

ONLINE DECISION-FOCUSED LEARNING

Aymeric Capitaine*

Maxime Haddouche †

Eric Moulines ‡

Michael I. Jordan

Etienne Boursier ¶

Alain Durmus*

§

ABSTRACT

Decision-focused learning (DFL) is an increasingly popular paradigm for training predictive models whose outputs are used in decision-making tasks. Instead of merely optimizing for predictive accuracy, DFL trains models to directly minimize the loss associated with downstream decisions. However, existing studies focus solely on scenarios where a fixed batch of data is available and the objective function does not change over time. We instead investigate DFL in dynamic environments where the objective function and data distribution evolve over time. This setting is challenging for online learning because the objective function has zero or undefined gradients, which prevents the use of standard first-order optimization methods, and is generally non-convex. To address these difficulties, we (i) regularize the objective to make it differentiable and (ii) use perturbation techniques along with a near-optimal oracle to overcome non-convexity. Combining those techniques yields two original online algorithms tailored for DFL, for which we establish respectively static and dynamic regret bounds. These are the first provable guarantees for the online decision-focused problem. Finally, we showcase the effectiveness of our algorithms on a knapsack experiment, where they outperform two standard benchmarks.

1 INTRODUCTION.

Many real-world decision problems involve uncertainty, which is commonly handled through the predict-then-optimize framework (Bertsimas and Kallus, 2020). First, a prediction model is trained on historical data; then, its output is fed into an optimization problem to guide decision-making. This natural strategy has been successfully applied in many operation research (OR) problems, ranging from supply chain management (Acimovic and Graves, 2015; Fisher et al., 2016; Ban and Rudin, 2019; Bertsimas and Kallus, 2020), revenue management (Farias et al., 2013; Ferreira et al., 2016; Cohen et al., 2017; Chen et al., 2022) to healthcare operation (Bertsimas et al., 2013; Aswani et al., 2019; Gupta et al., 2020; Rath et al., 2017), see Mivsić and Perakis (2020) for an extensive review. It is clear that this approach would yield optimal decisions if the predictions were perfectly accurate. However, in practice, prediction errors are inevitable, and even small inaccuracies can propagate through the optimization process, potentially leading to poor decisions.

To address this limitation, an approach known as *decision-focused learning* (Mandi et al., 2024), also referred to as *smart predict-then-optimize* (Elmachtoub and Grigas, 2022) or *integrated learning-optimization* (Sadana et al., 2025), has emerged. Instead of optimizing for prediction accuracy alone, this method trains the predictive model to directly minimize the downstream decision loss. By aligning the learning objective with the decision-making goal, it produces models that are more robust to prediction errors in practical applications. While decision-focused learning yields strong empirical performance (Donti et al., 2017; Verma et al., 2022; 2023; Wang et al., 2023), theoretical development has so far been limited to the batch setting, where models are trained on pre-collected,

*Centre de Mathématiques Appliquées - CNRS, École Polytechnique, Paris, France

†INRIA - CNRS, Ecole Normale Supérieure, PSL Research University, France

‡EPITA, Laboratoire de Recherche de l'EPITA, Mohamed Bin Zayed University of AI, Paris, France

§Inria Paris, Ecole Normale Supérieure PSL, UC Berkeley, Paris, France

¶Inria Saclay, Université Paris Saclay, LMO, Orsay, France

independently and identically distributed (i.i.d.) data (Wilder et al., 2019; Mandi et al., 2022; Shah et al., 2022; Schutte et al., 2024). This assumption breaks down in many real-world scenarios involving dynamic environments (Cheung et al., 2019; Padakandla et al., 2020) and shifting data distributions (Lu et al., 2018; Quiñero-Candela et al., 2022).

Online learning (Cesa-Bianchi and Lugosi, 2006; Hazan, 2023) provides a general way to cope with such non-stationarity of data-generating processes. This framework considers a learner who makes decisions sequentially, at each round leveraging data collected from previous rounds to inform its next decision. Crucially, the objective function is allowed to vary over time, either in a stochastic or adversarial way. This provides a natural framework for the work that we present here, which extends decision-focused learning beyond the i.i.d., batch setting.

Contributions. We develop a theoretical foundation for online decision-focused learning, enabling its application in non-stationary settings. This presents a significant technical challenge, as the inherent difficulties of decision-focused learning, such as the non-differentiability of the objective function or the lack of convexity of the losses due to the bi-level nature of the problem, compound those already present in online learning. Our contributions are as follows:

(i) We formalize the online decision-focused learning problem by assuming that at each round, a decision-maker seeks to solve a linear optimization problem over a polytope but does not have access to the true cost function. Thus the decision-maker has to predict the cost using whatever partial information that it has at hand. The cost function is then revealed, and the decision-maker updates its model in a decision-focused fashion. This results in a bi-level optimization problem, where the inner problem consists in making a decision and the outer problem involves optimizing the resulting decision cost.

(ii) We present two algorithms to tackle this problem, *Decision-Focused Follow-the-Perturbed-Leader* (DF-FTPL) and *Decision-Focused Online Gradient Descent* (DF-OGD). While both rely on regularizing the inner problem to make the resulting decision differentiable, they differ on the way they update the parameters of the prediction model. DF-FTPL uses the FTPL approach (Hutter and Poland, 2005), and DF-OGD leverages a variant of Online Gradient Descent (Zinkevich, 2003). We establish sublinear convergence guarantees for both procedures, in the form of a static regret bound for the former and a dynamic regret bound for the latter. To our knowledge, these are the first provable guarantees for the online decision-focused learning problem.

(iii) Finally, we assess the performance of our algorithms on a knapsack experiment inspired by Mandi et al. (2024). Our simulations demonstrate that our approach outperforms the online version of two popular baselines, namely prediction-focused learning and Smart-Predict-then-Optimize, in both static and dynamic environments.

Additional related work. Several decision-focused approaches have been proposed in the batch setting, among which differentiating the associated KKT conditions Gould et al. (2016); Amos and Kolter (2017); Donti et al. (2017); Wilder et al. (2019); Mandi and Guns (2020), smoothing via random perturbation Berthet et al. (2020), building surrogate losses via duality Elmachtoub and Grigas (2022) or directional gradients Huang and Gupta (2024) and relying on pairwise ranking techniques Mandi et al. (2022), see Mandi et al. (2024); Sadana et al. (2025) for additional references. However, the extension of decision-focused learning to the online setting remains unexplored. While several recent studies address online bi-level optimization Shen et al. (2023); Tarzanagh et al. (2024); Lin et al. (2023), none of them are applicable to decision-focused learning as they rely on restrictive smoothness assumptions, that are incompatible with the structure of decision-focused problems. In particular, the decision-focused objective is typically non-convex and features gradients that are either zero or undefined, owing to its underlying linear structure.

This motivates the use of different online methods, specifically tailored to address these problems. On the one hand, lack of differentiability is usually tackled through zero-th order methods Héliou et al. (2020); Frezat et al. (2023), sub-gradient Duchi et al. (2011), proximal Dixit et al. (2019) or smoothing approaches Abernethy et al. (2014). On the other hand, methods to address non-convexity in online learning often rely on the existence of a near optimal oracle Kalai and Vempala (2005); Agarwal et al. (2019); Suggala and Netrapalli (2020); Xu and Zhang (2024) or additional smoothness conditions Lesage-Landry et al. (2020); Ghai et al. (2022). All those results are derived for *static regret* (Zinkevich, 2003) which compares the learned predictors to the best static strategy

and is defined in Section 2. A more challenging criterion, also described in Section 2 is *dynamic regret*, introduced in Zinkevich (2003) and later developed in Hall and Willett (2013); Besbes et al. (2015); Zhao et al.; Zhao and Zhang (2021) among others, which takes into account the evolution of the environment. An alternative approach to obtain guarantees in the online non-convex setting is to weaken static regret to *local regret* Hazan et al. (2017); Aydore et al. (2019); Zhuang et al. (2020); Hallak et al. (2021), which is the sum of the objective gradient norms evaluated in the iterates over time. Minimizing local regret corresponds to encouraging convergence toward stationary points. However as we shall see, this approach is not sensible in our framework, as objective gradients are either zero or undefined.

Outline. In Section 2, we introduce the online decision-focused problem, our notions of regret and our assumptions. In Section 3, we present the DF-FTPL and DF-OGD algorithms, before deriving bounds on the static regret of the former and the dynamic regret of the latter. Finally, we present our experiment in Section 4.

Notation. For a differentiable map $\varphi : \mathbb{R}^m \rightarrow \mathcal{Y}$ such that $\mathcal{Y} \subseteq \mathbb{R}$, $\nabla\varphi(x)$ denotes the gradient of φ at $x \in \mathbb{R}^m$ and $\nabla^2\varphi(x)$ denotes the Hessian of φ . In the case where $\mathcal{Y} \subseteq \mathbb{R}^d$, $\nabla\varphi(x)$ denotes the Jacobian of φ . For two vectors $(v, w) \in \mathbb{R}^d \times \mathbb{R}^d$, $v \succcurlyeq w$ means that $v_i \geq w_i$ for any $i \in [d]$, and $\langle v, w \rangle = v^\top w$ refers to the standard Euclidian inner product. Also, $\|v\| = \sqrt{\langle v, v \rangle}$ is the standard euclidian norm. Given a compact convex set $\Theta \subseteq \mathbb{R}^d$, Π_Θ denotes the orthogonal projection onto Θ . For a matrix $M \in \mathbb{R}^{m \times d}$, $\|M\|_{\text{op}}$ refers to its L^2 -operator norm. In the case $d = m$, $\lambda_{\min}(M)$ refers to its lowest real eigenvalue and $\lambda_{\max}(M)$ it largest real eigenvalue. For $(x, y) \in \mathbb{R}^2$, $x \propto y$ means that there exists $\lambda \in \mathbb{R}^*$ such that $x = \lambda y$.

2 FRAMEWORK

Sequential decision-making. We consider an online decision-making problem over $T > 0$ periods, defined for any $t \in [T]$ as

$$\min_{w \in \mathcal{W}} \langle \bar{g}_t(X_t), w \rangle, \quad (1)$$

where $\mathcal{W} = \text{Conv}(v_1, \dots, v_K)$ is a bounded convex polytope of \mathbb{R}^d with non-empty interior and vertices $\{v_i\}_{i=1}^K$. This feasible set appears naturally in many problems such as shortest-path (Gallo and Pallottino, 1988), portfolio selection (Li and Hoi, 2014) or mixed strategy design in games (Syrkkanis et al., 2015). In (1), X_t are random covariates and $\bar{g}_t : \mathcal{X} \rightarrow \mathbb{R}^d$, is a deterministic cost function, which satisfies for instance $\bar{g}_t(X_t) = \mathbb{E}[Z_t \mid X_t]$ for some hidden state $Z_t \in \mathbb{R}^d$. At each period t , nature picks both a distribution for X_t and a cost function g_t . This situation corresponds to the stochastic adversary setting (Rakhlin et al., 2011). Importantly, X_t is revealed at the beginning of the round, but not $\bar{g}_t(X_t)$ which is only available at the end of the round. While $\bar{g}_t(X_t)$ is unknown to the decision-maker at the decision time, they have access to a family of models $g : \Theta \times \mathcal{X} \rightarrow \mathbb{R}^d$, parameterized by $\Theta \subset \mathbb{R}^m$, to predict it. The general form of the decision-making dynamics we consider can be described as follows: at each round $t \in [T]$, for a horizon $T \in \mathbb{N}$, given the current parameter θ_t ,

1. Nature picks a distribution for $X_t \in \mathcal{X}$ and a cost function $\bar{g}_t : \mathcal{X} \rightarrow \mathbb{R}^d$.
2. The decision-maker observes X_t and compute its prediction as $g(\theta_t, X_t)$.
3. Then, they take an action minimizing the resulting predicted cost

$$w_t = w_t^*(\theta_t) \in \underset{w \in \mathcal{W}}{\text{argmin}} \langle g(\theta_t, X_t), w \rangle \quad (2)$$

4. Finally, the decision-maker observes $\bar{g}_t(X_t)$ and update θ_{t+1} for the next round.

Formally, the decision-maker considers a joint process $(\theta_t, w_t)_{t \in [T]}$, starting from $\theta_1 \in \Theta$ and a history $\mathcal{H}_0 = \emptyset$, and defined by the following recursion. At round $t \in [T]$, the decision-maker takes the best action given the current estimate θ_t and a new feature X_t as in (2). Then, they observe $\bar{g}_t(X_t)$ and update their history $\mathcal{H}_t = \mathcal{H}_{t-1} \cup \{(X_t, \bar{g}_t(X_t), \theta_t, w_t^*(\theta_t))\}$. Finally, they update their prediction parameter θ_{t+1} based on \mathcal{H}_t through an algorithm Alg_t . A central question for the decision-maker is then the choice of algorithms $\{\text{Alg}_t\}_{t \in [T]}$ to fit their regression model, that is how to pick $\theta_t \in \Theta$ for each $t \in [T]$.

Online decision-focused learning. From an online perspective, *prediction-focused learning* consists in selecting at each $t \in [T]$ an algorithm Alg_t to estimate $\text{argmin}_{\theta \in \Theta} R_t(\theta)$, where R_t is a statistical risk based on the historical observations \mathcal{H}_t , typically chosen as the empirical risk $R_t(\theta) = \sum_{i=1}^t \ell(\bar{g}_i(X_i), g(\theta, X_i))$, for some loss function ℓ (e.g., cross-entropy or squared error).

Alternatively, we consider in this paper, the *decision-focused learning* approach, which selects parameters $\theta \in \Theta$ by directly minimizing the downstream objective instead of a risk function. Specifically, the decision-maker chooses Alg_t to solve $\text{argmin}_{\theta \in \Theta} \langle \bar{g}_t(X_t), w_t^*(\theta) \rangle$ where w_t^* is defined in (2). From the previous formulation, we see that the decision-focused learning formulation corresponds to a bilevel optimization problem (Colson et al., 2007; Sinha et al., 2017; Ji et al., 2021).

We emphasize that no stationary assumptions are made on the process $\{X_t, \bar{g}_t(X_t)\}_{t \in [T]}$. In other words, the adversary may select any distribution for X_t and any function \bar{g}_t (as long as **H2** below is satisfied). To fully accommodate this flexibility, we assess the optimality of $\{\theta_t\}_{t \in [T]}$ with two notions of regret. Let

$$f_t : \theta \in \Theta \mapsto \langle \bar{g}_t(X_t), w_t^*(\theta) \rangle, \quad (3)$$

denotes the loss incurred when taking an action based on the prediction parameter $\theta_t \in \Theta$. Following Zinkevich (2003), we consider the notions of static and dynamic regret to measure the effectiveness of a learning strategy $\{\theta_t\}_{t \in [T]}$, which are respectively:

$$\mathfrak{R}_T^s = \sum_{t \in [T]} f_t(\theta_t) - \inf_{\theta \in \Theta} \sum_{t \in [T]} \mathbb{E}[f_t(\theta)] \quad \text{and} \quad \mathfrak{R}_T^d = \sum_{t \in [T]} F_t(\theta_t) - \sum_{t \in [T]} \inf_{\theta \in \Theta} F_t(\theta), \quad (4)$$

where $F_t : \theta \mapsto \mathbb{E}[f_t(\theta) \mid \mathcal{H}_{t-1}]$. In our dynamic regret, the sequence of actions is compared against a sequence of oracles, each minimizing the instantaneous loss. Without taking the conditional expectation over \mathcal{H}_{t-1} , each comparator could overfit to the specific realization of X_t resulting in an unrealistically strong benchmark. By considering the conditional expectation of the loss, we effectively regularize the dynamic comparators, making them meaningful competitors. Note that our static regret compare to the best fixed strategy, with respect to the averaged losses. This notion makes sense here as we aim to control those regrets in expectation over the randomness of the process.

We make the following mild regularity assumptions for the rest of the analysis.

H1. (i) $\Theta \subset \mathbb{R}^m$ is a compact and convex set with diameter $D_\Theta < \infty$. (ii) for any $t \in [T]$, $\theta \mapsto g(\theta, X_t) \in \mathbb{R}^{m \times d}$ is continuously differentiable almost surely and $\|\nabla_\theta g(\theta, X_t)\|_{\text{op}} \leq G < \infty$ for any $\theta \in \Theta$ almost-surely. (iii) For any $t \in [T]$, $\|\bar{g}_t(X_t)\| \leq D_Z < \infty$ almost-surely.

It is common in online learning to assume boundedness of the parameter space, model gradient and prediction space (Boyd, 2004; Bishop and Nasrabadi, 2006). Note that in the well-specified setting, i.e., $\bar{g}_t = g(\bar{\theta}_t, \cdot)$ for some $\bar{\theta}_t$, these two former conditions automatically imply the latter. We emphasize however that we do not assume this realizability condition, in contrast to most of the literature on decision-focused learning (Bennouna et al., 2024).

Without further assumptions on the sequence of costs \bar{g}_t and models g , the problem can still be made arbitrarily hard. Therefore, we make an assumption coming from the classification (Mammen and Tsybakov, 1999; Tsybakov, 2004) and bandit literature (Zeevi and Goldenshluger, 2009; Perchet and Rigollet, 2013) about the *margins* of the cost function. Recall that we denote by $\{v_i\}_{i=1}^K$ the vertices of \mathcal{W} . We define for any $x \in \mathcal{X}$ and $i \in [K]$, $u_i(\theta, x) = \langle g(\theta, x), v_i \rangle$.

