_{\mathbf{y} \in \mathcal{Y}} F(\bar{\mathbf{x}}_T, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}, \bar{\mathbf{y}}_T) = O\left(\kappa L / \sqrt{T}\right).$$

### 3.1 Efficient implementations of CBA<sup>+</sup>

To obtain an implementation of CBA<sup>+</sup> and CBA, we need to efficiently resolve the functions  $\text{CHOOSEDECISION}_{\text{CBA}^+}$  and  $\text{UPDATEPAYOFF}_{\text{CBA}^+}$ . In particular, we need to compute  $\pi_{\mathcal{C}}(\mathbf{u})$ , the orthogonal projection of  $\mathbf{u}$  onto the cone  $\mathcal{C}$ , where  $\mathcal{C} = \text{cone}(\{\kappa\} \times \mathcal{X})$ :

$$\pi_{\mathcal{C}}(\mathbf{u}) \in \arg \min_{\mathbf{y} \in \mathcal{C}} \|\mathbf{y} - \mathbf{u}\|_2^2. \quad (4)$$

Even for CBA this problem must be resolved, since [ABH11] did not study whether (4) can be efficiently solved. It turns out that (4) can be computed in closed-form or quasi closed-form for many decision sets  $\mathcal{X}$  of interest. Interestingly, parts of the proofs rely on *Moreau's Decomposition Theorem* [CR13], which states that  $\pi_{\mathcal{C}}(\mathbf{u})$  can be recovered from  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  and vice-versa, because  $\pi_{\mathcal{C}}(\mathbf{u}) + \pi_{\mathcal{C}^\circ}(\mathbf{u}) = \mathbf{u}$ . We present the detailed complexity results and the proofs in Appendix D.

**Simplex**  $\mathcal{X} = \Delta(n)$  is the classical setting used for matrix games. Also, for extensive-form games, CFR decomposes the decision sets (treplexes) into a set of regret minimization problems over the simplex [FKS19a]. Here,  $n$  is the number of actions of a player and  $\mathbf{x} \in \Delta(n)$  represents a randomized strategy. In this case,  $\pi_{\mathcal{C}}(\mathbf{u})$  can be computed in  $O(n \log(n))$ . Note that RM and RM<sup>+</sup> are obtained by choosing a suboptimal solution to (4), avoiding the  $O(n \log(n))$  sorting operation, whereas CBA and CBA<sup>+</sup> choose optimally (see Appendix D). Thus, RM and RM<sup>+</sup> can be seen as approximate versions of CBA and CBA<sup>+</sup>, where (4) is solved approximatively at every iteration. In our numerical experiments, we will see that CBA<sup>+</sup> slightly outperforms RM<sup>+</sup> and CFR<sup>+</sup> in terms of iteration count.

**$\ell_p$  balls** This is when  $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\|_p \leq 1\}$  with  $p \geq 1$  or  $p = \infty$ . This is of interest for instance in distributionally robust optimization [BTHKM15, ND16],  $\ell_\infty$  regression [ST18] and saddle-point reformulation of Markov Decision Process [JS20]. For  $p = 2$ , we can compute  $\pi_{\mathcal{C}}(\mathbf{u})$  in closed-form, i.e., in  $O(n)$  arithmetic operations. For  $p \in \{1, \infty\}$ , we can compute  $\pi_{\mathcal{C}}(\mathbf{u})$  in  $O(n \log(n))$  arithmetic operations using a sorting algorithm.

**Ellipsoidal confidence region in the simplex** Here,  $\mathcal{X}$  is an *ellipsoidal subregion of the simplex*, defined as  $\mathcal{X} = \{\mathbf{x} \in \Delta(n) \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x\}$ . This type of decision set is widely used because it is associated with confidence regions when estimating a probability distribution from observed data [Iye05, BdHP19]. It can also be used in the Bellman update for robust Markov Decision Process [Iye05, WKR13, GGC18]. We also assume that the confidence region is “entirely contained in the simplex”:  $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}^\top \mathbf{e} = 1\} \cap \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x\} \subseteq \Delta(n)$ , to avoid degenerate components. In this case, using a change of basis we show that it is possible to compute  $\pi_{\mathcal{C}}(\mathbf{u})$  in closed-form, i.e., in  $O(n)$  arithmetic operations.

**Other potential sets of interests** Other important decision sets include sets based on Kullback-Leibler divergence  $\{\mathbf{x} \in \Delta(n) \mid KL(\mathbf{x}, \mathbf{x}_0) \leq \epsilon_x\}$ , or, more generally,  $\phi$ -divergence [BTDHDW<sup>+</sup>13]. For these sets, we did not find a closed-form solution to the projection problem (4). Still, as long as the domain  $\mathcal{X}$  is a convex set, computing  $\pi_{\mathcal{C}}(\mathbf{u})$  remains a convex problem, and it can be solved efficiently with solvers, although this results in a slower algorithm than with closed-form computations of  $\pi_{\mathcal{C}}(\mathbf{u})$ .

## 4 Numerical experiments

In this section we investigate the practical performances of our algorithms on several instances of saddle-point problems. We start by comparing CBA<sup>+</sup> with RM<sup>+</sup> in the matrix and extensive-form games setting. We then turn to comparing our algorithms on instances from the distributionally robust optimization literature. The code for all experiments is available in the supplemental material.

### 4.1 Matrix games on the simplex and Extensive-Form Games

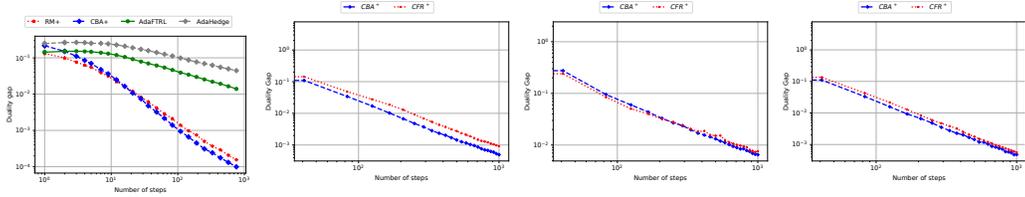
Since the motivation for CBA<sup>+</sup> is to obtain the strong empirical performances of RM<sup>+</sup> and CFR<sup>+</sup> on other decision sets than the simplex, we start by checking that CBA<sup>+</sup> indeed provides comparable

performance on simplex settings. We compare these methods on matrix games

$$\min_{\mathbf{x} \in \Delta(n)} \max_{\mathbf{y} \in \Delta(m)} \langle \mathbf{x}, \mathbf{A}\mathbf{y} \rangle,$$

where  $\mathbf{A}$  is the matrix of payoff, and on extensive-form games (EFGs). EFGs can also be written as SPPs with bilinear objective and  $\mathcal{X}, \mathcal{Y}$  polytopes encoding the players' space of sequential strategies [vS96]. EFGs can be solved via simplex-based regret minimization by using the counterfactual regret minimization (CFR) framework to decompose regrets into local regrets at each simplex. Explaining CFR is beyond the scope of this work; we point the reader to [ZJBP07] or newer explanations [FKS19c, FKS19a]. For matrix games, we generate 70 synthetic 10-dimensional matrix games with  $A_{ij} \sim U[0, 1]$  and compare the most efficient algorithms for matrix games with linear averaging:  $\text{CBA}^+$  and  $\text{RM}^+$ . We also compare with two other scale-free no-regret algorithms, AdaHedge [DRVEGK14] and AdaFTRL [OP15]. Figure 1a presents the duality gap of the current solutions vs. the number of steps. Here, both  $\text{CBA}^+$  and  $\text{RM}^+$  use *alternation*, which is a trick that is well-known to improve the performances of  $\text{RM}^+$  [TBJB15], where the repeated game framework is changed such that players take turns updating their strategies, rather than performing these updates simultaneously, see Appendix E.1 for details.<sup>2</sup>

For EFGs, we compare  $\text{CBA}^+$  and  $\text{CFR}^+$  on many poker AI benchmark instances, including Leduc, Kuhn, search games and sheriff (see [FKS21] for game descriptions). We present our results in Figures 1b-1d. Additional details and experiments for EFGs are presented in Appendix E.4. Overall, we see in Figure 1 that  $\text{CBA}^+$  may slightly outperform  $\text{RM}^+$  and  $\text{CFR}^+$ , two of the strongest algorithms for matrix games and EFGs, which were shown to achieve the best empirical performances compared to a wide range of algorithms, including Hedge and other first-order methods [Kro20, KFS18, FKS19b]. For matrix games, AdaHedge and AdaFTRL are both outperformed by  $\text{RM}^+$  and  $\text{CBA}^+$ ; we present more experiments to compare  $\text{RM}^+$  and  $\text{CBA}^+$  in Appendix E.2, and more experiments with matrix games in Appendix E.3. Recall that our goal is to generalize these strong performance to other settings: we present our numerical experiments for solving distributionally robust optimization problems in the next section.



(a) Matrix game. (b) EFG: Leduc, 2 players, 9 ranks. (c) EFG: Battleship, 3 turns, 1 ship. (d) EFG: Leduc, 2 players, 5 ranks.

Figure 1: Comparison of  $\text{CBA}^+$  with  $\text{RM}^+$  and  $\text{CFR}^+$  on matrix games and EFGs.

## 4.2 Distributionally Robust Optimization

**Problem setup** Broadly speaking, DRO attempts to exploit partial knowledge of the statistical properties of the model parameters to obtain risk-averse optimal solutions [RM19]. We focus on the following instance of distributionally robust classification with logistic losses [ND16, BTHKM15]. There are  $m$  observed feature-label pairs  $(\mathbf{a}_i, b_i) \in \mathbb{R}^n \times \{-1, 1\}$ , and we want to solve

$$\min_{\mathbf{x} \in \mathbb{R}^n, \|\mathbf{x} - \mathbf{x}_0\|_2 \leq R} \max_{\mathbf{y} \in \Delta(m), \|\mathbf{y} - \mathbf{y}_0\|_2 \leq \lambda} \sum_{i=1}^m y_i \ell_i(\mathbf{x}), \quad (5)$$

where  $\ell_i(\mathbf{x}) = \log(1 + \exp(-b_i \mathbf{a}_i^\top \mathbf{x}))$ . The formulation (5) takes a worst-case approach to put more weight on misclassified observations and provides some statistical guarantees, e.g., it can be seen as a convex regularization of standard empirical risk minimization instances [DGN21].

We compare  $\text{CBA}^+$  (with linear averaging and alternation) with Online Mirror Descent (OMD), Optimistic OMD (O-OMD), Follow-The-Regularized-Leader (FTRL) and Optimistic FTRL (O-FTRL). We provide a detailed presentation of our implementations of these algorithms in Appendix

<sup>2</sup>We note that  $\text{RM}^+$  is guaranteed to retain its convergence rate under alternation. In contrast, we leave resolving this property for  $\text{CBA}^+$  to future work.

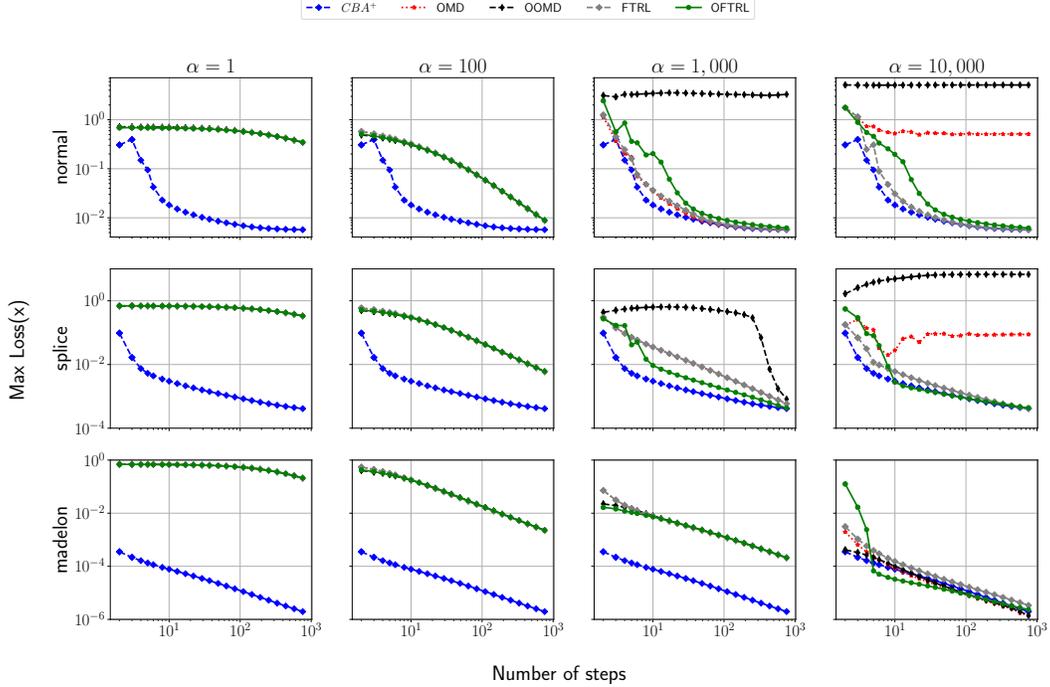


Figure 2: Comparisons of the performances of  $\text{CBA}^+$  with OMD, FTRL, O-OMD and O-FTRL with fixed step sizes, on synthetic (with *normal* distribution) and real data sets (*splice* and *madelon*).

F. We compare the performances of these algorithms with  $\text{CBA}^+$  on two synthetic data sets and four real data sets. We use linear averaging on decisions for all algorithms, and parameters  $\mathbf{x}_0 = \mathbf{0}$ ,  $R = 10$ ,  $\mathbf{y}_0 = (1, \dots, 1)/m$ ,  $\lambda = 1/2m$  in eq. (5).

