

Exact-Form Regret for Gradient Descent, Mirror Descent and Follow-the-Regularized-Leader

Ashkan Soleymani
MIT
ashkanso@mit.edu

Gabriele Farina
MIT
gfarina@mit.edu

Patrick Jaillet
MIT
jaillet@mit.edu

Abstract

Online gradient descent is usually studied through external regret, where the learner competes with fixed alternatives. Recent work shows that first-order methods control much richer action-dependent deviations. We ask for a geometric characterization of the deviations with respect to which online gradient descent, mirror descent, and follow-the-regularized-leader (FTRL) achieve no regret. We identify exactness as the common principle. Here, exactness means that the relevant displacement field is generated by a scalar potential, or equivalently that the associated one-form is exact in the geometry used by the algorithm. This geometry depends on the algorithm. For gradient descent it is ordinary Euclidean geometry, for mirror descent it is the geometry induced by the regularizer, and for FTRL it is the cumulative dual state. Under mild regularity conditions, exactness yields sublinear regret, while nonzero circulation provides the complementary obstruction and leads to linear regret. This gives a unified geometric framework for understanding the deviation classes controlled by these algorithms and reveals that different first-order methods can control genuinely different classes of deviations. These deviation classes have direct consequences for learning, particularly in games. We study the equilibrium notions induced by exact-form deviations and introduce conservative correlated equilibrium, reflecting both the conservative geometry of the underlying displacement fields and the restricted family of deviations available to the players. We characterize its relation to correlated equilibrium, determine when the resulting equilibrium notions coincide and when they separate, and show how these relationships depend on the geometry and the learning algorithm. Overall, this work gives a unified geometric account of what first-order online learning algorithms are no-regret with respect to, beyond fixed comparators.

1 Introduction

Online gradient descent is classically understood as an external-regret minimizer [Zinkevich, 2003, Hazan, 2016, Orabona, 2019]. Against a sequence of linear losses, its iterates compete with every fixed comparator $\mathbf{x}' \in K$. This guarantee is the basic engine behind the standard convergence of no-regret play to coarse correlated equilibrium [Hart and Mas-Colell, 2000, Cesa-Bianchi and Lugosi, 2006]. From the viewpoint of swap regret and the notion of hindsight rationality, however, external regret is only the simplest notion out of a wide spectrum. External regret compares $\mathbf{x}_t$ with fixed alternatives, while a general deviation may replace $\mathbf{x}_t$ by a feasible point $\phi(\mathbf{x}_t) \in K$ that depends on the current action [Greenwald and Jafari, 2003, Blum and Mansour, 2007, Gordon et al., 2008].

Recent work makes clear that gradient dynamics control more than fixed comparators, although the guarantees appear in different forms. In first-order equilibrium theory, Ahunbay [2026] studies infinitesimal strategy modifications generated by continuous vector fields. In the coarse adversarial setting, the fields certified by projected gradient dynamics are conservative gradient fields with the appropriate tangency to the action set. In normal-form games, Ahunbay and Bichler [2025] specialize this picture to a tractable linear-quadratic family and obtain semicoarse correlated equilibria through LP-based guarantees. In parallel, Cai, Daskalakis, Luo, Wei, and Zheng [2026] showed that online gradient descent and mirror descent control deviations generated by proximal operators of weakly convex functions. As discussed by Cai et al. [2026] and

Ahunbay [2026], these lines of work have substantial overlap but are not comparable in general, reflecting the different deviation classes and guarantees they study.

In this paper, we study Φ -regret against actual feasible deviation maps $\phi: K \rightarrow K$. Writing $\mathbf{D}_\phi(\mathbf{x}) := \mathbf{x} - \phi(\mathbf{x})$, we show that online gradient descent has $O(\sqrt{T})$ regret against every sufficiently regular feasible deviation whose displacement is exact, $\mathbf{D}_\phi = \nabla\Psi$, and obtain a uniform bound over normalized families with common smoothness and potential-oscillation parameters. Thus the comparator is the actual endpoint $\phi(\mathbf{x}_t)$, rather than an infinitesimal modification of the current action. The underlying Euclidean exactness principle can also be obtained by specializing the more general first-order framework of Ahunbay [2026] to displacement fields generated by feasible endomorphisms. Working directly with endomorphisms, however, removes the tangent-projection machinery and allows the boundary geometry to be handled directly through feasibility, convexity, and normal cones. This yields tighter and more explicit finite-time bounds for feasible endomorphisms.

A simple example reveals the geometry behind this formulation. Consider an affine deviation $\phi(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$, whose displacement is $\mathbf{D}_\phi(\mathbf{x}) = (I - A)\mathbf{x} - \mathbf{b}$. When $I - A$ is symmetric, this displacement is the gradient of a quadratic potential. Thus the symmetric affine deviations appearing in Ahunbay and Bichler [2025], Cai et al. [2026] are finite maps with exact displacement fields. The example suggests that the relevant property is not linearity itself, but exactness of the displacement.

The same structure appears in proximal deviations. If $F = f + \iota_K$ is weakly convex, the Moreau-envelope identity gives $\mathbf{x} - \text{prox}_F(\mathbf{x}) = \nabla M_F(\mathbf{x})$, where M_F is the Moreau envelope of F . Proximal deviations therefore also have exact displacement fields. From this viewpoint, proximality is one mechanism for producing exactness rather than the underlying geometric property itself. We show that the smooth exact-form class strictly contains the weakly convex proximal class, while every smooth Euclidean exact-form deviation admits a sufficiently small interpolation with the identity that is proximal.

As it turns out, this characterization is tight, as witnessed by circulations. Exact one-forms have zero circulation around closed loops, whereas nonzero circulation can be converted into a bounded cyclic adversary with linear regret. For feasible finite maps, this gives an exactness-versus-circulation criterion under the regularity and step-size conditions developed below. Together with the positive result, this identifies the geometric boundary of the finite deviation class controlled by online gradient descent.

We next extend this finite-map viewpoint to mirror descent. The relevant notion of exactness is determined by the mirror geometry, through the one-form $\mathbf{D}_\phi(\mathbf{x})^\top d\nabla R(\mathbf{x})$. We first develop the theory in the Legendre interior, where the geometry is most transparent, and recover Bregman proximal deviations as a special case. We then remove the Legendre assumption and obtain corresponding upper and lower bounds for standard constrained mirror descent with boundary iterates. This gives a finite-map exactness-versus-circulation characterization in mirror geometry as well.

Bregman proximality leads to a separate representation problem. Identity interpolation always transfers proximal regret to the radial closure of the Bregman proximal class, and this radial class is contained in the mirror-exact class. In Euclidean geometry the two classes coincide after radial closure, but this can fail for a general mirror regularizer. In the Legendre regime, we characterize proximal interpolation through the convexity of $R^* - \alpha\Psi$ and the curvature of the reconstructed proximal generator. A squared ℓ_p example shows that no positive interpolation may exist even when the dual potential is smooth with globally Lipschitz gradient.

We also treat FTRL separately. Mirror descent and FTRL coincide for constant-step linearized losses in the Legendre interior, but their finite deviation classes can separate once boundary effects are present. We develop the corresponding exactness-versus-circulation theory in the cumulative dual state of FTRL and show that FTRL exactness implies mirror exactness, while the converse can fail. A powered ℓ_p construction gives a deviation with $O(\sqrt{T})$ regret under mirror descent and linear regret under FTRL.

Finally, we study the equilibrium notions induced by these finite deviation classes. We call the equilibrium notion associated with Euclidean exact-form deviations *conservative correlated equilibrium* (ConCE). The name reflects the fact that these deviations are generated by conservative, or exact, displacement fields, and that players are correspondingly restricted to this conservative family of deviations rather than arbitrary measurable remappings. Although smooth exact-form maps strictly contain weakly convex proximal maps,

their unrestricted equilibrium constraints, in convex games with compact action sets and jointly continuous losses, coincide when the approximation error is zero, so $\text{ConCE} = \text{PCE}$. We further show that this common notion coincides with full correlated equilibrium on finite support and, under continuous losses, when every player’s action space is an interval. In contrast, with two-dimensional individual action spaces we construct an uncountably supported distribution with absolutely continuous marginals for which $\text{CE} \subsetneq \text{ConCE} = \text{PCE}$.

Contributions relative to prior work. Relative to Ahunbay [2026], our Euclidean analysis specializes the first-order geometric framework to finite feasible endomorphisms and studies the resulting finite Φ -regret and equilibrium classes. This specialization yields sharper finite-time bounds by exploiting the additional convex geometry available for endomorphisms and avoiding the tangent projections and auxiliary geometric quantities needed in the general infinitesimal formulation. Relative to Ahunbay and Bichler [2025], our exact-form class extends the tractable linear-quadratic family to nonlinear finite deviations while retaining the symmetric affine maps as a special case. Relative to Cai et al. [2026], we identify exactness as the broader finite-map mechanism containing weakly convex proximal deviations, prove strict map-level containment, and characterize the role of identity interpolation in relating the two classes.

Beyond the Euclidean setting, we develop the finite-map theory for standard constrained mirror descent without Legendre duality, separate mirror exactness from Bregman proximal representability, and show that the Euclidean interpolation phenomenon can fail in nonquadratic mirror geometry. We also develop a separate cumulative-dual theory for FTRL and exhibit a boundary separation between the finite deviation classes controlled by FTRL and mirror descent. These algorithmic results are complemented by the interpolation and equilibrium results above, which determine when distinctions between deviation classes persist at the level of correlated-equilibrium constraints.

2 Related Work

The notion of *regret* is one of the central concepts in the theory of online algorithms. In its simplest incarnation, regret measures competitiveness against the best *fixed* comparator in hindsight. In other words, letting K denote the set of possible decisions, $\mathbf{x}_t$ be the decision at time t , and $\langle \mathbf{g}_t, \cdot \rangle$ be the linear loss incurred at time t by the decision maker, regret takes the form

$$\sup_{\hat{\mathbf{x}} \in K} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \hat{\mathbf{x}} \rangle.$$

As is standard, we refer to this notion of regret as *external* regret. Since the seminal work of Zinkevich [2003], it is known that online gradient descent keeps external regret sublinear. Generalizations of this result to online mirror descent and follow-the-regularized-leader (FTRL) are by now standard [Shalev-Shwartz, 2012, Hazan, 2016, Orabona, 2019], in particular optimistic algorithms and their variants for the *self-play* setting in games [Rakhlin and Sridharan, 2013a,b, Syrgkanis et al., 2015, Chen and Peng, 2020, Daskalakis et al., 2021, Farina et al., 2022a, Piliouras et al., 2022a, Soleymani et al., 2025a,b].

Over the past few decades researchers have been interested in obtaining guarantees for stronger notions of online competitiveness than external regret. Within the constellation of alternative notions (which includes concepts such as interval regret, dynamic regret, policy regret, switching regret, among others), the notion of Φ -regret [Greenwald and Jafari, 2003] is the most related to our work.

Φ -regret is a generalization of external regret indexed by a set Φ of *deviation functions* $\phi : K \rightarrow K$, comparing the sequence of online choices $\mathbf{x}_t$ to the transformed sequence $\phi(\mathbf{x}_t)$ across all $\phi \in \Phi$:

$$\sup_{\phi \in \Phi} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle.$$

External regret is then understood as the special instantiation of Φ -regret in which the set Φ contains all *constant* functions. Partially motivated by game-theoretic applications and the celebrated wealth of

connections between regret and equilibrium notions [Foster and Vohra, 1997, Hart and Mas-Colell, 2000, Cesa-Bianchi and Lugosi, 2006], a long line of work has been concerned with characterizing the largest set of deviations Φ that admits efficient algorithms guaranteeing sublinear Φ -regret (e.g., [Gordon et al., 2008, Farina et al., 2022b, Morrill et al., 2021, Fujii, 2025, Peng and Rubinstein, 2024, Dagan et al., 2024] and references therein). Among the most general results in that space, it is now understood that Φ -regret minimization with respect to all linear or low-degree polynomial endomorphic transformations of K can be achieved efficiently for any convex and compact set K given via efficient oracle access [Daskalakis et al., 2025, Zhang et al., 2025]. These algorithms are quite involved, and make use of complex separation and fixed-point machinery to deal with the geometry of endomorphisms.

While the common wisdom had been for a long time that simple algorithms such as online gradient descent only minimize external regret—and the more complex algorithms mentioned above are thus necessary to target more challenging variants of Φ -regret—recent papers have challenged this view. Specifically, Ahunbay [2026], Ahunbay and Bichler [2025], Cai et al. [2024, 2026] demonstrated that online gradient descent keeps under control far more than just external regret: for example, it minimizes regret against all *symmetric* linear endomorphic transformations, and all proximal transformations of K . These results invite intriguing questions on exactly what notions of Φ -regret online gradient descent minimizes, for at least two reasons: (i) notions of Φ -regret induced by proximal transformations are not currently known to be achievable in other ways (for example, by low-degree polynomial transformations); and (ii) the extreme simplicity of online gradient descent compared to more advanced Φ -regret techniques makes a compelling case that the former should be used whenever possible. The main goal of the present paper is to develop a finite feasible-map theory that explains a broad class of Φ -regret guarantees given by these algorithms, including online gradient descent, mirror descent and FTRL.

Our Euclidean results are closely related to the first-order framework of Ahunbay [2026]. After specializing that framework to displacement fields generated by feasible endomorphisms, its qualitative conservativity and circulation conclusions yield the corresponding exactness picture for projected gradient descent. We focus instead on finite Φ -regret against the actual endpoints $\phi(\mathbf{x}_t)$ and on the equilibrium notions induced by these finite deviations. The endomorphism structure also permits a more direct quantitative analysis. Feasibility, convexity, and normal-cone geometry remove the tangent projections and auxiliary geometric quantities of the general infinitesimal formulation, leading to sharper explicit finite-time bounds. Relative to Ahunbay and Bichler [2025], our exact-form class extends the tractable linear–quadratic family to nonlinear feasible deviations while retaining symmetric affine endomorphisms as a special case.

Our results also address an open question raised by Cai et al. [2026]. After showing that gradient descent minimizes proximal regret, they ask for a tight characterization of the regret guarantees of gradient descent and note that the class controlled by the algorithm may be strictly broader than the proximal class. We characterize this larger finite-map class geometrically through exactness and circulation. For online gradient descent, smooth exact-form deviations strictly contain weakly convex proximal deviations, while the circulation converse identifies the corresponding obstruction to sublinear regret. We develop the analogous characterization for standard constrained mirror descent with differentiable strongly convex regularizers, where the relevant object is the mirror one-form $\mathbf{D}_\phi(\mathbf{x})^\top d\nabla R(\mathbf{x})$. Under our regularity assumptions, this identifies the full geometric deviation class controlled by mirror descent and goes beyond the Bregman proximal class studied by Cai et al. [2026]. We also study FTRL, which is not treated in Cai et al. [2026], and show that its controlled deviation class can differ from that of mirror descent once boundary effects are present. We further show that proximal representability is a separate question. In Euclidean geometry, identity interpolation recovers the smooth exact-form class, while in general mirror geometry the radial Bregman proximal class can be strictly smaller than the mirror-exact class. Finally, we derive the equilibrium consequences of these finite deviation classes. To our knowledge, the resulting equilibrium results are new, including the equality of ConCE and PCE under the unrestricted definitions, their coincidence with CE on finite support and on interval action spaces under continuous losses, and their strict separation from CE on a continuous support in higher dimension.

The gradient-equilibrium framework of Angelopoulos et al. [2025] is closely connected to our potential-based viewpoint. In the unconstrained setting, their gradient-equilibrium condition can be recovered from the linear

potentials $\Psi_{\mathbf{v}}(\mathbf{x}) = \langle \mathbf{v}, \mathbf{x} \rangle$, whose exact displacements are the constant vectors $\nabla \Psi_{\mathbf{v}} = \mathbf{v}$. Our framework replaces these constant directions by nonlinear finite displacements $\mathbf{D}_{\phi}(\mathbf{x}) = \nabla \Psi(\mathbf{x})$ generated by feasible endomorphisms $\phi: K \rightarrow K$. In this sense, both approaches are driven by the same telescoping mechanism, while our setting uses it to study finite Φ -regret and the equilibrium constraints induced by nonlinear action-dependent deviations. Their constrained formulation incorporates feasibility through boundary or regularization residuals, whereas ours builds feasibility directly into the deviation map.