H2. There exist $C_0 > 0$ and $\beta \in (0, 1]$, such that almost surely for any $t \in [T]$, $\theta \in \Theta$ and $\varepsilon \in [0, 1]$,

$$\mathbb{P} \left(\inf_{j \neq I_t(\theta)} \{u_j(\theta, X_t) - u_{I_t(\theta)}(\theta, X_t)\} \geq \varepsilon \mid \mathcal{H}_{t-1} \right) \geq 1 - C_0 \varepsilon^\beta,$$

where $I_t(\theta) \in \text{argmin}_{i \in [K]} u_i(\theta, X_t)$.

H2 controls how difficult it is to solve problem (2) since it determines the objective gap between the optimal vertex $v_{I_t(\theta)} \in \mathcal{W}$ and the other vertices. In other words, it quantifies how identifiable is the optimal vertex, $\varepsilon = 0$ meaning it is not distinguishable from the others. If it is satisfied for $\beta \geq 1$, then it is automatically satisfied for $\beta = 1$ thus we do not loose any generality restraining β to $[0, 1]$. This assumption is critical in our analysis, as it allows bounding the expected distance between the actual optimal decision $w_t^*(\theta)$ and the regularized approximation $\tilde{w}_t(\theta)$ introduced in (5) below.

We expect **H2** to hold in a wide variety of classical statistical settings. For example, in Appendix A we show that it is satisfied when $g : (\theta, X) \mapsto X\theta$ where X has i.i.d standard Gaussian columns.

3 ONLINE ALGORITHMS FOR DECISION-FOCUSED LEARNING .

On the need of regularization to get differentiation. From an online learning perspective, a natural approach to minimize the two regrets defined in (4) would be to update θ_{t+1} at each round t through an algorithm Alg_t being either a variant of either the Follow-The-Leader algorithm (FTL, Kalai and Vempala, 2005) or Online Gradient Descent (OGD, Zinkevich, 2003), applied to the objective function f_t defined in (3). However, these algorithms are designed for single-level optimization problems and require access to (sub-)gradients of the objective. In our setting, this requirement is problematic because the function f_t does not admit an informative gradient: by construction, computing the gradient of f_t involves differentiating the mapping $\theta \mapsto w_t^*(\theta)$, which is generally not (sub-)differentiable. Indeed, w_t^* minimizes a linear function over the convex polytope \mathcal{W} , thus $\theta \mapsto w_t^*(\theta)$ ranges in the finite set of vertices $\{v_1, \dots, v_K\}$.

To address this issue, following Wilder et al. (2019), we propose to add a regularizer \mathcal{R} to the objective function in (2). Accordingly, we define for any $\theta \in \Theta$ and some $\alpha_t > 0$

$$\tilde{w}_t(\theta) \in \underset{w \in \mathcal{W}}{\operatorname{argmin}} \{ \langle g(\theta, X_t), w \rangle + \alpha_t \mathcal{R}(w) \}, \quad (5)$$

which is a regularized approximation of w_t^* . When this surrogate is continuously differentiable, which is the case for our choices of \mathcal{R} in Section 3, we can use $\nabla \tilde{w}_t(\theta)$ in our algorithmic routine. This approach amounts to minimizing the surrogate

$$\tilde{f}_t : \theta \mapsto \langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle, \quad (6)$$

instead of f_t , which admits gradients of the form $\nabla \tilde{f}_t(\theta) = \nabla \tilde{w}_t(\theta)^\top \bar{g}_t(X_t)$. Note that there is a natural trade-off in the choice of the regularization parameter α_t . On the one hand, choosing a large α_t makes the function $\tilde{w}_t(\theta)$ smoother. On the other hand, the larger α_t is, the more \tilde{w}_t deviates from the true function w_t^* that we aim to approximate. As we will see later, it is possible to balance these two extremes by carefully tuning the parameter α_t .

On the choice of regularization on a general polytope. In what follows, we write \mathcal{W} as the intersection of n half-spaces, that is $\mathcal{W} = \{w \in \mathbb{R}^d : A^\top w - b \preceq 0\}$ where $n > 0$ is the number of faces, $A \in \mathbb{R}^{d \times n}$ and $b \in \mathbb{R}^n$. We assume that \mathcal{W} is not degenerated, i.e. $AA^\top \in \mathbb{R}^{d \times d}$ is full-rank.

It remains to determine what regularizer \mathcal{R} to choose in (6). We recall that our aim is to obtain a $\tilde{w}_t(\theta)$ in (5) which is differentiable for any $\theta \in \Theta$. A possible strategy is to choose \mathcal{R} so \tilde{w}_t remains in the strict interior of \mathcal{W} . In this case, $\tilde{w}_t(\theta)$ is differentiable in a neighborhood of any $\theta \in \Theta$ and admits a close-form Jacobian $\nabla \tilde{w}_t(\theta)$ by the implicit function theorem. A natural choice to force \tilde{w}_t to remain in the interior of \mathcal{W} is the corresponding log-barrier function

$$\mathcal{R} : w \mapsto - \sum_{i=1}^n \ln(b_i - A_i^\top w), \quad (7)$$

where $A_i \in \mathbb{R}^d$ is the i -th column of A . With this choice of regularization, we show in Lemma 4 that $\nabla \tilde{w}_t$ has an explicit formulation, allowing its use in practice.

Remark 1. We remark that in the special case where $\mathcal{W} = \{w \in \mathbb{R}^d : w \succcurlyeq 0, \mathbf{1}^\top w = 1\}$ is the simplex of \mathbb{R}^d , an alternative choice for \mathcal{R} in (5) is the negative entropy $\mathcal{R}_0 : w \mapsto \sum_{i \in [d]} w_i \ln(w_i)$. In this case, $\tilde{w}_t(\theta)$ in (5) reduces to the softmax mapping, which is differentiable and admits a close-form Jacobian. In Appendix C, we theoretically study the performances of our algorithms in this special case.

Approximate oracles to handle non-convexity. The bi-level structure of the problem yields unexpected properties. In particular, even when the regularizer \mathcal{R} is strongly-convex in (26), which the case with both log-barriers and negative entropy, and the model g is simple (such as $g(\theta, X) = X\theta$), $\tilde{f}_t : \theta \mapsto \langle \tilde{w}_t(\theta), \bar{g}_t(X_t) \rangle$ is lipschitz (see Lemmas 1 and 5) but not convex in general. This prevents us from directly using known online convex optimization algorithms (Hazan, 2023). However, a recent line of research (Agarwal et al., 2019; Suggala and Netrapalli, 2020; Xu and Zhang,

2024) has developed online algorithms in the non-convex case, provided the losses are Lipschitz-continuous. These studies combine near-optimal oracles with perturbation techniques to establish sublinear bounds on the expected regret. In line with this literature, we also assume to have access to an *approximate offline optimization oracle*, a notion that appeared in a slightly modified form in Suggala and Netrapalli (2020).

Definition 1. An ξ -approximate offline optimization oracle (or approximate oracle), adapted to a class \mathcal{C} is a mapping \mathbf{O}_ξ taking a function $f \in \mathcal{C}$ such that $f : \Theta \rightarrow \mathbb{R}$, and outputting $\vartheta = \mathbf{O}_\xi(f) \in \Theta$ satisfying:

$$f(\vartheta) \leq \inf_{\theta \in \Theta} f(\theta) + \xi. \quad (8)$$

This notion of an approximate oracle reflects the fact that, in non-convex settings, we cannot rely on subroutines that provably find a global minimizer of \tilde{f}_t at time t (as is possible in convex problems with offline gradient descent). Instead, we must be content with local minimizers. Their quality is measured by a parameter ξ , which becomes small in favorable loss landscapes. When \mathcal{C} contains differentiable and Lipschitz functions, the stochastic gradient descent (SGD) algorithm can be used for the oracle \mathbf{O}_ξ . Indeed, a large body of work has shown that, even with models as complicated as deep neural networks, SGD can converge to local minimizers (Ghadimi and Lan, 2013; Mertikopoulos et al., 2020; Patel and Zhang, 2021; Cutkosky et al., 2023), justifying its use as an approximate oracle. A more detailed discussion is provided in Appendix B.

We conclude this remark by mentioning that it is possible to conduct a non-convex analysis without an approximate oracle by considering the weaker notion of *local regret* defined as $\sum_{t \in [T]} \|\nabla F_t(\theta_t)\|$; see (Hazan et al., 2017; Aydore et al., 2019; Zhuang et al., 2020; Hallak et al., 2021). However, such a definition is not meaningful in the DFL framework, where gradients are zero or undefined due to the structure of the problem.

The combined use of approximate oracles and regularization allows us to derive original online algorithms. In Section 3.1, we focus on a variant of the FTL algorithm, enjoying a static regret bound, while in Section 3.2 we propose a version of OGD enjoying a dynamic regret bound, both results holding in the non-convex case.

3.1 STATIC REGRET: DECISION-FOCUSED FOLLOW THE PERTURBED LEADER

Our first algorithm is displayed in Algorithm 1, where $\text{Exp}(\eta)$ refers to an exponential distribution with parameter $\eta > 0$. It is inspired from the classic Follow-the-Leader approach (Kalai and Vempala, 2005), which consists in making a decision minimizing the sum of objective functions observed so far. When losses are non-convex, it is common to inject random noise at each step to regularize the total cost function, resulting in an approach known as Follow-the-Perturbed-Leader (Hutter and Poland, 2005; Suggala and Netrapalli, 2020). We employ this strategy as Alg_t to update θ_{t+1} in Algorithm 1. Note that the oracle in line 6 of Algorithm 1 crucially works with the *regularized* losses $\tilde{f}_1, \dots, \tilde{f}_t$. This allows to use in practice gradient-based methods to obtain an approximate minimizer as per our discussion in Section 3, since these surrogate losses are differentiable and Lipschitz.

Algorithm 1 Decision-Focused Follow The Perturbed Leader DF-FTPL

- 1: **Input:** horizon $T > 0$, initialization $\theta_1 \in \Theta$, ξ -approximate oracle \mathbf{O}_ξ and history $\mathcal{H}_0 = \emptyset$.
- 2: **for** each $t \in \{1, \dots, T\}$ **do**
- 3: Observe $X_t \in \mathcal{X}$ and play $w_t^*(\theta_t) = \operatorname{argmin}_{w \in \mathcal{W}} \langle g(\theta_t, X_t), w \rangle$.
- 4: Observe $\tilde{g}_t(X_t) \in \mathcal{Z}$ and update the history \mathcal{H}_t .
- 5: Draw $\sigma_t \in \mathbb{R}^d$ such that for all $j \in [d]$, the j -th component $\sigma_{j,t} \sim \text{Exp}(\eta)$.
- 6: Update

$$\theta_{t+1} = \mathbf{O}_\xi \left(\sum_{i=1}^t \tilde{f}_i - \langle \sigma_t, \cdot \rangle \right)$$

- 7: **end for**
-

Note that in line 3, the algorithm takes the best action $w_t^*(\theta_t)$ given the predicted cost $g(\theta_t, X_t)$ (the regularized action \tilde{w}_t is only needed to compute \tilde{f}_t). In most practical settings, w_t^* can be computed

efficiently with standard numerical solvers. However, we show in Appendix F that only being able to determine $\underline{w}_t(\theta_t) \in \mathcal{W}$ such that $\langle g(\theta_t, X_t), \underline{w}_t(\theta_t) \rangle - \langle g(\theta_t, X_t), w_t^*(\theta_t) \rangle \leq \kappa$ for some $\kappa > 0$ only shifts the regret bounds of the next section by κ .

We provide a theoretical guarantee on the convergence of Algorithm 1 with the following static regret bound.

Theorem 1. *Assume H1, H2 and having access to an ξ -approximate oracle \mathbf{O}_ξ adapted to $\{\sum_{i=1}^t \tilde{f}_i - \langle \sigma_t, \cdot \rangle\}_{t \in [T]}$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-FTPL (Algorithm 1) instantiated with learning step $\eta > 0$ and regularization coefficients $\alpha_t = \alpha > 0$ for any t . Then:*

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^s] = \tilde{\mathcal{O}}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha n\right),$$

where \mathbb{E} denotes the expectation on both data and the intrinsic randomness of DF-FTPL and $\tilde{\mathcal{O}}$ contains polynomial dependency in $\ln(1/\alpha), \ln(\ln(d))$.

Furthermore, taking $\eta \propto m^{1/4} T^{-3/4} n^{-1/2}$ and $\alpha \propto m^{3/4} n^{1/2} T^{-1/4}$ yields:

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^s] = \tilde{\mathcal{O}}\left(m^{3/4} n^{3/2} T^{-1/4} + \xi\right).$$

The proof of Theorem 1 can be found in Appendix G.2. In particular, Theorem 1 shows that DF-FTPL enjoys an average regret bound decaying in $T^{-1/4}$ as long as $\xi = \mathcal{O}(T^{-1/4})$. While it features a polynomial dependency $m^{3/4}$ on the dimension of Θ , it only depends on the dimension of the decision space \mathcal{W} through a $\ln \ln(d)$ term. This makes Algorithm 1 a particularly competitive approach when the decision space is of high dimension.

It is informative to compare our guarantee to existing bounds in the literature. Suggala and Netrapalli (2020), who also study non-convex online learning, achieves a rate of $\mathcal{O}(T^{-1/2})$ as long as their offline oracle satisfies $\xi = \mathcal{O}(T^{-1/2})$, at the cost of a degraded $m^{3/2}$ dependency on the dimension. Their faster rate in T comes from the fact that they tackle a simpler, single-level problem as compared to our bi-level, non differentiable setting. Indeed, we need to regularize the inner problem to overcome non-differentiability as discussed in Section 3. This introduces an additional trade-off on the regularization strength α on top of the usual trade-off in the learning rate η , which is reflected in our rate.

3.2 DECISION-FOCUSED ONLINE GRADIENT DESCENT

While DF-FTPL (Algorithm 1) benefits from a converging static regret bound, the techniques involved in Suggala and Netrapalli (2020) are not enough to reach dynamic regret guarantees, which is particularly relevant in highly non-stationary environments where the optimal decisions may change significantly from one round to another. To this end, we go beyond the Follow-the-Leader approach and propose an original algorithm based on the celebrated Online Gradient Descent (Zinkevich, 2003). This procedure, which we call DF-OGD, is presented in Algorithm 2.

In Algorithm 2, at each time $t \in [T]$, a decision $w_t^*(\theta_t) \in \mathcal{W}$ is made based on the current parameter $\theta_t \in \Theta$. The feedback $\bar{g}_t(X_t) \in \mathcal{Z}$ is then observed and used to define the surrogate objective \tilde{f}_t . Then, a near-minimizer of this objective is computed using the offline oracle from Definition 1. To address non-convexity, the gradient $\nabla \tilde{w}_t(u_t)$ is evaluated at a point u_t drawn uniformly at random from $[\theta_t, \vartheta_t]$. Finally, the parameter is updated to θ_{t+1} using the standard gradient step of OGD.

The main difference with Algorithm 1 is the update of θ_{t+1} through Alg_t . On the one hand, DF-FTPL invokes \mathbf{O}_ξ to minimize the cumulative loss observed so far and perturbs the entire objective function. On the other hand, DF-OGD relies only on a near-optimizer of the most recent regularized cost, $\vartheta_t = \mathbf{O}_\xi(\tilde{f}_t)$, and perturbs the point at which the descent direction (gradient) is evaluated.

Moreover, note that Algorithm 1 is instantiated with a single pair of parameters (α, η) whereas Algorithm 2 uses on a sequence $(\alpha_t, \eta_t)_{t \in [T]}$. This additional flexibility is crucial for the algorithm to adapt to the variation of the problem so as to maintain a low dynamic regret.

We provide a convergence guarantee for Algorithm 2 with the following dynamic regret bound.

Theorem 2. *Assume H1, H2, access to a ξ -approximate oracle adapted to $\{\tilde{f}_t\}_{t \in [T]}$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-OGD (Algorithm 2) instantiated with the non-increasing sequence $(\eta_t)_{t \in [T]}$ and*

Algorithm 2 Decision-Focused Online Gradient Descent (DF-OGD)

-
- 1: **Input:** horizon $T > 0$, initialization $\theta_1 \in \Theta$ and history $\mathcal{H}_0 = \emptyset$.
 - 2: **for** each $t \in \{1, \dots, T\}$ **do**
 - 3: Observe $X_t \in \mathcal{X}$ and play $w_t^*(\theta_t) = \operatorname{argmin}_{w \in \mathcal{W}} \langle g(\theta_t, X_t), w \rangle$.
 - 4: Observe $\bar{g}_t(X_t) \in \mathcal{Z}$ and update the history \mathcal{H}_t .
 - 5: Get $\vartheta_t = \mathbf{O}_\xi(\tilde{f}_t)$.
 - 6: Draw $\delta_t \sim \operatorname{Unif}([0, 1])$, compute $u_t = \vartheta_t + \delta_t(\theta_t - \vartheta_t)$ and $\tilde{\nabla}_t(u_t) = \nabla \tilde{w}_t(u_t)^\top \bar{g}_t(X_t)$.
 - 7: Update $\theta_{t+1} = \Pi_\Theta(\theta_t - \eta_t \tilde{\nabla}_t(u_t))$.
 - 8: **end for**
-

regularization coefficients $(\alpha_t)_{t \in [T]}$. Then:

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^d] = \tilde{\mathcal{O}} \left(\mathbb{E} \left[\frac{1 + P_T}{T\eta_T} + \frac{1}{T} \sum_{t \in [T]} \frac{\eta_t}{\alpha_t^2} + n\alpha_t \right] + \xi \right).$$

where $P_T = \sum_{t=1}^{T-1} \|\vartheta_{t+1} - \vartheta_t\|$, \mathbb{E} denotes the expectation on both data and the intrinsic randomness of DF-OGD and $\tilde{\mathcal{O}}$ contains polynomial dependency in $\ln(1/\alpha)$, $\ln(\ln(d))$.