**Synthetic and real instances** For the synthetic classification instances, we generate an optimal  $\mathbf{x}^* \in \mathbb{R}^n$ , sample  $\mathbf{a}_i \sim N(0, \mathbf{I})$  for  $i \in \{1, \dots, m\}$ , set labels  $b_i = \text{sign}(\mathbf{a}_i^\top \mathbf{x}^*)$ , and then we flip 10% of them. For the real classification instances, we use the following data sets from the *libsvm* website<sup>3</sup>: *adult*, *australian*, *splice*, *madelon*. Details about the empirical setting, the data sets and additional numerical experiments are presented in Appendix G.

**Choice of step sizes** One of the main motivation for  $\text{CBA}^+$  is to obtain a *parameter-free* algorithm. Choosing a *fixed* step size  $\eta$  for the other algorithms requires knowing a bound  $L$  on the norm of the instantaneous payoffs (see Appendix F.2 for our derivations of this upper bound). This is a major limitation in practice: these bounds may be very conservative, leading to small step sizes. We highlight this by showing the performance of all four algorithms, for various fixed step sizes  $\eta = \alpha \times \eta_{\text{th}}$ , where  $\alpha \in \{1, 100, 1,000, 10,000\}$  is a multiplier and  $\eta_{\text{th}}$  is the theoretical step size which guarantees the convergence of the algorithms for each instance. We detail the computation of  $\eta_{\text{th}}$  in Appendix F.2. We present the results of our numerical experiments on synthetic and real data sets in Figure 2. Additional simulations with adaptive step sizes  $\eta_t = 1/\sqrt{\sum_{\tau=1}^{t-1} \|\mathbf{f}_\tau\|_2^2}$  [Ora19] are presented in Figure 3 and in Appendix G.

**Results and discussion** In Figure 2, we present the worst-case loss of the current solution  $\bar{\mathbf{x}}_T$  in terms of the number of steps  $T$ . We see that when the step sizes is chosen as the theoretical step sizes guaranteeing the convergence of the non-parameter free algorithms ( $\alpha = 1$ ),  $\text{CBA}^+$  vastly outperforms all of the algorithms. When we take more aggressive step sizes, the non-parameter-free algorithms become more competitive. For instance, when  $\alpha = 1,000$ , OMD, FTRL and O-FTRL are competitive with  $\text{CBA}^+$  for the experiments on synthetic data sets. However, for this same instance and  $\alpha = 1,000$ , O-OMD diverges, because the step sizes are far greater than the theoretical step sizes guaranteeing convergence. At  $\alpha = 10,000$ , both OMD and O-OMD diverge. The same type of performances also hold for the *splice* data set. Finally, for the *madelon* data set, the non

<sup>3</sup><https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

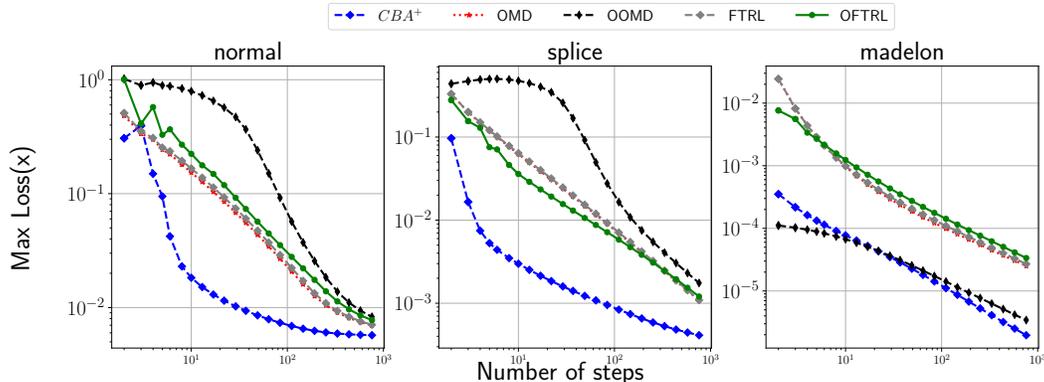


Figure 3: Comparisons of the performances of  $CBA^+$ , OMD, FTRL, O-OMD and O-FTRL with adaptive step sizes, on synthetic (with *normal* distribution) and real data sets (*splice* and *madelon*).

parameter-free algorithms start to be competitive with  $CBA^+$  only when  $\alpha = 10,000$ . Again, we note that this range of step sizes  $\eta$  is completely outside the values  $\eta_{th}$  that guarantee convergence of the algorithms, and fine-tuning the algorithms is time- and resource-consuming. In contrast,  $CBA^+$  can be used without wasting time on exploring and finding the best convergence rates, and with confidence in the convergence of the algorithm. Similar observations hold for adaptive step sizes (see Figure 3 and Appendix G). The overall poor performances of the optimistic methods (compared to their  $O(1/T)$  average regret guarantees) may reflect their sensibility to the choice of the step sizes. Additional experiments in Appendix G with other real and synthetic EFG and DRO instances show the robustness of the strong performances of  $CBA^+$  across additional problem instances.

**Running times compared to  $CBA^+$**  We would like to emphasize that all of our figures show the *number of steps* on the  $x$ -axis, and not the actual running times of the algorithms. Overall,  $CBA^+$  converges to an optimal solution to the DRO instance (5) vastly faster than the other algorithms. In particular, empirically,  $CBA^+$  is 2x-2.5x faster than OMD, FTRL and O-FTRL, and 3x-4x faster than O-OMD. This is because OMD, FTRL, O-OMD, and O-FTRL require binary searches at each step, see Appendix F. The functions used in the binary searches themselves require solving an optimization program (an orthogonal projection onto the simplex, see (33)) at each evaluation. Even though computing the orthogonal projection of a vector onto the simplex can be done in  $O(n \log(n))$ , this results in slower overall running time, compared to  $CBA^+$  with (quasi) closed-form updates at each step. The situation is even worse for O-OMD, which requires two proximal updates at each iteration. We acknowledge that the same holds for  $CBA^+$  compared to  $RM^+$ . In particular,  $CBA^+$  is slightly slower than  $RM^+$ , because of the computation of  $\pi_C(u)$  in  $O(n \log(n))$  operations at every iteration.

## 5 Conclusion

We have introduced  $CBA^+$ , a new algorithm for convex-concave saddle-point solving, that is 1) simple to implement for many practical decision sets, 2) completely parameter-free and does not attempt to learn any step sizes, and 3) competitive with, or even better than, state-of-the-art approaches for the best choices of parameters, both for matrix games, extensive-form games, and distributionally robust instances. Our paper is based on Blackwell approachability, which has been used to achieve important breakthroughs in poker AI in recent years, and we hope to generalize the use and implementation of this framework to other important problem instances. Interesting future directions of research include developing a theoretical understanding of the improvements related to alternation in our setting, designing efficient implementations for other widespread decision sets (e.g., based on Kullback-Leibler divergence or  $\phi$ -divergence), and novel accelerated versions based on strong convex-concavity or optimism.

**Societal impact** There is a priori no direct negative societal consequence to this work, since our methods simply return the same solutions as previous algorithms but in a simpler and more efficient way.

## References

- [ABH11] Jacob Abernethy, Peter L Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 27–46. JMLR Workshop and Conference Proceedings, 2011.
- [AHR09] Jacob D Abernethy, Elad Hazan, and Alexander Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. 2009.
- [BBJT15] Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.
- [BdHP19] Dimitris Bertsimas, Dick den Hertog, and Jean Pauphilet. Probabilistic guarantees in robust optimization. 2019.
- [Bla56] David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- [BMS19] Neil Burch, Matej Moravcik, and Martin Schmid. Revisiting cfr+ and alternating updates. *Journal of Artificial Intelligence Research*, 64:429–443, 2019.
- [BS18] Noam Brown and Tuomas Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.
- [BS19] Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365(6456):885–890, 2019.
- [BT03] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [BTDHDW<sup>+</sup>13] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.
- [BTHKM15] Aharon Ben-Tal, Elad Hazan, Tomer Koren, and Shie Mannor. Oracle-based robust optimization via online learning. *Operations Research*, 63(3):628–638, 2015.
- [BTN01] Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, volume 2. Siam, 2001.
- [CP11] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- [CP16] Antonin Chambolle and Thomas Pock. On the ergodic convergence rates of a first-order primal–dual algorithm. *Mathematical Programming*, 159(1-2):253–287, 2016.
- [CR13] Patrick L Combettes and Noli N Reyes. Moreau’s decomposition in banach spaces. *Mathematical Programming*, 139(1):103–114, 2013.
- [CYL<sup>+</sup>12] Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Conference on Learning Theory*, pages 6–1. JMLR Workshop and Conference Proceedings, 2012.
- [DGN21] John C Duchi, Peter W Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.
- [DRVEGK14] Steven De Rooij, Tim Van Erven, Peter D Grünwald, and Wouter M Koolen. Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316, 2014.
- [DSSSC08] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the L-1 ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279, 2008.
- [EPGMFBV03] Juan José Egozcue, Vera Pawlowsky-Glahn, Glòria Mateu-Figueras, and Carles Barcelo-Vidal. Isometric logratio transformations for compositional data analysis. *Mathematical Geology*, 35(3):279–300, 2003.

- [FKS19a] Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Online convex optimization for sequential decision processes and extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1917–1925, 2019.
- [FKS19b] Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Optimistic regret minimization for extensive-form games via dilated distance-generating functions. In *Advances in Neural Information Processing Systems*, pages 5222–5232, 2019.
- [FKS19c] Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Regret circuits: Composability of regret minimizers. In *International Conference on Machine Learning*, pages 1863–1872, 2019.
- [FKS21] Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Faster game solving via predictive blackwell approachability: Connecting regret matching and mirror descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI, 2021.
- [GCK20a] Julien Grand-Clément and Christian Kroer. First-order methods for Wasserstein distributionally robust MDP. *arXiv preprint arXiv:2009.06790*, 2020.
- [GCK20b] Julien Grand-Clément and Christian Kroer. Scalable first-order methods for robust mdps. *arXiv preprint arXiv:2005.05434*, 2020.
- [GGC18] Vineet Goyal and Julien Grand-Clément. Robust Markov decision process: Beyond rectangularity. *arXiv preprint arXiv:1811.00215*, 2018.
- [GKG19] Yuan Gao, Christian Kroer, and Donald Goldfarb. Increasing iterate averaging for solving saddle-point problems. *arXiv preprint arXiv:1903.10646*, 2019.
- [Gor07] Geoffrey J Gordon. No-regret algorithms for online convex programs. In *Advances in Neural Information Processing Systems*, pages 489–496. Citeseer, 2007.
- [HMC00] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [Iye05] Garud Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.
- [JS20] Yujia Jin and Aaron Sidford. Efficiently solving mdps with stochastic mirror descent. In *International Conference on Machine Learning*, pages 4890–4900. PMLR, 2020.
- [KFS18] Christian Kroer, Gabriele Farina, and Tuomas Sandholm. Solving large sequential games with the excessive gap technique. In *Advances in Neural Information Processing Systems*, pages 864–874, 2018.
- [Kro20] Christian Kroer. Ieor8100: Economics, ai, and optimization lecture note 5: Computing Nash equilibrium via regret minimization. 2020.
- [KWKKS20] Christian Kroer, Kevin Waugh, Fatma Kılınc-Karzan, and Tuomas Sandholm. Faster algorithms for extensive-form game solving via improved smoothing functions. *Mathematical Programming*, pages 1–33, 2020.
- [McM11] Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and  $\ell_1$  regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 525–533. JMLR Workshop and Conference Proceedings, 2011.
- [MSB<sup>+</sup>17] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017.
- [ND16] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with  $f$ -divergences. In *NIPS*, volume 29, pages 2208–2216, 2016.
- [Nem04] Arkadi Nemirovski. Prox-method with rate of convergence  $O(1/t)$  for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- [NY83] Arkadi Nemirovski and David Yudin. *Problem complexity and method efficiency in optimization*. 1983.

- [OP15] Francesco Orabona and Dávid Pál. Scale-free algorithms for online linear optimization. In *International Conference on Algorithmic Learning Theory*, pages 287–301. Springer, 2015.
- [Ora19] Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- [RM19] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.
- [RS13] Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pages 993–1019. PMLR, 2013.
- [SALS15] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. *arXiv preprint arXiv:1507.00407*, 2015.
- [Shi16] Nahum Shimkin. An online convex optimization approach to blackwell’s approachability. *The Journal of Machine Learning Research*, 17(1):4434–4456, 2016.
- [ST18] Aaron Sidford and Kevin Tian. Coordinate methods for accelerating l-infinity regression and faster approximate maximum flow. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 922–933. IEEE, 2018.
- [TBJB15] Oskari Tammelin, Neil Burch, Michael Johanson, and Michael Bowling. Solving heads-up limit Texas hold’em. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [Tse95] Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252, 1995.
- [vS96] Bernhard von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2):220–246, 1996.
- [WKR13] W. Wiesemann, D. Kuhn, and B. Rustem. Robust Markov decision processes. *Operations Research*, 38(1):153–183, 2013.
- [WLZL20] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations*, 2020.
- [ZJBP07] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. In *Advances in neural information processing systems*, pages 1729–1736, 2007.

## A Comparison to [Shi16] and [FKS21]

The focus of our paper is on developing new algorithms for convex-concave saddle-point solving via Blackwell approachability algorithms for sets beyond the simplex domain. This is what motivated the development of  $\text{CBA}^+$ , which attempts to generalize the ideas from  $\text{RM}^+$  and  $\text{CFR}^+$  beyond simplex settings. A complementary question is to consider the second construction of [ABH11], which is a way to convert a no-regret algorithm into an approachability algorithm. At a very level, that construction ends up using the no-regret algorithm to select which hyperplane to force, and performing regret minimization on the choice of hyperplane. An observation made by both [Shi16] and [FKS21] is that if one uses FTRL with the Euclidean regularizer, then it corresponds to Blackwell’s original approachability algorithm, and thus to regret matching on the simplex. This begs the question of what happens if one uses a different algorithm from FTRL. The most natural substitute would be OMD with a Bregman divergence derived from the Euclidean distance, which is also known as online gradient descent (OGD).