3 Setting and Notation

Let $K \subseteq \mathbb{R}^d$ be a nonempty closed convex action set. Throughout the Euclidean sections, gradients, projections, and normal cones are taken in the ambient space $\mathbb{R}^d$. When a statement is intrinsically lower-dimensional, such as the curl criterion on $\text{relint}(K)$, derivatives and differential forms are taken in affine coordinates on $\text{aff}(K)$. The mirror section introduces its own explicit affine-hull convention. Vectors are written in bold, e.g. $\mathbf{x}, \mathbf{g}$, and coordinates use brackets, e.g. $\mathbf{x}[j]$. We write $[T] := \{1, \dots, T\}$. We use $\|\cdot\|$ as shorthand for the Euclidean norm $\|\cdot\|_2$, and we denote projection onto K by Π_K . Define the ambient normal cone at any $\mathbf{x} \in K$ by

$$N_K(\mathbf{x}) := \{\mathbf{n} \in \mathbb{R}^d : \langle \mathbf{n}, \mathbf{y} - \mathbf{x} \rangle \leq 0 \text{ for every } \mathbf{y} \in K\}.$$

A learner plays $\mathbf{x}_t \in K$ and observes a vector loss $\mathbf{g}_t \in \mathbb{R}^d$. For a deviation $\phi: K \rightarrow K$, define

$$\mathbf{D}_{\phi}(\mathbf{x}) := \mathbf{x} - \phi(\mathbf{x}), \quad \text{Reg}_T(\phi) := \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_{\phi}(\mathbf{x}_t) \rangle$$

More generally, for a displacement field $\mathbf{D}: K \rightarrow \mathbb{R}^d$, define $\text{Reg}_T(\mathbf{D}) := \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}(\mathbf{x}_t) \rangle$. When losses $\ell_t: K \rightarrow \mathbb{R}$ are convex and differentiable, taking $\mathbf{g}_t = \nabla \ell_t(\mathbf{x}_t)$ gives

$$\ell_t(\mathbf{x}_t) - \ell_t(\phi(\mathbf{x}_t)) \leq \langle \nabla \ell_t(\mathbf{x}_t), \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle.$$

Thus every linearized regret bound below immediately implies the corresponding convex-loss deviation-regret bound.

Definition 3.1 (Exact-form deviation). *A deviation $\phi: K \rightarrow K$ is Euclidean exact-form if there exists a continuously differentiable potential Ψ , defined on a neighborhood of K , such that*

$$\mathbf{x} - \phi(\mathbf{x}) = \nabla \Psi(\mathbf{x}) \quad \text{for every } \mathbf{x} \in K.$$

When a result refers to a *smooth exact-form deviation*, the gradient of its potential is locally Lipschitz on a neighborhood of K . This is stronger than the bare C^1 regularity in the definition above and is the regularity used later to show that scaled exact deviations are proximal; see Theorem 6.2.

We denote the Euclidean exact-form class by

$$\Phi_{\text{exact}} := \{\phi: K \rightarrow K : \phi \text{ is a Euclidean exact-form deviation}\}.$$

Throughout the paper, the subscript identifies the deviation class. Thus Φ_{prox} denotes the Euclidean proximal class, $\Phi_{\text{exact}}^{\text{sm}}$ denotes the smooth exact-form class, and the superscript R is reserved for a regularizer-dependent mirror class. When a regularizer-dependent class is considered on more than one input domain, the domain is displayed explicitly in parentheses.

Finally, we use the symbol osc_K to denote the range of a function on the set K , that is,

$$\text{osc}_K(\Psi) := \sup_{\mathbf{x} \in K} \Psi(\mathbf{x}) - \inf_{\mathbf{x} \in K} \Psi(\mathbf{x}).$$

Bounds for a fixed deviation may depend on the parameters of that deviation and therefore need not yield a uniform Φ -regret bound. We therefore use the following normalized class whenever a supremum over deviations is taken.

Definition 3.2 (Normalized exact-form class). *For constants $B, L \geq 0$, let $\Phi_{\text{exact}}(B, L)$ consist of all feasible maps $\phi : K \rightarrow K$ for which there is a differentiable potential Ψ_ϕ , defined on a neighborhood of K , satisfying*

$$\mathbf{x} - \phi(\mathbf{x}) = \nabla \Psi_\phi(\mathbf{x}), \quad \text{osc}_K(\Psi_\phi) \leq B,$$

and whose gradient satisfies

$$\|\nabla \Psi_\phi(\mathbf{y}) - \nabla \Psi_\phi(\mathbf{x})\| \leq L \|\mathbf{y} - \mathbf{x}\| \quad \text{for all } \mathbf{x}, \mathbf{y} \in K.$$

Proofs not given in the main text are found in the appendices. Appendix A records the convex-analytic preliminaries; Appendix B proves the scaled proximal representation; Appendix C proves the finite interpolation theorem; and Appendix D contains the remaining proofs, organized in the same order as the main text.

4 Online Gradient Descent Minimizes Exact-form Regret

Online gradient descent with constant step size $\eta > 0$ is

$$\mathbf{x}_{t+1} := \Pi_K(\mathbf{x}_t - \eta \mathbf{g}_t). \quad (\text{OGD})$$

We prove a telescoping regret bound for every smooth feasible exact-form deviation ϕ . Under bounded gradients and bounded potential oscillation, this gives $\text{Reg}_T(\phi) = \mathcal{O}(\sqrt{T})$. Exactness of the displacement field, rather than proximality, is the geometric property that makes the bound telescope on Ψ . Compared with the unconstrained case (online gradient descent over an unbounded domain), the only additional issue is the projection step. The projection residual lies in the normal cone, and the feasibility condition $\phi(K) \subseteq K$ ensures that its contribution has the favorable sign.

Theorem 4.1. *Let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online gradient descent (OGD). Let $\phi : K \rightarrow K$ be an exact-form deviation $\mathbf{D}_\phi(\mathbf{x}) = \mathbf{x} - \phi(\mathbf{x}) = \nabla \Psi(\mathbf{x})$ where Ψ is L -smooth with respect to the Euclidean norm. Then, for every sequence $\mathbf{g}_1, \dots, \mathbf{g}_T$,*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{\Psi(\mathbf{x}_1) - \Psi(\mathbf{x}_{T+1})}{\eta} + \frac{3L\eta}{2} \sum_{t=1}^T \|\mathbf{g}_t\|^2. \quad (1)$$

The proof of Theorem 4.1 is given in Appendix D.1.

Corollary 4.2. *Let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online gradient descent (OGD) as in Theorem 4.1. If K is compact, $\|\mathbf{g}_t\| \leq G$, and $\text{osc}_K(\Psi) \leq B$ for constants $B, L, G > 0$, then*

$$\text{Reg}_T(\phi) \leq \frac{B}{\eta} + \frac{3L\eta G^2 T}{2}.$$

In particular, choosing $\eta = \sqrt{2B/(3LG^2T)}$ gives $\text{Reg}_T(\phi) \leq G\sqrt{6LBT} = \mathcal{O}(\sqrt{T})$.

Corollary 4.3. *Suppose $\|\mathbf{g}_t\| \leq G$ and fix $B, L > 0$. Then OGD with $\eta = \sqrt{2B/(3LG^2T)}$ satisfies*

$$\sup_{\phi \in \Phi_{\text{exact}}(B, L)} \text{Reg}_T(\phi) \leq G\sqrt{6LBT}.$$

Thus the algorithm controls the whole normalized class $\Phi_{\text{exact}}(B, L)$ with one learning rate.

Remark 4.4. *The proof of Theorem 4.1 and Corollary 4.2 uses the standard quadratic upper bound for Ψ to control the potential-difference term and L -Lipschitz continuity of $\nabla \Psi$ to control the projection-residual term. Consequently, the normalized class in Definition 3.2 imposes the full Lipschitz-gradient condition. On an unconstrained domain, the projection residual vanishes, so the quadratic upper bound alone is sufficient.*

5 Proximal Deviations are Exact-Form

This section proves that the exact-form deviation class includes the proximal deviations of Cai et al. [2026]. The key identity is the Moreau-envelope formula [Moreau, 1965]

$$\mathbf{x} - \text{prox}_F(\mathbf{x}) = \nabla M_F(\mathbf{x}). \quad (2)$$

These proximal-regret deviations are generated by ρ -weakly convex functions with $0 \leq \rho < 1$, so we use the corresponding weakly-convex Moreau-envelope identity rather than only its convex special case.

More precisely, let $F := f + \iota_K$, where ι_K is the indicator of K . For $0 \leq \rho < 1$, define

$$\text{prox}_F(\mathbf{x}) := \arg \min_{\mathbf{y} \in \mathbb{R}^d} \left\{ F(\mathbf{y}) + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2 \right\}, \quad M_F(\mathbf{x}) := \min_{\mathbf{y} \in \mathbb{R}^d} \left\{ F(\mathbf{y}) + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2 \right\}.$$

Because F includes ι_K , the minimization is over K .

Theorem 5.1. *Let $F := f + \iota_K$ be proper, lower semicontinuous, and ρ -weakly convex with $0 \leq \rho < 1$. Define $\phi_F(\mathbf{x}) := \text{prox}_F(\mathbf{x})$. Then $\phi_F : K \rightarrow K$ is feasible and Euclidean exact-form, with potential M_F , i.e., $\mathbf{x} - \phi_F(\mathbf{x}) = \nabla M_F(\mathbf{x})$, whose gradient is $L_F = (2 - \rho)/(1 - \rho)$ -Lipschitz. Consequently projected gradient descent controls every such proximal deviation by Theorem 4.1. In particular, if K is compact, $\|\mathbf{g}_t\| \leq G$, $\text{osc}_K(M_F) \leq B_F$, and $B_F, G > 0$, then choosing $\eta = \sqrt{2B_F/(3L_F G^2 T)}$ gives*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \text{prox}_F(\mathbf{x}_t) \rangle \leq \frac{B_F}{\eta} + \frac{3L_F \eta G^2 T}{2} = \mathcal{O}(\sqrt{T}).$$

If $B_F = 0$ or $G = 0$, the corresponding regret bound is trivial and the displayed optimizing step size need not be used.

The Moreau-envelope identity immediately recovers several standard deviation classes. The point of the following proposition is not to introduce new deviations, but to show that many of the deviations in Cai et al. [2026] have exact displacement fields, and therefore fall under our regret bound.

Proposition 5.2. *The exact-form class contains the following deviations.*

1. **External regret.** For $\mathbf{u} \in K$, take $F := \iota_{\{\mathbf{u}\}}$. Then $\text{prox}_F(\mathbf{x}) = \mathbf{u}$ and $\mathbf{x} - \mathbf{u} = \nabla \frac{1}{2} \|\mathbf{x} - \mathbf{u}\|^2$.
2. **Projection to a convex subset.** For a nonempty closed convex set $S \subseteq K$, take $F := \iota_S$. Then $\text{prox}_F(\mathbf{x}) = \Pi_S(\mathbf{x})$.
3. **Projection-based/no-move regret.** For a vector $\mathbf{v}$, take $F(\mathbf{y}) := \langle \mathbf{v}, \mathbf{y} \rangle + \iota_K(\mathbf{y})$. Then $\text{prox}_F(\mathbf{x}) = \Pi_K(\mathbf{x} - \mathbf{v})$.
4. **Interpolation deviations.** For $\mathbf{u} \in K$ and $\alpha \in (0, 1)$, the map $\phi(\mathbf{x}) := (1 - \alpha)\mathbf{x} + \alpha\mathbf{u}$ is the prox map of

$$F(\mathbf{y}) := \frac{\alpha}{2(1 - \alpha)} \|\mathbf{y} - \mathbf{u}\|^2 + \iota_K(\mathbf{y}).$$

The case $\alpha = 1$ is external regret.

5. **Local proximal deviations.** If $\|\mathbf{x} - \text{prox}_F(\mathbf{x})\| \leq \delta$ on K , then the deviation belongs to the local exact-form subclass $\|\nabla M_F(\mathbf{x})\| \leq \delta$.
6. **Symmetric affine swaps.** Let $\phi(\mathbf{x}) := A\mathbf{x} + \mathbf{b}$ map K into K . Let $T := \text{span}(K - K)$ be the direction space of $\text{aff}(K)$, fix any $\mathbf{a} \in K$, and assume $AT \subseteq T$ and the restriction of $I - A$ to T is self-adjoint. Then ϕ is exact-form in the ambient sense of our definition. On $\text{aff}(K)$, a potential is

$$\Psi(\mathbf{a} + \mathbf{h}) := \langle D_\phi(\mathbf{a}), \mathbf{h} \rangle + \frac{1}{2} \langle \mathbf{h}, (I - A)\mathbf{h} \rangle, \quad \mathbf{h} \in T. \quad (3)$$

An ambient extension of this potential is given in the proof.

Remark 5.3. *Theorem 5.1 and Proposition 5.2 cover the Euclidean exact-form deviation classes appearing in Cai et al. [2026] including point indicators for external regret, linear functions for projection-based/no-move regret, bounded-displacement local proximal deviations, interpolation deviations, and symmetric affine swaps. The symmetric affine case is even more direct in the exact-form view since the displacement of a symmetric affine map is already the gradient of a quadratic potential, without first representing the map as a proximal operator.*

6 The Exact-form Class Strictly Contains the Proximal Class

The inclusion from proximal deviations to exact-form deviations is strict, and the separation already appears in one dimension. Our proof strategy is to use the fact that every weakly convex proximal map with parameter $0 \leq \rho < 1$ is monotone, whereas feasible exact-form deviations need not be monotone.

Proposition 6.1. *There exists a smooth feasible exact-form deviation $\phi : [0, 1] \rightarrow [0, 1]$ that is not prox_F for any one-dimensional ρ -weakly convex function F with $0 \leq \rho < 1$.*

Proof. Let $k \geq 2$ be an even integer and let $a \in (0, 1]$ satisfy $a\pi k/2 > 1$. Define

$$D_\phi(x) := ax(1-x)\sin(2\pi kx), \quad \phi(x) := x - D_\phi(x).$$

Because $|D_\phi(x)| \leq ax(1-x) \leq \min\{x, 1-x\}$, we have $\phi(x) \in [0, 1]$. Also $D_\phi = \Psi'$, where

$$\Psi(x) := \int_0^x as(1-s)\sin(2\pi ks) \, ds,$$

so ϕ is feasible exact-form.

At $x = 1/2$, since k is even,

$$D'_\phi(1/2) = \frac{a\pi k}{2} > 1, \quad \phi'(1/2) = 1 - D'_\phi(1/2) < 0.$$

Thus ϕ is not monotone. By Lemma A.4, every one-dimensional weakly-convex proximal map with parameter $0 \leq \rho < 1$ is monotone. Hence ϕ cannot be represented as prox_F for any such F . $\square$

Thus while exact-form maps strictly contain weakly-convex proximal maps, we show that this strict inclusion does not carry over to the corresponding equilibrium sets, in particular when the approximation error is zero. The following scaling result is the key mechanism behind this equivalence. We show that scaled exact deviations are proximal.

Theorem 6.2. *Let $K \subseteq \mathbb{R}^d$ be compact and convex, and let $\phi(\mathbf{x}) := \mathbf{x} - \nabla\Psi(\mathbf{x})$ be a feasible exact-form deviation with $\Psi \in C_{\text{loc}}^{1,1}$ on a neighborhood of K . Then there exists $\alpha_0 > 0$ such that, for every $\alpha \in (0, \alpha_0]$, the interpolated map*

$$\phi_\alpha(\mathbf{x}) := (1-\alpha)\mathbf{x} + \alpha\phi(\mathbf{x}) = \mathbf{x} - \alpha\nabla\Psi(\mathbf{x})$$

coincides on K with $\text{prox}_{F_\alpha}(\mathbf{x})$ for some proper lower-semicontinuous $F_\alpha = f_\alpha + \iota_K$ that is ρ_α -weakly convex for some $\rho_\alpha < 1$.

The construction is closely related to the characterization of proximity operators in terms of convex potentials developed by Gribonval and Nikolova [2020]. The construction and its proof are given in Appendix B. It reconciles the strict inclusion

$$\Phi_{\text{prox}} \subsetneq \Phi_{\text{exact}}^{\text{sm}} \subseteq \Phi_{\text{exact}}.$$

Although a smooth exact-form deviation need not itself be proximal, every such deviation admits a sufficiently small interpolation with the identity that is proximal.

Write Φ_{prox} for the constrained weakly-convex proximal maps from Section 5, and write $\Phi_{\text{exact}}^{\text{sm}}$ for the feasible exact-form maps whose potentials are $C_{\text{loc}}^{1,1}$ near K . For any class $\mathcal{A}$ of feasible maps with a common domain, define its *identity-radial closure* by

$$\text{Rad}_I(\mathcal{A}) := \{\phi \text{ feasible} : \text{for some } \alpha \in (0, 1], (1 - \alpha)I + \alpha\phi \in \mathcal{A}\}.$$

Corollary 6.3. *For the radial closure of Euclidean proximal deviations, we have $\text{Rad}_I(\Phi_{\text{prox}}) = \Phi_{\text{exact}}^{\text{sm}}$.*

Although Corollary 6.3 does not identify the two map classes, it shows that their difference disappears after allowing interpolation with the identity. Indeed, every smooth exact-form deviation admits a sufficiently small positive interpolation that is proximal. Thus, while $\Phi_{\text{prox}} \subsetneq \Phi_{\text{exact}}^{\text{sm}}$, the strict separation between the two classes does not lead to a separation of the corresponding equilibrium sets when the approximation error is zero. The formal equilibrium statement is deferred to Section 10.