Furthermore, assume $(t^{-1}(1 + P_t))_{t \geq 1}$ is non-increasing almost surely with $P_t = \sum_{s=1}^{t-1} \|\vartheta_{s+1} - \vartheta_s\|$. Then, using $\alpha_t \propto n^{-1/2} t^{-1/4} (1 + P_t)^{1/4}$ and $\eta_t \propto n^{-1/2} t^{-3/4} (1 + P_t)^{3/4}$ for any $t \in [T]$ leads to:

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^d] = \tilde{\mathcal{O}} \left(\mathbb{E} \left[\sqrt{n} (1 + P_T)^{1/4} T^{-1/4} \right] + \xi \right).$$

Our dynamic regret bound naturally depends on P_T , which captures the problem’s variability by measuring the cumulative distance between approximate minimizers over time. When $\xi = \mathcal{O}((1 + P_T)^{1/4} T^{-1/4})$, the average dynamic regret decreases at the rate $\mathcal{O}((1 + P_T)^{1/4} T^{-1/4})$. Notably, the bound is independent of the dimension of Θ , and depends only mildly on the dimension of \mathcal{W} through a $\ln \ln(d)$ factor. This constitutes a key strength of the optimistic strategy underlying DF-OGD, making it particularly well-suited for high-dimensional settings.

It is instructive to compare our guarantee with those established in recent studies on dynamic regret. For instance, [Zhang et al. \(2018\)](#) derive a $\mathcal{O}(T^{-1/2} \sqrt{T(P_T + 1)})$ bound. However, their setting is not directly comparable to ours, as they consider a simpler single-level problem with differentiable and convex objectives. In contrast, our framework involves non-convex, non-differentiable losses due to the bi-level nature of decision-focused learning. More recently, [Huang and Wang \(2025\)](#) obtained a $\mathcal{O}((1 + P_T^\infty)^{1/3} T^{-1/3})$ bound, where $P_T^\infty := \sum_{t=1}^T \|f_{t+1} - f_t\|_\infty$. Yet, this rate is achieved under substantially more favorable conditions: a single-level problem with losses that are strongly convex or Lipschitz, an additional assumption of “quasi-stationary” and a more challenging path involving the full landscapes of the f_t s. By comparison, our losses are neither Lipschitz (indeed, they may even be discontinuous due to the linearity of the lower-level problem) nor required to be stationary over time—the only restriction being **H2** to hold. We view the ability of our algorithm to achieve efficient convergence despite these demanding conditions as a core contribution of our work.

4 EXPERIMENTS

In this section, we compare the performances of our algorithms DF-FTPL and DF-OGD to two important benchmarks, namely prediction-focused learning and SPO ([Elmachtoub and Grigas, 2022](#)).

Setting. Our experimental setup is inspired by the knapsack example from [Mandi et al. \(2024\)](#). More precisely, we consider a decision maker who must pick at each $t \in [T]$ an object $v_t \in \mathcal{V}$ among K items denoted $\mathcal{V} = \{1, 2, \dots, K\}$ with respective costs $\bar{g}_t(X) = (\bar{g}_{t,1}(X), \dots, \bar{g}_{t,K}(X)) \in [0, 1]^K$ depending on some covariates $X \in \mathcal{X}$. At the beginning of each period, the decision-maker only observe covariates $X_t \in \mathcal{X} \subset \mathbb{R}^p$, and have at their disposal a parametric model $g :$

$\Theta \times \mathcal{X} \rightarrow [0, 1]^K$ to predict \bar{g}_t . Given their current parameter $\theta_t \in \Theta$, they predict item costs as $g(\theta_t, X_t) = (g_1(\theta_t, X_t), \dots, g_K(\theta_t, X_t))$ and pick an item $v_t^*(\theta_t) = \operatorname{argmin}_{i \in [K]} g_i(\theta_t, X_t)$. This setting is depicted in Figure 2 for $K = 2$ in Appendix E. After having made their decision, the decision-maker observes the true item costs $\bar{g}_t(X_t)$, and update θ_{t+1} for the next round based on this feedback. Note that this setting can directly be mapped in the simplex example of Remark 1, since $v_t^*(\theta_t) = \operatorname{argmin}_{w \in \mathcal{W}_0} \langle w, g(\theta_t, X_t) \rangle$, \mathcal{W}_0 being the simplex of \mathbb{R}^K . For this reason, we instantiate our algorithms DF-OGD and DF-FTPL with the negative entropy regularizer, following Remark 1.

Synthetic data. We instantiate the previous problem with the following synthetic data. For any $t \in [T]$, we draw $X_t \in \mathbb{R}^{K \times p}$ with correlated rows, and generate a cost vector $c_t(X_t) \in \mathbb{R}^k$ as :

$$c_t(X_t) = A \sin^4((2X_t \theta_t^*)^{-1}) + \varepsilon_t, \quad (9)$$

where $A > 0$, $\varepsilon_t \sim \mathcal{N}(0, I_K)$ is a Gaussian noise and $\theta_t^* \in \mathbb{R}^p$ is a parameter which satisfies $\theta_t^* = 1/2 \theta^* + 1/2 \zeta_t$, where $\zeta_t \sim \mathcal{N}(0, I_p)$ for some $\theta^* \in \mathbb{R}^p$. This is a challenging data generating process, since $\bar{g}_t(X_t) = \sin^4((2X_t \theta_t^*)^{-1})$ is non-stationary and highly non-linear, and features are correlated. Equation (9) is discussed more in detail in Appendix E. To predict c_t from X_t , we assume that the decision-maker has access to a class of linear predictors of the form $g : (\theta, X) \mapsto X\theta$.

Benchmarks. In this setting, we compare the performances of DF-FTPL (Algorithm 1) and DF-OGD (Algorithm 2) to two benchmarks. First, we implement Prediction-Focused Online Gradient Descent (PF-OGD). This strategy consists in training in an online manner the model g at each timestep so it minimizes the statistical loss $\ell^{\text{mse}} : (v, X\theta) \mapsto \|v - X\theta\|^2$, irrespective of the downstream decision problem. Then, decisions are greedily made based on the predictions of the model. This approach is formally described in Algorithm 3 in Appendix E. Second, we compare our algorithms to an online version of the Smart Predict-then-Optimize (online SPO, see Algorithm 4) approach from Elmachtoub and Grigas (2022). This algorithm introduces a differentiable and convex loss which upper-bounds the actual decision-focused loss. Given its effectiveness and versatility, it is considered as a very important benchmark in the literature.

Results. In Figure 1, we plot on the left-hand panel the average cumulated cost $t \mapsto t^{-1} \sum_{s=1}^t \bar{g}_{s, v_s}(X_s)$ incurred by DF-OGD, DF-FTPL, PF-OGD and online SPO over 10 runs of $T = 5000$ timesteps, as well as the associated 95% confidence intervals. On the right-hand panel, we plot the average Mean Square Error (MSE) resulting from the sequence of prediction parameters $(\theta_t)_{t \in [T]}$. It appears that both DF-FTPL and DF-OGD outperform PF-OGD and online SPO from a decision point of view, while incurring a higher MSE. This is in line with decision-focused approach, which cares about decision cost rather than statistical loss. This experiment shows that (i) DFL outperforms PFL when models are clearly misspecified and (ii) our algorithms also outperforms the celebrated online SPO. We also mention the presence of additional numerical experiments in




Figure 1: Average cumulated cost (left) and prediction errors (right) of DF-OGD, PF-OGD, DF-FTPL and online SPO.

Appendix E, to study how model misspecification affects the relative performances of DFL and PFL.

5 CONCLUSION.

Decision-focused learning offers a promising way to integrate prediction into decision making. We extend its analysis from the batch to the online setting, enabling non-stationary data and varying objectives. Our algorithms DF-FTPL and DF-OGD comes with provable guarantees, and empirically outperform prediction-focused learning and online SPO in our experiment.

We believe that this work can be extended in several ways. First, it is plausible that faster rates (e.g., $T^{-1/2}$) could be obtained using more traditional online-learning techniques that bypass the bilevel framework, such as discretizing the parameter space and applying expert aggregation over the resulting bins. However, such approaches would likely incur a prohibitive dependence on the parameter dimension m (e.g., exponential in m). Exploring this direction remains however an interesting avenue in future work. Second, we believe that alternative smoothing techniques could be successfully employed, such as Moreau-Yosida transform and proximal operators, to adress the more challenging of a general convex set \mathcal{W} . Third, it would be valuable to investigate less adversarial environments, such as those involving i.i.d. data, to explore whether stronger theoretical guarantees could be obtained and whether novel algorithmic designs might emerge. Finally, implementing our method in more ambitious experimental settings would be of great interest from a practical perspective.

ACKNOWLEDGMENT

Funded by the European Union (ERC, Ocean, 101071601). M.H. is partially supported by the European Research Council Starting Grant DYNASTY – 101039676. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

- J. Abernethy, C. Lee, A. Sinha, and A. Tewari. Online linear optimization via smoothing. In *Conference on learning theory*, pages 807–823. PMLR, 2014.
- J. Acimovic and S. C. Graves. Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Operations Management*, 17(1):34–51, 2015.
- N. Agarwal, A. Gonen, and E. Hazan. Learning in non-convex games with an optimization oracle. In *Conference on Learning Theory*, pages 18–29. PMLR, 2019.
- B. Amos and J. Z. Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International conference on machine learning*, pages 136–145. PMLR, 2017.
- A. Aswani, Z.-J. M. Shen, and A. Siddiq. Data-driven incentive design in the medicare shared savings program. *Operations Research*, 67(4):1002–1026, 2019.
- S. Aydore, T. Zhu, and D. P. Foster. Dynamic local regret for non-convex online forecasting. *Advances in neural information processing systems*, 32, 2019.
- G.-Y. Ban and C. Rudin. The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108, 2019.
- O. Bennouna, J. Zhang, S. Amin, and A. Ozdaglar. Addressing misspecification in contextual optimization. *arXiv preprint arXiv:2409.10479*, 2024.
- Q. Berthet, M. Blondel, O. Teboul, M. Cuturi, J.-P. Vert, and F. Bach. Learning with differentiable perturbed optimizers. *Advances in neural information processing systems*, 33:9508–9519, 2020.
- D. Bertsimas and N. Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3): 1025–1044, 2020.
- D. Bertsimas, V. F. Farias, and N. Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.

- O. Besbes, Y. Gur, and A. Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- C. M. Bishop and N. M. Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- S. Boyd. Convex optimization. *Cambridge UP*, 2004.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- X. Chen, Z. Owen, C. Pixton, and D. Simchi-Levi. A statistical learning approach to personalization in revenue management. *Management Science*, 68(3):1923–1937, 2022.
- W. C. Cheung, D. Simchi-Levi, and R. Zhu. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1079–1087. PMLR, 2019.
- M. C. Cohen, N.-H. Z. Leung, K. Panchamgam, G. Perakis, and A. Smith. The impact of linear optimization on promotion planning. *Operations Research*, 65(2):446–468, 2017.
- B. Colson, P. Marcotte, and G. Savard. An overview of bilevel optimization. *Annals of operations research*, 153:235–256, 2007.
- A. Cutkosky, H. Mehta, and F. Orabona. Optimal Stochastic Non-smooth Non-convex Optimization through Online-to-Non-convex Conversion. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA, 2023*.
- O. de Oliveira. The implicit and the inverse function theorems: Easy proofs. *Real Analysis Exchange*, 39(1):207, 2014. ISSN 0147-1937. doi: 10.14321/realanalexch.39.1.0207. URL <http://dx.doi.org/10.14321/realanalexch.39.1.0207>.
- R. Dixit, A. S. Bedi, R. Tripathi, and K. Rajawat. Online learning with inexact proximal online gradient descent algorithms. *IEEE Transactions on Signal Processing*, 67(5):1338–1352, 2019.
- P. Donti, B. Amos, and J. Z. Kolter. Task-based end-to-end model learning in stochastic optimization. *Advances in neural information processing systems*, 30, 2017.
- J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.
- A. N. Elmachtoub and P. Grigas. Smart “predict, then optimize”. *Management Science*, 68(1):9–26, 2022.
- V. F. Farias, S. Jagabathula, and D. Shah. A nonparametric approach to modeling choice with limited data. *Management science*, 59(2):305–322, 2013.
- K. J. Ferreira, B. H. A. Lee, and D. Simchi-Levi. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & service operations management*, 18(1):69–88, 2016.
- M. Fisher, C. K. Glaeser, and X. Su. Optimal retail location: Empirical methodology and application to practice. *Available at SSRN*, 2016.
- H. Frezat, R. Fablet, G. Balarac, and J. L. Sommer. Gradient-free online learning of subgrid-scale dynamics with neural emulators. *arXiv preprint arXiv:2310.19385*, 2023.
- G. Gallo and S. Pallottino. Shortest path algorithms. *Annals of operations research*, 13(1):1–79, 1988.
- S. Ghadimi and G. Lan. Stochastic First- and Zeroth-Order Methods for Nonconvex Stochastic Programming. *SIAM J. Optim.*, 2013.
- U. Ghai, Z. Lu, and E. Hazan. Non-convex online learning via algorithmic equivalence. *Advances in Neural Information Processing Systems*, 35:22161–22172, 2022.

- S. Gould, B. Fernando, A. Cherian, P. Anderson, R. S. Cruz, and E. Guo. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *arXiv preprint arXiv:1607.05447*, 2016.
- V. Gupta, B. R. Han, S.-H. Kim, and H. Paek. Maximizing intervention effectiveness. *Management Science*, 66(12):5576–5598, 2020.
- E. Hall and R. Willett. Dynamical models and tracking regret in online convex programming. In *International Conference on Machine Learning*. PMLR, 2013.
- N. Hallak, P. Mertikopoulos, and V. Cevher. Regret minimization in stochastic non-convex learning via a proximal-gradient approach. In *International Conference on Machine Learning*, pages 4008–4017. PMLR, 2021.
- E. Hazan. Introduction to online convex optimization, 2023. URL <https://arxiv.org/abs/1909.05207>.
- E. Hazan, K. Singh, and C. Zhang. Efficient regret minimization in non-convex games. In *International Conference on Machine Learning*, pages 1433–1441. PMLR, 2017.
- A. Héliou, P. Mertikopoulos, and Z. Zhou. Gradient-free online learning in continuous games with delayed rewards. In *International conference on machine learning*, pages 4172–4181. PMLR, 2020.
- Y. Hsieh, M. R. K. Jaghargh, A. K. se, and P. Mertikopoulos. Riemannian stochastic optimization methods avoid strict saddle points. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.
- C. Huang and K. Wang. A stability principle for learning under nonstationarity. *Operations Research*, 2025.
- M. Huang and V. Gupta. Decision-focused learning with directional gradients. *Advances in Neural Information Processing Systems*, 37:79194–79220, 2024.
- M. Hutter and J. Poland. Adaptive online prediction by following the perturbed leader. 2005.
- K. Ji, J. Yang, and Y. Liang. Bilevel optimization: Convergence analysis and enhanced design. In *International conference on machine learning*, pages 4882–4892. PMLR, 2021.
- A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- A. Lesage-Landry, J. A. Taylor, and I. Shames. Second-order online nonconvex optimization. *IEEE Transactions on Automatic Control*, 66(10):4866–4872, 2020.
- B. Li and S. C. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3): 1–36, 2014.
- H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the Loss Landscape of Neural Nets. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 2018.
- S. Lin, D. Sow, K. Ji, Y. Liang, and N. Shroff. Non-convex bilevel optimization with time-varying objective functions. *Advances in Neural Information Processing Systems*, 36:29692–29717, 2023.
- J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang. Learning under concept drift: A review. *IEEE transactions on knowledge and data engineering*, 31(12):2346–2363, 2018.
- E. Mammen and A. B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6): 1808–1829, 1999.
- J. Mandi and T. Guns. Interior point solving for lp-based prediction+ optimisation. *Advances in Neural Information Processing Systems*, 33:7272–7282, 2020.