[Shi16] considers OGD on a variation of the original construction of [ABH11] with the simplex as the decision set. This setup (and the original setup from [ABH11]) uses (a subset of) the unit ball as the feasible set for the regret minimizer, and thus OGD ends up projecting the cumulated payoff vector onto the unit ball after every iteration. Therefore, the OGD setup by [Shi16] yields an algorithm that is reminiscent of  $\text{RM}^+$ , but where we repeatedly renormalize the cumulated payoff vector. One consequence of this is that the stepsize used to add new payoff vectors becomes important. [FKS21] show that the [ABH11] construction can be extended to allow the regret minimizer to use any decision set  $\mathcal{D}$  such that  $\mathcal{K} \subseteq \mathcal{D} \subseteq \mathcal{S}^\circ$ , where  $\mathcal{S}$  is the target set and  $\mathcal{S}^\circ$  its polar cone, and  $\mathcal{K}$  is  $\mathcal{S}^\circ$  intersected with the unit ball. Then, [FKS21] consider the simplex regret-minimization setting, and show that OGD instantiated on  $\mathbb{R}_+^n = \mathcal{S}^\circ$  is equivalent to  $\text{RM}^+$ .

$\text{CBA}^+$  does not generalize either of the two simplex approaches above. For [Shi16], this can be seen because of the projection and subsequent dependence on stepsize required by [Shi16]’s construction. For [FKS21], this can be seen by the fact that they obtain  $\text{RM}^+$  when applying their approach to the simplex setting, and  $\text{CBA}^+$  differs from  $\text{RM}^+$  (in fact, that paper does not attempt to derive new unaccelerated regret minimizers; the goal is to design accelerated, or predictive/optimistic, variants of  $\text{RM}$  and  $\text{RM}^+$ ). However, an alternative derivation of  $\text{CBA}^+$  can be accomplished by using the generalization from [FKS21] where OGD is run with the decision set  $\mathcal{S}^\circ$ . Instead of applying OGD on the no-regret formulation of the Blackwell formulation of regret-minimization on the simplex with target set  $\mathcal{S} = \mathbb{R}_+^n$  as in [FKS21], we can apply the OGD-setup of [FKS21] to the Blackwell formulation of a general no-regret problem with decision set  $\mathcal{X}$  and target set  $\mathcal{C}^\circ$  where  $\mathcal{C} = \text{cone}(\{\kappa\} \times \mathcal{X})$ . Then, we would get an indirect proof of correctness of the  $\text{CBA}^+$  algorithm through the correctness of OGD and the two layers of reduction from regret minimization to Blackwell approachability to regret minimization. However, we believe that this approach is significantly less intuitive than understanding  $\text{CBA}^+$  directly in terms of its properties as a Blackwell approachability algorithm.

## B Proofs of Theorem 3.1

**Notations and classical results in conic optimization** We make use of the following facts. We provide a proof here for completeness.

**Lemma B.1.** *Let  $\mathcal{C} \subset \mathbb{R}^{n+1}$  a closed convex cone and  $\mathcal{C}^\circ$  its polar.*

1. *If  $\mathbf{u} \in \mathbb{R}^{n+1}$ , then  $\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}) = \pi_{\mathcal{C}}(\mathbf{u}) \in \mathcal{C}$ ,  $\langle \mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}), \pi_{\mathcal{C}^\circ}(\mathbf{u}) \rangle = 0$ , and  $\|\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u})\|_2 \leq \|\mathbf{u}\|_2$ .*

2. *If  $\mathbf{u} \in \mathbb{R}^{n+1}$  then*

$$d(\mathbf{u}, \mathcal{C}) = \max_{\mathbf{w} \in \mathcal{C}^\circ \cap B_2(1)} \langle \mathbf{u}, \mathbf{w} \rangle,$$

*where  $B_2(1) = \{\mathbf{w} \in \mathbb{R}^{n+1} \mid \|\mathbf{w}\|_2 \leq 1\}$ .*

3. *If  $\mathbf{u} \in \mathcal{C}$ , then  $d(\mathbf{u}, \mathcal{C}^\circ) = \|\mathbf{u}\|_2$ .*

4. Assume that  $\mathcal{C} = \text{cone}(\{\kappa\} \times \mathcal{X})$  with  $\mathcal{X} \subset \mathbb{R}^n$  convex compact and  $\kappa = \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$ . Then  $\mathcal{C}^\circ$  is a closed convex cone. Additionally, if  $\mathbf{u} \in \mathcal{C}$  we have  $-\mathbf{u} \in \mathcal{C}^\circ$ .

5. Let us write  $\leq_{\mathcal{C}^\circ}$  the order induced by  $\mathcal{C}^\circ : \mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y} \iff \mathbf{y} - \mathbf{x} \in \mathcal{C}^\circ$ . Then

$$\mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y}, \mathbf{x}' \leq_{\mathcal{C}^\circ} \mathbf{y}' \Rightarrow \mathbf{x} + \mathbf{x}' \leq_{\mathcal{C}^\circ} \mathbf{y} + \mathbf{y}', \forall \mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}' \in \mathbb{R}^{n+1}, \quad (6)$$

$$\mathbf{x} + \mathbf{x}' \leq_{\mathcal{C}^\circ} \mathbf{y} \Rightarrow \mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y}, \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^{n+1}, \forall \mathbf{x}' \in \mathcal{C}^\circ, \quad (7)$$

6. Assume that  $\mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y}$  for  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{n+1}$ . Then  $d(\mathbf{y}, \mathcal{C}^\circ) \leq \|\mathbf{x}\|_2$ .

*Proof.* 1. The fact that  $\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}) = \pi_{\mathcal{C}}(\mathbf{u}) \in \mathcal{C}, \langle \mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}), \pi_{\mathcal{C}^\circ}(\mathbf{u}) \rangle = 0$  follows from Moreau's Decomposition Theorem [CR13]. The fact that  $\|\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u})\|_2 \leq \|\mathbf{u}\|_2$  is a straightforward consequence of  $\langle \mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}), \pi_{\mathcal{C}^\circ}(\mathbf{u}) \rangle = 0$ .

2. For any  $\mathbf{w} \in \mathcal{C}^\circ \cap B_2(1)$  we have

$$\langle \mathbf{u}, \mathbf{w} \rangle \leq \langle \mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u}), \mathbf{w} \rangle \leq \|\mathbf{w}\|_2 \|\mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u})\|_2 \leq \|\mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u})\|_2.$$

Conversely, since  $(\mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u})) / \|\mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u})\|_2 \in \mathcal{C}^\circ$ , we have

$$\max_{\mathbf{w} \in \mathcal{C}^\circ \cap B_2(1)} \langle \mathbf{u}, \mathbf{w} \rangle \geq \|\mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u})\|_2.$$

This shows that

$$\max_{\mathbf{w} \in \mathcal{C}^\circ \cap B_2(1)} \langle \mathbf{u}, \mathbf{w} \rangle = \|\mathbf{u} - \pi_{\mathcal{C}}(\mathbf{u})\|_2 = d(\mathbf{u}, \mathcal{C}).$$

3. For any  $\mathbf{u} \in \mathbb{R}^{n+1}$ , by definition we have  $d(\mathbf{u}, \mathcal{C}^\circ) = \|\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u})\|_2$ . Now if  $\mathbf{u} \in \mathcal{C}$  we have  $\pi_{\mathcal{C}^\circ}(\mathbf{u}) = 0$  so  $d(\mathbf{u}, \mathcal{C}^\circ) = \|\mathbf{u}\|_2$ .

4. Let  $\mathbf{u} \in \mathcal{C}$ . Then  $\mathbf{u} = \alpha(\kappa, \mathbf{x})$  for  $\alpha \geq 0, \mathbf{x} \in \mathcal{X}$ . We will show that  $-\mathbf{u} \in \mathcal{C}^\circ$ . We have

$$\begin{aligned} -\mathbf{u} \in \mathcal{C}^\circ &\iff \langle -\mathbf{u}, \mathbf{u}' \rangle \leq 0, \forall \mathbf{u}' \in \mathcal{C} \\ &\iff \langle -\alpha(\kappa, \mathbf{x}), \alpha'(\kappa, \mathbf{x}') \rangle \leq 0, \forall \alpha' \geq 0, \forall \mathbf{x}' \in \mathcal{X} \\ &\iff \kappa^2 + \langle \mathbf{x}, \mathbf{x}' \rangle \geq 0 \\ &\iff -\langle \mathbf{x}, \mathbf{x}' \rangle \leq \kappa^2, \end{aligned}$$

and  $-\langle \mathbf{x}, \mathbf{x}' \rangle \leq \kappa^2$  is true by Cauchy-Schwartz and the definition of  $\kappa = \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$ .

5. We start by proving (6). Let  $\mathbf{x}, \mathbf{x}', \mathbf{y}, \mathbf{y}' \in \mathbb{R}^{n+1}$ , and assume that  $\mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y}, \mathbf{x}' \leq_{\mathcal{C}^\circ} \mathbf{y}'$ . Then  $\mathbf{y} - \mathbf{x} \in \mathcal{C}^\circ, \mathbf{y}' - \mathbf{x}' \in \mathcal{C}^\circ$ . Because  $\mathcal{C}^\circ$  is a convex set, and a cone, we have  $2 \cdot \left( \frac{\mathbf{y} - \mathbf{x}}{2} + \frac{\mathbf{y}' - \mathbf{x}'}{2} \right) \in \mathcal{C}^\circ$ . Therefore,  $\mathbf{y} + \mathbf{y}' - \mathbf{x} - \mathbf{x}' \in \mathcal{C}^\circ$ , i.e.,  $\mathbf{x} + \mathbf{x}' \leq_{\mathcal{C}^\circ} \mathbf{y} + \mathbf{y}'$ .

We now prove (7). Let  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{n+1}, \mathbf{x}' \in \mathcal{C}^\circ$  and assume that  $\mathbf{x} + \mathbf{x}' \leq_{\mathcal{C}^\circ} \mathbf{y}$ . Then by definition  $\mathbf{y} - \mathbf{x} - \mathbf{x}' \in \mathcal{C}^\circ$ . Additionally,  $\mathbf{x}' \in \mathcal{C}^\circ$  by assumption. Since  $\mathcal{C}^\circ$  is convex, and is a cone,  $2 \cdot \left( \frac{\mathbf{y} - \mathbf{x} - \mathbf{x}'}{2} + \frac{\mathbf{x}'}{2} \right) \in \mathcal{C}^\circ$ , i.e.,  $\mathbf{y} - \mathbf{x} \in \mathcal{C}^\circ$ . Therefore,  $\mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y}$ .

6. Let  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{n+1}$  such that  $\mathbf{x} \leq_{\mathcal{C}^\circ} \mathbf{y}$ . Then  $\mathbf{y} - \mathbf{x} \in \mathcal{C}^\circ$ . We have

$$d(\mathbf{y}, \mathcal{C}^\circ) = \min_{\mathbf{z} \in \mathcal{C}^\circ} \|\mathbf{y} - \mathbf{z}\|_2 \leq \|\mathbf{y} - (\mathbf{y} - \mathbf{x})\|_2 = \|\mathbf{x}\|_2.$$

□

Based on Moreau's Decomposition Theorem, we will use  $\pi_{\mathcal{C}}(\mathbf{u})$  and  $\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u})$  interchangeably.

**Results for various linear averaging schemes** We now present our convergence results for various linear averaging schemes. As a warm-up, we start with two theorems, Theorem B.2 and Theorem B.3, which show that CBA and CBA<sup>+</sup> are compatible with weighted average schemes, when *both* the decisions and the payoffs are weighted. The proofs for these theorems will be used in the proof of our main theorem, Theorem 3.1. For the sake of concision, in all the proofs of this section we will always write  $L = \max\{\|\mathbf{f}_t\|_2 \mid t \geq 1\}, \kappa = \max\{\|\mathbf{x}\|_2 \mid t \geq 1\}, \mathbf{v}_t = (\langle \mathbf{f}_t, \mathbf{x}_t \rangle / \kappa, -\mathbf{f}_t)$ . We start with the following theorem.

**Theorem B.2.** Let  $(\mathbf{x}_t)_{t \geq 0}$  the sequence of decisions generated by CBA with weights  $(\omega_t)_{t \geq 0}$  and let  $S_t = \sum_{\tau=1}^t \omega_\tau$  for any  $t \geq 1$ . Then

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O(\kappa \cdot d(\mathbf{u}_T, \mathcal{C}^\circ)).$$

Additionally,

$$d(\mathbf{u}_T, \mathcal{C}^\circ)^2 = O\left(L^2 \cdot \frac{\sum_{t=1}^T \omega_t^2}{\left(\sum_{t=1}^T \omega_t\right)^2}\right).$$

Overall,

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O\left(\kappa L \frac{\sqrt{\sum_{t=1}^T \omega_t^2}}{\sum_{t=1}^T \omega_t}\right).$$

*Proof.* The proof proceeds in two steps. We start by proving

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O(\kappa \cdot d(\mathbf{u}_T, \mathcal{C}^\circ)).$$

We have

$$d(\mathbf{u}_T, \mathcal{C}^\circ) = \max_{\mathbf{w} \in \text{cone}(\{\kappa\} \times \mathcal{X}) \cap B_2(1)} \left\langle \frac{1}{S_T} \sum_{t=1}^T \omega_t \mathbf{v}_t, \mathbf{w} \right\rangle \quad (8)$$

$$\begin{aligned} &\geq \max_{\mathbf{x} \in \mathcal{X}} \left\langle \frac{1}{S_T} \sum_{t=1}^T \omega_t \mathbf{v}_t, \frac{(\kappa, \mathbf{x})}{\|(\kappa, \mathbf{x})\|_2} \right\rangle \\ &\geq \frac{1}{S_T} \max_{\mathbf{x} \in \mathcal{X}} \frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{\|(\kappa, \mathbf{x})\|_2}, \end{aligned} \quad (9)$$

where (8) follows from Statement 1 in Lemma B.1, and (9) follows from CBA maintaining

$$\mathbf{u}_t = \left( \frac{1}{S_t} \sum_{\tau=1}^t \omega_\tau \frac{\langle \mathbf{f}_\tau, \mathbf{x}_\tau \rangle}{\kappa}, -\frac{1}{S_t} \sum_{\tau=1}^t \omega_\tau \mathbf{f}_\tau \right), \forall t \geq 1.$$