Remark 6.4. *For $D_{\phi_\alpha} = \alpha D_\phi$, linearized regret scales exactly as $\text{Reg}_T(\phi_\alpha) = \alpha \text{Reg}_T(\phi)$. Moreover, convexity of a loss ℓ gives $\ell(\mathbf{x}) - \ell(\phi_\alpha(\mathbf{x})) \geq \alpha(\ell(\mathbf{x}) - \ell(\phi(\mathbf{x})))$. Thus a profitable smooth exact deviation remains profitable under every positive interpolation with the identity, while Theorem 6.2 ensures that a sufficiently small such interpolation is proximal. The resulting equivalence of the equilibrium constraints, and in particular $\text{ConCE} = \text{PCE}$ for convex games with compact action sets and jointly continuous losses, is stated formally in Theorem 10.4.*

7 Curl Induces Linear Regret for Online Gradient Descent

The upper bound telescopes because exact one-forms integrate to zero around closed loops. We now show that nonzero circulation gives the corresponding obstruction for finite deviation maps. The construction is the finite-map analogue of the nonconservative vector-field lower bound of Ahunbay [2026, Proposition 5.4]; here we formulate it using a globally feasible endpoint map and the regret convention of our setting.

Definition 7.1 (Circulation). *Let $D : K \rightarrow \mathbb{R}^d$ be continuous and let $\gamma : [0, 1] \rightarrow K$ be a closed piecewise C^1 loop. The circulation of D along γ is*

$$\oint_\gamma D(\mathbf{x})^\top d\mathbf{x} := \int_0^1 \langle D(\gamma(s)), \gamma'(s) \rangle ds.$$

First, we show that exact fields have zero circulation, which is why this obstruction does not arise for exact-form regret.

Lemma 7.2. *If $D = \nabla\Psi$ on a neighborhood of K , then we have $\oint_\gamma D(\mathbf{x})^\top d\mathbf{x} = 0$ for every closed loop $\gamma \subseteq K$.*

Thus exactness rules out circulation. The next theorem proves the converse algorithmic statement. Nonzero circulation can be converted into a bounded cyclic adversary and hence into linear regret.

Theorem 7.3. *Let $D : K \rightarrow \mathbb{R}^d$ be continuous. Suppose there exists a closed piecewise C^1 loop $\gamma : [0, 1] \rightarrow K$ such that, after possibly reversing orientation,*

$$-\oint_\gamma D(\mathbf{x})^\top d\mathbf{x} := c_0 > 0.$$

Then for every gradient bound $G > 0$ and every sufficiently small $\eta > 0$, there exists a sequence of losses with $\|g_t\| \leq G$ such that online gradient descent, initialized at $\gamma(0)$, suffers

$$\sum_{t=1}^T \langle g_t, D(\mathbf{x}_t) \rangle \geq cT - O(1/\eta)$$

for a constant $c > 0$ depending only on D , γ , and G . In particular the regret is $\Omega(T)$. If the algorithm starts elsewhere in the same path-connected component of K , the adversary can first steer it to the loop along any fixed polygonal path; this changes only the $O(1/\eta)$ transient term.

The same lower-bound mechanism also extends to standard anytime learning-rate schedules.

Remark 7.4. *Theorem 7.3 is stated for a sufficiently small constant step size, which may depend on the witnessing loop. This is slightly different from the usual $O(\sqrt{T})$ upper bound, which uses the known-horizon choice $\eta \asymp T^{-1/2}$. The lower-bound construction can nevertheless be made anytime by using geometric epochs: during an epoch of length 2^k , use a constant step size of order $2^{-k/2}$. Within each epoch, the loop construction applies after its $O(1/\eta)$ steering transient. Across epochs, the positive main terms sum to $\Omega(T)$, whereas the transient terms sum only to $O(\sqrt{T})$. Thus the linear lower bound is not an artifact of knowing the horizon in advance.*

Together, the upper and lower bounds give the following Euclidean curl criterion for finite deviation maps. Curl-free fields admit a potential and, under the feasibility and smoothness conditions below, yield sublinear regret. By contrast, nonzero curl produces a local loop along which gradient descent can be forced to incur linear regret. The two directions involve slightly different learning-rate assumptions. The lower bound uses a sufficiently small constant step size that may depend on the witnessing loop, while the uniform upper bound over a normalized family is given in Corollary 4.3.

Corollary 7.5 (Euclidean curl criterion). *Let $E_K := \text{span}(K - K)$, let $U := \text{relint}(K)$ be open and simply connected in $\text{aff}(K)$, and let $\mathbf{D} : U \rightarrow E_K$ be C^1 . In any orthonormal affine coordinate system on E_K , $\partial_j \mathbf{D}_i = \partial_i \mathbf{D}_j$ throughout U if and only if $\mathbf{D} = \nabla_{E_K} \Psi$ on U for a C^2 potential Ψ . If a skew derivative is nonzero somewhere, online gradient descent can be forced to incur linear regret against $\mathbf{D}$ along a sufficiently small feasible loop. When the field is curl-free, assume that Ψ admits an ambient extension $\tilde{\Psi}$, defined on a neighborhood of K , that satisfies the smoothness and bounded-oscillation assumptions of Theorem 4.1 and Corollary 4.2. If $\tilde{\phi}(\mathbf{x}) := \mathbf{x} - \nabla \tilde{\Psi}(\mathbf{x})$ lies in K for every $\mathbf{x} \in K$, then projected GD has $O(\sqrt{T})$ regret against $\tilde{\phi}$.*

Remark 7.6. *Ahunbay [2026, Theorem 1.2] identifies two requirements for first-order vector-field tests: the field must be conservative and tangent to the action set. In the finite-map setting, the second requirement is already built into feasibility. Indeed, if $\phi(\mathbf{x}) \in K$, convexity of K ensures that the segment from $\mathbf{x}$ to $\phi(\mathbf{x})$ remains in K . Exactness is then the finite-map counterpart of the conservative-field condition. The circulation obstruction used above is therefore closely related to the nonconservative-field lower bound of Ahunbay [2026, Proposition 5.4], while here the deviation is specified by a feasible endpoint map.*

8 Online Mirror Descent and Exact One-forms

Our Euclidean result in Section 4 says that online gradient descent controls deviations whose displacement is a gradient. Mirror descent changes the coordinates in which the same statement is true. Its natural variable is not the primal point $\mathbf{x}$, but the dual point

$$\boldsymbol{\theta} := \nabla R(\mathbf{x}),$$

for a distance-generating function R . At an interior iterate, where the normal-cone term vanishes, mirror descent is the dual update

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \mathbf{g}_t.$$

Therefore, as expected we show that the same telescoping proof works whenever the primal displacement $\mathbf{D}_\phi(\mathbf{x})$ is a gradient as a function of the dual coordinate $\boldsymbol{\theta}$. This way, we are able to generalize our results to mirror-exact deviations.

To make the mirror generalization transparent, we first work in the Legendre interior. In this setting, ∇R provides a smooth change of coordinates between the primal and dual spaces, and interior mirror-descent iterates satisfy an additive update in the dual variable. This removes boundary normal-cone terms and makes the connection with the Euclidean argument particularly clear. The first three subsections use this setting to

introduce mirror-exact deviations, prove the corresponding dual-potential regret bound, and recover Bregman proximal maps of Cai et al. [2026] as a special case.

We then return to standard constrained mirror descent in Section 8.4 and remove the Legendre assumption. The iterates may now reach the boundary, and the mirror map need not provide a globally invertible dual coordinate system. A normal-cone residual therefore appears in the mirror-descent optimality condition. The underlying geometric condition, however, remains the same. Exactness of the mirror one-form gives the upper bound after the boundary residual is controlled, while nonzero circulation gives the corresponding lower bound. This shows that the exactness principle is not a consequence of Legendre duality, even though the Legendre setting provides the cleanest way to see it initially.

With the regret characterization in place, Section 8.5 returns to the identity interpolation of Theorem 6.2. We keep the original weakly convex Bregman proximal class and use the scaling identity to transfer proximal regret to its radial closure. The Euclidean conclusion does not extend in full generality. The radial proximal class is always contained in the mirror-exact class, but the containment can be strict. In the Legendre setting, convexity of $R^* - \alpha\Psi$ gives the first obstruction to proximal representation, and an actual weakly convex proximal generator requires an additional curvature condition. This separates the part of the Euclidean interpolation argument that is purely algebraic from the part that depends on the curvature of the regularizer.

There is a further distinction once one leaves the Legendre interior. For constant-step linearized losses, Legendre mirror descent has an additive dual recursion, and unrolling that recursion gives the corresponding follow-the-regularized-leader (FTRL) update. This relationship between mirror descent, FTRL, and closely related dual-averaging methods is well known in online optimization [McMahan, 2011, 2017, Xiao, 2010]. At the boundary, however, standard constrained mirror descent and FTRL need not generate the same dynamics. Mirror descent performs a one-step Bregman update and may carry a normal-cone residual, whereas FTRL retains the accumulated linear losses in a single regularized minimization. The distinction also affects which deviations the two algorithms control. In Section 9, we show that their exact deviation classes can differ and give a smooth powered ℓ_p example for which mirror descent has $O(\sqrt{T})$ regret while FTRL can incur linear regret. For this reason, we study FTRL separately in Section 9.

Throughout the first three subsections of this section, all convex-analytic objects are taken in the affine hull of K . After fixing an origin, we identify this affine hull with a Euclidean vector space E , equipped with the induced inner product. These three subsections work in the standard Legendre regime. Let $R : E \rightarrow (-\infty, +\infty]$ be a proper closed Legendre function satisfying

$$\overline{\text{dom } R} = K, \quad \mathcal{D} := \text{int}_E(\text{dom } R) = \text{relint}(K).$$

Assume that R is C^2 on $\mathcal{D}$ and that $\nabla^2 R(\mathbf{x})$ is positive definite there. Let

$$R^*(\boldsymbol{\theta}) := \sup_{\mathbf{x} \in E} \{\langle \boldsymbol{\theta}, \mathbf{x} \rangle - R(\mathbf{x})\}, \quad \Theta := \text{int}_E(\text{dom } R^*) = \nabla R(\mathcal{D}).$$

Legendre duality gives mutually inverse C^1 diffeomorphisms $\nabla R : \mathcal{D} \rightarrow \Theta$ and $\nabla R^* : \Theta \rightarrow \mathcal{D}$. In particular,

$$D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) = D_R(\mathbf{x}, \mathbf{x}') \quad \text{when } \boldsymbol{\theta} = \nabla R(\mathbf{x}), \quad \boldsymbol{\theta}' = \nabla R(\mathbf{x}').$$

We initialize mirror descent in $\mathcal{D}$. Under the standing Legendre assumptions, attained mirror-descent subproblems remain in the Legendre interior, so the iterates stay in $\mathcal{D}$. We write $\boldsymbol{\theta} := \nabla R(\mathbf{x})$.

We use the Bregman divergence generated by R , defined as

$$D_R(\mathbf{y}, \mathbf{x}) := R(\mathbf{y}) - R(\mathbf{x}) - \langle \nabla R(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

With this notation, online mirror descent takes the form

$$\mathbf{x}_{t+1} := \arg \min_{\mathbf{x} \in K} \{\eta \langle \mathbf{g}_t, \mathbf{x} \rangle + D_R(\mathbf{x}, \mathbf{x}_t)\}. \quad (4)$$

Definition 8.1 (Mirror-exact deviation). *A deviation $\phi : \mathcal{D} \rightarrow K$, with displacement $\mathbf{D}_\phi(\mathbf{x}) := \mathbf{x} - \phi(\mathbf{x})$, is R -mirror-exact if there is a differentiable potential $\Psi : \Theta \rightarrow \mathbb{R}$ such that*

$$\mathbf{D}_\phi(\mathbf{x}) = \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\mathbf{x})) \quad \forall \mathbf{x} \in \mathcal{D}.$$

Equivalently, the one-form $\mathbf{D}_\phi(\mathbf{x})^\top d\nabla R(\mathbf{x})$ is exact.

We write $\Phi_{\text{exact}}^R(\mathcal{D})$ for the class of mirror-exact deviations in the Legendre interior.

Remark 8.2. If $K = E$ and $R(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2$, then $\nabla R(\mathbf{x}) = \mathbf{x}$. Mirror-exactness becomes $\mathbf{D}_\phi(\mathbf{x}) = \nabla \Psi(\mathbf{x})$, exactly the Euclidean exact-form condition from Section 4 on the unconstrained action space.

The next proposition gives a useful primal-coordinate test for this condition. It shows that mirror-exactness is equivalent to ordinary exactness after multiplying the displacement by the mirror metric $\nabla^2 R(\mathbf{x})$.

Proposition 8.3. A field $\mathbf{D} : \mathcal{D} \rightarrow E$ is R -mirror-exact if and only if there exists a differentiable scalar potential $V : \mathcal{D} \rightarrow \mathbb{R}$ such that

$$\nabla V(\mathbf{x}) = \nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x}). \quad (5)$$

Equivalently, $\mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x}) = dV(\mathbf{x})$.

8.1 Examples of Mirror-exact Deviations

Definition 8.1 can look abstract because it is written in dual coordinates. The quickest way to build examples is simple: “choose a scalar potential $\Psi(\boldsymbol{\theta})$ in dual space, differentiate it to get $\nabla_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta})$, and then pull the resulting field back through $\boldsymbol{\theta} = \nabla R(\mathbf{x})$ to get $\phi(\mathbf{x}) = \mathbf{x} - \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\mathbf{x}))$.”

In Appendix D.5, we provide several examples of mirror-exact deviations. They show that the mirror-exact deviation family contains every fixed-comparator deviation. They also give nontrivial deviations for different instances of online mirror descent, such as, for the multiplicative weights update algorithm, multiplicative deviations with arbitrary temperature vector $\mathbf{a} \in \mathbb{R}^d$,

$$\phi_{\mathbf{a}}(\mathbf{x})[i] := \frac{\mathbf{x}[i] \exp(-\mathbf{a}[i])}{\sum_{j=1}^d \mathbf{x}[j] \exp(-\mathbf{a}[j])},$$

and geometric interpolation deviations anchored at $\mathbf{u} \in \text{relint}(\Delta_d)$, for $\lambda \in [0, 1]$,

$$\phi_{\lambda, \mathbf{u}}(\mathbf{x})[i] = \frac{\mathbf{x}[i]^{1-\lambda} \mathbf{u}[i]^\lambda}{\sum_{j=1}^d \mathbf{x}[j]^{1-\lambda} \mathbf{u}[j]^\lambda};$$

and there exist mirror-exact deviations that are not Euclidean-exact.

8.2 Online Mirror Descent Controls Mirror-exact Regret

We now prove no-regret results for mirror-exact deviations under online mirror descent in the Legendre interior. For a differentiable function A on dual space, we write its Bregman divergence as

$$D_A(\boldsymbol{\theta}', \boldsymbol{\theta}) := A(\boldsymbol{\theta}') - A(\boldsymbol{\theta}) - \langle \nabla A(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle,$$

and, for $L \geq 0$, we say that the potential Ψ is L -smooth relative to R^* on a set of ordered dual pairs if

$$D_\Psi(\boldsymbol{\theta}', \boldsymbol{\theta}) \leq LD_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) \quad (6)$$

for every such pair $\boldsymbol{\theta}', \boldsymbol{\theta} \in \Theta$. When $K = E$ and $R(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2$, this is the usual upper quadratic inequality. For a general mirror map, it measures the one-step Taylor error of Ψ in the same geometry used by mirror descent.

Theorem 8.4. Assume the standing Legendre-domain conditions and that R is m -strongly convex on $\mathcal{D}$. Let $\phi : \mathcal{D} \rightarrow K$ be feasible and R -mirror-exact with dual potential Ψ . Let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online mirror descent (4), initialized in $\mathcal{D}$, with iterates remaining in $\mathcal{D}$. Write

$$\boldsymbol{\theta}_t := \nabla R(\mathbf{x}_t), \quad \mathbf{D}_\phi(\mathbf{x}_t) = \nabla_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta}_t).$$

Assume that (6) holds for every consecutive pair $(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t)$, $t \in [T]$. A sufficient trajectory-independent condition is that it hold for every ordered pair in $\Theta \times \Theta$. Then

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \frac{\Psi(\boldsymbol{\theta}_1) - \Psi(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{L\eta}{m} \sum_{t=1}^T \|\mathbf{g}_t\|^2. \quad (7)$$

Consequently, suppose $\|\mathbf{g}_t\| \leq G$ and that, for the horizon-dependent choices of constant step size considered below, there is a constant B , independent of the horizon, such that

$$\max_{1 \leq t \leq T+1} \Psi(\boldsymbol{\theta}_t) - \min_{1 \leq t \leq T+1} \Psi(\boldsymbol{\theta}_t) \leq B,$$

Then choosing $\eta \asymp 1/\sqrt{T}$ gives $O(\sqrt{T})$ regret against ϕ . More precisely, when $B, L, G > 0$, the choice $\eta = \sqrt{mB/(LG^2T)}$ gives the bound $2G\sqrt{LBT}/m$. The global condition $\text{osc}_\Theta(\Psi) \leq B$ is a sufficient trajectory-independent assumption that avoids any dependence of the oscillation bound on the chosen step size.