- J. Mandi, V. Bucarey, M. M. K. Tchomba, and T. Guns. Decision-focused learning: Through the lens of learning to rank. In *International conference on machine learning*, pages 14935–14947. PMLR, 2022.
- J. Mandi, J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns, and F. Fioretto. Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*, 80:1623–1701, 2024.
- P. Mertikopoulos, N. Hallak, A. Kavis, and V. Cevher. On the Almost Sure Convergence of Stochastic Gradient Descent in Non-Convex Problems. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- V. V. Mivsić and G. Perakis. Data analytics in operations management: A review. *Manufacturing & Service Operations Management*, 22(1):158–169, 2020.
- S. Padakandla, P. KJ, and S. Bhatnagar. Reinforcement learning algorithm for non-stationary environments. *Applied Intelligence*, 50(11):3590–3606, 2020.
- V. Patel and S. Zhang. Stochastic Gradient Descent on Nonconvex Functions with General Noise Models, 2021.
- V. Perchet and P. Rigollet. The multi-armed bandit problem with covariates. 2013.
- J. Quiñero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence. *Dataset shift in machine learning*. Mit Press, 2022.
- A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Stochastic, constrained, and smoothed adversaries. *Advances in neural information processing systems*, 24, 2011.
- S. Rath, K. Rajaram, and A. Mahajan. Integrated anesthesiologist and room scheduling for surgeries: Methodology and application. *Operations Research*, 65(6):1460–1478, 2017.
- U. Sadana, A. Chenreddy, E. Delage, A. Forel, E. Frejinger, and T. Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025.
- N. Schutte, K. Postek, and N. Yorke-Smith. Robust losses for decision-focused learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 4868–4875, 2024.
- S. Shah, K. Wang, B. Wilder, A. Perrault, and M. Tambe. Decision-focused learning without decision-making: Learning locally optimized decision losses. *Advances in Neural Information Processing Systems*, 35:1320–1332, 2022.
- L. Shen, N. Ho-Nguyen, and F. Kilinc-Karzan. An online convex optimization-based framework for convex bilevel optimization. *Mathematical Programming*, 198(2):1519–1582, 2023.
- A. Sinha, P. Malo, and K. Deb. A review on bilevel optimization: From classical to evolutionary approaches and applications. *IEEE transactions on evolutionary computation*, 22(2):276–295, 2017.
- A. S. Suggala and P. Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. In *Algorithmic Learning Theory*, pages 845–861. PMLR, 2020.
- V. Syrgkanis, A. Agarwal, H. Luo, and R. E. Schapire. Fast convergence of regularized learning in games. *Advances in Neural Information Processing Systems*, 28, 2015.
- D. A. Tarzanagh, P. Nazari, B. Hou, L. Shen, and L. Balzano. Online bilevel optimization: Regret analysis of online alternating gradient methods. In *International Conference on Artificial Intelligence and Statistics*, pages 2854–2862. PMLR, 2024.
- A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.

- S. Verma, A. Mate, K. Wang, A. Taneja, and M. Tambe. Case study: Applying decision focused learning in the real world. In *Workshop on Trustworthy and Socially Responsible Machine Learning, NeurIPS 2022*, 2022.
- S. Verma, A. Mate, K. Wang, N. Madhiwalla, A. Hegde, A. Taneja, and M. Tambe. Restless multi-armed bandits for maternal and child health: Results from decision-focused learning. In *AAMAS*, pages 1312–1320, 2023.
- K. Wang, S. Verma, A. Mate, S. Shah, A. Taneja, N. Madhiwalla, A. Hegde, and M. Tambe. Scalable decision-focused learning in restless multi-armed bandits with application to maternal and child health. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 12138–12146, 2023.
- B. Wilder, B. Dilkina, and M. Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1658–1665, 2019.
- Z. Xu and L. Zhang. Online non-convex learning in dynamic environments. *Advances in Neural Information Processing Systems*, 37:51930–51962, 2024.
- A. Zeevi and A. Goldenshluger. Woodrooffe’s one-armed bandit problem revisited. *Annals of Applied Probability*, 19(4):1603–1633, 2009.
- L. Zhang, S. Lu, and Z.-H. Zhou. Adaptive online learning in dynamic environments. *Advances in neural information processing systems*, 31, 2018.
- P. Zhao and L. Zhang. Improved analysis for dynamic regret of strongly convex and smooth functions. In *Learning for Dynamics and Control*, pages 48–59. PMLR, 2021.
- P. Zhao, Y.-J. Zhang, L. Zhang, and Z.-H. Zhou. Dynamic Regret of Convex and Smooth Functions. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.
- Z. Zhuang, Y. Wang, K. Yu, and S. Lu. No-regret non-convex online meta-learning. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3942–3946. IEEE, 2020.
- M. Zinkevich. Online Convex Programming and Generalized Infinitesimal Gradient Ascent. In *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, 2003.

A **H2** IS SATISFIED IN THE GAUSSIAN LINEAR CASE.

In this section, we prove that **H2** holds true when \mathcal{W} is the simplex of \mathbb{R}^d , g is a linear model and X has columns distributed according to a Gaussian distribution. In what follows, we denote by $X_j \in \mathbb{R}^d$ the j -th column of X .

Proposition 1. *Assume that $g : (\theta, X) \in \Theta \times \mathcal{X} \mapsto X\theta$ where $\Theta = \{\theta \in \mathbb{R}^m : \|\theta\|_2 = 1\}$ and $X = (X_1 | \dots | X_m) \in \mathbb{R}^{d \times m}$ where $X_j \stackrel{i.i.d.}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I})$. Then, **H2** is true for any $\theta \in \Theta$ with*

$$C_0 = \frac{d(d-1)}{2\sqrt{\pi}} \quad \text{and} \quad \beta = 1.$$

Proof. Let $\theta \in \Theta$. In what follows, we denote $i = I(\theta)$ where $I(\theta)$ is defined as in **H2**. Since \mathcal{W} is the simplex of \mathbb{R}^d , v_j is the j -th vector of the canonical basis of \mathbb{R}^d . For any $\varepsilon > 0$ we have:

$$\begin{aligned} \mathbb{P}\left(\inf_{j \neq i} |\langle X\theta, v_j \rangle - \langle X\theta, v_i \rangle| \leq \varepsilon\right) &\leq \mathbb{P}\left(\bigcup_{j \neq i} \{|\langle X\theta, v_j \rangle - \langle X\theta, v_i \rangle| \leq \varepsilon\}\right) \\ &\leq \sum_{j \neq i} \mathbb{P}(|\langle \theta, X^\top(v_j - v_i) \rangle| \leq \varepsilon) \\ &= \sum_{j \neq i} \mathbb{P}(|\langle \theta, X_j - X_i \rangle| \leq \varepsilon). \end{aligned} \quad (10)$$

Since for any $j \neq i$, $\langle \theta, X_j - X_i \rangle \sim \mathcal{N}(0, 2\|\theta\|_2^2)$ with density f_j , we have:

$$\mathbb{P}(|\langle \theta, X_j - X_i \rangle| \leq \varepsilon) = \int_{-\varepsilon}^{\varepsilon} f_j(x) dx = (2\sqrt{\pi}\|\theta\|_2)^{-1} \int_{-\varepsilon}^{\varepsilon} \exp\left(-\frac{x^2}{4\|\theta\|_2^2}\right) dx$$

and since the integrand reaches its maximum in $x = 0$,

$$\leq \varepsilon(\sqrt{\pi}\|\theta\|_2)^{-1}. \quad (11)$$

Then, plugging (11) in (10) and summing over all pairs (i, j) such that $j \neq i$ yields the desired result. \square

Note that the d^2 factor in C_0 comes from the use of an union bound, and could be reduced with a refined analysis.

B EXTENDED DISCUSSION ON THE RELEVANCE OF APPROXIMATE ORACLES.

The notion of an approximate oracle reflects the fact that, in non-convex settings, we cannot rely on subroutines that provably reach a global minimizer of f_t at time t (as offline gradient descent would in convex problems). Instead, we must settle for local minimizers, whose quality is governed by a parameter ξ —which vanishes in favorable loss landscapes.

As a concrete example, consider **O** as the stochastic gradient descent (SGD) algorithm. Even in the context of deep networks, it is plausible that **O** converges to a local minimizer. Indeed, a large body of work has analyzed SGD in non-convex settings, establishing convergence to stationary points either in expectation (Ghadimi and Lan, 2013) or almost surely (Mertikopoulos et al., 2020; Patel and Zhang, 2021; Cutkosky et al., 2023). Such stationary points may correspond to local/global minima or saddle points.

Recent studies further show that SGD avoids saddle points. For instance, Mertikopoulos et al. (2020) proved that the trajectories of SGD almost surely avoid all strict saddle manifolds—i.e., sets of critical points x where the Hessian has at least one negative eigenvalue. These manifolds include connected families of non-isolated saddle points, a phenomenon common in the loss landscapes of overparametrized neural networks (Li et al., 2018).

Beyond SGD, similar guarantees extend to more general methods. In particular, the stochastic Riemannian Robbins–Monro method (a broad template encompassing various algorithms) has been shown to converge almost surely to local or global minima (Hsieh et al., 2023).

Finally, we highlight that, in most practical industrial settings, solving the downstream optimization problem is typically far more computationally demanding than updating the prediction model. It is therefore reasonable to assume that the oracle call represents only a small fraction of the overall computational cost.

C SUPPLEMENTARY RESULTS ON THE SIMPLEX.

We now provide convergence guarantees for Algorithms 1 and 2 when \mathcal{W} is the simplex. To do so we choose another regularizer than done in the main document.

C.1 ADDITIONAL FRAMEWORK

We focus on the case where \mathcal{W} is the simplex of \mathbb{R}^d , that is

$$\mathcal{W}_0 = \{w \in \mathbb{R}^d : w \succcurlyeq 0, \mathbf{1}^\top w = 1\}.$$

This setting encompasses various decision-making scenarios such as portfolio selection Li and Hoi (2014) or mixed strategy design in multiplayer games Syrgkanis et al. (2015). In this case, we choose the negative entropy

$$\mathcal{R}_0 : w \mapsto \sum_{i \in [d]} w_i \ln(w_i),$$

as the regularizer in (5). The minimizer $\tilde{w}_t(\theta) \in \mathcal{W}_0$ in (5) with $\mathcal{R} = \mathcal{R}_0$ can easily be shown to be the softmax mapping, that is it satisfies for any $i \in [d]$:

$$\tilde{w}_{t,i}(\theta) = \frac{\exp(-\alpha^{-1}g_i(\theta, X_t))}{\sum_{k \in [d]} \exp(-\alpha^{-1}g_k(\theta, X_t))}.$$

It is clear from this expression that \tilde{w}_t is differentiable, and that for any $\theta \in \Theta$:

$$\nabla \tilde{w}_t(\theta) = -\frac{1}{\alpha} [\text{diag}[\tilde{w}_t(\theta)] - \tilde{w}_t(\theta) \tilde{w}_t(\theta)^\top] \nabla_\theta g(\theta, X_t). \quad (12)$$

C.2 RESULTS

DF-FTPL

Proposition 2. *Assume **H1**, **H2** and having access to an ξ -approximate optimization oracle \mathbf{O}_ξ adapted to $\left\{ \sum_{i=1}^t \tilde{f}_i - \langle \sigma_t, \cdot \rangle \right\}_{t \in [T]}$. Fix $\mathcal{W} = \mathcal{W}_0$, $\mathcal{R} = \mathcal{R}_0$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-FTPL (Algorithm 1) instantiated with learning step $\eta > 0$ and regularization coefficients $\alpha_t = \alpha > 0$ for any t . Then:*

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^s] \leq \tilde{O} \left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha \ln(d) \right),$$

where \mathbb{E} denotes the expectation on both data and the intrinsic randomness of DF-FTPL and \tilde{O} contains polynomial dependency in $\ln(1/\alpha)$, $\ln(\ln(d))$.

The proof of this result is postponed to Appendix G.4.

DF-OGD

Proposition 3. *Assume **H1**, **H2**, and having access to an ξ -approximate optimization oracle \mathbf{O}_ξ adapted to $\left\{ \tilde{f}_i \right\}_{i \in [T]}$. Fix $\mathcal{W} = \mathcal{W}_0$, $\mathcal{R} = \mathcal{R}_0$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-OGD (Algorithm 2)*

instantiated with the non-increasing sequence $(\eta_t)_{t \in [T]}$ and regularization coefficients $(\alpha_t)_{t \in [T]}$. Then:

$$\begin{aligned} T^{-1} \mathbb{E}[\mathfrak{R}_T^d] &\leq \frac{D_\Theta^2 + D_\Theta \mathbb{E}[P_T]}{2T\eta_T} + \sum_{t \in [T]} \frac{25D_Z G \eta_t}{32T\alpha_t^2} \\ &\quad + 2D_Z \sum_{t \in [T]} \alpha_t [1 + 2\ln(d)C_0 \{\ln(2\alpha_t^{-1}) + (1 - \beta) \ln^2(\alpha_t \ln d)\}] + \xi, \end{aligned}$$

where $P_T = \sum_{t=1}^{T-1} \|\vartheta_{t+1} - \vartheta_t\|$.

The proof of this result is postponed to Appendix G.5.

Note that the principal gain, compared to the general polytope case, is to avoid the dependency in n the number of faces and enjoys a nice $\ln(d)$ dependency in the dimension of the decision space.

D BACKGROUND ON THE FTPL ALGORITHM

Approximate Optimization Oracle. The recent work of [Suggala and Netrapalli \(2020\)](#) proposed an online algorithm for nonconvex losses $(\ell_t)_{t \in [T]}$ with static regret guarantees. They rely on an *approximate optimization oracle* \mathcal{O} which takes as input any function ℓ , a d -dimensional vector σ and returns an approximate minimizer of $\ell - \langle \sigma, \cdot \rangle$. An optimization oracle is called " (ξ, χ) -approximate optimization oracle" if it returns $\mu^* \in \mathcal{K}$ such that

$$\ell(\mu^*) - \langle \sigma, \mu^* \rangle \leq \inf_{\mu \in \mathcal{K}} [\ell(\mu) - \langle \sigma, \mu \rangle] + (\xi + \chi \|\sigma\|_1). \quad (13)$$

We denote the output of such an optimization oracle by $\mu^* = \mathcal{O}_{\xi, \chi}(\ell - \langle \sigma, \cdot \rangle)$. Note the notion of oracle described in (13) is very close from the Definition 1 we made in Section 3 on the expert sequence $(\vartheta_t)_{t \in [T]}$ (in this case $\chi = 0$ as no σ is involved).

Remark 2. Note that, in our work, we made the choice to fix $\chi = 0$. This stronger assumption is due to the will of having a unifying framework encompassing both the setup for DF-FTPL and DF-OGD.

Follow The Perturbed Leader. Given access to an (ξ, χ) -approximate optimization oracle, [Suggala and Netrapalli \(2020\)](#) study the FTPL algorithm which is described by the following recursion. Starting from $\hat{\mu}_1$, at each time steps t , $\hat{\mu}_t \in \mathcal{K}$ is updated as follows:

$$\hat{\mu}_{t+1} = \mathcal{O}_{\xi, \chi} \left(\sum_{i=1}^{t-1} \ell_i - \langle \sigma_t, \cdot \rangle \right) \quad (14)$$

where $\sigma_t \in \mathbb{R}^d$ is a random perturbation such that $\sigma_{t,j}$, the j -th coordinate of σ_t , is sampled from $\text{Exp}(\eta)$, the exponential distribution with parameter η .

Following this update route, the addition of an exponential noise allowing to handle the non-convexity of the losses, they reach the following static regret bound.

Proposition 4 (Theorem 1 of [Suggala and Netrapalli \(2020\)](#)). *Let D be the diameter of \mathcal{K} . Suppose that ℓ_t is L -Lipschitz w.r.t L_1 norm, for all $t \in [T]$. For any fixed η , FTPL (Equation (14)) with access to a (ξ, χ) -approximate" optimization oracle satisfies the following static regret bound:*

$$\frac{1}{T} \left(\sum_{t=1}^T \ell_t(\hat{\mu}_t) - \inf_{\mu} \sum_{t=1}^T \ell_t(\mu) \right) \leq \mathcal{O} \left(\eta m^2 D L^2 + \frac{m(\chi T + D)}{\eta T} + \xi + \chi m L \right),$$

where \mathbb{E} denotes the expectation over $\sigma_1 \cdots \sigma_T$.

E ADDITIONAL EXPERIMENTAL MATERIAL

The motivating example from [Mandi et al. \(2024\)](#). In this paragraph, we recall the example that motivates our experiment. [Mandi et al. \(2024\)](#) illustrates the interest of decision-focused learning

with a simple problem where a decision maker seeks to pick, between two objects, the one with the lowest cost. Before picking an object, they do not know costs but have at their disposal a model to predict it. This situation is depicted in Figure 2. For instance, if the true cost $\bar{g}_t = (\bar{g}_{t,1}, \bar{g}_{t,K})$ is $(*)$, any prediction $\hat{g}_t = (\hat{g}_{t,1}, \hat{g}_{t,K})$ lying in the blue-shaded area, such as $(+)$, induces the optimal decision $v_t = 1$. As observed in Mandi et al. (2024), a prediction-focused approach may underperform compared to a decision-focused one in this setting. This is because generating a prediction \hat{v}_t that closely approximates the true value v_t —in the sense of minimizing statistical loss—does not necessarily ensure that \hat{v}_t falls on the same side of the 45° line as v_t . For example, the prediction (\times) is satisfactory from a predictive point of view, but induces a sub-optimal action. On the contrary, one would expect the decision-focused approach to produce predictions that lie further away from $(*)$ since it does not minimize prediction error, but on the right side of the 45° line; see for instance (\blacktriangle) or (\blacktriangledown) on Figure 2.