We can conclude that

$$2\kappa d(\mathbf{u}_T, \mathcal{C}^\circ) \geq \frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T}.$$

We now prove that

$$d(\mathbf{u}_T, \mathcal{C}^\circ)^2 = O\left(L^2 \frac{\sum_{\tau=1}^T \omega_\tau^2}{\left(\sum_{\tau=1}^T \omega_\tau\right)^2}\right).$$

We have

$$d(\mathbf{u}_{t+1}, \mathcal{C}^\circ)^2 = \min_{\mathbf{z} \in \mathcal{C}^\circ} \|\mathbf{u}_{t+1} - \mathbf{z}\|_2^2 \quad (10)$$

$$\begin{aligned} &\leq \|\mathbf{u}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2 \\ &\leq \left\| \frac{S_t}{S_t + \omega_{t+1}} \mathbf{u}_t + \frac{\omega_{t+1}}{S_t + \omega_{t+1}} \mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t) \right\|_2^2 \\ &\leq \left\| \frac{S_t}{S_t + \omega_{t+1}} (\mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)) + \frac{\omega_{t+1}}{S_t + \omega_{t+1}} (\mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)) \right\|_2^2 \\ &\leq \frac{1}{S_{t+1}^2} (S_t^2 \|\mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2 + \omega_{t+1}^2 \|\mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2) \\ &\quad + 2S_t \omega_{t+1} \langle \mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t), \mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t) \rangle \\ &\leq \frac{1}{S_{t+1}^2} (S_t^2 \|\mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2 + \omega_{t+1}^2 \|\mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2), \end{aligned} \quad (11)$$

where (11) follows from

$$\langle \mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t), \mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t) \rangle = 0. \quad (12)$$

This is because:

- $\langle \mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t), \mathbf{v}_{t+1} \rangle = 0$ . This is one of the crucial component of Blackwell's approachability framework: the current decision is chosen to force a hyperplane on the aggregate payoff. To see this, first note that  $\mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t) = \pi_{\mathcal{C}}(\mathbf{u}_t)$ . Let us write  $\boldsymbol{\pi} = (\tilde{\pi}, \hat{\pi}) = \pi_{\mathcal{C}}(\mathbf{u}_t)$ . Note that by definition,  $\mathbf{x}_{t+1} = (\kappa/\tilde{\pi})\hat{\pi}$ , and  $\mathbf{v}_{t+1} = (\langle \mathbf{f}_{t+1}, \mathbf{x}_{t+1} \rangle / \kappa, -\mathbf{f}_{t+1})$ . Therefore,

$$\begin{aligned} \langle \mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t), \mathbf{v}_{t+1} \rangle &= \langle \boldsymbol{\pi}, \mathbf{v}_{t+1} \rangle \\ &= \langle (\tilde{\pi}, \hat{\pi}), (\langle \mathbf{f}_{t+1}, \mathbf{x}_{t+1} \rangle / \kappa, -\mathbf{f}_{t+1}) \rangle \\ &= \langle (\tilde{\pi}, \hat{\pi}), (\langle \mathbf{f}_{t+1}, (\kappa/\tilde{\pi})\hat{\pi} \rangle / \kappa, -\mathbf{f}_{t+1}) \rangle \\ &= \langle \hat{\pi}, \mathbf{f}_{t+1} \rangle - \langle \hat{\pi}, \mathbf{f}_{t+1} \rangle \\ &= 0. \end{aligned}$$

- $\langle \mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t), \pi_{\mathcal{C}^\circ}(\mathbf{u}_t) \rangle = 0$  from Statement 3 of Lemma B.1 and  $\mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t) = \pi_{\mathcal{C}}(\mathbf{u}_t) \in \mathcal{C}$ .

We therefore have

$$d(\mathbf{u}_{t+1}, \mathcal{C}^\circ)^2 \leq \frac{1}{S_{t+1}^2} (S_t^2 \|\mathbf{u}_t - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2 + \omega_{t+1}^2 \|\mathbf{v}_{t+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_t)\|_2^2).$$

This recursion directly gives

$$d(\mathbf{u}_{t+1}, \mathcal{C}^\circ)^2 \leq \frac{1}{S_{t+1}^2} \sum_{\tau=1}^t \omega_{\tau+1}^2 \|\mathbf{v}_{\tau+1} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_\tau)\|_2^2 \leq O\left(L^2 \cdot \frac{\sum_{\tau=1}^{t+1} \omega_\tau^2}{S_{t+1}^2}\right),$$

where the last inequality follows from the definition of  $\mathbf{v}_t$  and  $L$ .  $\square$

**Theorem B.3.** Let  $(\mathbf{x}_t)_{t \geq 0}$  the sequence of decisions generated by CBA<sup>+</sup> with weights  $(\omega_t)_{t \geq 0}$  and let  $S_t = \sum_{\tau=1}^t \omega_\tau$  for any  $t \geq 1$ . Then

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O(\kappa \cdot d(\mathbf{u}_T, \mathcal{C}^\circ)).$$

Additionally,

$$d(\mathbf{u}_T, \mathcal{C}^\circ)^2 = O\left(L^2 \cdot \frac{\sum_{t=1}^T \omega_t^2}{\left(\sum_{t=1}^T \omega_t\right)^2}\right).$$

Overall,

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O\left(\kappa L \frac{\sqrt{\sum_{t=1}^T \omega_t^2}}{\sum_{t=1}^T \omega_t}\right).$$

*Proof of Theorem B.3.* The proof proceeds in two steps. We start by proving

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O(\kappa \cdot d(\mathbf{u}_T, \mathcal{C}^\circ)).$$

Recall that  $\mathbf{v}_t = (\langle \mathbf{f}_t, \mathbf{x}_t \rangle / \kappa, -\mathbf{f}_t)$ , and let us consider  $\mathbf{R}_t = \frac{1}{S_t} \sum_{\tau=1}^t \omega_\tau \mathbf{v}_\tau$ . By definition of  $\mathbf{R}_t$ , similarly as in the proof of Theorem B.2, we have

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O(\kappa \cdot d(\mathbf{R}_T, \mathcal{C}^\circ)).$$

Note that at any period  $t$ , we have

$$S_{t+1} \mathbf{u}_{t+1} - S_t \mathbf{u}_t \leq_{\mathcal{C}^\circ} S_{t+1} \mathbf{R}_{t+1} - S_t \mathbf{R}_t. \quad (13)$$

This is simply because  $\mathbf{u}_{t+1} = \pi_{\mathcal{C}}(\mathbf{u}_{t+1/2}) = \mathbf{u}_{t+1/2} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_{t+1/2})$  with

$$\mathbf{u}_{t+1/2} = \text{UPDATEPAYOFF}_{\text{CBA}}(\mathbf{u}_t) = \frac{S_t}{S_t + \omega_{t+1}} \mathbf{u}_t + \frac{\omega_{t+1}}{S_t + \omega_{t+1}} \mathbf{v}_{t+1}.$$

Now we have

$$\begin{aligned} S_{t+1} \mathbf{R}_{t+1} - S_t \mathbf{R}_t - (S_{t+1} \mathbf{u}_{t+1} - S_t \mathbf{u}_t) &= \omega_{t+1} \mathbf{v}_{t+1} + S_t \mathbf{u}_t - S_{t+1} \mathbf{u}_{t+1/2} + S_{t+1} \pi_{\mathcal{C}^\circ}(\mathbf{u}_{t+1/2}) \\ &= S_{t+1} \mathbf{u}_{t+1/2} - S_{t+1} \mathbf{u}_{t+1/2} + S_{t+1} \pi_{\mathcal{C}^\circ}(\mathbf{u}_{t+1/2}) \\ &= S_{t+1} \pi_{\mathcal{C}^\circ}(\mathbf{u}_{t+1/2}) \in \mathcal{C}^\circ. \end{aligned}$$

From (6) in Lemma B.1, we can sum the inequalities (13). Noticing that  $\mathbf{u}_1 = \mathbf{R}_1$ , we can conclude that

$$\mathbf{u}_t \leq_{\mathcal{C}^\circ} \mathbf{R}_t.$$

From  $\mathbf{u}_t \in \mathcal{C}$  and Statement 6 in Lemma B.1, we have  $d(\mathbf{R}_t, \mathcal{C}^\circ) \leq \|\mathbf{u}_t\|_2$ . This implies

$$\frac{\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle}{S_T} = O(\kappa \|\mathbf{u}_T\|_2).$$

We now turn to proving

$$\|\mathbf{u}_T\|_2^2 = O\left(L^2 \frac{\sum_{t=1}^T \omega_t^2}{\left(\sum_{t=1}^T \omega_t\right)^2}\right).$$

We have

$$\|\mathbf{u}_{t+1}\|_2^2 = \|\mathbf{u}_{t+1/2} - \pi_{\mathcal{C}^\circ}(\mathbf{u}_{t+1/2})\|_2^2 \tag{14}$$

$$\leq \|\mathbf{u}_{t+1/2}\|_2^2 \tag{15}$$

$$\leq \left\| \frac{S_t}{S_t + \omega_{t+1}} \mathbf{u}_t + \frac{\omega_{t+1}}{S_t + \omega_{t+1}} \mathbf{v}_{t+1} \right\|_2^2, \tag{16}$$

where (15) follows from Statement 1 in Lemma B.1. Therefore,

$$\|\mathbf{u}_{t+1}\|_2^2 \leq \frac{1}{(S_t + \omega_{t+1})^2} (S_t^2 \|\mathbf{u}_t\|_2^2 + \omega_{t+1}^2 \|\mathbf{v}_{t+1}\|_2^2 + 2S_t \omega_{t+1} \langle \mathbf{u}_t, \mathbf{v}_{t+1} \rangle).$$

By construction and for the same reason as for (12),  $\langle \mathbf{u}_t, \mathbf{v}_{t+1} \rangle = 0$ . Therefore, we have the recursion

$$\|\mathbf{u}_{t+1}\|_2^2 \leq \frac{1}{S_{t+1}^2} (S_t^2 \|\mathbf{u}_t\|_2^2 + \omega_{t+1}^2 \|\mathbf{v}_{t+1}\|_2^2).$$

By telescoping the inequality above we obtain

$$d(\mathbf{u}_{t+1}, \mathcal{C}^\circ)^2 \leq \frac{1}{S_{t+1}^2} \left( \sum_{\tau=1}^{t+1} \omega_\tau^2 \|\mathbf{v}_\tau\|_2^2 \right).$$

By definition of  $L$ ,

$$\|\mathbf{u}_{t+1}\|_2^2 = O\left(L^2 \cdot \frac{\sum_{\tau=1}^{t+1} \omega_\tau^2}{S_{t+1}^2}\right).$$

□

**Linear averaging only on decisions** We are now ready to prove our main convergence result, Theorem 3.1. Our proof heavily relies on the sequence of payoffs belonging to the cone  $\mathcal{C}$  at every iteration ( $\mathbf{u}_t \in \mathcal{C}, \forall t \geq 1$ ), and for this reason it does not extend to CBA. We also note that the use of conic optimization somewhat simplifies the argument compared to the proof that RM<sup>+</sup> is compatible with linear averaging [TBJB15].

*Proof of Theorem 3.1.* Recall that  $\mathbf{v}_t = (\langle \mathbf{f}_t, \mathbf{x}_t \rangle / \kappa, -\mathbf{f}_t)$ . By construction and following the same argument as for the proof of Theorem B.3, we have

$$\sum_{t=1}^T t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T t \langle \mathbf{f}_t, \mathbf{x} \rangle = O \left( \kappa \cdot d \left( \sum_{t=1}^T t \mathbf{v}_t, \mathcal{C}^\circ \right) \right). \quad (17)$$

Additionally, Equation (13) for uniform weights ( $\omega_\tau = 1, S_\tau = \tau$ ) yields

$$\mathbf{v}_{t+1} \geq_{\mathcal{C}^\circ} (t+1) \mathbf{u}_{t+1} - t \mathbf{u}_t.$$

Therefore,

$$(t+1) \mathbf{v}_{t+1} \geq_{\mathcal{C}^\circ} (t+1)^2 \mathbf{u}_{t+1} - t^2 \mathbf{u}_t - t \mathbf{u}_t.$$

Summing up the previous inequalities from  $t = 1$  to  $t = T - 1$  and using  $\mathbf{u}_1 = \mathbf{v}_1$  we obtain

$$\sum_{t=1}^T t \mathbf{v}_t \geq_{\mathcal{C}^\circ} T^2 \mathbf{u}_T - \sum_{t=1}^{T-1} t \mathbf{u}_t.$$

Note that since  $\sum_{t=1}^{T-1} t \mathbf{u}_t \in \mathcal{C}$ , Statement 4 in Lemma B.1 shows that  $-\sum_{t=1}^{T-1} t \mathbf{u}_t \in \mathcal{C}^\circ$ . Now, by applying (7) in Lemma B.1, we have

$$\sum_{t=1}^T t \mathbf{v}_t \geq_{\mathcal{C}^\circ} T^2 \mathbf{u}_T - \sum_{t=1}^{T-1} t \mathbf{u}_t \Rightarrow \sum_{t=1}^T t \mathbf{v}_t \geq_{\mathcal{C}^\circ} T^2 \mathbf{u}_T.$$

Since  $T^2 \mathbf{u}_T \in \mathcal{C}$ , Statement 6 shows that

$$d \left( \sum_{t=1}^T t \mathbf{v}_t, \mathcal{C}^\circ \right) \leq \|T^2 \mathbf{u}_T\|_2.$$

By construction  $\mathbf{u}_T$  is the output of CBA<sup>+</sup> with uniform weight, so that  $d(\mathbf{u}_T, \mathcal{C}^\circ) = \|\mathbf{u}_T\|_2 = O(L/\sqrt{T})$  (see Theorem B.3). Therefore,  $d(\sum_{t=1}^T t \mathbf{v}_t, \mathcal{C}^\circ) = O(L \cdot T^{3/2})$ . This shows that

$$\frac{\sum_{t=1}^T t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T t \langle \mathbf{f}_t, \mathbf{x} \rangle}{T(T+1)} = O \left( \kappa \frac{d \left( \sum_{t=1}^T t \mathbf{v}_t, \mathcal{C}^\circ \right)}{T(T+1)} \right) = O \left( \kappa L / \sqrt{T} \right).$$

□

**Comparisons of different weighted average schemes** We conclude this section with an empirical comparisons of the different weighted average schemes (Theorem B.2, Theorem B.3, and Theorem 3.1). We also compare these algorithms with RM<sup>+</sup>. We present our numerical experiments on sets of random matrix game instances in Figure 4. The setting is the same as in our simulation section, Section 4. We note that CBA<sup>+</sup> with linear averaging only on decisions outperforms both CBA<sup>+</sup> and CBA with linear averaging on both decisions and payoffs, as well as RM<sup>+</sup> with linear averaging on decisions.