The coefficient L/m in this bound uses the standard Legendre-regime assumption that every iterate remains in $\text{relint}(K)$. The relative normal cone then vanishes, so the proof works directly with the unconstrained dual update and introduces no additional projection term. Boundary iterates are treated separately in Section 8.4, where we remove the Legendre assumption and analyze standard constrained mirror descent directly.

The circulation converse is also developed in Section 8.4. In particular, Theorem 8.11 applies directly to standard constrained mirror descent, allows boundary iterates, and gives the corresponding lower bound without relying on the interior dual-coordinate recursion.

Remark 8.5. For the regret proof, it is enough that (6) hold along the realized trajectory, namely for $(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t)$ at each round. Imposing the condition on all of $\Theta \times \Theta$ gives a trajectory-independent sufficient condition, but may be considerably stronger when Θ is unbounded.

8.3 Bregman Proximal Regret is Mirror-exact

We next show that the Bregman proximal deviations of Cai et al. [2026] are mirror-exact. A related Hessian-weighted integrability observation appears in Ahunbay [2026, Appendix B.2.2], but there the argument requires a steep, smooth regularizer and does not recover the full Bregman proximal guarantee of Cai et al. [2026]. Here we first work in the Legendre interior and later, in Section 8.4, remove the Legendre assumption and allow boundary iterates. The proof is the mirror version of the Moreau-envelope identity (2).

Let $F := f + \iota_K$ be proper, lower semicontinuous, and ρ -weakly convex for some $\rho \geq 0$. The indicator is included so that all proximal outputs lie in K , giving pointwise feasibility automatically. For $\mathbf{x} \in \mathcal{D}$, define

$$\text{prox}_F^R(\mathbf{x}) := \arg \min_{\mathbf{y} \in E} \{F(\mathbf{y}) + D_R(\mathbf{y}, \mathbf{x})\}, \quad M_F^R(\mathbf{x}) := \min_{\mathbf{y} \in E} \{F(\mathbf{y}) + D_R(\mathbf{y}, \mathbf{x})\}.$$

Because F contains ι_K , the minimization is equivalently over $\mathbf{y} \in K$.

Theorem 8.6. Assume the standing Legendre-domain conditions. Let $F := f + \iota_K$ be proper and lower semicontinuous, assume that every Bregman proximal minimum above is attained, and suppose that prox_F^R is single-valued and continuous on $\mathcal{D}$. Then the Bregman proximal deviation $\phi_F^R : \mathcal{D} \rightarrow K$ defined by $\phi_F^R(\mathbf{x}) := \text{prox}_F^R(\mathbf{x})$ is feasible and R -mirror-exact. Equivalently,

$$(\mathbf{x} - \text{prox}_F^R(\mathbf{x}))^\top \text{d}\nabla R(\mathbf{x}) = \text{d}M_F^R(\mathbf{x}). \quad (8)$$

Thus Bregman proximal regret is a special case of mirror-exact regret.

We can now directly invoke Theorem 8.4 for this class of deviations in the following proposition.

Proposition 8.7. *Assume the standing Legendre-domain conditions. Assume, in the extended-valued sense on E , that R is m -strongly convex and that $F := f + \iota_K$ is proper, closed, and ρ -weakly convex, where $0 \leq \rho < m$ and $H := F + R$ is proper. Then the Bregman proximal subproblem has a unique solution on $\mathcal{D}$, and its solution map is continuous there. Let $\{\mathbf{x}_t\}_{t=1}^{T+1} \subset \mathcal{D}$ be generated by online mirror descent (4). Then these iterates control Bregman proximal regret with the bound*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \text{prox}_F^R(\mathbf{x}_t) \rangle \leq \frac{\Psi_F(\boldsymbol{\theta}_1) - \Psi_F(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{\eta}{m} \sum_{t=1}^T \|\mathbf{g}_t\|^2, \quad (9)$$

where $\Psi_F(\boldsymbol{\theta}_t) = M_F^R(\mathbf{x}_t)$ for all $t \in [T+1]$.

The key idea behind this result is to show that weakly convex Bregman proximal potentials are ($L = 1$)-relatively smooth with respect to R^* , as defined in (6), and hence Theorem 8.4 applies.

8.4 Full Mirror Characterization Beyond Legendre Regularizers

The preceding subsections use Legendre duality to regard $\boldsymbol{\theta} := \nabla R(\mathbf{x})$ as a global coordinate and to eliminate the normal cone from the mirror-descent optimality condition. The geometric condition itself does not depend on either property. It is exactness of the one-form

$$\mathbf{D}_\phi(\mathbf{x})^\top d\nabla R(\mathbf{x}).$$

We now keep the standard constrained mirror-descent update (4), allow boundary iterates, and prove matching upper and lower bounds without assuming that ∇R is onto or that R is Legendre.

Let K be a compact convex subset of a finite-dimensional normed space $(E, \|\cdot\|)$, with dual norm $\|\cdot\|_*$. Throughout this subsection, R is differentiable on a neighborhood of K and is m -strongly convex on K :

$$R(\mathbf{y}) \geq R(\mathbf{x}) + \langle \nabla R(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{m}{2} \|\mathbf{y} - \mathbf{x}\|^2 \quad (\mathbf{x}, \mathbf{y} \in K).$$

The gradient is uniformly continuous on K . Its modulus of continuity is

$$\omega_R(\delta) := \sup \{ \|\nabla R(\mathbf{x}) - \nabla R(\mathbf{y})\|_* : \mathbf{x}, \mathbf{y} \in K, \|\mathbf{x} - \mathbf{y}\| \leq \delta \}.$$

Thus $\omega_R(\delta) \rightarrow 0$ as $\delta \downarrow 0$.

The same exactness condition extends beyond the Legendre setting by defining the potential directly on a neighborhood of $\nabla R(K)$.

Definition 8.8. *A feasible deviation $\phi : K \rightarrow K$ is R -mirror-exact if there are an open convex set $\Omega \supseteq \nabla R(K)$ and a differentiable potential $\Psi : \Omega \rightarrow \mathbb{R}$ such that*

$$\mathbf{D}_\phi(\mathbf{x}) = \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\mathbf{x})) \quad (\mathbf{x} \in K).$$

We denote this class by $\Phi_{\text{exact}}^R(K)$. Under the Legendre assumptions of the preceding subsections, the definition reduces to Definition 8.1 on $\mathcal{D}$. If R is twice differentiable and $V = \Psi \circ \nabla R$, then

$$\nabla V(\mathbf{x}) = \nabla^2 R(\mathbf{x}) \mathbf{D}_\phi(\mathbf{x}),$$

so Definition 8.8 is equivalent to exactness of $\mathbf{D}_\phi(\mathbf{x})^\top d\nabla R(\mathbf{x})$.

The discrete potential argument requires a one-sided curvature bound. We assume that, for some $L \geq 0$,

$$D_\Psi(\nabla R(\mathbf{y}), \nabla R(\mathbf{x})) \leq LD_R(\mathbf{x}, \mathbf{y}) \quad (\mathbf{x}, \mathbf{y} \in K). \quad (10)$$

Equivalently,

$$\Psi(\nabla R(\mathbf{y})) \leq \Psi(\nabla R(\mathbf{x})) + \langle \mathbf{D}_\phi(\mathbf{x}), \nabla R(\mathbf{y}) - \nabla R(\mathbf{x}) \rangle + LD_R(\mathbf{x}, \mathbf{y}).$$

In the Legendre interior, $D_R(\mathbf{x}, \mathbf{y}) = D_{R^*}(\nabla R(\mathbf{y}), \nabla R(\mathbf{x}))$, and this condition becomes the relative-smoothness assumption in Theorem 8.4.

These ingredients yield the corresponding exact-form regret bound for standard constrained mirror descent.

Theorem 8.9. Assume R is differentiable on a neighborhood of K , so that $\mathcal{X} = K$, and let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be generated by mirror descent (4) with a constant step size $\eta > 0$. Let $\phi : K \rightarrow K$ be R -mirror-exact with potential Ψ . Suppose (10) holds and $\mathbf{D}_\phi$ is L_ϕ -Lipschitz on K . Then

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle &\leq \frac{\Psi(\nabla R(\mathbf{x}_1)) - \Psi(\nabla R(\mathbf{x}_{T+1}))}{\eta} + \frac{L\eta}{m} \sum_{t=1}^T \|\mathbf{g}_t\|_*^2 \\ &\quad + \frac{L_\phi}{m} \sum_{t=1}^T \|\mathbf{g}_t\|_* \left[\eta \|\mathbf{g}_t\|_* + \omega_R\left(\frac{\eta}{m} \|\mathbf{g}_t\|_*\right) \right]. \end{aligned}$$

The first line is the same potential estimate as in the Legendre interior. The second line controls the boundary normal-cone residual. Feasibility gives the correct sign at $\mathbf{x}_{t+1}$, and Lipschitz continuity of the displacement transfers that sign to the displacement evaluated at $\mathbf{x}_t$.

The preceding bound immediately gives sublinear regret under bounded gradients and bounded potential oscillation. When ∇R is Lipschitz on K , it also recovers the usual $O(\sqrt{T})$ rate.

Corollary 8.10. Assume the hypotheses of Theorem 8.9, $\|\mathbf{g}_t\|_* \leq G$, and $\text{osc}_{\nabla R(K)}(\Psi) \leq B$. For $\eta = T^{-1/2}$,

$$\frac{1}{T} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \frac{B}{\sqrt{T}} + \frac{(L + L_\phi)G^2}{m\sqrt{T}} + \frac{L_\phi G}{m} \omega_R\left(\frac{G}{m\sqrt{T}}\right).$$

Consequently, mirror descent has $o(T)$ regret against ϕ . If ∇R is M -Lipschitz on K , then

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \frac{B}{\eta} + C_{R,\phi} \eta G^2 T, \quad C_{R,\phi} := \frac{L + L_\phi}{m} + \frac{L_\phi M}{m^2}.$$

The optimized bound is $2G\sqrt{BC_{R,\phi}T}$ whenever $BC_{R,\phi} > 0$.

We next establish necessity of exactness at the geometric level and show how nonzero circulation obstructs sublinear regret. Let $\mathbf{D} : K \rightarrow E$ be continuous and let $\gamma : [0, 1] \rightarrow K$ be a closed piecewise C^1 loop such that $\nabla R \circ \gamma$ is piecewise C^1 . The mirror circulation of $\mathbf{D}$ along γ is

$$\oint_\gamma \mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x}) := \int_0^1 \left\langle \mathbf{D}(\gamma(s)), \frac{d}{ds} \nabla R(\gamma(s)) \right\rangle ds. \quad (11)$$

By the chain rule, every R -mirror-exact field has zero mirror circulation. Conversely, nonzero circulation along a feasible loop can be exploited by an adversary to make standard constrained mirror descent repeatedly traverse the loop and accumulate linear regret.

Theorem 8.11. Suppose that, after reversing the orientation of γ if necessary,

$$-\oint_\gamma \mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x}) := c_0 > 0.$$

For every $G > 0$ and every sufficiently small constant step size $\eta > 0$, there is a sequence with $\|\mathbf{g}_t\|_* \leq G$ such that standard constrained mirror descent, initialized at $\gamma(0)$, satisfies

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}(\mathbf{x}_t) \rangle \geq cT - O(1/\eta)$$

for a constant $c > 0$ depending only on $\mathbf{D}$, γ , R , and G .

The adversary discretizes the loop in the variables $\nabla R(\mathbf{x})$. For two consecutive points $\mathbf{z}_j, \mathbf{z}_{j+1}$, it chooses

$$\mathbf{g}_{j+1} = \frac{\nabla R(\mathbf{z}_j) - \nabla R(\mathbf{z}_{j+1})}{\eta}.$$

The derivative of the mirror-descent objective vanishes at $\mathbf{z}_{j+1}$. Strong convexity therefore makes $\mathbf{z}_{j+1}$ the unique constrained minimizer, including when it lies on the boundary. One traversal accumulates a Riemann sum for the negative circulation divided by η . The learning-rate quantifiers are the same as in Theorem 7.3 and the geometric-epoch construction in Remark 7.4 gives the corresponding anytime statement.

Combining the exactness, cross-derivative, and circulation viewpoints gives the following characterization for standard constrained mirror descent.

Corollary 8.12. *Let $U := \text{relint}(K)$. The set U is open and convex in $\text{aff}(K)$, and hence simply connected there. Suppose that $R \in C^2(U)$ and $\nabla^2 R(\mathbf{x})$ is positive definite for every $\mathbf{x} \in U$, and the coefficient field $\mathbf{x} \mapsto \nabla^2 R(\mathbf{x})\mathbf{D}(\mathbf{x})$ is C^1 . All derivatives and coordinate components below are taken in affine coordinates on $\text{aff}(K)$. Then the following conditions are equivalent.*

(i) *The one-form $\mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x})$ is exact on U .*

(ii) *For every coordinate pair i, j ,*

$$\partial_j[\nabla^2 R(\mathbf{x})\mathbf{D}(\mathbf{x})]_i = \partial_i[\nabla^2 R(\mathbf{x})\mathbf{D}(\mathbf{x})]_j \quad (\mathbf{x} \in U).$$

(iii) *The mirror circulation vanishes on every closed piecewise C^1 loop in U .*

If these conditions fail, standard mirror descent can be forced to incur linear regret along a sufficiently small loop contained in U . Suppose instead that they hold, that $\mathbf{D}$ extends to the displacement $\mathbf{D}_\phi$ of a feasible map on K , and that the resulting potential and displacement satisfy the regularity hypotheses of Theorem 8.9. Then mirror descent has sublinear regret. Under the Lipschitz-gradient assumptions in Corollary 8.10, the regret is $O(\sqrt{T})$.

Exactness is therefore the sharp geometric boundary for standard constrained mirror descent under the stated regularity. Smoothness and potential oscillation determine the quantitative rate, while nonzero circulation rules out no regret. Proximality plays no role in this characterization. It is addressed next as a separate representation problem.

8.5 Identity Interpolation of Bregman Proximal Deviations

Theorem 6.2 has two consequences in Euclidean geometry. First, interpolation with the identity scales the displacement exactly. Second, every smooth exact displacement has a sufficiently small interpolation that is proximal. The first fact is purely algebraic and survives for every mirror map. The second depends on the curvature of the regularizer and can fail even for a smooth Legendre geometry. This subsection separates these two statements.

The representation and interpolation arguments do not require compactness. Let K be a nonempty closed convex set, and let $R : E \rightarrow (-\infty, +\infty]$ be proper, closed, and m -strongly convex in the extended-valued sense. Let $\mathcal{X} \subseteq K \cap \text{dom } R$ be a convex input domain on which R is differentiable. We use $\mathcal{X} = K$ when R is differentiable on a neighborhood of K , as in standard constrained mirror descent, and $\mathcal{X} = \mathcal{D}$ in the Legendre specialization below.

For a proper lower-semicontinuous function $F := f + \iota_K$, assume that $F + R$ is proper (equivalently, $\text{dom } F \cap \text{dom } R \neq \emptyset$), and define

$$\text{prox}_F^R(\mathbf{x}) := \arg \min_{\mathbf{y} \in E} \{F(\mathbf{y}) + D_R(\mathbf{y}, \mathbf{x})\}, \quad \mathbf{x} \in \mathcal{X}.$$

The indicator in F restricts the minimization to K . If F is ρ -weakly convex for some $0 \leq \rho < m$, the objective is proper and lower semicontinuous as well as $(m - \rho)$ -strongly convex and the minimizer exists and is unique.

Definition 8.13. For a specified input domain $\mathcal{X} \subseteq K \cap \text{dom } R$, let

$$\Phi_{\text{prox}}^R(\mathcal{X}) := \left\{ \text{prox}_F^R|_{\mathcal{X}} : \begin{array}{l} F = f + \iota_K \text{ is proper, lower semicontinuous,} \\ \text{and } \rho\text{-weakly convex for some } 0 \leq \rho < m \end{array} \right\},$$

$$\Phi_{\text{rad}}^R(\mathcal{X}) := \text{Rad}_I(\Phi_{\text{prox}}^R(\mathcal{X})).$$

Thus $\phi \in \Phi_{\text{rad}}^R(\mathcal{X})$ when $\phi_\alpha := (1 - \alpha)I + \alpha\phi$ belongs to $\Phi_{\text{prox}}^R(\mathcal{X})$ for at least one $\alpha \in (0, 1]$.

For the constrained regularizer, we write $R_K := R + \iota_K$. This notation is used only when the constraint is absorbed into the regularizer and the Bregman divergence itself continues to be generated by R .

The Bregman proximal-regret guarantee of Cai et al. [2026] applies to standard constrained mirror descent and does not require Legendre duality. We therefore take that guarantee as given and ask how the corresponding deviation class relates to mirror exactness. In short, we show that radial proximal deviations are mirror-exact.

Proposition 8.14. Assume $\mathcal{X} = K$. For every admissible generator F in Definition 8.13, the function $\Psi_F(\theta) := R_K^*(\theta) - (F + R)^*(\theta)$ is differentiable and satisfies

$$\mathbf{x} - \text{prox}_F^R(\mathbf{x}) = \nabla_{\theta} \Psi_F(\nabla R(\mathbf{x})) \quad (\mathbf{x} \in K).$$

Consequently, $\Phi_{\text{prox}}^R(K) \subseteq \Phi_{\text{rad}}^R(K) \subseteq \Phi_{\text{exact}}^R(K)$, where $\Phi_{\text{exact}}^R(K)$ is the class from Definition 8.8. More precisely, if $\phi_\alpha = \text{prox}_F^R$, then ϕ is mirror-exact with potential Ψ_F/α .