Figure 2: Figure 2 from Mandi et al. (2024).

Details on the experimental setup. In this paragraph, we provide more detail about the experimental setup in Section 4 and perform additional numerical simulations. We start by explaining more precisely how the data used in our experiment are drawn. First, for any $t \in [T]$, $X_t \in \mathbb{R}^{K \times p}$ is constructed as follows: (i) we generate a Toeplitz covariance matrix $\Sigma = (\rho^{|i-j|})_{(i,j) \in [K]^2}$ for some $\rho \in (0, 1)$, (ii) apply a Cholesky decomposition to it: $\Sigma = LL^T$, and (iii) generate a matrix \bar{X}_t with standard Gaussian entries. Then, X_t is defined as $X_t = L\bar{X}_t$. This introduces correlation between features, which makes convergence harder for an online linear model. Second, we generate $v_t \in \mathbb{R}^K$ as follows:

$$c_t = \min(\max(\tilde{c}_t, 0), 1) \quad \text{with} \quad \tilde{c}_t = A \sin^4((2X_t \theta_t^*)^{-1}) + \varepsilon.$$

In the above equation, $A > 0$ is a constant, θ_t^* satisfies:

$$\theta_t^* = \frac{1}{2}\theta^* + \frac{1}{2}\zeta_t, \quad \text{where} \quad \zeta_t \sim \mathcal{N}(0, I_K) \quad \text{and} \quad \theta^* \in \mathbb{R}^K,$$

and $\varepsilon_t \sim \mathcal{N}(0, I_K)$ is Gaussian noise. The fact that θ_t^* varies throughout time and that the relationship between v_t and X_t is highly linear makes learning hard for a linear model.

We now present the values used for the different parameters in our experiment. We consider $K = 5$ items and $p = 10$ features. The horizon is set to $T = 5000$. For each plot, we run $N = 10$ times DF-OGD and PF-OGD. The reported error bars are 95% Gaussian confidence intervals (from the CLT). In (9), $A = 45$ and the correlation coefficient for X_t is $\rho = 0.8$. Our algorithms are instantiated with the following parameters. First, DF-OGD runs with (α_t) and (η_t) as suggested by our theoretical analysis. The oracle from H2 is obtained through a SGD algorithm performing $t_\vartheta = 10$ steps at each iteration with a learning rate set to $\eta_\vartheta = 0.01$, while being initialized at ϑ_{t-1} . On the other hand, PF-OGD runs with a learning rate $\eta = 10$ determined by grid search. Our code is implemented with pytorch.

Benchmark algorithms.

Algorithm 3 Prediction-Focused Online Gradient Descent (PF-OGD).

-
- 1: **Input:** horizon $T > 0$, initialization $\theta_1 \in \Theta$.
 - 2: **for** each $t \in \{1, \dots, T\}$ **do**
 - 3: Observe X_t , predict $c_t = g(\theta_t, X_t)$
 - 4: Play $v_t^*(\theta_t) = \operatorname{argmin}_{i \in [K]} c_{t,i}$, observe $\bar{g}_t(X_t)$.
 - 5: Update $\theta_{t+1} = \Pi_{\Theta}(\theta_t - \eta \nabla \ell^{\text{mse}}(\bar{g}_t(X_t), X_t \theta_t))$.
 - 6: **end for**
-

Deviation from realizability. As argued in Section 1, prediction-focused learning would be optimal if the prediction model made no mistake, since in this case problem (2) would perfectly align with (1). As a consequence, one would expect prediction-focused learning to perform very well in the realizable setting, that is when for any $t \in [T]$, there exists $\bar{\theta}_t \in \Theta$ such that $\bar{g}_t(X_t) = g(\bar{\theta}_t, X_t)$. On the flip side, it is likely to struggle in a misspecified setting. This is in contrast with DF-OGD, whose theoretical performance analysis does not highlight any peculiar dependence on realizability. We therefore hypothesize that PFL should outperform DFL in the realizable case, and DFL should take the upper hand when the prediction model is highly misspecified. To test this hypothesis, we run an alternative experiment where we interpolate between the well-specified, linear case, and ill-specified, non-linear one. More precisely, for $\gamma \in [0, 1]$, we generate $(X_t, v_t^{(\gamma)})_{t \in [T]}$ as follows:

$$v_t^{(\gamma)} = (1 - \gamma)X_t \theta_t^* + \gamma \sin^4((2X_t \theta_t^*)^{-1}) + \varepsilon_t.$$

In Figure 3, we make γ vary from 0 to 1. At each value, we plot the average cumulated cost gap between DF-OGD and PF-OGD, averaged over 50 runs with horizon $T = 1000$. This experiment provides empirical support to the previous conjecture, namely, DFL becomes a competitive option when the prediction model is ill-specified.




Figure 3: Average cumulated cost gap PFL - DFL as a function of γ .

Algorithm 4 Online Smart Predict-then-Optimize Online Gradient Descent (Online SPO).

-
- 1: **Input:** horizon $T > 0$, initialization $\theta_1 \in \Theta$.
 - 2: **for** each $t \in \{1, \dots, T\}$ **do**
 - 3: Observe X_t , predict $\hat{c}_t = g(\theta_t, X_t)$.
 - 4: Play $v_t^*(\theta_t) = \operatorname{argmin}_{v \in \mathcal{V}} \hat{c}_t^\top v$, observe c_t .
 - 5: Compute $v^*(\hat{c}_t) = \operatorname{argmin}_{v \in \mathcal{V}} \hat{c}_t^\top v$ and $v^*(c_t) = \operatorname{argmin}_{v \in \mathcal{V}} c_t^\top v$.
 - 6: Define SPO⁺ surrogate:

$$\ell^{\text{SPO}^+}(\hat{c}_t, c_t) = 2c_t^\top v^*(\hat{c}_t) - c_t^\top v^*(c_t) - \hat{c}_t^\top v^*(\hat{c}_t).$$

- 7: Update $\theta_{t+1} = \Pi_{\Theta}(\theta_t - \eta \nabla_{\theta} \ell^{\text{SPO}^+}(\hat{c}_t, c_t))$, where the gradient is computed via chain rule through $\hat{c}_t = g(\theta_t, X_t)$.
 - 8: **end for**
-




Figure 4: Experiment results with 80 items.

Higher dimension problem. Finally, the following figure shows the performance of our algorithms in a higher dimension problem, with $K = 80$ items. We see that our algorithms still significantly beat the baselines, suggesting that they remain effective in moderately high dimension problems.

F REGRET BOUND WITH AN APPROXIMATE OPTIMAL SOLUTION.

In this section, we show that when we cannot access $w_t^*(\theta_t) \in \mathcal{W}$ but only $\underline{w}_t(\theta_t) \in \mathcal{W}$ such that $\langle g(\theta, X_t), \underline{w}_t(\theta_t) - w_t^*(\theta_t) \rangle \leq \kappa$ for $\kappa > 0$, our regret bounds are only shifted by κ . The following proof is for the general convex polytope case. The simplex case is established in the exact same way.

Proposition 5. Assume the assumptions from Theorem 2 and that at each step, $w_t^*(\theta_t) \in \mathcal{W}$ is replaced by $\underline{w}_t(\theta_t) \in \mathcal{W}$ which satisfies

$$\langle g(\theta, X_t), \underline{w}_t(\theta_t) - w_t^*(\theta_t) \rangle \leq \kappa,$$

where $\kappa > 0$. Then,

$$\begin{aligned} \mathbb{E}[\mathfrak{R}_T^d] &\leq \frac{D_\Theta^2 - D_\Theta \mathbb{E}[P_T]}{2\eta_T} + \sum_{t \in [T]} \frac{\eta_t}{2} \left(\frac{GD_{\mathcal{Z}} \sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty}{\alpha_t \lambda_{\min}(AA^\top)} \right)^2 \\ &\quad + 2D_{\mathcal{Z}} \sum_{t \in [T]} \alpha_t [1 + 2n \max_{k \in [K]} \|v_k\|_1 C_0 (\ln(1/\alpha_t) + (1 - \beta)^2 \ln^2(\alpha_t \ln d))] + \xi + \kappa. \end{aligned}$$

The proof Proposition 5 is deferred to Appendix G.6.

G PROOFS

G.1 PROOFS OF PRELIMINARY LEMMAS.

G.1.1 SIMPLEX

Lemma 1. Assume **H1**, $\mathcal{W} = \mathcal{W}_0$ and $\mathcal{R} = \mathcal{R}_0$. Then for any $t \in [T]$, \tilde{f}_t is K_t -Lipchitz almost surely, with $K_t = 5D_{\mathcal{Z}}G(4\alpha_t)^{-1}$.

Proof. Let $(\theta, \theta') \in \Theta^2$. We have:

$$\begin{aligned} \left| \tilde{f}_t(\theta) - \tilde{f}_t(\theta') \right| &= \langle \bar{g}_t(X_t), \tilde{w}_t(\theta) - \tilde{w}_t(\theta') \rangle \\ &\leq \|\bar{g}_t(X_t)\| \|\tilde{w}_t(\theta) - \tilde{w}_t(\theta')\| \\ &= \|\bar{g}_t(X_t)\| \int_{t=0}^1 \|\nabla \tilde{w}_t(\theta + t(\theta' - \theta))(\theta - \theta')\| dt \\ &\leq D_{\mathcal{Z}} \int_{t=0}^1 \|\nabla \tilde{w}_t(\theta + t(\theta' - \theta))\|_{\text{op}} \|\theta - \theta'\| dt \quad (\text{by H1-(ii)}) \quad (15) \end{aligned}$$

Now, denoting $\zeta = \theta + t(\theta' - \theta)$, we have:

$$\begin{aligned} \|\nabla \tilde{w}_t(\zeta)\|_{\text{op}} &= \alpha_t^{-1} \|(\text{diag}[\tilde{w}_t(\zeta)] - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top)\nabla_{\theta}g(\zeta, X_t)\|_{\text{op}} \\ &\leq \alpha_t^{-1} \|\text{diag}[\tilde{w}_t(\zeta)] - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top\|_{\text{op}} \|\nabla_{\theta}g(\zeta, X_t)\|_{\text{op}} \\ &\leq G\alpha_t^{-1} \|\text{diag}[\tilde{w}_t(\zeta)] - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top\|_{\text{op}} \quad (\text{by H1-(iii)}) \end{aligned} \quad (16)$$

Since $\text{diag}[\tilde{w}_t(\zeta)] - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top$ is symmetric, its operator norm equals its largest eigenvalue, denoted λ_{\max} . It then follows that:

$$\begin{aligned} \|\text{diag}[\tilde{w}_t(\zeta)] - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top\|_{\text{op}} &= \lambda_{\max}(\text{diag}[\tilde{w}_t(\zeta)] - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top) \\ &\leq \max_{i \in [d]} \sum_{j \in [d]} |\text{diag}[\tilde{w}_t(\zeta)]_{i,j} - \tilde{w}_t(\zeta)\tilde{w}_t(\zeta)^\top_{i,j}| \\ &\leq \max_{i \in [d]} \{ \tilde{w}_t(\zeta)_{i,i}(1 - \tilde{w}_t(\zeta)_{i,i}) + \tilde{w}_t(\theta)_{i,i} \sum_{j \neq i} \tilde{w}_t(\zeta)_{i,j} \} \\ &\leq \frac{1}{4} + 1 = \frac{5}{4}, \end{aligned}$$

where the two last inequalities hold because $\tilde{w}_t \in [0, 1]^d$. Therefore, we obtain from (16) that $\|\nabla \tilde{w}_t(\zeta)\|_{\text{op}} \leq 5G(4\alpha_t)^{-1}$, and plugging this in (15) yields:

$$\left| \tilde{f}_t(\theta) - \tilde{f}_t(\theta') \right| \leq \frac{5GD_Z}{4\alpha_t} \int_0^1 \|\theta - \theta'\| dt = \frac{5GD_Z}{4\alpha_t} \|\theta - \theta'\|.$$

□

Lemma 2. Assume H2, $\mathcal{W} = \mathcal{W}_0$ and $\mathcal{R} = \mathcal{R}_0$. Then for any $\theta \in \Theta$ and $t \in [T]$, for any $\theta \in \theta$

$$\mathbb{E}[\|\tilde{w}_t(\theta) - w_t^*(\theta)\|_1 \mid \mathcal{H}_{t-1}] \leq \alpha_t + 2(\alpha_t \ln(d))^\beta C_0 \min\left\{\ln\left(\frac{2}{\alpha_t}\right), (1 - \beta)^{-1}\right\}.$$

Proof. Let $\theta \in \Theta$ and $t \in [T]$. Since $\mathcal{W} = \mathcal{W}_0$, denoting for any $j \in [d]$ e_j the j -th element of the canonical basis, we have $u_j(\theta, X_t) = \langle g(\theta, X_t), e_j \rangle$. We recall that we denote $u_{I_t(\theta)}(\theta, X_t) = \min_{j \in [d]} u_j(\theta, X_t)$. We first prove the following lemma.

Lemma 3. Assume that there exists $\varepsilon > 0$ such that $u_{I_t(\theta)}(\theta, X_t) + \varepsilon \leq u_j(\theta, X_t)$ for any $j \in \{1, \dots, K\} \setminus \{I_t(\theta)\}$. Then,

$$\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \leq \frac{2\alpha_t \ln(d)}{\varepsilon}.$$

Proof. In this proof, $I_t(\theta)$ from H2 is denoted I_t to lighten notation. We have by assumption:

$$\begin{aligned} \langle g(\theta, X_t), \tilde{w}_t(\theta) \rangle &= \left\langle g(\theta, X_t), \sum_{j=1}^d \tilde{w}_{t,j}(\theta) e_j \right\rangle \\ &= u_{I_t}(\theta, X_t) \tilde{w}_{t,I_t}(\theta) + \sum_{j \neq I_t} u_j(\theta, X_t) \tilde{w}_{t,j}(\theta) \\ &\geq u_{I_t}(\theta, X_t) \tilde{w}_{t,I_t}(\theta) + (u_{I_t}(\theta, X_t) + \varepsilon) \sum_{j \neq I_t} \tilde{w}_{t,j}(\theta) \\ &= u_{I_t}(\theta, X_t) \tilde{w}_{t,I_t}(\theta) + (u_{I_t}(\theta, X_t) + \varepsilon)(1 - \tilde{w}_{t,I_t}(\theta)). \end{aligned} \quad (17)$$

The last equality holds because $\|\tilde{w}_t(\theta)\|_1 = 1$. On the other hand, using the facts that $\mathcal{R}_0 \geq 0$, $\tilde{w}_t(\theta) = \text{argmin} \langle g(\theta, X_t), w \rangle + \alpha_t \mathcal{R}_0(w)$ and $u_{I_t}(\theta, X_t) = \langle g(\theta, X_t), w_t^*(\theta) \rangle$ yields:

$$\begin{aligned} \langle g(\theta, X_t), \tilde{w}_t(\theta) \rangle &\leq \langle g(\theta, X_t), \tilde{w}_t(\theta) \rangle + \alpha_t \mathcal{R}_0(\tilde{w}_t(\theta)) \\ &\leq \langle g(\theta, X_t), w_t^*(\theta) \rangle + \alpha_t \mathcal{R}_0(w_t^*(\theta)) \\ &= u_{I_t}(\theta, X_t) + \alpha_t \mathcal{R}_0(w_t^*(\theta)) \\ &\leq u_{I_t}(\theta, X_t) + \alpha_t \ln(d). \end{aligned} \quad (18)$$

The last line holds because $\mathcal{R}_0(w) \leq \ln(d)$ for any $w \in \mathcal{W}_0$. Combining (17) and (18) gives:

$$u_{I_t}(\theta, X_t) \tilde{w}_{t,I_t}(\theta) + (u_{I_t} + \varepsilon)(1 - \tilde{w}_{t,I_t}(\theta)) \leq u_{I_t}(\theta, X_t) + \alpha_t \ln(d).$$

Re-organizing the terms then yields $\tilde{w}_{t,I_t}(\theta) \geq 1 - \frac{\alpha_t \ln(d)}{\varepsilon}$ and again, as $\|\tilde{w}_t(\theta)\|_1 = 1$, we have: $\sum_{j \neq I_t} \tilde{w}_{t,j}(\theta) \leq \frac{\alpha_t \ln(d)}{\varepsilon}$.