## B.1 Geometric intuition on the projection step of CBA<sup>+</sup>

Figure 5 illustrates the projection step  $\pi_{\mathcal{C}}(\cdot)$  of CBA<sup>+</sup>. At a high level, from  $\mathbf{u}_t$  to  $\mathbf{u}_{t+1}$ , an instantaneous payoff  $\mathbf{v}_t$  is first added to  $\mathbf{u}_t$  (where  $\mathbf{v}_t = (\langle \mathbf{f}_t, \mathbf{x}_t \rangle / \kappa, -\mathbf{f}_t)$ ), and then the resulting vector  $\mathbf{u}_t^+ = \mathbf{u}_t + \mathbf{v}_t$  is projected onto  $\mathcal{C}$ . The projection  $\pi_{\mathcal{C}}(\cdot)$  moves the vector  $\mathbf{u}_t^+$  along the edges of the cone  $\mathcal{C}^\circ$ , preserving the (orthogonal) distance  $d$  to  $\mathcal{C}^\circ$ .

## C Proof of Theorem 3.2

Let  $\omega_t = t, S_T = \sum_{t=1}^T \omega_t = T(T+1)/2$ , and

$$\bar{\mathbf{x}}_T = \frac{1}{S_T} \sum_{t=1}^T \omega_t \mathbf{x}_t, \bar{\mathbf{y}}_T = \frac{1}{S_T} \sum_{t=1}^T \omega_t \mathbf{y}_t.$$

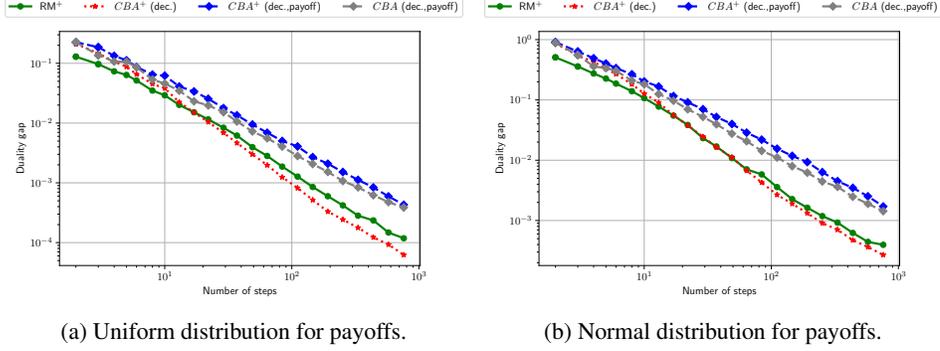


Figure 4: Comparison of RM<sup>+</sup> vs. CBA<sup>+</sup> and CBA with different linear averaging schemes: only on decisions (CBA<sup>+</sup> (dec.)), or on both the decisions and the payoffs  $u$  (CBA<sup>+</sup> (dec.,payoff), CBA (dec.,payoff)).

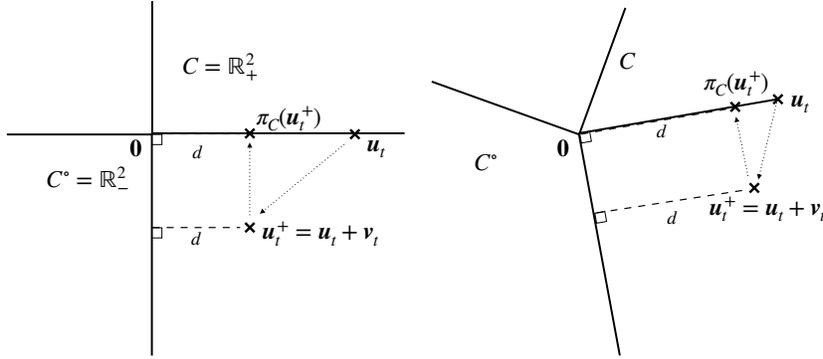


Figure 5: Illustration of  $\pi_C(\cdot)$  for  $C = \mathbb{R}_+^2$  (left-hand side) and  $C$  any cone in  $\mathbb{R}^2$  (right-hand side).

Since  $F$  is convex-concave, we first have

$$\max_{\mathbf{y} \in \mathcal{Y}} F(\bar{\mathbf{x}}_T, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}, \bar{\mathbf{y}}_T) \leq \frac{2}{T(T+1)} \left( \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t F(\mathbf{x}, \mathbf{y}_t) \right).$$

Now,

$$\begin{aligned} \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t F(\mathbf{x}, \mathbf{y}_t) &= \left( \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}) - \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}_t) \right) \\ &\quad + \left( \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}_t) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t F(\mathbf{x}, \mathbf{y}_t) \right). \end{aligned}$$

Now since  $F$  is convex-concave, we can use the following upper bound:

$$\begin{aligned} \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}) - \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}_t) &\leq \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \omega_t \langle \mathbf{g}_t, \mathbf{y} \rangle - \sum_{t=1}^T \omega_t \langle \mathbf{g}_t, \mathbf{y}_t \rangle, \\ \sum_{t=1}^T \omega_t F(\mathbf{x}_t, \mathbf{y}_t) - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t F(\mathbf{x}, \mathbf{y}_t) &\leq \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle, \end{aligned}$$

where  $\mathbf{f}_t \in \partial_{\mathbf{x}} F(\mathbf{x}_t, \mathbf{y}_t)$ ,  $\mathbf{g}_t \in \partial_{\mathbf{y}} F(\mathbf{x}_t, \mathbf{y}_t)$  (recall the repeated game framework presented at the beginning of Section 2).

Now we have proved in Theorem 3.1 that

$$\begin{aligned} \frac{2}{T(T+1)} \max_{\mathbf{y} \in \mathcal{Y}} \sum_{t=1}^T \omega_t \langle \mathbf{g}_t, \mathbf{y} \rangle - \sum_{t=1}^T \langle \omega_t \mathbf{g}_t, \mathbf{y}_t \rangle &= O\left(\kappa L / \sqrt{T}\right), \\ \frac{2}{T(T+1)} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle &= O\left(\kappa L / \sqrt{T}\right). \end{aligned}$$

Therefore, we can conclude that

$$\max_{\mathbf{y} \in \mathcal{Y}} F(\bar{\mathbf{x}}_T, \mathbf{y}) - \min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}, \bar{\mathbf{y}}_T) = O\left(\kappa L / \sqrt{T}\right).$$

**Remark C.1.** Note that we essentially reprove the folk theorem, except that we consider *weighted* average for the decisions of both players. This is because Theorem 3.2 uses linear averaging on decisions, whereas Theorem 2.1 is written with uniform averaging on decisions.

## D Proofs of the projections of Section 3.1

We will extensively use *Moreau's Decomposition Theorem* [CR13]: for any convex cone  $\mathcal{C} \subset \mathbb{R}^{n+1}$  and  $\mathbf{u} \in \mathbb{R}^{n+1}$ , we can decompose  $\mathbf{u} = \pi_{\mathcal{C}}(\mathbf{u}) + \pi_{\mathcal{C}^\circ}(\mathbf{u})$ , where  $\mathcal{C}^\circ$  is the *polar cone* of  $\mathcal{C}$ . Therefore, to compute  $\pi_{\mathcal{C}}(\mathbf{u})$ , it is sufficient to compute  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$ , the orthogonal projection of  $\mathbf{u}$  onto  $\mathcal{C}^\circ$ . We will see that in some cases, it is simpler to compute  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  and then use  $\pi_{\mathcal{C}}(\mathbf{u}) = \mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u})$  than directly computing  $\pi_{\mathcal{C}}(\mathbf{u})$  via solving (4).

### D.1 The case of the simplex

We consider  $\mathcal{X} = \Delta(n)$ . Note that in this case,  $\kappa = \max_{\mathbf{x} \in \Delta(n)} \|\mathbf{x}\|_2 = 1$ . The next lemma gives a closed-form expression of  $\mathcal{C}^\circ$ .

**Lemma D.1.** *Let  $\mathcal{C} = \text{cone}(\{1\}) \times \Delta(n)$ . Then  $\mathcal{C}^\circ = \{(\tilde{y}, \hat{\mathbf{y}}) \in \mathbb{R}^{n+1} \mid \max_{i=1, \dots, n} \hat{y}_i \leq -\tilde{y}\}$ .*

*Proof of Lemma D.1.* Note that for  $\mathbf{y} = (\tilde{y}, \hat{\mathbf{y}}) \in \mathbb{R}^{n+1}$  we have

$$\begin{aligned} \mathbf{y} \in \mathcal{C}^\circ &\iff \langle \mathbf{y}, \mathbf{z} \rangle \leq 0, \forall \mathbf{z} \in \mathcal{C} \\ &\iff \langle (\tilde{y}, \hat{\mathbf{y}}), \alpha(1, \mathbf{x}) \rangle \leq 0, \forall \mathbf{x} \in \Delta(n), \forall \alpha \geq 0 \\ &\iff \tilde{y} + \langle \hat{\mathbf{y}}, \mathbf{x} \rangle \leq 0, \forall \mathbf{x} \in \Delta(n) \\ &\iff \max_{\mathbf{x} \in \Delta(n)} \langle \hat{\mathbf{y}}, \mathbf{x} \rangle \leq -\tilde{y} \\ &\iff \max_{i=1, \dots, n} \hat{y}_i \leq -\tilde{y}. \end{aligned}$$

□

For a given  $\mathbf{u} = (\tilde{u}, \hat{\mathbf{u}})$ , computing  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  is now equivalent to solving

$$\min\{(\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \mid (\tilde{y}, \hat{\mathbf{y}}) \in \mathbb{R}^{n+1}, \max_{i=1, \dots, n} \hat{y}_i \leq -\tilde{y}\}. \quad (18)$$

Using the reformulation (18), we show that for a fixed  $\tilde{y}$ , the optimal  $\hat{\mathbf{y}}(\tilde{y})$  can be computed in closed-form. It is then possible to avoid a binary search over  $\tilde{y}$  and to simply use a sorting algorithm to obtain the optimal  $\tilde{y}$ . The next proposition summarizes our complexity result for  $\mathcal{X} = \Delta(n)$ .

**Proposition D.2.** *An optimal solution  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  to (18) can be computed in  $O(n \log(n))$  time.*

*Proof.* Computing  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  is equivalent to computing

$$\begin{aligned} \min & (\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \\ & \tilde{y} \in \mathbb{R}, \hat{\mathbf{y}} \in \mathbb{R}^n, \\ & \max_{i=1, \dots, n} \hat{y}_i \leq -\tilde{y}. \end{aligned}$$

Let us fix  $\tilde{y} \in \mathbb{R}$  and let us first solve

$$\begin{aligned} \min \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \\ \hat{\mathbf{y}} \in \mathbb{R}^n, \\ \max_{i=1, \dots, n} \hat{y}_i \leq -\tilde{y}. \end{aligned} \quad (19)$$

This is essentially the projection of  $\hat{\mathbf{u}}$  on  $(-\infty, -\tilde{y}]^n$ . So a solution to (19) is  $\hat{y}_i(\tilde{y}) = \min\{-\tilde{y}, \hat{u}_i\}, \forall i = 1, \dots, n$ . Note that in this case we have  $\hat{\mathbf{u}} - \hat{\mathbf{y}}(\tilde{y}) = (\hat{\mathbf{u}} + \tilde{y}\mathbf{e})^+$ . So overall the projection brings down to the optimization of  $\phi : \mathbb{R} \mapsto \mathbb{R}_+$  such that

$$\phi : \tilde{y} \mapsto (\tilde{y} - \tilde{u})^2 + \|(\hat{\mathbf{u}} + \tilde{y}\mathbf{e})^+\|_2^2. \quad (20)$$

In principle, we could use binary search with a doubling trick to compute a  $\epsilon$ -minimizer of the convex function  $\phi$  in  $O(\log(\epsilon^{-1}))$  calls to  $\phi$ . However, it is possible to find a minimizer  $\tilde{y}^*$  of  $\phi$  using the following remark. By construction, we know that  $\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}) \in \mathcal{C}$ . Here,  $\mathcal{C} = \text{cone}(\{1\} \times \Delta(n))$ , and  $\mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u}) = (\tilde{u} - \tilde{y}^*, (\hat{\mathbf{u}} + \tilde{y}^*\mathbf{e})^+)$ . This proves that

$$\frac{(\hat{\mathbf{u}} + \tilde{y}^*\mathbf{e})^+}{\tilde{u} - \tilde{y}^*} \in \Delta(n),$$

which in turns imply that

$$\tilde{y}^* + \sum_{i=1}^n \max\{\hat{u}_i + \tilde{y}^*, 0\} = \tilde{u}. \quad (21)$$

We can use (21) to efficiently compute  $\tilde{y}^*$  without using any binary search. In particular, we can sort the coefficients of  $\hat{\mathbf{u}}$  in  $O(n \log(n))$  operations, and use (21) to find  $\tilde{y}^*$ .  $\square$

Having obtained  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$ , we can obtain  $\pi_{\mathcal{C}}(\mathbf{u})$  by using the identity  $\pi_{\mathcal{C}}(\mathbf{u}) = \mathbf{u} - \pi_{\mathcal{C}^\circ}(\mathbf{u})$ . Note that  $\text{RM}$  and  $\text{RM}^+$  are obtained by choosing the closed-form feasible point corresponding to  $\tilde{y} = \tilde{u}$  in (18).