The first inclusion follows by taking $\alpha = 1$. For the second, if $\phi_\alpha = \text{prox}_F^R$, then $\mathbf{x} - \phi_\alpha(\mathbf{x}) = \alpha(\mathbf{x} - \phi(\mathbf{x}))$. Thus exactness of ϕ_α implies exactness of ϕ . The same identity also transfers the Bregman proximal-regret guarantee of Cai et al. [2026]. For linearized regret, $\text{Reg}_T(\phi) = \alpha^{-1} \text{Reg}_T(\phi_\alpha)$, and for convex losses, $\ell(\mathbf{x}) - \ell(\phi_\alpha(\mathbf{x})) \geq \alpha[\ell(\mathbf{x}) - \ell(\phi(\mathbf{x}))]$. Hence, for each fixed deviation, the proximal-regret guarantee extends to $\Phi_{\text{rad}}^R(K)$.

Thus standard mirror descent controls the identity-radial closure of the Bregman proximal class. The remaining question is whether this radial class exhausts the mirror-exact class. Theorem 6.2 gives an affirmative answer for smooth Euclidean exact deviations. We now show that the corresponding statement can fail even for Legendre regularizers. It is therefore enough to return temporarily to the Legendre setting, where duality provides a convenient test for when a mirror-exact deviation can admit a Bregman proximal interpolation.

The Legendre test. We now return to the Legendre setting to test when a mirror-exact deviation admits a Bregman proximal interpolation. We use the assumptions of the first three subsections, except that no C^2 Hessian assumption is needed. Thus R is a proper closed Legendre function with $\text{dom } R = K$, and it is m -strongly convex in the norm under consideration. Write $\mathcal{D} := \text{int}_E(\text{dom } R)$ and $\Theta := \text{int}_E(\text{dom } R^*)$. Legendre duality identifies these two sets through the inverse maps ∇R and ∇R^* .

Let $\phi : \mathcal{D} \rightarrow K$ be R -mirror-exact with dual potential Ψ . For $\alpha \in (0, 1]$, define

$$u_\alpha(\theta) := R^*(\theta) - \alpha\Psi(\theta).$$

If $\mathbf{x} = \nabla R^*(\theta)$, mirror exactness and the definition of ϕ_α give

$$\phi_\alpha(\nabla R^*(\theta)) = \nabla u_\alpha(\theta).$$

Thus, in dual coordinates, the interpolated map ϕ_α is the gradient map of u_α . If ϕ_α is Bregman proximal, this gradient must also arise from a convex conjugate. This gives the first necessary condition for proximal representability.

Proposition 8.15. Suppose $\phi_\alpha = \text{prox}_{F_\alpha}^R$ on $\mathcal{D}$ for an admissible Bregman proximal generator F_α . Then the dual endpoint potential u_α agrees, up to an additive constant, with the convex conjugate of $F_\alpha + R$. More

precisely, there is a constant $c_\alpha \in \mathbb{R}$ such that $u_\alpha = (F_\alpha + R)^* + c_\alpha$, on Θ . Consequently, u_α is convex, and hence we have

$$D_\Psi(\theta', \theta) \leq \frac{1}{\alpha} D_{R^*}(\theta', \theta) \quad (\theta, \theta' \in \Theta).$$

Under the domain assumptions of Proposition 8.16, convexity of u_α yields a variational Bregman representation. It also reveals the connection with the relative smoothness condition of (6) used in the mirror-exact regret bound, since any proximal interpolation necessarily satisfies such a condition with constant $1/\alpha$. To obtain an admissible proximal map, however, the reconstructed generator F_α below must additionally satisfy the required weak-convexity condition.

Proposition 8.16. *Suppose u_α is convex on Θ , and let*

$$\bar{u}_\alpha := \text{cl}(u_\alpha + \iota_\Theta).$$

Assume $\bar{u}_\alpha$ is proper and that $\phi_\alpha(\mathcal{D}) \subseteq \text{dom } R$. Define the candidate proximal generator

$$F_\alpha(\mathbf{y}) := \begin{cases} \bar{u}_\alpha^*(\mathbf{y}) - R(\mathbf{y}), & \mathbf{y} \in \text{dom } R, \\ +\infty, & \mathbf{y} \notin \text{dom } R. \end{cases}$$

Then, for every $\mathbf{x} \in \mathcal{D}$,

$$\phi_\alpha(\mathbf{x}) \in \arg \min_{\mathbf{y}} \{F_\alpha(\mathbf{y}) + D_R(\mathbf{y}, \mathbf{x})\}.$$

Under these domain assumptions, convexity of u_α yields a variational Bregman representation of ϕ_α . If, in addition, F_α is proper, lower semicontinuous, and ρ_α -weakly convex for some $\rho_\alpha < m$, the minimizer is unique and $\phi_\alpha = \text{prox}_{F_\alpha}^R$. In that case $\phi \in \Phi_{\text{rad}}^R(\mathcal{D})$.

Propositions 8.15 and 8.16 separate proximal representability into two requirements. First, $R^* - \alpha\Psi$ must be convex. Second, the reconstruction must satisfy the stated domain conditions and produce a generator in the admissible weakly convex proximal class. The Euclidean construction in Theorem 6.2 satisfies both requirements for a sufficiently small interpolation scale. In a general mirror geometry, either requirement may fail.

We begin with the first requirement. Define

$$L_R(\Psi) := \inf \{L \geq 0 : D_\Psi(\theta', \theta) \leq L D_{R^*}(\theta', \theta) \text{ for all } \theta, \theta' \in \Theta\}.$$

This quantity compares the curvature of Ψ with that of R^* .

Corollary 8.17. *For $0 < \alpha \leq 1$, the function $u_\alpha = R^* - \alpha\Psi$ is convex exactly when $\alpha L_R(\Psi) \leq 1$. Define the convexity threshold*

$$\alpha_{\text{cvx}} := \min \left\{ 1, \frac{1}{L_R(\Psi)} \right\},$$

with the conventions $1/0 := +\infty$ and $1/(+\infty) := 0$; a zero threshold means that no positive scale passes the convexity test. Every Bregman proximal interpolation must satisfy $0 < \alpha \leq \alpha_{\text{cvx}}$. Conversely, every such α for which $\bar{u}_\alpha$ is proper and $\phi_\alpha(\mathcal{D}) \subseteq \text{dom } R$ yields the variational representation in Proposition 8.16. It is an admissible Bregman proximal interpolation if the resulting F_α also satisfies the weak-convexity and regularity conditions in that proposition.

The endpoint-domain condition is automatic for $0 < \alpha < 1$, since $\mathcal{D} = \text{int}_E(\text{dom } R)$, $K = \overline{\text{dom } R}$, and a strict interpolation of an interior point with a point of K lies in $\mathcal{D}$. At $\alpha = 1$, one must additionally require $\phi(\mathcal{D}) \subseteq \text{dom } R$.

When R^* and Ψ are twice differentiable, the same condition becomes

$$\alpha \nabla^2 \Psi(\boldsymbol{\theta}) \preceq \nabla^2 R^*(\boldsymbol{\theta}) \quad (\boldsymbol{\theta} \in \Theta).$$

Thus the curvature contributed by $\alpha \Psi$ cannot exceed that of R^* in any direction. This is a relative curvature condition and can be much stronger than an ordinary Euclidean Lipschitz bound on $\nabla \Psi$.

We next consider the second requirement. Suppose R is M -smooth on the relevant primal domain and $\bar{u}_\alpha$ is a finite closed convex function that is L_α -smooth on its affine hull. Then $\bar{u}_\alpha^*$ is $1/L_\alpha$ -strongly convex. Comparing this with the M -smoothness of R shows that the reconstructed F_α is

$$\rho_\alpha := \left(M - \frac{1}{L_\alpha} \right)_+$$

weakly convex. Hence it is admissible when $\rho_\alpha < m$, provided the reconstruction domain assumptions hold and F_α is proper and lower semicontinuous.

In the Euclidean case $M = m = 1$, a sufficiently small interpolation scale provides both the convexity of u_α and the required weak-convexity margin for F_α , as shown in Theorem 6.2. For a nonquadratic regularizer, these two curvature requirements need not follow from one another.

The next example shows that failure can already occur at the first step. Although the dual potential is smooth with globally Lipschitz gradient, $R^* - \alpha \Psi$ is nonconvex for every positive interpolation scale.

Proposition 8.18. *Let $1 < p < 2$, let $q := p/(p-1) > 2$, and take the full-domain Legendre regularizer*

$$R(\mathbf{x}) := \frac{1}{2} \|\mathbf{x}\|_p^2 \quad (\mathbf{x} \in \mathbb{R}^2).$$

There is a feasible R -mirror-exact deviation whose dual potential is C^∞ with globally Lipschitz gradient, but which does not belong to $\Phi_{\text{rad}}^R(\mathbb{R}^2)$.

To see why, note that $R^*(\boldsymbol{\theta}) = \frac{1}{2} \|\boldsymbol{\theta}\|_q^2$. Take

$$\Psi(\boldsymbol{\theta}) := \frac{1}{2} \boldsymbol{\theta}[2]^2, \quad \phi(\mathbf{x}) := \mathbf{x} - \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\mathbf{x})).$$

The action space is all of $\mathbb{R}^2$, so feasibility is automatic, and $\nabla \Psi$ is globally 1-Lipschitz. Along the line $\boldsymbol{\theta} = \mathbf{e}_1 + t\mathbf{e}_2$,

$$u_\alpha(\mathbf{e}_1 + t\mathbf{e}_2) = \frac{1}{2}(1 + |t|^q)^{2/q} - \frac{\alpha}{2}t^2.$$

As $t \rightarrow 0$,

$$u_\alpha(\mathbf{e}_1 + t\mathbf{e}_2) - u_\alpha(\mathbf{e}_1) = \frac{1}{q}|t|^q - \frac{\alpha}{2}t^2 + O(|t|^{2q}).$$

Because $q > 2$, the negative quadratic term dominates for every $\alpha > 0$ and every sufficiently small nonzero t . Hence $R^* - \alpha \Psi$ is nonconvex for every positive α .

The obstruction comes from a mismatch in curvature. At $\mathbf{e}_1$, the conjugate regularizer has no quadratic curvature in the transverse $\mathbf{e}_2$ direction, and its first increase is of order $|t|^q$. The potential Ψ has quadratic curvature in that same direction. Scaling Ψ changes only the coefficient of this quadratic term, so no positive scale can restore convexity. In the Euclidean case, $\frac{1}{2} \|\boldsymbol{\theta}\|_2^2$ has quadratic curvature in every direction, which is why the small-scale construction in Theorem 6.2 succeeds. Hence

$$\Phi_{\text{prox}}^R(\mathbb{R}^2) \subseteq \Phi_{\text{rad}}^R(\mathbb{R}^2) \subsetneq \Phi_{\text{exact}}^R(\mathbb{R}^2)$$

can hold even for a Legendre regularizer with a C^∞ dual potential whose gradient is globally Lipschitz. The regularizer is differentiable but not C^2 on the coordinate axes, which is why the interpolation argument is formulated through conjugacy rather than Hessians.

9 Follow-the-Regularized-Leader Beyond Legendre Regularizers

In the Legendre interior, online mirror descent and follow-the-regularized-leader (FTRL) coincide for constant-step linearized losses [McMahan, 2011, 2017]. Indeed, mirror descent satisfies

$$\nabla R(\mathbf{x}_{t+1}) = \nabla R(\mathbf{x}_t) - \eta \mathbf{g}_t,$$

which unrolls to the FTRL update

$$\mathbf{x}_t := \arg \min_{\mathbf{x} \in E} \left\{ \eta \sum_{s=1}^{t-1} \langle \mathbf{g}_s, \mathbf{x} \rangle + R(\mathbf{x}) - \langle \nabla R(\mathbf{x}_1), \mathbf{x} \rangle \right\}.$$

Beyond the Legendre interior, this equivalence need not persist. Boundary constraints introduce normal-cone terms into mirror descent, while FTRL continues to use the cumulative linear losses in a single regularized minimization; related distinctions between constrained mirror descent and cumulative-dual methods are discussed by Fang et al. [2022]. In Section 8.4, we characterized the full deviation class controlled by standard constrained mirror descent beyond the Legendre setting. For FTRL, the relevant geometry is different because the algorithm evolves through its cumulative dual state rather than the one-step constrained mirror update. We therefore study FTRL separately and characterize the deviation class for which it has sublinear regret, together with the corresponding dual-circulation obstruction.

9.1 Coincidence with Mirror Descent in the Legendre Interior

The equivalence between mirror descent and follow-the-regularized-leader (FTRL) for constant-step linearized losses in the Legendre interior is standard; see, for example, McMahan [2011, 2017]. We record the precise form used here for completeness.

Proposition 9.1. *Under the standing Legendre assumptions of Section 8, let $\boldsymbol{\theta}_1 := \nabla R(\mathbf{x}_1)$ and use a constant step size η . The mirror-descent iterates satisfy*

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in E} \left\{ \eta \sum_{s=1}^{t-1} \langle \mathbf{g}_s, \mathbf{x} \rangle + R(\mathbf{x}) - \langle \boldsymbol{\theta}_1, \mathbf{x} \rangle \right\} = \arg \min_{\mathbf{x} \in K} \left\{ \eta \sum_{s=1}^{t-1} \langle \mathbf{g}_s, \mathbf{x} \rangle + D_R(\mathbf{x}, \mathbf{x}_1) \right\}.$$

Thus they are the FTRL iterates for the linearized losses and the tilted regularizer $R - \langle \boldsymbol{\theta}_1, \cdot \rangle$.

Since the two algorithms generate the same iterates in this setting, the Legendre-interior regret guarantees established above for mirror descent, including the mirror-exact and Bregman proximal guarantees of Sections 8.2 and 8.3, apply verbatim to FTRL. The situation changes beyond the Legendre interior. Standard constrained mirror descent can acquire a boundary normal-cone residual, whereas FTRL continues to evolve through the cumulative linear losses. We analyzed the resulting deviation class for mirror descent in Section 8.4. We now carry out the corresponding analysis for FTRL and show that, once boundary effects are present, the two algorithms can control different deviation classes.

9.2 Exactness and Circulation in the Cumulative Dual State

Unlike constrained mirror descent, FTRL continues to evolve additively through its cumulative dual state beyond the Legendre setting. This suggests imposing exactness directly in that variable.

Let $R : E \rightarrow (-\infty, +\infty]$ be proper, closed, and convex, set $K := \overline{\text{dom } R}$, and assume that R^* is differentiable on the nonempty open convex set $\Theta := \text{int}_E(\text{dom } R^*)$. Given $\boldsymbol{\theta}_1 \in \Theta$, define FTRL by

$$\boldsymbol{\theta}_t := \boldsymbol{\theta}_1 - \eta \sum_{s=1}^{t-1} \mathbf{g}_s, \quad \mathbf{x}_t := \nabla R^*(\boldsymbol{\theta}_t). \quad (12)$$

Equivalently,

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in E} \left\{ \eta \sum_{s=1}^{t-1} \langle \mathbf{g}_s, \mathbf{x} \rangle + R(\mathbf{x}) - \langle \boldsymbol{\theta}_1, \mathbf{x} \rangle \right\}.$$

Because R may be extended-valued, it can encode constraints and the resulting iterates may lie on the boundary of K .

The relevant displacement is therefore the displacement viewed as a function of the cumulative dual state.

Definition 9.2 (FTRL-exact deviation). *A feasible map $\phi : K \rightarrow K$ is R -FTRL-exact if there is a differentiable potential $\Psi : \Theta \rightarrow \mathbb{R}$ such that*

$$\nabla R^*(\boldsymbol{\theta}) - \phi(\nabla R^*(\boldsymbol{\theta})) = \nabla_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta}) \quad (\boldsymbol{\theta} \in \Theta).$$

When R is Legendre, this is the mirror-exact condition on the interior choice range.

We denote this class by $\Phi_{\text{FTRL}}^R(K)$. Its dual displacement is

$$\tilde{D}_{\phi}(\boldsymbol{\theta}) := \nabla R^*(\boldsymbol{\theta}) - \phi(\nabla R^*(\boldsymbol{\theta})).$$

Thus FTRL-exactness is ordinary exactness of $\tilde{D}_{\phi}$ in the cumulative dual state.

To obtain a quantitative regret bound, we use uniform convexity of the regularizer. Let $\|\cdot\|$ be a norm on E , with dual norm $\|\cdot\|_*$. We say that R is (σ, r) -uniformly convex if, for every $\mathbf{x}, \mathbf{y} \in \text{dom } R$ and every $\boldsymbol{\theta} \in \partial R(\mathbf{x})$,

$$R(\mathbf{y}) \geq R(\mathbf{x}) + \langle \boldsymbol{\theta}, \mathbf{y} - \mathbf{x} \rangle + \frac{\sigma}{r} \|\mathbf{y} - \mathbf{x}\|^r. \quad (13)$$

Here $\sigma > 0$, $r > 1$, and $q := r/(r-1)$.