Combining those inequalities yields:

$$\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 = \|e_{I_t} - \tilde{w}_t(\theta)\|_1 = 1 - \tilde{w}_{t,I_t}(\theta) + \sum_{j \neq I_t} \tilde{w}_{t,j}(\theta) \leq \frac{2\alpha_t \ln(d)}{\varepsilon}.$$

□

To conclude the proof, we first write the expected distance as follows:

$$\mathbb{E}[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \mid \mathcal{H}_{t-1}] = \int_0^{+\infty} \mathbb{P}(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > y \mid \mathcal{H}_{t-1}) dy,$$

Since both $w_t^*(\theta)$ and $\tilde{w}_t(\theta)$ belong to the simplex, we can restrict the integral to:

$$= \int_0^2 \mathbb{P}(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > y \mid \mathcal{H}_{t-1}) dy \quad (19)$$

$$= \int_0^{\alpha_t} \mathbb{P}(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > y \mid \mathcal{H}_{t-1}) dy \quad (20)$$

$$+ \int_{\alpha_t}^2 \mathbb{P}(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > y \mid \mathcal{H}_{t-1}) dy \\ \leq \alpha_t + \int_{\alpha_t}^2 \mathbb{P}(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > y \mid \mathcal{H}_{t-1}) dy \quad (21)$$

We now apply the change of variable $y = 2\alpha_t \ln(d)\varepsilon^{-1}$ to obtain:

$$\mathbb{E}[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \mid \mathcal{H}_{t-1}] \\ = \alpha_t + 2\alpha_t \ln(d) \int_{(\alpha_t \ln(d))}^{(2 \ln(d))} \mathbb{P}\left(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > 2 \frac{\alpha_t \ln(d)}{\varepsilon} \mid \mathcal{H}_{t-1}\right) \frac{d\varepsilon}{\varepsilon^2}. \quad (22)$$

Moreover, for any $\varepsilon > 0$ we have by **H2** and Lemma 3 that:

$$1 - C_0 \varepsilon^\beta \leq \mathbb{P}\left(\inf_{j \neq I_t} \{u_j(\theta, X_t) - u_{I_t}(\theta, X_t)\} \geq \varepsilon \mid \mathcal{H}_{t-1}\right) \\ \leq \mathbb{P}(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \leq 2\alpha_t \ln(d)\varepsilon^{-1} \mid \mathcal{H}_{t-1}),$$

so it follows that:

$$\mathbb{P}_t\left(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 > 2 \frac{\alpha_t \ln(d)}{\varepsilon} \mid \mathcal{H}_{t-1}\right) = 1 - \mathbb{P}_t\left(\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \leq 2 \frac{\alpha_t \ln(d)}{\varepsilon} \mid \mathcal{H}_{t-1}\right) \\ \leq C_0 \varepsilon^\beta, \quad (23)$$

Hence, plugging (23) in (22) gives:

$$\mathbb{E}[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \mid \mathcal{H}_{t-1}] \leq \alpha_t + 2\alpha_t \ln(d) C_0 \int_{\alpha_t \ln(d)}^{2 \ln(d)} \varepsilon^{\beta-2} d\varepsilon.$$

We now proceed to bound $\varphi(\beta) = \int_{\alpha_t \ln(d)}^{2 \ln(d)} \varepsilon^{\beta-2} d\varepsilon$ for $\beta \in (0, 1)$. With $a = \alpha_t \ln(d)$ and $b = 2 \ln(d)$, we have

$$\varphi(\beta) = (1 - \beta)^{-1} (a^{\beta-1} - b^{\beta-1}) = (1 - \beta)^{-1} [e^{-(1-\beta) \ln(a)} - e^{-(1-\beta) \ln(b)}].$$

Factoring out $e^{-(1-\beta)\ln(a)} = a^{\beta-1}$ yields

$$\varphi(\beta) = a^{\beta-1} (1-\beta)^{-1} \left[1 - e^{-(1-\beta)\ln(b/a)} \right].$$

Since for any $x \in \mathbb{R}$ one has $1 - e^{-x} \leq \min\{x, 1\}$, we obtain

$$\varphi(\beta) \leq a^{\beta-1} (1-\beta)^{-1} \min\{(1-\beta)\ln(b/a), 1\} = a^{\beta-1} \min\{\ln(b/a), (1-\beta)^{-1}\}.$$

Substituting the values of a and b finally gives

$$\varphi(\beta) \leq (\alpha_t \ln d)^{\beta-1} \min\left\{\ln\left(\frac{2}{\alpha_t}\right), (1-\beta)^{-1}\right\},$$

which concludes the proof. \square

G.1.2 GENERAL POLYTOPE

In this section, we denote by \mathcal{W}_1 the general polytope described in ?? and \mathcal{R}_1 the log-barrier regularization described in Equation (7).

Lemma 4. Assume **H1**, $\mathcal{W} = \mathcal{W}_1$ and $\mathcal{R} = \mathcal{R}_1$. Then for any $t \in [T]$, \tilde{f}_t is differentiable and for any $\theta \in \Theta$, $\nabla \tilde{f}_t(\theta) = \nabla \tilde{w}_t(\theta) \tilde{g}_t(X_t)$ where:

$$\nabla \tilde{w}_t(\theta) = -\alpha_t^{-1} \left(\sum_{i=1}^n (b_i - A_i^\top \tilde{w}_t(\theta))^{-2} A_i A_i^\top \right) \nabla_\theta g(\theta, X_t).$$

Proof. Let $\theta \in \mathbb{R}$ and $t \in [T]$. With

$$h_t : (w, \theta) \mapsto \langle g(\theta, X_t), w \rangle + \alpha_t \mathcal{R}_1(w),$$

we have $h_t(w) \rightarrow +\infty$ as $w \rightarrow \text{bdry}(\mathcal{W}_1)$. Since $h_t(\tilde{w}_t(\theta), \theta) = \min_{w \in \mathcal{W}_1} h_t(w, \theta)$, we deduce that $\tilde{w}_t(\theta) \in \text{int}(\mathcal{W}_1)$. It follows from the first order condition and the implicit function theorem (de Oliveira, 2014, Theorem 2) that $\tilde{w}_t : \theta \mapsto \tilde{w}_t(\theta)$ is differentiable, and

$$\begin{aligned} \nabla \tilde{w}_t(\theta) &= -(\nabla_{ww}^2 h_t(\tilde{w}_t(\theta), \theta))^{-1} \nabla_{\theta w}^2 h_t(\tilde{w}_t(\theta), \theta) \\ &= -\frac{1}{\alpha_t} [\nabla^2 \mathcal{R}_1(\tilde{w}_t(\theta))]^{-1} \nabla_\theta g(\theta, X_t). \end{aligned} \quad (24)$$

Since $\mathcal{R}_1(w) = -\sum_{i=1}^n \ln(b_i - A_i w)$, simple computations give:

$$\nabla^2 \mathcal{R}_1(\tilde{w}_t(\theta)) = \sum_{i=1}^n (b_i - A_i^\top \tilde{w}_t(\theta))^{-2} A_i A_i^\top, \quad (25)$$

Moreover, since $\text{rank}(AA^\top) = d$ by assumption, $\nabla^2 \mathcal{R}_1(\tilde{w}_t(\theta))$ is invertible. Plugging (25) in (24) gives the result. \square

Lemma 5. Assume **H1** and $\mathcal{R} = \mathcal{R}_1$. For any $t \in [T]$, \tilde{f}_t is K_t -Lipschitz almost-surely, with

$$K_t = \frac{GD_{\mathcal{Z}} \sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty}{\alpha_t \lambda_{\min}(AA^\top)}.$$

Proof. Let $(\theta, \theta') \in \Theta^2$. We proved in (15) that:

$$\left| \tilde{f}_t(\theta) - \tilde{f}_t(\theta') \right| \leq D_{\mathcal{Z}} \int_{t=0}^1 \|\nabla \tilde{w}_t(\theta + t(\theta' - \theta))\|_{\text{op}} \|\theta - \theta'\| dt.$$

Denoting $\zeta = \theta + t(\theta' - \theta)$, by (24), since $\|\cdot\|_{\text{op}}$ is sub-multiplicative:

$$\begin{aligned} \|\nabla \tilde{w}_t(\zeta)\|_{\text{op}} &= -\alpha_t^{-1} \left\| [\nabla^2 \mathcal{R}_1(\tilde{w}_t(\zeta))]^{-1} \nabla_\theta g(\theta, X_t) \right\|_{\text{op}} \\ &\leq -\alpha_t^{-1} \left\| [\nabla^2 \mathcal{R}_1(\tilde{w}_t(\zeta))]^{-1} \right\|_{\text{op}} G \quad (\text{by H1-(iii)}). \end{aligned}$$

Moreover, we have:

$$\begin{aligned} \|\nabla^2 \mathcal{R}_1(\zeta)\|_{\text{op}}^{-1} &= \left\| \left(\sum_{i=1}^n (b_i - A_i^\top \tilde{w}_t(\zeta))^{-2} A_i A_i^\top \right)^{-1} \right\|_{\text{op}} \\ &\leq \sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty^2 \|(AA^\top)^{-1}\|_{\text{op}}, \end{aligned}$$

Since AA^\top is symmetric, $\|(AA^\top)^{-1}\|_{\text{op}} = \lambda_{\min}(AA^\top)^{-1}$, where $\lambda_{\min}(M)$ is the lowest eigenvalue of M . We therefore obtain by (15) and the previous inequalities that almost-surely:

$$\begin{aligned} \left| \tilde{f}_t(\theta) - \tilde{f}_t(\theta') \right| &\leq \frac{GD_{\mathcal{Z}}}{\alpha_t} \frac{\sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty^2}{\lambda_{\min}(AA^\top)} \int_{t=0}^1 \|\theta - \theta'\| dt \\ &= \frac{GD_{\mathcal{Z}}}{\alpha_t} \frac{\sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty^2}{\lambda_{\min}(AA^\top)} \|\theta - \theta'\|. \end{aligned}$$

□

Lemma 6. Assume H2, $\mathcal{W} = \mathcal{W}_1$ and $\mathcal{R} = \mathcal{R}_1$. Then for any $t \in [T]$ and $\theta \in \Theta$,

$$\mathbb{E}[\|\tilde{w}_t(\theta) - w_t^*(\theta)\|_1 \mid \mathcal{H}_{t-1}] \leq \alpha_t + 2n \max_{k \in [K]} \|v_k\|_1 \alpha_t (\alpha_t \ln d)^{\beta-1} \min\left\{ \ln\left(\frac{2}{\alpha_t}\right), (1-\beta)^{-1} \right\}$$

Proof. Let $\theta \in \Theta$ and $t \in [T]$. Since $\mathcal{W}_1 = \text{Conv}(v_1, \dots, v_K)$, there exists $(\lambda_{t,I_t}(\theta), \dots, \lambda_{t,I_t}(\theta)) \in [0, 1]^K$ such that

$$\tilde{w}_t(\theta) = \sum_{i=1}^K \lambda_{t,i}(\theta) v_i \quad \text{and} \quad \sum_{i=1}^K \lambda_{t,i}(\theta) = 1.$$

We recall that $u_j(\theta, X_t) = \langle g(\theta, X_t), v_j \rangle$ denotes the value of the objective function on the vertex $v_j \in \mathcal{W}$, and $u_{I_t}(\theta, X_t) = \min_{j \in [K]} u_j(\theta, X_t)$ so $v_{I_t} = w_t^*(\theta)$. We start by proving the following lemma:

Lemma 7. If there exists $\varepsilon > 0$ such that $u_{I_t} + \varepsilon \leq u_j$ for any $j \in \{1, \dots, K\} \setminus \{I_t\}$, then,

$$\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 \leq \frac{2n\alpha_t}{\varepsilon} \max_{k \in [K]} \|v_k\|_1.$$

Proof. On the one hand, $\tilde{w}_t(\theta)$ is by definition the solution to the problem:

$$\tilde{w}_t(\theta) = \underset{w \in \mathcal{W}}{\text{argmin}} \langle g(\theta, X_t), w \rangle - \alpha_t \sum_{i=1}^d \ln(b_i - A_i w), \quad (26)$$

which is uniquely determined by strong convexity of the objective. We know by (Boyd, 2004, page 566) that:

$$\langle g(\theta, X_t), \tilde{w}_t(\theta) \rangle - \inf_{w \in \mathcal{W}_1} \langle g(\theta, X_t), w \rangle \leq n\alpha_t,$$

that is

$$\langle g(\theta, X_t), \tilde{w}_t(\theta) \rangle \leq u_{I_t} + n\alpha_t. \quad (27)$$

On the other hand, by assumption we know that for any $j \neq I_t$, $u_j \geq u_{I_t} + \varepsilon$ so:

$$\langle g(\theta, X_t), \tilde{w}_t(\theta) \rangle = \left\langle g(\theta, X_t), \sum_{j=1}^K \lambda_{t,j}(\theta) v_j \right\rangle \quad (28)$$

$$= \lambda_{t,I_t}(\theta) u_{I_t} + \sum_{j \neq I_t} \lambda_{t,j}(\theta) u_j$$

$$\geq \lambda_{t,I_t}(\theta) u_{I_t} + \sum_{j \neq I_t} \lambda_{t,j}(\theta) (u_{I_t} + \varepsilon)$$

$$= \lambda_{t,I_t}(\theta) u_{I_t} + (u_{I_t} + \varepsilon) (1 - \lambda_{t,I_t}(\theta)). \quad (29)$$

Combining (27) and (29) yields:

$$\lambda_{t, I_t}(\theta) \geq 1 - \frac{n\alpha_t}{\varepsilon}, \quad (30)$$

and it follows from (30) that:

$$\begin{aligned} \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1 &= \|v_{I_t} - \tilde{w}_t(\theta)\|_1 = \left\| \sum_{j=1}^K (\mathbb{1}_{\{j=I_t\}} - \lambda_{t,j}(\theta)) v_j \right\|_1 \\ &\leq [(1 - \lambda_{t, I_t}(\theta)) + \sum_{j \neq I_t} \lambda_{t,j}(\theta)] \max_{k \in [K]} \|v_k\|_1 \\ &\leq \frac{2n\alpha_t}{\varepsilon} \max_{k \in [K]} \|v_k\|_1. \end{aligned}$$

□

The result follows from combining **H2** and Lemma 7 according to the same lines of computation as in the proof of Lemma 2. □

G.1.3 OTHER LEMMAS

Lemma 8. *Assume that for any $t \in [T]$, \tilde{f}_t is K_t -Lipschitz almost-surely and that the sequence of steps $(\eta_t)_{t \geq 1}$ is non-increasing. Denote $\tilde{\mathfrak{R}}_T^d = \sum_{t \in [T]} \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t)$. Then it holds almost surely that:*

$$\mathbb{E}[\tilde{\mathfrak{R}}_T^d] \leq \frac{1}{2\eta_T} \left(D_\Theta^2 + 2D_\Theta \sum_{t=1}^T \|\vartheta_{t+1} - \vartheta_t\| \right) + \sum_{t=1}^T \frac{\eta_t}{2} K_t^2,$$

where the expectation is taken over the sequence (u_1, \dots, u_T) in Algorithm 2.

Proof. In what follows, we write \mathbb{E}_{u_t} the expectation under the distribution $\text{Unif}([0, 1])$ of u_t , and $\mathbb{E}_{u_{1:T}}$ the expectation under the joint distribution $\text{Unif}([0, 1])^{\otimes T}$ of (u_1, \dots, u_T) . For any $t \in [T]$, we have:

$$\begin{aligned} \sum_{t=1}^T \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) &= \sum_{t=1}^T \int_{u=0}^1 \langle \nabla \tilde{f}_t(\vartheta_t + u(\theta_t - \vartheta_t)), \theta_t - \vartheta_t \rangle du \\ &= \sum_{t=1}^T \langle \mathbb{E}_{u_t} [\nabla \tilde{f}_t(\vartheta_t + u_t(\theta_t - \vartheta_t))], \theta_t - \vartheta_t \rangle \\ &= \sum_{t=1}^T \langle \mathbb{E}_{u_t} [\tilde{\nabla}_t(u_t)], \theta_t - \vartheta_t \rangle, \end{aligned} \quad (31)$$

where $\tilde{\nabla}_t(u_t) = \nabla \tilde{f}_t(\vartheta_t + u_t(\theta_t - \vartheta_t))$ as in Algorithm 2. To control (31) we note that by definition of θ_{t+1} in Algorithm 2,

$$\begin{aligned} \mathbb{E}_{u_t} [\|\theta_{t+1} - \vartheta_t\|^2] &= \mathbb{E}_{u_t} \left[\left\| \Pi_\Theta[\theta_t - \eta_t \tilde{\nabla}_t(u_t)] - \vartheta_t \right\|^2 \right] \\ &\leq \mathbb{E}_{u_t} \left[\left\| \theta_t - \vartheta_t - \eta_t \tilde{\nabla}_t(u_t) \right\|^2 \right] \\ &\leq \|\theta_t - \vartheta_t\|^2 - 2\eta_t \langle \mathbb{E}_{u_t} [\tilde{\nabla}_t(u_t)], \theta_t - \vartheta_t \rangle + \eta_t^2 K_t^2 \end{aligned} \quad (32)$$

where we have used that $\|\tilde{\nabla}_t(u_t)\|^2 \leq K_t^2$ because \tilde{f}_t is K_t -Lipschitz. Re-arranging (32) and taking the expectation over (u_1, \dots, u_T) yields:

$$\mathbb{E}_{u_{1:T}} \left[\langle \tilde{\nabla}_t(u_t), \theta_t - \vartheta_t \rangle \right] \leq \mathbb{E}_{u_{1:T}} \left[\frac{1}{2\eta_t} (\|\theta_t - \vartheta_t\|^2 - \|\theta_{t+1} - \vartheta_t\|^2) + \frac{\eta_t K_t^2}{2} \right], \quad (33)$$