## D.2 The case of an $\ell_p$ ball

In this section we assume that  $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\|_p \leq 1\}$  with  $p \geq 1$  or  $p = \infty$ . The next lemma provides a closed-form reformulation of the polar cone  $\mathcal{C}^\circ$ .

**Lemma D.3.** *Let  $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\|_p \leq 1\}$ , with  $p \geq 1$  or  $p = \infty$ . Then  $\mathcal{C}^\circ = \{(\tilde{y}, \hat{\mathbf{y}}) \in \mathbb{R} \times \mathbb{R}^n \mid \|\hat{\mathbf{y}}\|_q \leq -\kappa\tilde{y}\}$ , with  $q$  such that  $1/p + 1/q = 1$ .*

*Proof of Lemma D.3.* Let us write  $B_p(1) = \{\mathbf{z} \in \mathbb{R}^n \mid \|\mathbf{z}\|_p \leq 1\}$ . Note that for  $\mathbf{y} = (\tilde{y}, \hat{\mathbf{y}}) \in \mathbb{R}^{n+1}$  we have

$$\begin{aligned} \mathbf{y} \in \mathcal{C}^\circ &\iff \langle \mathbf{y}, \mathbf{z} \rangle \leq 0, \forall \mathbf{z} \in \mathcal{C} \\ &\iff \langle (\tilde{y}, \hat{\mathbf{y}}), \alpha(\kappa, \mathbf{x}) \rangle \leq 0, \forall \mathbf{x} \in B_p(1), \forall \alpha \geq 0 \\ &\iff \kappa\tilde{y} + \langle \hat{\mathbf{y}}, \mathbf{x} \rangle \leq 0, \forall \mathbf{x} \in B_p(1), \\ &\iff \max_{\mathbf{x} \in B_p(1)} \langle \hat{\mathbf{y}}, \mathbf{x} \rangle \leq -\kappa\tilde{y} \\ &\iff \|\hat{\mathbf{y}}\|_q \leq -\kappa\tilde{y}, \end{aligned}$$

since  $\|\cdot\|_q$  is the dual norm of  $\|\cdot\|_p$ .  $\square$

The orthogonal projection problem onto  $\mathcal{C}^\circ$  becomes

$$\min\{(\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \mid (\tilde{y}, \hat{\mathbf{y}}) \in \mathbb{R}^{n+1}, \|\hat{\mathbf{y}}\|_q \leq -\kappa\tilde{y}\}. \quad (22)$$

For  $p = 2$ , (22) has a closed-form solution. For  $p = 1$ , a quasi-closed-form solution to (22) can be obtained efficiently using sorting. For  $p = \infty$ , it is more efficient to directly compute  $\pi_{\mathcal{C}}(\mathbf{u})$ . This is because the dual norm of  $\|\cdot\|_\infty$  is  $\|\cdot\|_1$ .

**Proposition D.4.** • For  $p = 1$ ,  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  can be computed in  $O(n \log(n))$  arithmetic operations.

- For  $p = \infty$ ,  $\pi_{\mathcal{C}}(\mathbf{u})$  can be computed in  $O(n \log(n))$  arithmetic operations.
- For  $p = 2$ ,  $\pi_{\mathcal{C}}(\mathbf{u})$  can be computed in closed-form.

*Proof.* **The case  $p = 1$ .** Assume that  $p = 1$ . Then  $\|\cdot\|_q = \|\cdot\|_\infty$  and  $\kappa = 1$ . We want to compute the projection of  $(\tilde{u}, \hat{\mathbf{u}})$  on  $\mathcal{C}^\circ$ :

$$\begin{aligned} \min_{\mathbf{y} \in \mathcal{C}^\circ} \|\mathbf{y} - \mathbf{u}\|_2^2 &= \min (\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \\ \tilde{y} \in \mathbb{R}, \hat{\mathbf{y}} \in \mathbb{R}^n, \\ \|\hat{\mathbf{y}}\|_\infty &\leq -\tilde{y}. \end{aligned} \quad (23)$$

For a fixed  $\tilde{y}$ , we want to compute

$$\begin{aligned} \min \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \\ \hat{\mathbf{y}} \in \mathbb{R}^n, \\ \|\hat{\mathbf{y}}\|_\infty &\leq -\tilde{y}. \end{aligned} \quad (24)$$

The projection (24) can be computed in closed-form as

$$\hat{\mathbf{y}}^*(\tilde{y}) = \min\{-\tilde{y}, \max\{\tilde{y}, \hat{\mathbf{u}}\}\} \quad (25)$$

since this is simply the orthogonal projection of  $\hat{\mathbf{u}}$  onto the  $\ell_\infty$  ball of radius  $-\tilde{y}$ . Let us call  $\phi : \mathbb{R} \mapsto \mathbb{R}$  such that

$$\phi(\tilde{y}) = (\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}}^*(\tilde{y}) - \hat{\mathbf{u}}\|_2^2.$$

Because of the closed-form expression for  $\hat{\mathbf{y}}^*(\tilde{y})$  as in (25), we have

$$\phi : \tilde{y} \mapsto (\tilde{y} - \tilde{u})^2 + \|(\hat{\mathbf{u}} + \tilde{y}\mathbf{e})^+\|_2^2.$$

Finding a minimizer of  $\phi$  can be done in  $O(n \log(n))$ , with the same methods as in the proof in the previous section (Appendix D.1).

**The case  $p = \infty$ .** Let  $p = \infty$ . The problem of computing  $\pi_{\mathcal{C}}(\mathbf{u})$ , the orthogonal projection onto the cone  $\mathcal{C}$ , is equivalent to

$$\begin{aligned} \min_{\mathbf{y} \in \mathcal{C}} \|\mathbf{y} - \mathbf{u}\|_2^2 &= \min (\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \\ \tilde{y} \in \mathbb{R}, \hat{\mathbf{y}} \in \mathbb{R}^n, \\ \|\hat{\mathbf{y}}\|_\infty &\leq \tilde{y}. \end{aligned} \quad (26)$$

Note the similarity between (26) (computing the orthogonal projection onto  $\mathcal{C}$  when  $p = \infty$ ), and (23) (computing the orthogonal projection onto  $\mathcal{C}^\circ$  when  $p = 1$ ). From Lemma D.3, we know that this is the case because  $\|\cdot\|_1$  and  $\|\cdot\|_\infty$  are dual norms to each other.

Therefore, the methods described for computing  $\pi_{\mathcal{C}^\circ}(\mathbf{u})$  for  $p = 1$  can be applied to the case  $p = \infty$  for directly computing  $\pi_{\mathcal{C}}(\mathbf{u})$ . This gives the complexity results as stated in Proposition D.4:  $\pi_{\mathcal{C}}(\mathbf{u})$  can be computed in  $O(n \log(n))$  operations.

**The case  $p = 2$ .** Let  $\|\cdot\|_p = \|\cdot\|_2$ , then  $\|\cdot\|_q = \|\cdot\|_2$ . Let us fix  $\tilde{y}$  and consider solving

$$\begin{aligned} \min \|\hat{\mathbf{y}} - \hat{\mathbf{u}}\|_2^2 \\ \hat{\mathbf{y}} \in \mathbb{R}^n, \\ \|\hat{\mathbf{y}}\|_2 &\leq -\tilde{y}. \end{aligned} \quad (27)$$

The projection (27) can be computed in closed-form as

$$\hat{\mathbf{y}}^*(\tilde{y}) = (-\tilde{y}) \frac{\hat{\mathbf{u}}}{\|\hat{\mathbf{u}}\|_2},$$

since this is just the orthogonal projection of the vector  $\hat{\mathbf{u}}$  onto the  $\ell_2$ -ball of radius  $-\tilde{y}$ . Let us call  $\phi : \mathbb{R} \mapsto \mathbb{R}$  such that

$$\phi(\tilde{y}) = (\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}}^*(\tilde{y}) - \hat{\mathbf{u}}\|_2^2.$$

Note that here,  $\tilde{y} \mapsto \hat{\mathbf{y}}^*(\tilde{y})$  is differentiable. Therefore  $\phi : \tilde{y} \mapsto (\tilde{y} - \tilde{u})^2 + \|\hat{\mathbf{y}}^*(\tilde{y}) - \hat{\mathbf{u}}\|_2^2$  is also differentiable. First-order optimality conditions yield a closed-form solution for computing  $(\tilde{y}^*, \hat{\mathbf{y}}^*) = \pi_{\mathcal{C}^\circ}(\mathbf{u})$ , as

$$\tilde{y}^* = \frac{\tilde{u} - \|\hat{\mathbf{u}}\|_2}{2}, \hat{\mathbf{y}}^* = -\frac{1}{2}(\tilde{u} - \|\hat{\mathbf{u}}\|_2) \frac{\hat{\mathbf{u}}}{\|\hat{\mathbf{u}}\|_2}. \quad (28)$$

□

### D.3 The case of an ellipsoidal confidence region in the simplex

In this section we assume that  $\mathcal{X}$  is  $\mathcal{X} = \{\mathbf{x} \in \Delta(n) \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x\}$ . We also assume that  $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x}^\top \mathbf{e} = 1\} \cap \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x\} \subseteq \Delta(n)$ , so that we can write  $\mathcal{X} = \mathbf{x}_0 + \epsilon_x \tilde{B}$ , where  $\tilde{B} = \{\mathbf{z} \in \mathbb{R}^n \mid \mathbf{z}^\top \mathbf{e} = 0, \|\mathbf{z}\|_2 \leq 1\}$ .

Suppose we took a sequence of decisions  $\mathbf{x}_1, \dots, \mathbf{x}_T$ , which can be written as  $\mathbf{x}_t = \mathbf{x}_0 + \epsilon_x \mathbf{z}_t$  for  $\mathbf{z}_t \in \tilde{B}$ . Then it is clear that for any sequence of payoffs  $\mathbf{f}_1, \dots, \mathbf{f}_T$ , we have

$$\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle = \epsilon_x \left( \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{z}_t \rangle - \min_{\mathbf{z} \in \tilde{B}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{z} \rangle \right). \quad (29)$$

Therefore, if we run CBA<sup>+</sup> on the set  $\tilde{B}$  to obtain  $O(\sqrt{T})$  growth of the right-hand side of (29), we obtain a no-regret algorithm for  $\mathcal{X}$ . We now show how to run CBA<sup>+</sup> for the set  $\tilde{B}$ . Let  $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^n \mid \mathbf{v}^\top \mathbf{e} = 0\}$ . We use the following orthonormal basis of  $\mathcal{V}$ : let  $\mathbf{v}_1, \dots, \mathbf{v}_{n-1} \in \mathbb{R}^n$  be the vectors  $\mathbf{v}_i = \sqrt{i/(i+1)}(1/i, \dots, 1/i, -1, 0, \dots, 0), \forall i = 1, \dots, n-1$ , where the component  $1/i$  is repeated  $i$  times. The vectors  $\mathbf{v}_1, \dots, \mathbf{v}_{n-1}$  are orthonormal and constitute a basis of  $\mathcal{V}$  [EPGMFBV03]. Writing  $\mathbf{V} = (\mathbf{v}_1, \dots, \mathbf{v}_{n-1}) \in \mathbb{R}^{n \times (n-1)}$ , and noting that  $\mathbf{V}^\top \mathbf{V} = \mathbf{I}$ , we can write  $\tilde{B} = \{\mathbf{V}\mathbf{s} \mid \mathbf{s} \in \mathbb{R}^{n-1}, \|\mathbf{s}\|_2 \leq 1\}$ . Now, if  $\mathbf{x} = \mathbf{x}_0 + \epsilon_x \mathbf{z}_t$  with  $\mathbf{z}_t \in \mathcal{V}$ , we have  $\mathbf{z}_t = \mathbf{V}\mathbf{s}_t$ , for  $\mathbf{s}_t \in \mathbb{R}^{n-1}$  and  $\|\mathbf{s}_t\|_2 \leq 1$ . Finally,

$$\sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x}_t \rangle - \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^T \omega_t \langle \mathbf{f}_t, \mathbf{x} \rangle = \epsilon_x \left( \sum_{t=1}^T \omega_t \langle \mathbf{V}^\top \mathbf{f}_t, \mathbf{s}_t \rangle - \min_{\mathbf{s} \in \mathbb{R}^{n-1}, \|\mathbf{s}\|_2 \leq 1} \sum_{t=1}^T \omega_t \langle \mathbf{V}^\top \mathbf{f}_t, \mathbf{s} \rangle \right). \quad (30)$$

Therefore, to obtain a regret minimizer for the left-hand side of (30) with observed payoffs  $(\mathbf{f}_t)_{t \geq 0}$ , we can run CBA<sup>+</sup> on the right-hand side, where the decision set is an  $\ell_2$  ball and the sequence of observed payoffs is  $(\mathbf{V}^\top \mathbf{f}_t)_{t \geq 0}$ . In the previous section we showed how to efficiently instantiate CBA<sup>+</sup> in this setting (see Proposition D.4).

**Remark D.5.** In this section we have highlighted a sequence of reformulations of the regret, from (29) to (30). We essentially showed how to instantiate CBA<sup>+</sup> for settings where the decision set  $\mathcal{X}$  is the intersection of an  $\ell_2$  ball with a hyperplane for which we have an orthonormal basis.

## E Additional details and numerical experiments for matrix games and EFGs

### E.1 Numerical setup

**Numerical setup for matrix games** For the experiments on matrix games, we sample at random the matrix of payoffs  $\mathbf{A} \in \mathbb{R}^{n \times m}$  and we let  $n, m = 10$ , where  $n, m$  represent the number of actions of each player. We average our results over 70 instances. The decision sets  $\mathcal{X}$  and  $\mathcal{Y}$  are given as  $\mathcal{X} = \Delta(n)$  and  $\mathcal{Y} = \Delta(m)$ .