Exactness makes the potential differences telescope, while relative smoothness and uniform convexity control the Bregman error incurred by each dual update. We are now ready to state and prove our exact-form regret result for FTRL.

Theorem 9.3. *Assume that R is proper, closed, (σ, r) -uniformly convex, and that R^* is finite on E . Let the iterates follow (12). If ϕ is R -FTRL-exact with potential Ψ and*

$$D_{\Psi}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \leq LD_{R^*}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \quad (t \in [T]),$$

then

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{\Psi(\boldsymbol{\theta}_1) - \Psi(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{L\eta^{q-1}}{q\sigma^{q-1}} \sum_{t=1}^T \|\mathbf{g}_t\|_*^q.$$

As in the mirror-descent analysis and Condition (6) in Section 8.2, the relative-smoothness condition is needed only along the realized trajectory, namely on the consecutive pairs $(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t)$. Imposing it on all of $\Theta \times \Theta$ gives a trajectory-independent sufficient condition.

Under bounded gradients and bounded oscillation of the potential, optimizing the preceding bound over the step size gives the following rate.

Corollary 9.4. *Suppose $\|\mathbf{g}_t\|_* \leq G$ and the oscillation of Ψ along the dual trajectory is at most B . When $B, L, G > 0$, choosing $\eta = \left(\frac{rB\sigma^{q-1}}{LG^qT} \right)^{1/q}$ gives*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq r^{1/r} \frac{GL^{1/q} B^{1/r}}{\sigma^{1/r}} T^{1/q}.$$

For $r = 2$, this is $G\sqrt{2LBT/\sigma}$.

As in the mirror-descent converse of Theorem 8.11, the obstruction is nonzero circulation. Here the relevant loop lives in the cumulative dual state of FTRL. If the dual displacement has nonzero circulation, an adversary can force the iterates to repeatedly traverse such a loop and incur linear regret.

Theorem 9.5. *Suppose $\tilde{D}_\phi$ is continuous and there is a closed piecewise C^1 loop $\gamma \subset \Theta$ with*

$$\oint_\gamma \langle \tilde{D}_\phi(\boldsymbol{\theta}), d\boldsymbol{\theta} \rangle \neq 0.$$

For every $G > 0$ and every sufficiently small constant $\eta > 0$, bounded gradients $\|\mathbf{g}_t\|_ \leq G$ can force FTRL (12) to incur $cT - O(1/\eta)$ regret against ϕ , for some $c > 0$. The statement holds from every initial dual state in Θ since steering to the loop changes only the transient term.*

Combining the upper and lower bounds gives the corresponding exactness-versus- circulation characterization for FTRL.

Corollary 9.6. *Assume that $\tilde{D}_\phi$ is C^1 on the open convex set Θ . It is R -FTRL-exact if and only if its Jacobian is symmetric. If a skew derivative is nonzero, Theorem 9.5 gives linear regret along a small dual loop. Under the regularity and normalization assumptions of Theorem 9.3, an exact field has sublinear regret.*

9.3 Mirror Descent and FTRL Control Different Deviation Classes

The equivalence in Proposition 9.1 relies on the Legendre interior. Once boundary points are allowed, the two algorithms use different dual information. Standard constrained mirror descent is governed by the primal mirror coordinate $\nabla R(\mathbf{x})$ together with a normal-cone residual, whereas FTRL with the extended regularizer $R_K := R + \iota_K$ evolves through a cumulative dual state that may also contain normal components. As a result, FTRL exactness imposes conditions in dual directions that are not present in mirror exactness. In particular, every FTRL-exact deviation is mirror-exact.

Proposition 9.7. *Let K be compact and convex, and let R be differentiable and strongly convex on a neighborhood of K . If a feasible map $\phi : K \rightarrow K$ is R_K -FTRL-exact, then it is R -mirror-exact in the sense of Definition 8.8. Consequently,*

$$\Phi_{\text{FTRL}}^{R_K}(K) \subseteq \Phi_{\text{exact}}^R(K).$$

We next show that this inclusion can be strict, even for a smooth power-type regularizer. Fix $1 < p < \infty$ and $0 < a < b$, and set

$$K := [a, b]^2, \quad R_p(\mathbf{x}) := \frac{1}{p} \|\mathbf{x}\|_p^p = \frac{1}{p} (\mathbf{x}[1]^p + \mathbf{x}[2]^p), \quad R_{p,K} := R_p + \iota_K.$$

Let $c := (a + b)/2$, choose $\beta \in (0, 1)$, and write

$$S_\beta := \begin{pmatrix} 1 & \beta \\ \beta & 1 \end{pmatrix}, \quad m_p := (p - 1) \min_{s \in [a, b]} s^{p-2}.$$

Choose $0 < \varepsilon \leq m_p(1 - \beta)/(1 + \beta)$ and define

$$\mathbf{D}_{p,\varepsilon}(\mathbf{x}) := \varepsilon [\nabla^2 R_p(\mathbf{x})]^{-1} S_\beta (\mathbf{x} - c\mathbf{1}), \quad \phi_{p,\varepsilon}(\mathbf{x}) := \mathbf{x} - \mathbf{D}_{p,\varepsilon}(\mathbf{x}). \quad (14)$$

The same deviation has opposite regret behavior under the two algorithms.

Theorem 9.8. *The map $\phi_{p,\varepsilon}$ is a smooth feasible deviation. Standard constrained mirror descent with distance generator R_p has $O(\sqrt{T})$ regret against it for bounded gradients and the usual $T^{-1/2}$ step size. In contrast, FTRL with the extended regularizer $R_{p,K}$ can be forced, for every sufficiently small constant step size, to incur $cT - O(1/\eta)$ regret against the same deviation. Hence the inclusion in Proposition 9.7 is strict and $\Phi_{\text{FTRL}}^{R_{p,K}}(K) \subsetneq \Phi_{\text{exact}}^{R_p}(K)$.*

For mirror descent, exactness is visible directly in primal coordinates. The one-form satisfies

$$[\nabla^2 R_p(\mathbf{x}) \mathbf{D}_{p,\varepsilon}(\mathbf{x})]^\top d\mathbf{x} = d \left[\frac{\varepsilon}{2} (\mathbf{x} - c\mathbf{1})^\top S_\beta (\mathbf{x} - c\mathbf{1}) \right].$$

All coefficients are smooth and bounded on the positive box, so the relative-curvature and Lipschitz conditions in Corollary 8.10 hold with finite constants. The mirror-descent regret bound therefore applies.

For FTRL, the obstruction appears when the cumulative dual state reaches a region where one coordinate is saturated at the boundary. For the extended regularizer, the choice map is coordinatewise

$$\nabla R_{p,K}^*(\boldsymbol{\theta})[i] = \begin{cases} a, & \boldsymbol{\theta}[i] \leq a^{p-1}, \\ \boldsymbol{\theta}[i]^{1/(p-1)}, & a^{p-1} < \boldsymbol{\theta}[i] < b^{p-1}, \\ b, & \boldsymbol{\theta}[i] \geq b^{p-1}. \end{cases}$$

On the open strip, we have

$$\boldsymbol{\theta}[1] > b^{p-1}, \quad a^{p-1} < \boldsymbol{\theta}[2] < b^{p-1}.$$

The dual displacement $\tilde{\mathbf{D}}_{p,\varepsilon}$ satisfies

$$\frac{\partial \tilde{\mathbf{D}}_{p,\varepsilon}[1]}{\partial \boldsymbol{\theta}[2]} = \frac{\varepsilon \beta}{(p-1)^2} b^{2-p} \boldsymbol{\theta}[2]^{(2-p)/(p-1)} > 0, \quad \frac{\partial \tilde{\mathbf{D}}_{p,\varepsilon}[2]}{\partial \boldsymbol{\theta}[1]} = 0.$$

The unequal cross derivatives give nonzero dual circulation on a small rectangle in this strip. Hence Theorem 9.5 gives linear regret for FTRL, even though the same deviation is mirror-exact and has sublinear regret under standard constrained mirror descent. At $p = 2$, the displacement in (14) is affine and the example reduces to a symmetric cross-coordinate deviation on a box.

9.4 Examples and Applications of the FTRL Bound

We conclude the FTRL analysis with two standard power-type regularizers that illustrate how Theorem 9.3 specializes in concrete geometries. In both examples, $K \subseteq \mathbb{R}^d$ is nonempty, closed, and convex, the constraint is encoded in R , ϕ is FTRL-exact, the relative-smoothness condition holds with constant L , and the oscillation of Ψ along the dual trajectory is at most B .

For squared ℓ_p geometry with $1 < p \leq 2$, take

$$R(\mathbf{x}) := \frac{1}{2} \|\mathbf{x}\|_p^2 + \iota_K(\mathbf{x}).$$

Since this regularizer is $(p-1)$ -strongly convex with respect to $\|\cdot\|_p$, Theorem 9.3 gives the following specialization.

Corollary 9.9. *Let $1 < p \leq 2$ and let $q := p/(p-1)$. Then*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{B}{\eta} + \frac{L\eta}{2(p-1)} \sum_{t=1}^T \|\mathbf{g}_t\|_q^2.$$

If $\|\mathbf{g}_t\|_q \leq G$, the optimized bound is $G\sqrt{2LBT/(p-1)}$.

For $p \geq 2$, the natural power-type regularizer is instead

$$R(\mathbf{x}) := \frac{1}{p} \|\mathbf{x}\|_p^p + \iota_K(\mathbf{x}).$$

Its uniform convexity has exponent p , which leads to a different dependence on the horizon.

Corollary 9.10. *Let $2 \leq p < \infty$ and let $q := p/(p-1)$. Then*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{B}{\eta} + \frac{L\eta^{q-1}}{q(2^{2-p})^{q-1}} \sum_{t=1}^T \|\mathbf{g}_t\|_q^q.$$

If $\|\mathbf{g}_t\|_q \leq G$, the optimized bound is $p^{1/p} 2^{(p-2)/p} GL^{1/q} B^{1/p} T^{1/q}$, which is $O(T^{1-1/p})$.

These examples show how the geometry of the regularizer determines the quantitative FTRL rate through its uniform-convexity exponent. The one-homogeneous norm $\|\mathbf{x}\|_p$ itself is not suitable for Theorem 9.3, since it is affine along positive rays and its conjugate is an indicator of the dual unit ball. At $p = 1$ or $p = \infty$, one can instead use an entropic regularizer, a strongly convex perturbation, or a nearby ℓ_p geometry with $1 < p < \infty$.

Sections 8.4 and 9 give parallel characterizations for different algorithms. For standard constrained mirror descent, exactness of $\mathbf{D}_\phi(\mathbf{x})^\top \nabla R(\mathbf{x})$ is the sharp geometric condition, and the boundary residual determines the additional analytic term in the upper bound. For FTRL, exactness and circulation live directly in the cumulative dual state, which also permits nonsmooth and extended-valued regularizers. The theories coincide on Legendre interior trajectories. Once normal directions appear, FTRL exactness is a stricter extension requirement, and Theorem 9.8 shows that the corresponding no-regret deviation classes can differ even for the same powered ℓ_p geometry.

10 Conservative Correlated Equilibria in Convex Games

The regret guarantees developed above translate directly into equilibrium constraints in convex games. It is useful to distinguish three questions. The first concerns inclusion between deviation classes, the second concerns quantitative approximate equilibria for normalized families, and the third concerns equilibrium constraints when the approximation error is zero. In particular, a strict inclusion between deviation classes, such as the one in Proposition 6.1, does not by itself imply a strict inclusion between the corresponding equilibrium sets.

Consider an N -player game. Player i has compact convex action set K_i and a loss $\ell_i(\mathbf{x}_i, \mathbf{x}_{-i})$ convex and differentiable in its own action. Let $K := \prod_{i=1}^N K_i$. Assume also that each loss $\ell_i : K \rightarrow \mathbb{R}$ is jointly continuous. Throughout this section, every distribution is a Borel probability measure on K , and every feasible deviation is Borel measurable. Joint continuity and compactness make each loss bounded, so all the expected gains below are well defined and finite, including those of identity interpolations. For such a distribution μ and a feasible deviation ϕ_i , define its gain

$$\Gamma_i(\mu, \phi_i) := \mathbb{E}_{\mathbf{x} \sim \mu} [\ell_i(\mathbf{x}_i, \mathbf{x}_{-i}) - \ell_i(\phi_i(\mathbf{x}_i), \mathbf{x}_{-i})].$$

The interpolation results in Sections 6 and 8.5 have the same consequence when the approximation error is zero. Passing from a deviation class to its identity-radial closure does not change the resulting equilibrium constraints.

Proposition 10.1. *Let $\mathcal{A}_i$ be any family of feasible Borel-measurable deviations for player i . A distribution satisfies $\Gamma_i(\mu, \phi_i) \leq 0$ for every $\phi_i \in \mathcal{A}_i$ and every player i if and only if it satisfies the same inequalities for every $\phi_i \in \text{Rad}_I(\mathcal{A}_i)$.*

This equivalence concerns the equilibrium constraints themselves. Quantitative bounds over a radial closure additionally require a common lower bound on the admissible interpolation scales.

For each player i , let $\Phi_{i,\text{prox}}$ be the unrestricted union of all constrained proximal maps $\text{prox}_F|_{K_i}$, where $F := f + \iota_{K_i}$ is proper, lower-semicontinuous, and ρ -weakly convex for some map-dependent $0 \leq \rho < 1$. We use the following equilibrium notions.

- A distribution μ is a *coarse correlated equilibrium* (CCE) if $\Gamma_i(\mu, \tau_{i,\mathbf{u}_i}) \leq 0$ for every player i and every constant deviation $\tau_{i,\mathbf{u}_i}(\mathbf{x}_i) \equiv \mathbf{u}_i \in K_i$.

- It is a *proximal correlated equilibrium* (PCE) if $\Gamma_i(\mu, \tau_i) \leq 0$ for every player i and every $\tau_i \in \Phi_{i,\text{prox}}$.
- It is a *correlated equilibrium* (CE) if $\Gamma_i(\mu, \tau_i) \leq 0$ for every player i and every Borel-measurable deviation $\tau_i : K_i \rightarrow K_i$.

Thus, under the unrestricted proximal convention used here,

$$\text{CE} \subseteq \text{PCE} \subseteq \text{CCE}.$$

The CE notion originates with [Aumann \[1974\]](#); here we use its standard measurable-deviation extension to compact convex action spaces. The original PCE work introduced PCE as a tractable refinement of CCE [\[Cai et al., 2026\]](#). We determine below how PCE compares with measurable-deviation CE. The two notions coincide on finite support, while $\text{CE} \subsetneq \text{PCE}$ can occur on infinite support.

We next define the equilibrium notion associated with exact-form deviations. The quantitative formulation and the case of zero approximation error require slightly different conventions because uniform regret bounds require common bounds on the potentials.

Definition 10.2 (ConCE for a specified deviation family). *Fix for every player a specified family Φ_i of feasible Euclidean exact-form deviations. A distribution μ is an ε -conservative correlated equilibrium relative to $(\Phi_i)_i$ if*

$$\sup_{\phi_i \in \Phi_i} \Gamma_i(\mu, \phi_i) \leq \varepsilon \quad \text{for every player } i.$$

For approximate, quantitative conclusions, we use a family such as $\Phi_i := \Phi_{i,\text{exact}}(B_i, L_i)$, where $\Phi_{i,\text{exact}}(B_i, L_i)$ denotes $\Phi_{\text{exact}}(B_i, L_i)$ on the action set K_i . In this family all potentials satisfy the same oscillation bound B_i and smoothness bound L_i . By contrast, when the approximation error is zero, ConCE tests every smooth feasible exact-form deviation in the sense of Definition 3.1 (equivalently, every feasible exact-form deviation admitting a $C_{\text{loc}}^{1,1}$ potential on a neighborhood of K_i). Thus each such deviation must have nonpositive gain. Since this larger class has no common bounds on the size or smoothness of its potentials, we do not claim a single uniform quantitative estimate over the entire class. Unless stated otherwise, ConCE in this section refers to this Euclidean exact-form notion.

The standard regret-to-equilibrium argument gives the following quantitative connection.

Theorem 10.3. *Let $\bar{\mu}_T := T^{-1} \sum_{t=1}^T \delta_{\mathbf{x}_t}$ and define $\mathbf{g}_{i,t} := \nabla_{\mathbf{x}_i} \ell_i(\mathbf{x}_{i,t}, \mathbf{x}_{-i,t})$. If every player i guarantees*

$$\sup_{\phi_i \in \Phi_i} \sum_{t=1}^T \langle \mathbf{g}_{i,t}, \mathbf{x}_{i,t} - \phi_i(\mathbf{x}_{i,t}) \rangle \leq R_{i,T},$$

then $\bar{\mu}_T$ is an ε_T -ConCE relative to $(\Phi_i)_i$, where $\varepsilon_T = \max_i R_{i,T}/T$. In particular, OGD on $\Phi_{i,\text{exact}}(B_i, L_i)$ with $\|\mathbf{g}_{i,t}\| \leq G_i$ gives $\varepsilon_T \leq \max_i G_i \sqrt{6L_i B_i / T}$.