Now notice that, for any $t \in \{1, \dots, T-1\}$,

$$\begin{aligned} \|\theta_{t+1} - \vartheta_{t+1}\|^2 &= \|\theta_{t+1} - \vartheta_t\|^2 + \|\theta_{t+1} - \vartheta_{t+1}\|^2 - \|\theta_{t+1} - \vartheta_t\|^2 \\ &= \|\theta_{t+1} - \vartheta_t\|^2 \\ &\quad + (\|\theta_{t+1} - \vartheta_{t+1}\| + \|\theta_{t+1} - \vartheta_t\|)(\|\theta_{t+1} - \vartheta_{t+1}\| - \|\theta_{t+1} - \vartheta_t\|) \\ &\leq \|\theta_{t+1} - \vartheta_t\|^2 + 2D_\Theta \|\vartheta_{t+1} - \vartheta_t\|, \end{aligned} \quad (34)$$

where we have used **H1**-(i) and the reversed triangular inequality in the last line. It follows that:

$$-\|\theta_{t+1} - \vartheta_t\|^2 \leq -\|\theta_{t+1} - \vartheta_{t+1}\|^2 + 2D_\Theta \|\vartheta_{t+1} - \vartheta_t\| \quad (35)$$

Plugging (35) into (33) then gives:

$$\begin{aligned} \mathbb{E}_{u_{1:T}} \left[\left\langle \tilde{\nabla}_t(u_t), \theta_t - \vartheta_t \right\rangle \right] &\leq \mathbb{E}_{u_{1:T}} \left[\frac{1}{2\eta_t} \left(\|\theta_t - \vartheta_t\|^2 - \|\theta_{t+1} - \vartheta_{t+1}\|^2 \right) \right. \\ &\quad \left. + \frac{1}{\eta_t} D_\Theta \|\vartheta_{t+1} - \vartheta_t\| + \frac{\eta_t K_t^2}{2} \right]. \end{aligned}$$

Thus, summing over T and plugging the resulting sum in (31) gives:

$$\begin{aligned} \mathbb{E}_{u_{1:T}} \left[\sum_{t=1}^T \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) \right] &\leq \mathbb{E}_{u_{1:T}} \left[\sum_{t=1}^T \frac{\|\theta_t - \vartheta_t\|^2}{2} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \right. \\ &\quad \left. + \sum_{t=1}^T \frac{1}{\eta_t} D_\Theta \|\vartheta_{t+1} - \vartheta_t\| + \sum_{t=1}^T \frac{\eta_t K_t^2}{2} \right]. \\ &\leq \frac{D_\Theta^2}{2} \sum_{t=1}^T \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) + \frac{D_\Theta}{\eta_T} \sum_{t=1}^T \|\vartheta_{t+1} - \vartheta_t\| + \sum_{t=1}^T \frac{\eta_t K_t^2}{2}. \end{aligned}$$

where we used that for all t , $\|\theta_t - \vartheta_t\| \leq D_\Theta$ by **H1**-(i) and that for all $t \in [T]$, $\eta_{t-1} \geq \eta_t \geq \eta_T$. Then, we have by telescoping the first sum:

$$\mathbb{E}_{u_{1:T}} \left[\sum_{t=1}^T \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) \right] \leq \frac{D_\Theta^2 + 2D_\Theta}{2\eta_T} \sum_{t=1}^T \|\vartheta_{t+1} - \vartheta_t\| + \sum_{t=1}^T \frac{\eta_t K_t^2}{2}.$$

□

G.2 PROOF OF THEOREM 1.

Theorem 1. Assume **H1**, **H2** and having access to an ξ -approximate oracle \mathbf{O}_ξ adapted to $\{\sum_{i=1}^t \tilde{f}_i - \langle \sigma_t, \cdot \rangle\}_{t \in [T]}$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-FTPL (Algorithm 1) instantiated with learning step $\eta > 0$ and regularization coefficients $\alpha_t = \alpha > 0$ for any t . Then:

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^s] = \tilde{\mathcal{O}} \left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha n \right),$$

where \mathbb{E} denotes the expectation on both data and the intrinsic randomness of DF-FTPL and $\tilde{\mathcal{O}}$ contains polynomial dependency in $\ln(1/\alpha)$, $\ln(\ln(d))$.

Furthermore, taking $\eta \propto m^{1/4} T^{-3/4} n^{-1/2}$ and $\alpha \propto m^{3/4} n^{1/2} T^{-1/4}$ yields:

$$T^{-1} \mathbb{E}[\mathfrak{R}_T^s] = \tilde{\mathcal{O}} \left(m^{3/4} n^{3/2} T^{-1/4} + \xi \right).$$

Proof. For the sake of conciseness we use the notation \mathbb{E}_t to denote $\mathbb{E}[\cdot \mid \mathcal{H}_{t-1}]$. In what follows, we consider, for any θ , the intermediary regret: $R_T^s(\theta) := \sum_{t=1}^T F_t(\theta_t) - F_t(\theta)$, where $F_t : \theta \mapsto$

$\mathbb{E}[f_t(\theta) \mid \mathcal{H}_{t-1}]$. We have, for any θ , the following decomposition of the static regret

$$R_T^s(\theta) = \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta_t) - \tilde{w}_t(\theta_t) \rangle] \quad (36)$$

$$+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) \rangle] - \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] \quad (37)$$

$$+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] - \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle] \quad (38)$$

Then, taking the sup over Θ in Equations (36) and (38), and defining

$$\text{Reg}_T^S := \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) \rangle] - \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle],$$

in Equation (37) yields for any θ :

$$R_T^s(\theta) \leq \text{Reg}_T^S + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) - \tilde{w}_t(\theta) \rangle]$$

Using the fact that $\|\cdot\|_\infty$ is dual to $\|\cdot\|_1$:

$$R_T^s(\theta) \leq \text{Reg}_T^S + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\|\bar{g}_t(X_t)\|_\infty \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1]$$

Since for any $t \in [T]$, $\|\bar{g}_t(X_t)\|_\infty \leq \|\bar{g}_t(X_t)\|_2 \leq D_Z$ by **H1**-(ii),

$$R_T^s(\theta) \leq \text{Reg}_T^S + 2D_Z \sum_{t \in [T]} \mathbb{E}_t \sup_{\theta \in \Theta} \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1$$

First, notice that, by Lemma 6, and because for any t , $\alpha_t = \alpha$ we have for all t :

$$\mathbb{E}[\|\tilde{w}_t(\theta) - w_t^*(\theta)\|_1 \mid \mathcal{H}_{t-1}] \leq \alpha + 2n \max_{k \in [K]} \|v_k\|_1 \alpha (\alpha \ln d)^{\beta-1} \min\left\{\ln\left(\frac{2}{\alpha}\right), (1-\beta)^{-1}\right\}. \quad (39)$$

Second, taking the expectation (denoted by \mathbb{E}) over $(\sigma_t)_{t \in [T]}$ and X_1, \dots, X_T on both sides gives:

$$\mathbb{E}[R_T^s(\theta)] \leq \mathbb{E}[\text{Reg}_T^S] + 2D_Z T \left(\alpha + 2n \max_{k \in [K]} \|v_k\|_1 \alpha (\alpha \ln d)^{\beta-1} \min\left\{\ln\left(\frac{2}{\alpha}\right), (1-\beta)^{-1}\right\} \right). \quad (40)$$

We now bound the first term in the right-hand-side of Equation (40). By the definition of conditional expectation, we have:

$$\mathbb{E}[\text{Reg}_T^S] = \mathbb{E}_{X_1 \dots X_T} \mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\sum_{t \in [T]} \langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) \rangle - \langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle \right]$$

and then,

$$\mathbb{E}[\text{Reg}_T^S] \leq \mathbb{E}_{X_1 \dots X_T} \mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\sum_{t \in [T]} \tilde{f}_t(\theta_t) - \inf_{\theta \in \Theta} \sum_{t \in [T]} \tilde{f}_t(\theta) \right].$$

One recognises (up to a factor T) the left-hand side of Proposition 4 on the loss sequence $(\tilde{f}_t)_{t \in [T]}$. Furthermore, we can use this proposition as, given our choice of \mathcal{R} , for any $t \in [T]$, \tilde{f}_t is L -Lipschitz with $L = 5D_Z G(4\alpha)^{-1}$ almost surely by Lemma 1 (with $\chi = 0$). We then have:

$$\mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\frac{1}{T} \sum_{t \in [T]} \tilde{f}_t(\theta_t) - \inf_{\theta \in \Theta} \frac{1}{T} \sum_{t \in [T]} \tilde{f}_t(\theta) \right] \leq \mathcal{O}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi\right). \quad (41)$$

Dividing Equation (40) by T , and plugging Equation (41) gives: for all $\theta \in \Theta$:

$$\begin{aligned} T^{-1}\mathbb{E}[\mathcal{R}_T^s(\theta)] &\leq \mathcal{O}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi\right) \\ &\quad + 2D_{\mathcal{Z}} T \left(\alpha + 2n \max_{k \in [K]} \|v_k\|_1 \alpha (\alpha \ln d)^{\beta-1} \min\left\{\ln\left(\frac{2}{\alpha}\right), (1-\beta)^{-1}\right\}\right) \\ &= \tilde{\mathcal{O}}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha n\right). \end{aligned}$$

Finally remark that, by the definition of the conditional expectation (thus of F_t), and because θ_t is \mathcal{F}_{t-1} -measurable, we have for any θ :

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T^s(\theta)] &= \mathbb{E}\left[\sum_{t=1}^T f_t(\theta_t) - f_t(\theta)\right] \\ &= \mathbb{E}\left[\sum_{t=1}^T f_t(\theta_t)\right] - \mathbb{E}[f_t(\theta)] \end{aligned}$$

Then taking the infimum over θ yields:

$$T^{-1}\mathbb{E}[\mathfrak{R}_T^s] \leq \tilde{\mathcal{O}}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha n\right).$$

This concludes the proof. The second equation consists in simply plugging the proposed value of η, α in this bound. \square

G.3 PROOF OF THEOREM 2

Theorem 2. Assume **H1**, **H2**, access to a ξ -approximate oracle adapted to $\{\tilde{f}_t\}_{t \in [T]}$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-OGD (Algorithm 2) instantiated with the non-increasing sequence $(\eta_t)_{t \in [T]}$ and regularization coefficients $(\alpha_t)_{t \in [T]}$. Then:

$$T^{-1}\mathbb{E}[\mathfrak{R}_T^d] = \tilde{\mathcal{O}}\left(\mathbb{E}\left[\frac{1 + P_T}{T\eta_T} + \frac{1}{T} \sum_{t \in [T]} \frac{\eta_t}{\alpha_t^2} + n\alpha_t\right] + \xi\right).$$

where $P_T = \sum_{t=1}^{T-1} \|\vartheta_{t+1} - \vartheta_t\|$, \mathbb{E} denotes the expectation on both data and the intrinsic randomness of DF-OGD and $\tilde{\mathcal{O}}$ contains polynomial dependency in $\ln(1/\alpha), \ln(\ln(d))$.

Furthermore, assume $(t^{-1}(1 + P_t))_{t \geq 1}$ is non-increasing almost surely with $P_t = \sum_{s=1}^{t-1} \|\vartheta_{s+1} - \vartheta_s\|$. Then, using $\alpha_t \propto n^{-1/2} t^{-1/4} (1 + P_t)^{1/4}$ and $\eta_t \propto n^{-1/2} t^{-3/4} (1 + P_t)^{3/4}$ for any $t \in [T]$ leads to:

$$T^{-1}\mathbb{E}[\mathfrak{R}_T^d] = \tilde{\mathcal{O}}\left(\mathbb{E}\left[\sqrt{n}(1 + P_T)^{1/4} T^{-1/4}\right] + \xi\right).$$

Proof. In this proof, for the sake of conciseness, we use the notation \mathbb{E}_t to denote $\mathbb{E}[\cdot | \mathcal{H}_{t-1}]$. Observe that the dynamic regret can be decomposed as follows:

$$\mathfrak{R}_T^d = \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta_t) - \tilde{w}_t(\theta_t) \rangle] + \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) - \tilde{w}_t(\vartheta_t) \rangle] \quad (42)$$

$$+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\vartheta_t) \rangle] - \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] \quad (43)$$

$$+ \sum_{t \in [T]} \left[\inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] - \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle] \right] \quad (44)$$

First, we remark that for any t , given $\vartheta_t = \mathbf{O}(\tilde{f}_t)$, we control (43) by Jensen:

$$\begin{aligned} \mathbb{E}_t[\langle \tilde{g}_t(X_t), \tilde{w}_t(\vartheta_t) \rangle] - \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \tilde{g}_t(X_t), \tilde{w}_t(\theta) \rangle] &\leq \mathbb{E}_t \left[\tilde{f}_t(\vartheta_t) - \inf_{\theta \in \Theta} \tilde{f}_t(\theta) \right] \\ &\leq \xi. \end{aligned}$$

Then, taking the sup over Θ for each summand of the first sum of (42) (which is valid as θ_t is \mathcal{F}_{t-1} -measurable), defining by $\text{Reg}_T := \sum_{t \in [T]} \mathbb{E}_t[\langle \tilde{g}_t(X_t), \tilde{w}_t(\theta_t) - \tilde{w}_t(\vartheta_t) \rangle]$, and noticing for (44) that $\inf_{\theta \in \Theta} \mathbb{E}_t[f_t(\theta)] - \inf_{\theta \in \Theta} \mathbb{E}_t[\tilde{f}_t(\theta)] \leq |\inf_{\theta \in \Theta} \mathbb{E}_t[f_t(\theta)] - \inf_{\theta \in \Theta} \mathbb{E}_t[\tilde{f}_t(\theta)]| \leq \sup_{\theta \in \Theta} |\mathbb{E}_t[f_t(\theta) - \tilde{f}_t(\theta)]| \leq \sup_{\theta \in \Theta} \mathbb{E}_t[|f_t(\theta) - \tilde{f}_t(\theta)|]$ leads to:

$$\mathfrak{R}_T^d \leq \text{Reg}_T + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[|\langle \tilde{g}_t(X_t), w_t^*(\theta) - \tilde{w}_t(\theta) \rangle|] + \xi T$$

Using the fact that $\|\cdot\|_\infty$ is dual to $\|\cdot\|_1$:

$$\mathfrak{R}_T^d \leq \text{Reg}_T + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\|\tilde{g}_t(X_t)\|_\infty \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1] + \xi T$$

Since for any $t \in [T]$, $\|\tilde{g}_t(X_t)\|_\infty \leq \|\tilde{g}_t(X_t)\|_2 \leq D_Z$ almost surely by **H1**-(ii),

$$\mathfrak{R}_T^d \leq \text{Reg}_T + 2D_Z \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1] + \xi T \quad (45)$$

First, by Lemma 6, we know have for any t :

$$\sup_{\theta \in \Theta} \mathbb{E}_t[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1] \leq \alpha_t + 2n \max_{k \in [K]} \|v_k\|_1 \alpha_t (\alpha_t \ln d)^{\beta-1} \min \left\{ \ln \left(\frac{2}{\alpha_t} \right), (1-\beta)^{-1} \right\}.$$

Thus combining the two last equations and taking the expectation (denoted by \mathbb{E}) over $(u_t)_{t \in [T]}$ and X_1, \dots, X_T on both sides gives:

$$\begin{aligned} \mathbb{E}[\mathfrak{R}_T^d] &\leq \mathbb{E}[\text{Reg}_T] \\ &+ 2D_Z \sum_{t \in [T]} \left(\alpha_t + 2n \max_{k \in [K]} \|v_k\|_1 \alpha_t (\alpha_t \ln d)^{\beta-1} \min \left\{ \ln \left(\frac{2}{\alpha_t} \right), (1-\beta)^{-1} \right\} \right) + \xi T. \quad (46) \end{aligned}$$

Now, to bound the first term in Equation (46) note that the definition of conditional expectation implies:

$$\mathbb{E}[\text{Reg}_T] = \mathbb{E}_{X_1 \dots X_T} \mathbb{E}_{u_1, \dots, u_T} \left[\sum_{t \in [T]} \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) \right].$$

One recognizes the definition of $\tilde{\mathfrak{R}}_T^d$ of Lemma 8. Since \tilde{f}_t is K_t -Lipschitz with $K_t = \frac{GD_Z \sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty}{\alpha_t \lambda_{\min}(AA^\top)}$ almost surely for any t by Lemma 5, we deduce from Lemma 8 that:

$$\mathbb{E}_{u_1, \dots, u_T} \left[\sum_{t \in [T]} \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) \right] \leq \frac{D_\Theta^2 - D_\Theta P_T}{2\eta_T} + \sum_{t \in [T]} \frac{\eta_t}{2} \left(\frac{GD_Z \sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty^2}{\alpha_t \lambda_{\min}(AA^\top)} \right)^2. \quad (47)$$

Plugging Equation (47) into Equation (46) and then dividing by $T > 0$ on both sides gives:

$$\begin{aligned}
T^{-1}\mathbb{E}[\mathfrak{R}_T^d] &\leq \frac{D_\Theta^2 + D_\Theta\mathbb{E}[P_T]}{2T\eta_T} + \sum_{t \in [T]} \frac{\eta_t}{2} \left(\frac{GD_Z \sup_{w \in \mathcal{W}_1} \|b - Aw\|_\infty^2}{\alpha_t \lambda_{\min}(AA^\top)} \right)^2 \\
&\quad + 2D_Z \sum_{t \in [T]} \left(\alpha_t + 2n \max_{k \in [K]} \|v_k\|_1 \alpha_t (\alpha_t \ln d)^{\beta-1} \min \left\{ \ln \left(\frac{2}{\alpha_t} \right), (1-\beta)^{-1} \right\} \right) + \xi.
\end{aligned} \tag{48}$$

Rewriting Equation (48) using $\tilde{\mathcal{O}}$ concludes the proof.