**Alternation** Alternation is a method which improves the performances of RM and RM<sup>+</sup> [BMS19]. We leave proving this for CBA and CBA<sup>+</sup> to future works. Using alternation, the players play in turn, instead of playing at the same time. In particular, the  $y$ -player may observe the current decision  $\mathbf{x}_t$  of the  $x$ -player at period  $t$ , before choosing its own decision  $\mathbf{y}_t$ . For CBA and CBA<sup>+</sup>, it is implemented as follows. At period  $t \geq 2$ ,

1. The  $x$ -player chooses  $\mathbf{x}_t$  using its payoff  $\mathbf{u}_{t-1}^x$  :  $\mathbf{x}_t = \text{CHOOSEDECISION}(\mathbf{u}_{t-1}^x)$ .
2. The  $y$ -player observes  $\mathbf{g} \in \partial_{\mathbf{y}}F(\mathbf{x}_t, \mathbf{y}_{t-1})$  and updates  $\mathbf{u}_t^y$ :

$$\mathbf{u}_t^y = \text{UPDATEPAYOFF}_{\text{CBA}^+}(\mathbf{u}_{t-1}^y, \mathbf{y}_{t-1}, \mathbf{g}, \omega_t, \sum_{\tau=1}^{t-1} \omega_\tau).$$

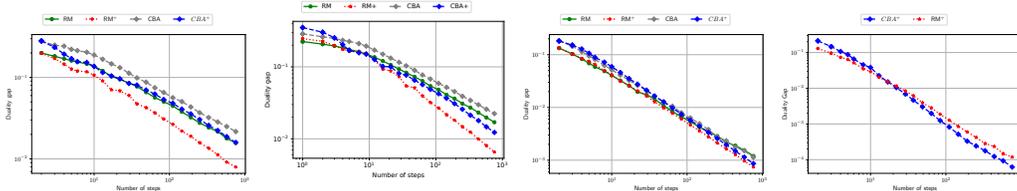
3. The  $y$ -player chooses  $\mathbf{y}_t$  using  $\mathbf{u}_{t-1}^y$  :  $\mathbf{y}_t = \text{CHOOSEDECISION}(\mathbf{u}_{t-1}^y)$ .
4. The  $x$ -player observes  $\mathbf{f} \in \partial_{\mathbf{x}}F(\mathbf{x}_t, \mathbf{y}_t)$  and updates  $\mathbf{u}_t^x$ :

$$\mathbf{u}_t^x = \text{UPDATEPAYOFF}_{\text{CBA}^+}(\mathbf{u}_t^x, \mathbf{x}_t, \mathbf{f}, \omega_t, \sum_{\tau=1}^{t-1} \omega_\tau).$$

## E.2 Comparing RM, RM<sup>+</sup>, CBA, and CBA<sup>+</sup> on matrix games

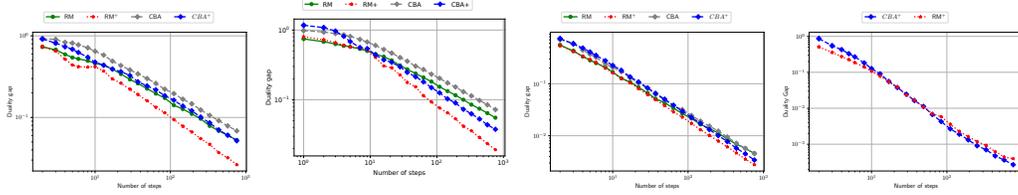
In Figure 6 and Figure 7, we show the performances of RM, RM<sup>+</sup>, CBA and CBA<sup>+</sup> with and without alternation, and with and without linear averaging. On the  $y$ -axis we show the duality gap of the current averaged decisions  $(\bar{\mathbf{x}}_T, \bar{\mathbf{y}}_T)$ . On the  $x$ -axis we show the number of iterations.

- In Figure 6a and Figure 7a, the four algorithms do not use alternation nor linear averaging, i.e., the four algorithms use uniform weights on the sequence of decisions. We note that RM<sup>+</sup> is the best algorithm in this setting, outperforming CBA<sup>+</sup>, which performs similarly as RM.
- In Figure 6b and Figure 7b, we note that linear averaging (only on the decisions) slightly improves the performances of RM<sup>+</sup> and CBA<sup>+</sup>. Note that RM and CBA are not known to be compatible with linear averaging (on decisions only), so we use uniform weights for RM and CBA here.
- In Figures 6c and Figure 7c, the four algorithms use alternation, but not linear averaging. The performances of the four algorithms are very similar.
- Finally, in Figure 6d and Figure 7d, RM<sup>+</sup> and CBA<sup>+</sup> use linear averaging on decisions *and* alternation. We see that the strongest performances are achieved by CBA<sup>+</sup>. Recall that RM and CBA are not known to be compatible with linear averaging (on decisions only), so we do not show their performances here.
- **Conclusion of our experiments.** We see that it is both alternation *and* linear averaging that enable the strong empirical performances of CBA<sup>+</sup>, rendering it capable to outperform RM<sup>+</sup>. Crucially, it is the “+ operation” that enables CBA<sup>+</sup> (and RM<sup>+</sup>) to be compatible with linear averaging on the decisions only and to outperform CBA and RM.



(a) Without alternation nor linear averaging. (b) Without alternation but with linear averaging (only for RM<sup>+</sup> and CBA<sup>+</sup>). (c) With alternation but no linear averaging. (d) With alternation and linear averaging.

Figure 6: Comparison of RM, RM<sup>+</sup>, CBA and CBA<sup>+</sup> for  $\mathcal{X} = \Delta(n), \mathcal{Y} = \Delta(m)$  and random matrices, with and without alternations, and with and without linear averaging. We choose  $n, m = 10$  and  $A_{ij} \sim U[0, 1]$  over 70 instances.



(a) Without alternation nor linear averaging. (b) Without alternation but with linear averaging (only for  $RM^+$  and  $CBA^+$ ). (c) With alternation but no linear averaging. (d) With alternation and linear averaging.

Figure 7: Comparison of  $RM$ ,  $RM^+$ ,  $CBA$  and  $CBA^+$  for  $\mathcal{X} = \Delta(n)$ ,  $\mathcal{Y} = \Delta(m)$  and random matrices, with and without alternations, and with and without linear averaging. We choose  $n, m = 10$  and  $A_{ij} \sim N(0, 1)$  over 70 instances.

### E.3 Additional numerical experiments for matrix games

We have seen in Appendix E.2 that  $RM^+$  and  $CBA^+$  with alternation and linear averaging are outperforming  $RM$  and  $CBA$ . In Figure 1a, we have compared both  $RM^+$  and  $CBA^+$  with AdaFTRL [OP15] and AdaHedge [DRVEGK14], two algorithms that also enjoy the desirable *scale-free* property, i.e., their sequences of decisions remain invariant when the losses are scaled by a constant factor. In the next figure, we provide additional comparisons of  $RM^+$ ,  $CBA^+$ , AdaHedge and AdaFTRL when the coefficients of the matrix of payoff are normally distributed. We found that  $RM^+$  and  $CBA^+$  are both outperforming AdaHedge and AdaFTRL, a situation similar to the case of uniform payoffs (Figure 1a.)

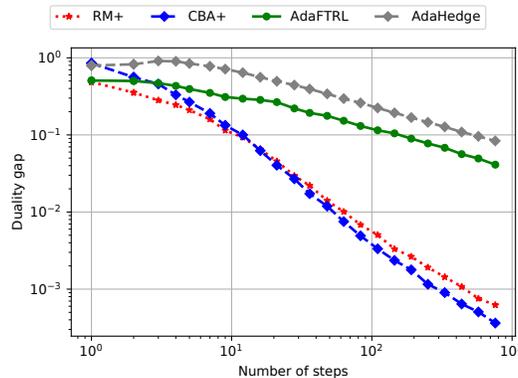


Figure 8: Comparison of  $RM^+$ ,  $CBA^+$ , AdaHedge and AdaFTRL for  $\mathcal{X} = \Delta(n)$ ,  $\mathcal{Y} = \Delta(m)$  and random matrices. We choose  $n, m = 10$  and  $A_{ij} \sim N(0, 1)$  over 70 instances.

### E.4 Additional numerical experiments for EFGs

In Section 4.1, we have compared  $CBA^+$  (using alternation and linear averaging) and  $CFR^+$  on various EFGs instances. We present in Figure 9 additional simulations where  $CBA^+$  and  $CFR^+$  performs similarly. A description of the games can be found in [FKS21]. On the  $y$ -axis we show the duality gap of the current averaged decisions  $(\bar{x}_T, \bar{y}_T)$ . On the  $x$ -axis we show the number of iterations.

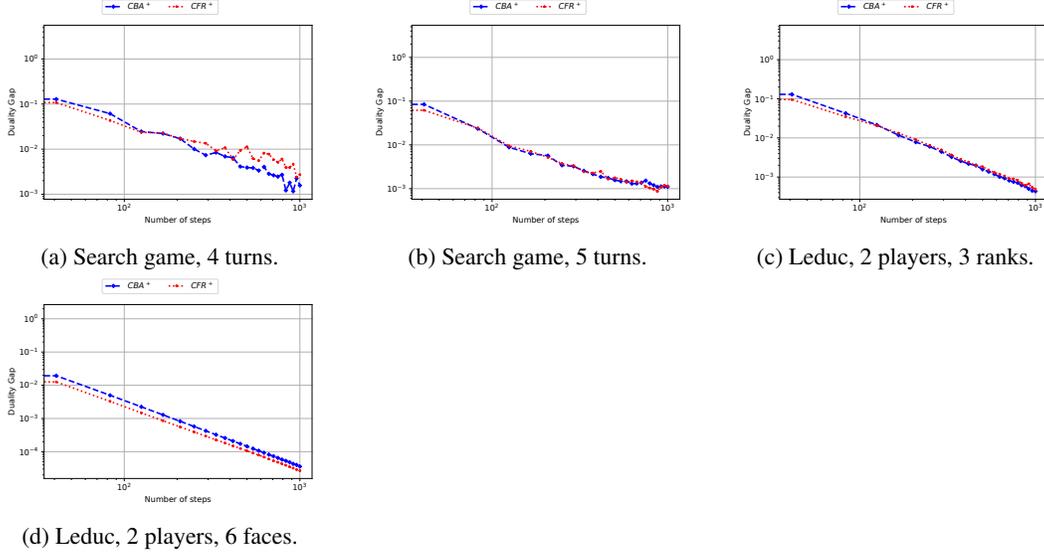


Figure 9: Comparison of  $\text{CBA}^+$  and  $\text{CFR}^+$  for various Extensive-Form Games (EFG) instances.

## F OMD, FTRL and optimistic variants

### F.1 Algorithms

Let us fix some step size  $\eta > 0$ . For solving our instances of distributionally robust optimization, we compare Algorithm  $\text{CBA}^+$  with the following four state-of-the-art algorithms:

1. Follow-The-Regularized-Leader (FTRL) [AHR09, McM11]:

$$\mathbf{x}_{t+1} \in \arg \min_{\mathbf{x} \in \mathcal{X}} \left\langle \sum_{\tau=1}^t \mathbf{f}_{\tau}, \mathbf{x} \right\rangle + \frac{1}{\eta} \|\mathbf{x}\|_2^2. \quad (\text{FTRL})$$

Optimistic FTRL [RS13]: given an estimation  $\mathbf{m}^{t+1}$  of loss at period  $t+1$ , choose

$$\mathbf{x}_{t+1} \in \arg \min_{\mathbf{x} \in \mathcal{X}} \left\langle \sum_{\tau=1}^t \mathbf{f}_{\tau} + \mathbf{m}^{t+1}, \mathbf{x} \right\rangle + \frac{1}{\eta} \|\mathbf{x}\|_2^2. \quad (\text{O-FTRL})$$

2. Online Mirror Descent (OMD) [NY83, BT03]:

$$\mathbf{x}_{t+1} \in \min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{f}_t, \mathbf{x} \rangle + \frac{1}{\eta} \|\mathbf{x} - \mathbf{x}_t\|_2^2. \quad (\text{OMD})$$

Optimistic OMD [CYL<sup>+</sup>12]: given an estimation  $\mathbf{m}^{t+1}$  of loss at period  $t+1$ ,

$$\mathbf{z}_{t+1} \in \min_{\mathbf{z} \in \mathcal{X}} \langle \mathbf{m}_{t+1}, \mathbf{z} \rangle + \frac{1}{\eta} \|\mathbf{z} - \mathbf{x}_t\|_2^2,$$

Observe the loss  $\mathbf{f}_{t+1}$  related to  $\mathbf{z}_{t+1}$ , (O-OMD)

$$\mathbf{x}_{t+1} \in \min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{f}_{t+1}, \mathbf{x} \rangle + \frac{1}{\eta} \|\mathbf{x} - \mathbf{x}_t\|_2^2.$$

Note that a priori these algorithms can be written more generally using Bregman divergence (e.g., [BTN01]). We choose to work with  $\|\cdot\|_2$  instead of *Kullback-Leibler divergence* as this  $\ell_2$ -setup is usually associated with faster empirical convergence rates [CP16, GKG19]. Additionally, following [CYL<sup>+</sup>12, RS13], we use the last observed loss as the predictor for the next loss, i.e., we set  $\mathbf{m}^{t+1} = \mathbf{f}_t$ .

## F.2 Implementations

When  $\mathcal{X}$  is the simplex or a ball based on the  $\ell_2$ -distance and centered at  $\mathbf{0}$ , there is a closed-form solution to the proximal updates for FTRL, OMD, O-FTRL and O-OMD. However, it is not clear how to compute these proximal updates for different settings, e.g., when  $\mathcal{X}$  is a subset of the simplex or an  $\ell_p$ -ball. We present the details of our implementation below. The results in the rest of this section are reminiscent to the novel tractable proximal setups presented in [GCK20a, GCK20b].