Concerning the second bound, We know by Equation (48) that the regret of Algorithm 2 satisfies the following bound:

$$\mathbb{E}[\mathfrak{R}_T^d] = \tilde{\mathcal{O}} \left(\mathbb{E} \left[\frac{1+P_T}{\eta_T} + \sum_{t \in [T]} \frac{\eta_t}{\alpha_t^2} + n \sum_{t \in [T]} \alpha_t + \xi T \right] \right).$$

With $\alpha_t \propto n^{-1/2} t^{-1/4} (1+P_t)^{1/4}$ and $\eta_t \propto \alpha_t t^{-1/2} (1+P_t)^{1/2} \propto n^{-1/2} t^{-3/4} (1+P_t)^{3/4}$, the sequence $(\eta_t)_{t \in [T]}$ is non-increasing because so is $(t^{-1}(1+P_t))_{t \in [T]}$ by assumption. Moreover, the first term satisfies:

$$\frac{1+P_T}{\eta_T} = \mathcal{O} \left(n^{1/2} \frac{T^{3/4}}{(1+P_T)^{3/4}} (1+P_T) \right) \quad \text{so} \quad \frac{1+P_T}{T\eta_T} = \mathcal{O} \left(n^{1/2} \left(\frac{1+P_T}{T} \right)^{1/4} \right).$$

For the second term,

$$\sum_{t=1}^T \frac{\eta_t}{\alpha_t^2} = \sum_{t=1}^T \frac{\sqrt{1+P_t}}{\sqrt{t}\alpha_t} = \mathcal{O} \left(n^{1/2} \sum_{t=1}^T \left(\frac{1+P_t}{t} \right)^{1/4} \right).$$

Moreover, since $1+P_t \leq 1+P_T$ for any $t \in [T]$, we have:

$$\sum_{t=1}^T \frac{\eta_t}{\alpha_t^2} = \mathcal{O} \left(n^{1/2} (1+P_T)^{1/4} \sum_{t=1}^T \left(\frac{1}{t} \right)^{1/4} \right) = \mathcal{O} \left(n^{1/2} (1+P_T)^{1/4} T^{3/4} \right),$$

where we used in the last line that $\sum_{t=1}^T t^{-\gamma} = \mathcal{O}(T^{1-\gamma})$ for $\gamma \in (0, 1)$. Dividing by T again gives a rate of $\mathcal{O} \left(\left(\frac{1+P_T}{T} \right)^{1/4} \right)$ for the second term. Finally, we obtain by the same reasoning that

$$n \sum_{t \in [T]} \alpha_t \leq n^{1/2} (1+P_T)^{1/4} \sum_{t \in [T]} t^{-1/4} = \mathcal{O} \left(n^{1/2} (1+P_T)^{1/4} T^{3/4} \right),$$

and dividing by $T > 0$ yields the desired rate. \square

G.4 PROOF OF PROPOSITION 2

Proposition 2. Assume **H1**, **H2** and having access to an ξ -approximate optimization oracle \mathbf{O}_ξ adapted to $\left\{ \sum_{i=1}^t \tilde{f}_i - \langle \sigma_t, \cdot \rangle \right\}_{t \in [T]}$. Fix $\mathcal{W} = \mathcal{W}_0$, $\mathcal{R} = \mathcal{R}_0$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-FTPL (Algorithm 1) instantiated with learning step $\eta > 0$ and regularization coefficients $\alpha_t = \alpha > 0$ for any t . Then:

$$T^{-1}\mathbb{E}[\mathfrak{R}_T^s] \leq \tilde{\mathcal{O}} \left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha \ln(d) \right),$$

where \mathbb{E} denotes the expectation on both data and the intrinsic randomness of DF-FTPL and $\tilde{\mathcal{O}}$ contains polynomial dependency in $\ln(1/\alpha)$, $\ln(\ln(d))$.

Proof. For the sake of conciseness we use the notation \mathbb{E}_t to denote $\mathbb{E}[\cdot \mid \mathcal{H}_{t-1}]$. In what follows, we consider, for any θ , the intermediary regret: $R_T^s(\theta) := \sum_{t=1}^T F_t(\theta_t) - F_t(\theta)$, where $F_t : \theta \mapsto \mathbb{E}[f_t(\theta) \mid \mathcal{H}_{t-1}]$. We have, for any θ , the following decomposition of the static regret

$$R_T^s(\theta) = \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta_t) - \tilde{w}_t(\theta_t) \rangle] \quad (49)$$

$$+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) \rangle] - \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] \quad (50)$$

$$+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] - \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle] \quad (51)$$

Then, taking absolute values, using the triangular inequality and taking the sup over Θ in Equations (49) and (51), as well as defining

$$\text{Reg}_T^S := \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) \rangle] - \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle],$$

in Equation (50) yields for any θ :

$$R_T^s(\theta) \leq \text{Reg}_T^S + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) - \tilde{w}_t(\theta) \rangle]$$

Using the fact that $\|\cdot\|_\infty$ is dual to $\|\cdot\|_1$:

$$R_T^s(\theta) \leq \text{Reg}_T^S + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\|\bar{g}_t(X_t)\|_\infty \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1]$$

Since for any $t \in [T]$, $\|\bar{g}_t(X_t)\|_\infty \leq \|\bar{g}_t(X_t)\|_2 \leq D_Z$ by **H1**-(ii),

$$R_T^s(\theta) \leq \text{Reg}_T^S + 2D_Z \sum_{t \in [T]} \mathbb{E}_t \sup_{\theta \in \Theta} \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1$$

First, notice that, by Lemma 2, and because for any t , $\alpha_t = \alpha$ we have for all t :

$$\mathbb{E}[\|\tilde{w}_t(\theta) - w_t^*(\theta)\|_1 \mid \mathcal{H}_{t-1}] \leq \alpha + 2(\alpha \ln(d))^\beta C_0 \min\left\{\ln\left(\frac{2}{\alpha}\right), (1-\beta)^{-1}\right\}, \quad (52)$$

Second, taking the expectation (denoted by \mathbb{E}) over $(\sigma_t)_{t \in [T]}$ and X_1, \dots, X_T on both sides gives:

$$\mathbb{E}[R_T^s(\theta)] \leq \mathbb{E}[\text{Reg}_T^S] + 2D_Z T \left(\alpha + 2(\alpha \ln(d))^\beta C_0 \min\left\{\ln\left(\frac{2}{\alpha}\right), (1-\beta)^{-1}\right\} \right). \quad (53)$$

To bound the first term in the right-hand-side of Equation (53). by the definition of conditional expectation, we have:

$$\mathbb{E}[\text{Reg}_T^S] = \mathbb{E}_{X_1 \dots X_T} \mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\sum_{t \in [T]} \langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) \rangle - \langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle \right]$$

Then,

$$\mathbb{E}[\text{Reg}_T^S] \leq \mathbb{E}_{X_1 \dots X_T} \mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\sum_{t \in [T]} \tilde{f}_t(\theta_t) - \inf_{\theta \in \Theta} \sum_{t \in [T]} \tilde{f}_t(\theta) \right].$$

One recognises (up to a factor T) the left-hand side of Proposition 4 on the loss sequence $(\tilde{f}_t)_{t \in [T]}$. Furthermore, we can use this proposition as, given our choice of \mathcal{R} , for any $t \in [T]$, \tilde{f}_t is L -Lipschitz with $L = 5D_Z G(4\alpha)^{-1}$ almost surely by Lemma 1 (with $\chi = 0$). We then have:

$$\mathbb{E}_{\sigma_1, \dots, \sigma_T} \left[\frac{1}{T} \sum_{t \in [T]} \tilde{f}_t(\theta_t) - \inf_{\theta \in \Theta} \frac{1}{T} \sum_{t \in [T]} \tilde{f}_t(\theta) \right] \leq \mathcal{O}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi\right). \quad (54)$$

Dividing Equation (53) by T , and plugging Equation (54) gives: for all $\theta \in \Theta$:

$$\begin{aligned} T^{-1}\mathbb{E}[\mathcal{R}_T^s(\theta)] &\leq \mathcal{O}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi\right) \\ &\quad + 2D_{\mathcal{Z}}\alpha\left(1 + 2\ln(d)C_0\left(\ln\left(\frac{2}{\alpha}\right) + (1-\beta)\ln^2(\alpha\ln(d))\right)\right), \\ &= \tilde{\mathcal{O}}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha\ln(d)\right). \end{aligned}$$

Finally remark that, by the definition of the conditional expectation (thus of F_t), and because θ_t is \mathcal{F}_{t-1} -measurable, we have for any θ :

$$\begin{aligned} \mathbb{E}[\mathcal{R}_T^s(\theta)] &= \mathbb{E}\left[\sum_{t=1}^T f_t(\theta_t) - f_t(\theta)\right] \\ &= \mathbb{E}\left[\sum_{t=1}^T f_t(\theta_t)\right] - \mathbb{E}[f_t(\theta)] \end{aligned}$$

Then taking the infimum over θ yields:

$$T^{-1}\mathbb{E}[\mathfrak{R}_T^s] \leq \tilde{\mathcal{O}}\left(\eta m^2 D \frac{1}{\alpha^2} + \frac{mD}{\eta T} + \xi + \alpha\ln(d)\right).$$

This concludes the proof. The second equation consists in simply plugging the proposed value of η, α in this bound. \square

G.5 PROOF OF PROPOSITION 3

Proposition 3. Assume **H1**, **H2**, and having access to an ξ -approximate optimization oracle \mathbf{O}_ξ adapted to $\{\tilde{f}_i\}_{i \in [T]}$. Fix $\mathcal{W} = \mathcal{W}_0$, $\mathcal{R} = \mathcal{R}_0$. Let $\{\theta_t\}_{t \in [T]}$ be the output of DF-OGD (Algorithm 2) instantiated with the non-increasing sequence $(\eta_t)_{t \in [T]}$ and regularization coefficients $(\alpha_t)_{t \in [T]}$. Then:

$$\begin{aligned} T^{-1}\mathbb{E}[\mathfrak{R}_T^d] &\leq \frac{D_\Theta^2 + D_\Theta \mathbb{E}[P_T]}{2T\eta_T} + \sum_{t \in [T]} \frac{25D_{\mathcal{Z}}G\eta_t}{32T\alpha_t^2} \\ &\quad + 2D_{\mathcal{Z}} \sum_{t \in [T]} \alpha_t [1 + 2\ln(d)C_0\{\ln(2\alpha_t^{-1}) + (1-\beta)\ln^2(\alpha_t \ln d)\}] + \xi, \end{aligned}$$

where $P_T = \sum_{t=1}^{T-1} \|\vartheta_{t+1} - \vartheta_t\|$.

Proof. In this proof, for the sake of conciseness, we denote by \mathbb{E}_t the conditional expectation $\mathbb{E}[\cdot | \mathcal{H}_{t-1}]$.

Observe that the dynamic regret can be decomposed as follows:

$$\mathfrak{R}_T^d = \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta_t) - \tilde{w}_t(\theta_t) \rangle] + \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta_t) - \tilde{w}_t(\vartheta_t) \rangle] \quad (55)$$

$$+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\vartheta_t) \rangle] - \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] \quad (56)$$

$$+ \sum_{t \in [T]} \left[\inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), \tilde{w}_t(\theta) \rangle] - \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle] \right] \quad (57)$$

First, we remark that for any t , given $\vartheta_t = \mathbf{O}(\tilde{f}_t)$, we control (56) by Jensen:

$$\begin{aligned} \mathbb{E}_t[\langle \tilde{g}_t(X_t), \tilde{w}_t(\vartheta_t) \rangle] - \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \tilde{g}_t(X_t), \tilde{w}_t(\theta) \rangle] &\leq \mathbb{E}_t \left[\tilde{f}_t(\vartheta_t) - \inf_{\theta \in \Theta} \tilde{f}_t(\theta) \right] \\ &\leq \xi. \end{aligned}$$

Then, taking the sup over Θ for each summand of the first sum of (55) (which is valid as θ_t is \mathcal{F}_{t-1} -measurable), defining by $\text{Reg}_T := \sum_{t \in [T]} \mathbb{E}_t[\langle \tilde{g}_t(X_t), \tilde{w}_t(\theta_t) - \tilde{w}_t(\vartheta_t) \rangle]$, and noticing for (57) that $\inf_{\theta \in \Theta} \mathbb{E}_t[f_t(\theta)] - \inf_{\theta \in \Theta} \mathbb{E}_t[\tilde{f}_t(\theta)] \leq \sup_{\theta \in \Theta} \mathbb{E}_t[f_t(\theta) - \tilde{f}_t(\theta)]$ leads to:

$$\mathfrak{R}_T^d \leq \text{Reg}_T + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\langle \tilde{g}_t(X_t), w_t^*(\theta) - \tilde{w}_t(\theta) \rangle] + \xi T$$

Using the fact that $\|\cdot\|_\infty$ is dual to $\|\cdot\|_1$:

$$\mathfrak{R}_T^d \leq \text{Reg}_T + 2 \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\|\tilde{g}_t(X_t)\|_\infty \|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1] + \xi T$$

Since for any $t \in [T]$, $\|\tilde{g}_t(X_t)\|_\infty \leq \|\tilde{g}_t(X_t)\|_2 \leq D_Z$ almost surely by **H1**-(ii),

$$\mathfrak{R}_T^d \leq \text{Reg}_T + 2D_Z \sum_{t \in [T]} \sup_{\theta \in \Theta} \mathbb{E}_t[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1] + \xi T \quad (58)$$

First, notice that, by Lemma 2, for all t :

$$\sup_{\theta \in \Theta} \mathbb{E}_t[\|w_t^*(\theta) - \tilde{w}_t(\theta)\|_1] \leq \alpha_t \left(1 + 2 \ln(d) C_0 \left(\ln \left(\frac{2}{\alpha_t} \right) + (1 - \beta) \ln^2(\alpha_t \ln(d)) \right) \right). \quad (59)$$

Second, taking the expectation (denoted by \mathbb{E}) over $(u_t)_{t \in [T]}$ and X_1, \dots, X_T on both sides gives:

$$\begin{aligned} \mathbb{E}[\mathfrak{R}_T^d] &\leq \mathbb{E}[\text{Reg}_T] \\ &\quad + 2D_Z \sum_{t \in [T]} \alpha_t \left(1 + 2 \ln(d) C_0 \left(\ln \left(\frac{2}{\alpha_t} \right) + (1 - \beta) \ln^2(\alpha_t \ln(d)) \right) \right) + \xi T. \quad (60) \end{aligned}$$

To bound the first term in the right-hand side of Equation (60). Remark that, by the definition of conditional expectation:

$$\mathbb{E}[\text{Reg}_T] = \mathbb{E}_{X_1 \dots X_T} \mathbb{E}_{u_1, \dots, u_T} \left[\sum_{t \in [T]} \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) \right].$$

One recognises the definition of $\tilde{\mathfrak{R}}_T^d$ of Lemma 8. Thus, by this lemma, we know that for any $t \in [T]$, because \tilde{f}_t is K_t -Lipschitz with $K_t = 5D_Z G(4\alpha_t)^{-1}$ almost surely by Lemma 1, we have:

$$\mathbb{E}_{u_1, \dots, u_T} \left[\sum_{t \in [T]} \tilde{f}_t(\theta_t) - \tilde{f}_t(\vartheta_t) \right] \leq \frac{1}{2\eta_T} (D_\Theta^2 + D_\Theta P_T) + \sum_{t \in [T]} \frac{25D_Z^2 G^2 \eta_t}{32\alpha_t^2}. \quad (61)$$

Plugging Equation (61) into Equation (60) and then dividing by $T > 0$ on both sides gives the desired result. \square

G.6 PROOFS OF APPENDIX F

Proof of Proposition 5. The regret simply decomposes as follows:

$$\begin{aligned}
\mathfrak{R}_T^d &= \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \underline{w}_t(\theta_t) \rangle] - \sum_{t \in [T]} \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle] \\
&= \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), \underline{w}_t(\theta_t) - w_t^*(\theta_t) \rangle] \\
&+ \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta_t) \rangle] - \sum_{t \in [T]} \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle] \\
&\leq \kappa + \sum_{t \in [T]} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta_t) \rangle] - \sum_{t \in [T]} \inf_{\theta \in \Theta} \mathbb{E}_t[\langle \bar{g}_t(X_t), w_t^*(\theta) \rangle],
\end{aligned}$$

and the proof continues as in Theorem 2. □