**Computing the projection steps for min-player** For  $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x\}$ ,  $\mathbf{c}, \mathbf{x}' \in \mathbb{R}^n$  and a step size  $\eta > 0$ , the proximal update becomes

$$\min_{\|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x} \langle \mathbf{c}, \mathbf{x} \rangle + \frac{1}{2\eta} \|\mathbf{x} - \mathbf{x}'\|_2^2. \quad (31)$$

This is the same arg min as

$$\min_{\|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x} \|\mathbf{x} - (\mathbf{x}' - \eta\mathbf{c})\|_2^2.$$

We can change  $\mathbf{x}$  by  $\mathbf{z} = (\mathbf{x} - \mathbf{x}_0) / \epsilon_x$  to solve the equivalent program

$$\min_{\|\mathbf{z}\|_2 \leq 1} \|\mathbf{z} - \frac{1}{\epsilon_x} (\mathbf{x}' - \eta\mathbf{c} - \mathbf{x}_0)\|_2^2.$$

The solution to the above program is

$$\mathbf{z} = \frac{\mathbf{x}' - \eta\mathbf{c} - \mathbf{x}_0}{\max\{\epsilon_x, \|\mathbf{x}' - \eta\mathbf{c} - \mathbf{x}_0\|_2\}}.$$

From  $\mathbf{x} = \mathbf{x}_0 + \epsilon_x \mathbf{z}$  we obtain  $\mathbf{x}^*$  the solution to (31)

$$\mathbf{x}^* = \mathbf{x}_0 + \epsilon_x \frac{\mathbf{x}' - \eta\mathbf{c} - \mathbf{x}_0}{\max\{\epsilon_x, \|\mathbf{x}' - \eta\mathbf{c} - \mathbf{x}_0\|_2\}}.$$

**Computing the projection steps for max-player** For  $\mathcal{Y} = \{\mathbf{y} \in \Delta(m) \mid \|\mathbf{y} - \mathbf{y}_0\|_2 \leq \epsilon_y\}$ , the proximal update of the max-player from a previous point  $\mathbf{y}'$  and a step size of  $\eta > 0$  becomes

$$\min_{\|\mathbf{y} - \mathbf{y}_0\|_2 \leq \epsilon_y, \mathbf{y} \in \Delta(m)} \langle \mathbf{c}, \mathbf{y} \rangle + \frac{1}{2\eta} \|\mathbf{y} - \mathbf{y}'\|_2^2. \quad (32)$$

If we dualize the  $\ell_2$  constraint with a Lagrangian multiplier  $\mu \geq 0$  we obtain the relaxed problem  $q(\mu)$  where

$$q(\mu) = -(1/2)\epsilon_y^2\mu + \min_{\mathbf{y} \in \Delta(m)} \langle \mathbf{c}, \mathbf{y} \rangle + \frac{1}{2\eta} \|\mathbf{y} - \mathbf{y}'\|_2^2 + \frac{\mu}{2} \|\mathbf{y} - \mathbf{y}_0\|_2^2. \quad (33)$$

Note that the arg min in

$$\min_{\mathbf{y} \in \Delta(m)} \langle \mathbf{c}, \mathbf{y} \rangle + \frac{1}{2\eta} \|\mathbf{y} - \mathbf{y}'\|_2^2 + \frac{\mu}{2} \|\mathbf{y} - \mathbf{y}_0\|_2^2$$

is the same arg min as in

$$\min_{\mathbf{y} \in \Delta(m)} \left\| \mathbf{y} - \frac{\eta}{\eta\mu + 1} \left( \frac{1}{\eta} \mathbf{y}' + \mu \mathbf{y}_0 - \mathbf{c} \right) \right\|_2^2. \quad (34)$$

Note that (34) is an orthogonal projection onto the simplex. Therefore, it can be solved efficiently [DSSSC08]. We call  $\mathbf{y}(\mu)$  an optimal solution of (34). Then  $q(\mu)$  can be rewritten

$$q(\mu) = -(1/2)\epsilon_y^2\mu + \langle \mathbf{c}, \mathbf{y}(\mu) \rangle + \frac{1}{2\eta} \|\mathbf{y}(\mu) - \mathbf{y}'\|_2^2 + \frac{\mu}{2} \|\mathbf{y}(\mu) - \mathbf{y}_0\|_2^2.$$

We can therefore binary search  $q(\mu)$  as in the previous expression. An upper bound  $\bar{\mu}$  for  $\mu^*$  can be computed as follows. Note that

$$q(\mu) \leq -(1/2)\epsilon_y^2\mu + \langle \mathbf{c}, \mathbf{y}_0 \rangle + \frac{1}{2\eta} \|\mathbf{y}_0 - \mathbf{y}'\|_2^2.$$

Since  $\mu \mapsto q(\mu)$  is concave we can choose  $\bar{\mu}$  such that  $q(\bar{\mu}) \leq q(0)$ . Using the previous inequality this yields

$$\bar{\mu} = \frac{2}{\epsilon_y^2} \left( \langle \mathbf{c}, \mathbf{y}_0 \rangle + \frac{1}{2\eta} \|\mathbf{y}_0 - \mathbf{y}'\|_2^2 - q(0) \right).$$

We choose a precision of  $\epsilon = 0.001$  in our simulations. Note that these binary searches make OMD, FTRL, O-FTRL and O-OMD slower than CBA<sup>+</sup> in terms of running times, since the updates in CBA<sup>+</sup> only requires to compute the projection  $\pi_{\mathcal{C}}(\mathbf{u})$ , and we have shown in Proposition D.4 and Appendix D.2 how to compute this in  $O(n)$  when  $\mathcal{X}$  is an  $\ell_2$  ball  $\mathcal{X} = \{\mathbf{x} \in \Delta(n) \mid \|\mathbf{x} - \mathbf{x}_0\|_2 \leq \epsilon_x\}$ .

**Computing the theoretical step sizes** We now give details about the choice of choice of theoretical step sizes. In theory (e.g., [BTN01]), for a player with decision set  $\mathcal{X}$ , we can choose  $\eta_{\text{th}} = \sqrt{2}\Omega/L\sqrt{T}$  with  $\Omega = \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{X}} \|\mathbf{x} - \mathbf{x}'\|_2$ , and  $L$  an upper bound on the norm of any observed loss  $\mathbf{f}_t$ :  $\|\mathbf{f}_t\|_2 \leq L, \forall t \geq 1$ . Note that this requires to know 1) the number of steps  $T$ , and 2) the upper bound  $L$  on the norm of any observed loss  $\mathbf{f}_t$ , before the losses are generated. We now show how to compute  $L_x$  and  $L_y$  (for the  $x$ -player and the  $y$ -player) for an instance of the distributionally robust optimization problem (5).

1. For the  $y$ -player, the loss  $\mathbf{f}_t$  is  $\mathbf{f}_t = (\ell_i(\mathbf{x}_t))_{i \in [1, m]}$ , with  $\ell_i(\mathbf{x}) = \log(1 + \exp(-b_i \mathbf{a}_i^\top \mathbf{x}))$ . For each  $i \in [1, m]$  we have  $|\ell_i| \leq \log(1 + \exp(|b_i|R\|\mathbf{a}_i\|_2))$  so that

$$L_y = \sqrt{\sum_{i=1}^m \log(1 + \exp(|b_i|R\|\mathbf{a}_i\|_2))^2}.$$

2. For the  $x$ -player we have  $\mathbf{f}_t = \mathbf{A}^t \mathbf{y}_t$ , where  $\mathbf{A}^t$  is the matrix of subgradients of  $\mathbf{x} \mapsto F(\mathbf{x}, \mathbf{y}_t)$  at  $\mathbf{x}_t$ :

$$A_{ij}^t = \frac{-1}{1 + \exp(b_i \mathbf{a}_i^\top \mathbf{x}_t)} b_i a_{i,j}, \forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}.$$

Therefore,  $\|\mathbf{f}_t\|_2 \leq \|\mathbf{A}^t\|_2 \|\mathbf{y}_t\|_2 \leq \|\mathbf{A}^t\|_2$ , because  $\mathbf{y} \in \Delta(m)$ . Now we have  $\|\mathbf{A}^t\|_2 \leq \|\mathbf{A}^t\|_F = \sqrt{\sum_{i,j} |A_{ij}^t|^2}$ . From  $|A_{ij}^t| \leq |b_i a_{i,j}|$  we use

$$L_x = \sqrt{\sum_{i,j} |b_i a_{i,j}|^2}.$$

## G Additional details and numerical experiments for distributionally robust optimization

We compare CBA<sup>+</sup> with alternation and linear averaging, OMD, FTRL, O-OMD and O-OMD for various step sizes  $\eta$  where  $\eta = \alpha \eta_{\text{th}}$  for  $\alpha \in \{1, 100, 1,000, 10,000\}$ , on additional synthetic and real data sets. We also add a comparison with adaptive step sizes.

**Data sets** We present here the characteristics of the data sets that we use in our DRO simulations. All data sets can be downloaded from the `libsvm` classification libraries<sup>4</sup>

- *Adult* data set: two classes,  $m = 1,605$  samples with  $n = 123$  features.
- *Australian* data set: two classes,  $m = 690$  samples with  $n = 14$  features.
- *Madelon* data set: two classes,  $m = 2,000$  samples with  $n = 500$  features.
- *Splice* data set: two classes,  $m = 1,000$  samples with  $n = 60$  features.

**Additional experiments with fixed step sizes** In this section we present additional numerical experiments for solving distributionally robust optimization instances in Figure 10. We use a synthetic data set, where we sample the features  $a_{ij}$  as uniform random variables in  $[0, 1]$ . We also present results for the *adult* and the *australian* data sets from `libsvm`. We vary the aggressiveness of the step sizes  $\eta = \alpha \eta_{\text{th}}$  by multiplying the theoretical step sizes  $\eta_{\text{th}}$  by a multiplicative step factor  $\alpha$ . The empirical setting is the same as in Section 4. We note that our algorithm still outperforms or performs on par with the classical approaches after  $10^2$  iterations, without requiring a single choice of parameter.

<sup>4</sup><https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

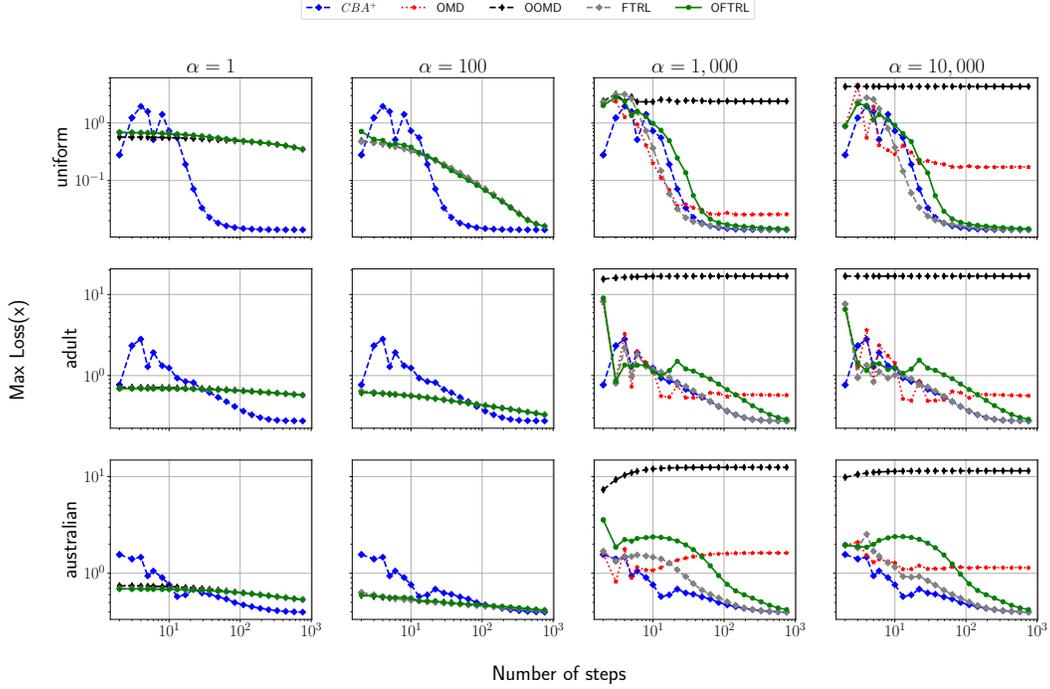


Figure 10: Comparisons of the performances of  $\text{CBA}^+$  with alternation and linear averaging, OMD, FTRL, O-OMD and O-FTRL on synthetic (with *uniform* distribution) and real data sets (*adult* and *australian*). We use fixed step sizes  $\eta = \alpha\eta_{\text{th}}$ , where  $\eta_{\text{th}}$  is the theoretical step size that guarantees convergence.

**Additional experiments with adaptive step sizes** We present our additional results with adaptive step sizes in Figure 11. Given  $v_t$  the payoff observed by the player at period  $t$ , and following [Ora19], we choose the step sizes  $(\eta_t)_{t \geq 1}$  as

$$\eta_t = 1 / \sqrt{\sum_{\tau=1}^t \|v_\tau\|_2^2}. \quad (35)$$

We note that  $\text{CBA}^+$  still outperforms, or performs on par, with the state-of-the-art approaches.

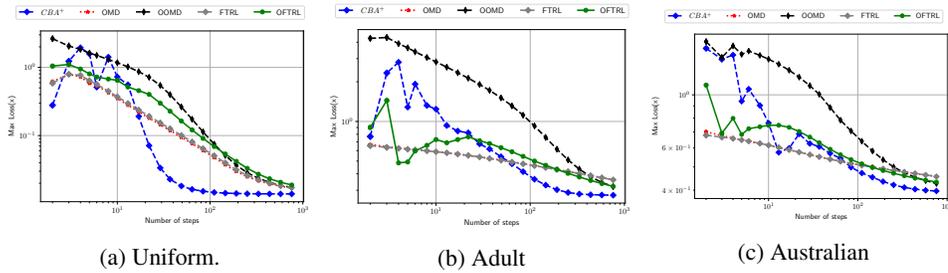


Figure 11: Comparisons of the performances of  $\text{CBA}^+$  with alternation and linear averaging, OMD, FTRL, O-OMD and O-FTRL on synthetic (with *uniform* distribution) and real data sets (*adult* and *australian*). For the non-parameter free algorithms, we use the adaptive step sizes as in (35).