The proof uses only convexity of the losses and the linearized regret bound. The same argument therefore applies to the regularizer-dependent mirror and FTRL deviation families. We introduce their corresponding equilibrium notions in Section 10.2. The distinction between them matters at the boundary, as shown by Theorem 9.8.

10.1 Exact and Proximal Equilibrium Constraints

The class $\Phi_{i,\text{prox}}$ is the unrestricted proximal class used in [Cai et al. \[2026\]](#). We impose no fixed ρ , normalization, locality, or computational restriction. Let $\Phi_{i,\text{exact}}^{\text{sm}}$ denote the smooth exact-form class tested by ConCE. By Corollary 6.3,

$$\text{Rad}_I(\Phi_{i,\text{prox}}) = \Phi_{i,\text{exact}}^{\text{sm}}.$$

Together with the radial invariance in Proposition 10.1, this immediately gives the equality of the corresponding equilibrium constraints.

Theorem 10.4. *Under the unrestricted conventions above, every convex game with compact convex action sets and jointly continuous losses satisfies*

$$\text{ConCE} = \text{PCE}.$$

The proof is given in Appendix D.14. This equality relies on the unrestricted deviation classes and does not directly extend to classes with a fixed weak-convexity parameter, normalization, locality restriction, or computational restriction. For normalized ε -equilibria, the admissible interpolation scale may depend on the deviation, so the strict difference between the map classes can still affect quantitative guarantees.

10.2 Mirror-exact, Mirror-radial, and FTRL Equilibrium Constraints

We now record the equilibrium consequences of the regularizer-dependent deviation classes studied in Sections 8 and 9. We begin with standard constrained mirror descent. Fix for every player i a differentiable m_i -strongly convex regularizer R_i on K_i , as in Section 8.4. The full geometric class controlled by the exactness characterization is $\Phi_{\text{exact}}^{R_i}(K_i)$. For quantitative equilibrium guarantees, we restrict to a subfamily with common regularity parameters. Let

$$\Phi_i^{\text{MD}} \subseteq \Phi_{\text{exact}}^{R_i}(K_i)$$

be such a family. Assume that every $\phi_i \in \Phi_i^{\text{MD}}$ has a potential with oscillation at most B_i^{MD} , satisfies (10) with the same constant L_i^{MD} , and has an S_i -Lipschitz displacement. Suppose also that ∇R_i is M_i -Lipschitz on K_i and $\|\mathbf{g}_{i,t}\|_* \leq G_i$. Set

$$C_i^{\text{MD}} := \frac{L_i^{\text{MD}} + S_i}{m_i} + \frac{S_i M_i}{m_i^2}.$$

If player i uses standard mirror descent with

$$\eta_i^{\text{MD}} := \sqrt{\frac{B_i^{\text{MD}}}{C_i^{\text{MD}} G_i^2 T}},$$

whenever $B_i^{\text{MD}} C_i^{\text{MD}} > 0$, then Corollary 8.10 and Theorem 10.3 give

$$\sup_{\phi_i \in \Phi_i^{\text{MD}}} \Gamma_i(\bar{\mu}_T, \phi_i) \leq 2G_i \sqrt{\frac{B_i^{\text{MD}} C_i^{\text{MD}}}{T}}.$$

The degenerate cases follow from the corresponding unoptimized bound. If ∇R_i is only uniformly continuous, the modulus estimate in Corollary 8.10 still gives an $o(1)$ equilibrium error for every family with common values of B_i^{MD} , L_i^{MD} , and S_i .

The identity-interpolation results of Section 8.5 give a second equilibrium notion associated with the same regularizer. Let $\Phi_{\text{prox}}^{R_i}(K_i)$ and $\Phi_{\text{rad}}^{R_i}(K_i)$ be the Bregman proximal and radial classes from Definition 8.13. A distribution is a *Bregman R -PCE* when every deviation in $\Phi_{\text{prox}}^{R_i}(K_i)$ has nonpositive gain for every player. It is a *radial R -equilibrium* when the same condition holds for every deviation in $\Phi_{\text{rad}}^{R_i}(K_i)$. Denote the two sets by $\text{PCE}_{(R_i)_i}^{\text{Breg}}$ and $\text{Eq}_{(R_i)_i}^{\text{rad}}$, respectively.

By radial invariance, these two equilibrium notions coincide. We record this result in the following corollary.

Corollary 10.5. *For any set of regularizers R_i for $i \in [N]$, we have*

$$\text{PCE}_{(R_i)_i}^{\text{Breg}} = \text{Eq}_{(R_i)_i}^{\text{rad}}.$$

This equality concerns the proximal class and its identity-radial closure; it does not identify the radial class with the full mirror-exact class. To make this distinction explicit, define

$$\text{Eq}_{(R_i)_i}^{\text{exact}} := \left\{ \mu : \Gamma_i(\mu, \phi_i) \leq 0 \text{ for every player } i \text{ and every } \phi_i \in \Phi_{\text{exact}}^{R_i}(K_i) \right\}.$$

Proposition 8.14 gives

$$\text{Eq}_{(R_i)_i}^{\text{exact}} \subseteq \text{Eq}_{(R_i)_i}^{\text{rad}} = \text{PCE}_{(R_i)_i}^{\text{Breg}}.$$

The first inclusion can be strict at the level of deviation maps. In the Legendre regime, Proposition 8.15 shows that a mirror-exact deviation can belong to the radial class only if it passes the relative curvature test in Corollary 8.17. Proposition 8.16 gives the converse when the reconstructed generator is an admissible weakly convex Bregman proximal function, while Proposition 8.18 shows that no positive interpolation may exist even for a smooth dual potential with globally Lipschitz gradient. As in the Euclidean setting, strict inclusion of deviation classes does not by itself imply strict inclusion of the corresponding equilibrium sets.

FTRL produces a third, distinct family of equilibrium constraints. For each player, let $\mathcal{R}_i : \mathbb{R}^{d_i} \rightarrow (-\infty, +\infty]$ be a proper closed (σ_i, r_i) -uniformly convex FTRL regularizer with $K_i = \overline{\text{dom } \mathcal{R}_i}$, and let

$$\Phi_i^{\text{FTRL}} \subseteq \Phi_{\text{FTRL}}^{\mathcal{R}_i}(K_i)$$

be a specified family. Player i uses (12) with R replaced by $\mathcal{R}_i$. When FTRL and mirror descent use the same base geometry and constraint, one takes $\mathcal{R}_i := R_i + \iota_{K_i}$. Let $q_i := r_i/(r_i - 1)$. Assume that every potential in Φ_i^{FTRL} has trajectory oscillation at most B_i^{FTRL} , satisfies the consecutive relative bound with the same constant L_i^{FTRL} , and that $\|g_{i,t}\|_* \leq G_i$. With the learning rate from Corollary 9.4, the empirical distribution is an $\varepsilon_T^{\text{FTRL}}$ -equilibrium relative to $(\Phi_i^{\text{FTRL}})_i$, where

$$\varepsilon_T^{\text{FTRL}} \leq \max_i \left\{ r_i^{1/r_i} \frac{G_i (L_i^{\text{FTRL}})^{1/q_i} (B_i^{\text{FTRL}})^{1/r_i}}{\sigma_i^{1/r_i}} T^{-1/r_i} \right\}.$$

For $r_i = 2$, the rate is $O(T^{-1/2})$. This guarantee applies to the FTRL-exact deviation family rather than the mirror-exact family. As shown in Theorem 9.8, the two classes can differ once boundary effects are present, even when the algorithms use the same underlying geometry and action set.

The remaining subsections return to the Euclidean ConCE and PCE notions. Their comparison with unrestricted measurable CE does not depend on a selected mirror regularizer.

10.3 Finite Support and Boundary-compatible Interpolation

On a finite support, only finitely many recommended actions and deviation values matter. The following result shows that any such finite collection can be realized, after a sufficiently small common interpolation with the identity, by a single proximal map. Its proof appears in Appendix C.

Theorem 10.6. *Let $K \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Let $\mathbf{x}^1, \dots, \mathbf{x}^m \in K$ be distinct and choose arbitrary targets $\mathbf{y}^1, \dots, \mathbf{y}^m \in K$. Fix $\rho \in (0, 1)$. There is $\alpha_0 > 0$ such that for every $\alpha \in (0, \alpha_0]$ one can find a proper lower-semicontinuous $F_\alpha = f_\alpha + \iota_K$ that is ρ -weakly convex and satisfies*

$$\text{prox}_{F_\alpha}(\mathbf{x}^j) = (1 - \alpha)\mathbf{x}^j + \alpha\mathbf{y}^j, \quad j = 1, \dots, m.$$

The proximal minimizer is unique.

It follows that the three equilibrium notions cannot be separated on a finite support.

Corollary 10.7. *For every finitely supported distribution in a convex game,*

$$\mu \in \text{CE} \iff \mu \in \text{ConCE} \iff \mu \in \text{PCE}.$$

In particular the three notions coincide in every finite game.

For the finite-game statement, we use the standard mixed extension, with each pure-action set embedded as the vertices of its mixed-strategy simplex and the distribution supported on pure-action profiles. Since the extended losses are affine in a player’s own mixed action, any profitable deviation to a mixed action has a profitable pure action in its support. The proof of Corollary 10.7 is given in Appendix D.14.

Indeed, Monnot and Piliouras [2017, Proposition 6.1] show that $\text{CE} = \text{CCE}$ when every player has two actions: in the CCE constraint for the constant deviation to one action, the terms corresponding to recommendations of that same action cancel, leaving exactly the CE constraint for switching from the other action. The remaining equalities follow from Corollary 10.7. This discrete binary-action statement is distinct from Theorem 10.8 below: with a continuum of actions on an interval, the same cancellation does not identify CCE with CE.

10.4 Full Correlated Equilibrium on Intervals

On an interval, every smooth scalar displacement has an antiderivative. Under the joint continuity assumption on the losses, smooth feasible deviations are also sufficiently rich to approximate all measurable CE deviations.

Theorem 10.8 (Full CE on one-dimensional action spaces). *Suppose that every player has a compact interval action set $K_i := [a_i, b_i] \subseteq \mathbb{R}$ and, in addition to the standing assumptions, every loss $\ell_i : K \rightarrow \mathbb{R}$ is continuous on $K := \prod_{j=1}^N K_j$. Then*

$$\text{CE} = \text{ConCE} = \text{PCE}$$

under the unrestricted definitions used in this section when the approximation error is zero.

The proof of Theorem 10.8 is given in Appendix D.14. In particular, infinite support alone does not separate these equilibrium notions. Under continuous losses, such a separation requires at least one player’s individual action set to have affine dimension at least two.

Two related results help place the statement above. In finite two-action games, Monnot and Piliouras [2017, Proposition 6.1] show that $\text{CE} = \text{CCE}$ through a cancellation specific to the two available actions. Although mixed strategies in that setting admit a one-dimensional parameterization as $(q, 1 - q)$, the argument relies on the underlying binary action structure and does not extend to interval action spaces, where each player has a continuum of possible deviations. In particular, we do not obtain an analogous equality with CCE. Separately, Piliouras et al. [2022b, Theorem 5.2] study continuous-time replicator dynamics with two actions and obtain a localized regret guarantee under a condition on how the resulting one-dimensional trajectory crosses an interval. Their main game-theoretic results likewise focus on two-player 2×2 games. The result below differs in both scope and method. It applies to any number of players with continuous interval action spaces and does not depend on continuous-time dynamics or on any particular learning trajectory. Instead, it is an equilibrium-set statement for arbitrary Borel distributions, proved by approximating measurable deviations with smooth feasible exact-form maps.

10.5 Strict Separation from CE on Continuous Support

Finite interpolation shows that any separation from CE must involve infinite support. The following result establishes such a separation and gives it a simple variational interpretation.

The construction uses the diagonal law (X, X) with X uniform on a disk. Thus the marginals are absolutely continuous with respect to two-dimensional Lebesgue measure and the joint support is uncountable, while the joint law itself is singular in the product action space. Accordingly, the result does not establish a separation for a joint distribution with a full-dimensional density.

Theorem 10.9. *There is a two-player convex game and a distribution μ with uncountable support—indeed, with marginals absolutely continuous with respect to two-dimensional Lebesgue measure—for which*

$$\mu \in \text{ConCE} = \text{PCE} \quad \text{but} \quad \mu \notin \text{CE}.$$

Consequently $\text{CE} \subsetneq \text{ConCE} = \text{PCE}$ under the definitions above.

The proof of Theorem 10.9 is given in Appendix D.14. It constructs a conditional loss-gradient field that is invisible to exact-form deviations but can be exploited by unrestricted deviations. The construction also illustrates why a simple localization argument cannot establish CE equality on general continuous supports. We defer the question of whether a similar strict separation can be obtained with an absolutely continuous, full-dimensional joint law to future work.

11 Conclusion

This paper develops a geometric theory of the action-dependent comparators controlled by projected first-order methods. Its central principle is that exactness determines the interior behavior of a deviation, while feasibility controls its interaction with the boundary. In Euclidean geometry, exact feasible displacement fields yield explicit regret bounds for projected gradient descent, whereas nonzero circulation can be exploited to force linear regret. The same exactness-versus-circulation principle extends to mirror geometry, where the regularizer determines the relevant one-form and the additional regularity needed to control boundary residuals.

This viewpoint also clarifies several distinctions among deviation classes, algorithms, and equilibrium notions. Proximal deviations and their identity interpolations form a subclass of exact-form deviations, but the Euclidean radial-closure representation does not extend to general Bregman geometry: convexity of the reconstructed dual potential provides an additional obstruction. Mirror descent and FTRL likewise agree in the Legendre interior but can control different deviation classes once constraints create nontrivial boundary effects. At the equilibrium level, strict inclusion of deviation maps need not produce strict inclusion of zero-tolerance equilibrium sets. In particular, we showed that unrestricted ConCE and PCE coincide, and they coincide with CE on finite supports and on one-dimensional action spaces with continuous losses, while a two-dimensional continuous-support example yields

$$\text{CE} \subsetneq \text{ConCE} = \text{PCE}.$$

Two questions appear particularly natural. First, can the relative-curvature, Lipschitz, and ambient regularity assumptions in the positive mirror-descent result be weakened while retaining control of the boundary residual and hence sublinear regret? A sharp answer would bring the sufficient conditions closer to the circulation-based obstruction, which requires only continuity of the field. Second, can the strict continuous-support equilibrium separation be realized by a full-dimensional absolutely continuous joint distribution? The present construction has absolutely continuous marginals but is supported on a lower-dimensional diagonal, so resolving this question would determine whether the separation persists for genuinely diffuse outcome distributions.

More broadly, the exact-form viewpoint suggests connections with Blackwell approachability [Blackwell, 1956, Abernethy et al., 2011] and gradient equilibrium [Angelopoulos et al., 2025, Lee et al., 2026]. All three frameworks exploit geometric structure to identify quantities that first-order dynamics can drive toward a target set or make telescope over time. Understanding these connections may lead to alternative characterizations of exact-form regret and to broader classes of deviations controlled by simple first-order methods. A related direction is to combine exactness with lifting [Farina et al., 2022a, Soleymani et al., 2025b]. Deviations that are not exact in the original action space may admit an exact representation in a higher-dimensional feature space, potentially allowing gradient-based algorithms to control richer families such as general linear endomorphisms. Developing such exact lifts, and relating them to approachability and gradient equilibrium, could provide a bridge between the geometric simplicity of the present framework and more general Φ -regret guarantees.

Acknowledgments

This research was partially supported by the Office of Naval Research (ONR) grants N00014-24-1-2470 and N00014-25-1-2296, National Science Foundation awards CCF-2443068 and IIS-2552046, and a Schmidt

Sciences AI2050 Early Career Fellowship. Part of this research was performed while A.S. was visiting the Institute for Mathematical and Statistical Innovation (IMSI), which is supported by the National Science Foundation (Grant No. DMS-2425650).

References

- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, 2003.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4): 157–325, 2016.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Amy Greenwald and Amir Jafari. A general class of no-regret learning algorithms and game-theoretic equilibria. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003. Proceedings*, pages 2–12. Springer, 2003.
- Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8(6), 2007.
- Geoffrey J Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of the 25th international conference on Machine learning*, pages 360–367, 2008.
- Mete Şeref Ahunbay. First-order (coarse) correlated equilibria in non-concave games. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, 2026.
- Mete Şeref Ahunbay and Martin Bichler. Semicoarse correlated equilibria and lp-based guarantees for gradient dynamics in normal-form games. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 91–91, 2025.
- Yang Cai, Constantinos Daskalakis, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. Proximal regret and proximal correlated equilibria: A new tractable solution concept for online learning and games. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pages 1639–1650, 2026.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pages 993–1019. PMLR, 2013a.
- Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. *Advances in Neural Information Processing Systems*, 26, 2013b.
- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. *Advances in Neural Information Processing Systems*, 28, 2015.
- Xi Chen and Binghui Peng. Hedging in games: Faster convergence of external and swap regrets. *Advances in Neural Information Processing Systems*, 33:18990–18999, 2020.

- Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. Near-optimal no-regret learning in general games. *Advances in Neural Information Processing Systems*, 34:27604–27616, 2021.
- Gabriele Farina, Ioannis Anagnostides, Haipeng Luo, Chung-Wei Lee, Christian Kroer, and Tuomas Sandholm. Near-optimal no-regret learning dynamics for general convex games. *Advances in Neural Information Processing Systems*, 35:39076–39089, 2022a.
- Georgios Piliouras, Ryann Sim, and Stratis Skoulakis. Beyond time-average convergence: Near-optimal uncoupled online learning via clairvoyant multiplicative weights update. *Advances in Neural Information Processing Systems*, 35:22258–22269, 2022a.
- Ashkan Soleymani, Georgios Piliouras, and Gabriele Farina. Faster rates for no-regret learning in general games via cautious optimism. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 518–529, 2025a.
- Ashkan Soleymani, Georgios Piliouras, and Gabriele Farina. Cautious optimism: A meta-algorithm for near-constant regret in general games. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 870–870, 2025b.
- Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1–2):40–55, 1997.
- Gabriele Farina, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Simple uncoupled no-regret learning dynamics for extensive-form correlated equilibrium. *Journal of the ACM*, 69(6):41:1–41:41, 2022b. doi: 10.1145/3563772.
- Dustin Morrill, Ryan D’Orazio, Marc Lanctot, James R. Wright, Michael Bowling, and Amy R. Greenwald. Efficient deviation types and learning for hindsight rationality in extensive-form games. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- Kaito Fujii. Bayes correlated equilibria, no-regret dynamics in bayesian games, and the price of anarchy. In *The Thirty Eighth Annual Conference on Learning Theory*, pages 2190–2191. PMLR, 2025.
- Binghui Peng and Aviad Rubinstein. Fast swap regret minimization and applications to approximate correlated equilibria. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1223–1234, 2024.
- Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. From external to swap regret 2.0: An efficient reduction for large action spaces. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1216–1222, 2024.
- Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Charilaos Pipis, and Jon Schneider. Efficient learning and computation of linear correlated equilibrium in general convex games. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 542–553, 2025.
- Brian Hu Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, and Tuomas Sandholm. Learning and computation of ϕ -equilibria at the frontier of tractability. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 1105–1105, 2025.
- Yang Cai, Constantinos Daskalakis, Haipeng Luo, Chen-Yu Wei, and Weiqiang Zheng. On tractable Φ -equilibria in non-concave games. *Advances in Neural Information Processing Systems*, 37:140366–140404, 2024.
- Anastasios N. Angelopoulos, Michael I. Jordan, and Ryan J. Tibshirani. Gradient equilibrium in online learning: Theory and applications. *Journal of Machine Learning Research*, 26(305):1–68, 2025.

- Jean Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France*, 93:273–299, 1965.
- Rémi Gribonval and Mila Nikolova. A characterization of proximity operators. *Journal of Mathematical Imaging and Vision*, 62(6):773–789, 2020.
- Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and l1 regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 525–533. JMLR Workshop and Conference Proceedings, 2011.
- H Brendan McMahan. A survey of algorithms and analysis for adaptive online learning. *Journal of Machine Learning Research*, 18(90):1–50, 2017.
- Lin Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 2010.
- Huang Fang, Nicholas JA Harvey, Victor S Portella, and Michael P Friedlander. Online mirror descent and dual averaging: keeping pace in the dynamic case. *Journal of Machine Learning Research*, 23(121):1–38, 2022.
- Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974. doi: 10.1016/0304-4068(74)90037-8.
- Barnabé Monnot and Georgios Piliouras. Limits and limitations of no-regret learning in games. *The Knowledge Engineering Review*, 32:e21, 2017. doi: 10.1017/S0269888917000133.
- Georgios Piliouras, Mark Rowland, Shayegan Omidshafiei, Romuald Elie, Daniel Hennes, Jerome Connor, and Karl Tuyls. Evolutionary dynamics and Φ -regret minimization in games. *Journal of Artificial Intelligence Research*, 74:1125–1158, 2022b.
- David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1): 1–8, 1956.
- Jacob Abernethy, Peter L Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of the 24th Annual Conference on Learning Theory*, 2011.
- Brian W. Lee, Nika Haghtalab, Michael I. Jordan, and Ryan J. Tibshirani. Blackwell approachability and gradient equilibrium are equivalent. In *Proceedings of Thirty Ninth Conference on Learning Theory*, 2026.

A Preliminaries

Lemma A.1. *Let $\mathbf{u} \in \mathbb{R}^d$ and $\mathbf{x}^+ := \Pi_K(\mathbf{u})$. Then*

$$\mathbf{u} - \mathbf{x}^+ \in N_K(\mathbf{x}^+),$$

Proof. The point $\mathbf{x}^+$ minimizes $\mathbf{y} \mapsto \frac{1}{2} \|\mathbf{y} - \mathbf{u}\|^2$ over K . The first-order optimality condition is

$$\langle \mathbf{x}^+ - \mathbf{u}, \mathbf{y} - \mathbf{x}^+ \rangle \geq 0 \quad \forall \mathbf{y} \in K.$$

Equivalently, $\langle \mathbf{u} - \mathbf{x}^+, \mathbf{y} - \mathbf{x}^+ \rangle \leq 0$ for all $\mathbf{y} \in K$, which is the normal-cone condition. $\square$

Definition A.2. *For $\rho \geq 0$, a proper lower-semicontinuous function $F : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ is ρ -weakly convex if*

$$\mathbf{x} \mapsto F(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{x}\|^2$$

is convex. We use the weakly convex subdifferential convention

$$\partial F(\mathbf{x}) := \partial_{\text{cvx}} \left(F + \frac{\rho}{2} \|\cdot\|^2 \right) (\mathbf{x}) - \rho \mathbf{x},$$

where ∂_{cvx} is the convex subdifferential. With this convention, ∂F is ρ -hypomonotone: for all $\mathbf{u} \in \partial F(\mathbf{p})$ and $\mathbf{v} \in \partial F(\mathbf{q})$,

$$\langle \mathbf{u} - \mathbf{v}, \mathbf{p} - \mathbf{q} \rangle \geq -\rho \|\mathbf{p} - \mathbf{q}\|^2. \quad (15)$$

Lemma A.3. *Assume $F := f + \iota_K$ is proper, lower semicontinuous, and ρ -weakly convex with $0 \leq \rho < 1$. Then, for every $\mathbf{x} \in \mathbb{R}^d$, the proximal objective attains its minimum at a unique point. Moreover, M_F is differentiable, and*

$$\nabla M_F(\mathbf{x}) = \mathbf{x} - \text{prox}_F(\mathbf{x}). \quad (16)$$

Moreover,

$$\|\text{prox}_F(\mathbf{x}) - \text{prox}_F(\mathbf{y})\| \leq \frac{1}{1 - \rho} \|\mathbf{x} - \mathbf{y}\|, \quad (17)$$

so ∇M_F is Lipschitz with the crude bound

$$\|\nabla M_F(\mathbf{x}) - \nabla M_F(\mathbf{y})\| \leq \frac{2 - \rho}{1 - \rho} \|\mathbf{x} - \mathbf{y}\|. \quad (18)$$

Proof. For fixed $\mathbf{x}$, the function $\mathbf{y} \mapsto F(\mathbf{y}) + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2$ is proper, lower semicontinuous, coercive, and $(1 - \rho)$ -strongly convex. It therefore attains its minimum at a unique point. Let

$$\mathbf{p} = \text{prox}_F(\mathbf{x}), \quad \mathbf{q} = \text{prox}_F(\mathbf{y}).$$

The first-order optimality conditions are

$$\mathbf{x} - \mathbf{p} \in \partial F(\mathbf{p}), \quad \mathbf{y} - \mathbf{q} \in \partial F(\mathbf{q}).$$

Applying hypomonotonicity (15) gives

$$\begin{aligned} \langle (\mathbf{x} - \mathbf{p}) - (\mathbf{y} - \mathbf{q}), \mathbf{p} - \mathbf{q} \rangle &\geq -\rho \|\mathbf{p} - \mathbf{q}\|^2, \\ \langle \mathbf{x} - \mathbf{y}, \mathbf{p} - \mathbf{q} \rangle - \|\mathbf{p} - \mathbf{q}\|^2 &\geq -\rho \|\mathbf{p} - \mathbf{q}\|^2. \end{aligned}$$

Thus

$$(1 - \rho) \|p - q\|^2 \leq \langle x - y, p - q \rangle \leq \|x - y\| \|p - q\|,$$

which proves (17).

For fixed y , the derivative with respect to x of $F(y) + \frac{1}{2} \|y - x\|^2$ is $x - y$. Since the minimizer is unique, Danskin's theorem gives $\nabla M_F(x) = x - \text{prox}_F(x)$. Finally,

$$\begin{aligned} \|\nabla M_F(x) - \nabla M_F(y)\| &= \|(x - y) - (\text{prox}_F(x) - \text{prox}_F(y))\| \\ &\leq \|x - y\| + \|\text{prox}_F(x) - \text{prox}_F(y)\| \\ &\leq \left(1 + \frac{1}{1 - \rho}\right) \|x - y\|. \end{aligned}$$

This is (18). □

In turn, we recall that weakly-convex proximal maps are monotone.

Lemma A.4. *Let $F : \mathbb{R}^d \rightarrow (-\infty, +\infty]$ be a proper lower-semicontinuous ρ -weakly convex function with $0 \leq \rho < 1$. Then prox_F is monotone in the sense that, for all $x, y \in \mathbb{R}^d$,*

$$\langle x - y, \text{prox}_F(x) - \text{prox}_F(y) \rangle \geq 0.$$

Proof. Fix $x, y \in \mathbb{R}^d$. The proximal optimality conditions give

$$x - \text{prox}_F(x) \in \partial F(\text{prox}_F(x)), \quad y - \text{prox}_F(y) \in \partial F(\text{prox}_F(y)).$$

Since F is ρ -weakly convex, its subdifferential is ρ -hypomonotone. Hence

$$\langle (x - \text{prox}_F(x)) - (y - \text{prox}_F(y)), \text{prox}_F(x) - \text{prox}_F(y) \rangle \geq -\rho \|\text{prox}_F(x) - \text{prox}_F(y)\|^2.$$

Rearranging gives

$$\langle x - y, \text{prox}_F(x) - \text{prox}_F(y) \rangle \geq (1 - \rho) \|\text{prox}_F(x) - \text{prox}_F(y)\|^2 \geq 0.$$

Thus prox_F is monotone. □

B Scaled Exact Deviations and Proximal Representation

Lemma B.1. *Let K be compact and let Ψ have locally Lipschitz gradient on an open neighborhood of K . Then Ψ agrees on a neighborhood of K with a globally defined differentiable function $\tilde{\Psi}$ whose gradient is globally Lipschitz.*

Proof. Choose open sets $K \subset V \Subset W$ with $\overline{W}$ inside the domain of Ψ , and a smooth cutoff χ equal to one on V and supported in W . The function $\tilde{\Psi} := \chi\Psi$ on the original domain, extended by zero outside it, has the required properties. Compactness of $\overline{W}$ gives a finite Lipschitz constant for its gradient. □

Theorem 6.2. *Let $K \subseteq \mathbb{R}^d$ be compact and convex, and let $\phi(x) := x - \nabla\Psi(x)$ be a feasible exact-form deviation with $\Psi \in C_{\text{loc}}^{1,1}$ on a neighborhood of K . Then there exists $\alpha_0 > 0$ such that, for every $\alpha \in (0, \alpha_0]$, the interpolated map*

$$\phi_\alpha(x) := (1 - \alpha)x + \alpha\phi(x) = x - \alpha\nabla\Psi(x)$$

coincides on K with $\text{prox}_{F_\alpha}(x)$ for some proper lower-semicontinuous $F_\alpha = f_\alpha + \iota_K$ that is ρ_α -weakly convex for some $\rho_\alpha < 1$.

Proof. By Lemma B.1, replace Ψ outside a neighborhood of K so that $\nabla\Psi$ is globally L -Lipschitz. If $L = 0$, set $\alpha_0 := 1/2$; otherwise set $\alpha_0 := \frac{1}{2} \min\{1, 1/L\}$. Fix any $\alpha \in (0, \alpha_0]$ and set

$$h_\alpha(\mathbf{x}) := \frac{1}{2} \|\mathbf{x}\|^2 - \alpha\Psi(\mathbf{x}), \quad f_\alpha(\mathbf{y}) := h_\alpha^*(\mathbf{y}) - \frac{1}{2} \|\mathbf{y}\|^2.$$

The function h_α is $(1 - \alpha L)$ -strongly convex and $(1 + \alpha L)$ -smooth. Hence h_α^* is $1/(1 + \alpha L)$ -strongly convex and f_α is ρ_α -weakly convex for $\rho_\alpha := \alpha L/(1 + \alpha L) < 1$. Put $F_\alpha := f_\alpha + \iota_K$.

For fixed $\mathbf{x}$, the unconstrained proximal objective equals, up to the constant $\|\mathbf{x}\|^2/2$,

$$h_\alpha^*(\mathbf{y}) - \langle \mathbf{x}, \mathbf{y} \rangle.$$

Its unique minimizer is $\mathbf{y} = \nabla h_\alpha(\mathbf{x}) = \mathbf{x} - \alpha \nabla \Psi(\mathbf{x}) = \phi_\alpha(\mathbf{x})$. If $\mathbf{x} \in K$, convexity of K and feasibility of ϕ give $\phi_\alpha(\mathbf{x}) := (1 - \alpha)\mathbf{x} + \alpha\phi(\mathbf{x}) \in K$. The constraint therefore does not change the minimizer, proving $\text{prox}_{F_\alpha} = \phi_\alpha$ on K . $\square$

Corollary 6.3. *For the radial closure of Euclidean proximal deviations, we have $\text{Rad}_I(\Phi_{\text{prox}}) = \Phi_{\text{exact}}^{\text{sm}}$.*

Proof. Let $\phi \in \Phi_{\text{exact}}^{\text{sm}}$. By Theorem 6.2, there is $\alpha > 0$ for which $\phi_\alpha = (1 - \alpha)I + \alpha\phi$ is proximal. Hence $\phi \in \text{Rad}_I(\Phi_{\text{prox}})$.

Conversely, let $\phi \in \text{Rad}_I(\Phi_{\text{prox}})$. Choose $\alpha \in (0, 1]$ and a proper lower-semicontinuous weakly convex function F with $\phi_\alpha = \text{prox}_F$ on K . By Theorem 5.1,

$$\mathbf{x} - \phi_\alpha(\mathbf{x}) = \nabla M_F(\mathbf{x}),$$

and ∇M_F is Lipschitz. Since $\mathbf{x} - \phi_\alpha(\mathbf{x}) = \alpha(\mathbf{x} - \phi(\mathbf{x}))$, we have

$$\mathbf{x} - \phi(\mathbf{x}) = \nabla \left(\frac{1}{\alpha} M_F \right) (\mathbf{x}).$$

Thus ϕ is smooth exact-form, proving the reverse inclusion. $\square$

C Boundary-Compatible Finite Interpolation

Lemma C.1. *For finite data $(\mathbf{z}^j, \mathbf{s}^j)_{j=1}^m$, there is a proper closed convex function H with $\mathbf{s}^j \in \partial H(\mathbf{z}^j)$ for every j if and only if the data are cyclically monotone, namely*

$$\sum_{r=1}^q \langle \mathbf{s}^{j_r}, \mathbf{z}^{j_{r+1}} - \mathbf{z}^{j_r} \rangle \leq 0 \quad \text{for every cycle } j_1, \dots, j_q, j_{q+1} = j_1.$$

It is enough to check nontrivial simple cycles.

Proof. Necessity follows by summing the subgradient inequalities around a cycle. For sufficiency, form the complete directed graph on $[m]$ with edge weight $w(j, k) = \langle \mathbf{s}^j, \mathbf{z}^k - \mathbf{z}^j \rangle$. Add a source vertex 0 with a zero-weight edge to every j , and let

$$a_k = \max \left\{ \sum_{r=0}^{q-1} w(j_r, j_{r+1}) : j_q = k, j_0 \in [m], \text{ and the path is simple} \right\}.$$

The maximum is finite because there are finitely many simple paths. Any walk can be reduced to a simple path by deleting closed subwalks; cyclic monotonicity says each deleted closed subwalk has nonpositive weight. Hence a_k is also the supremum over all directed walks ending at k . Appending the edge (j, k) to a walk ending at j therefore gives

$$a_k \geq a_j + w(j, k) = a_j + \langle \mathbf{s}^j, \mathbf{z}^k - \mathbf{z}^j \rangle \quad (j, k \in [m]).$$

Define the finite polyhedral function

$$H(\mathbf{z}) := \max_{j \in [m]} \{a_j + \langle \mathbf{s}^j, \mathbf{z} - \mathbf{z}^j \rangle\}.$$

It is proper, closed, and convex. At $\mathbf{z}^k$, the preceding inequality shows that the k th affine function has value a_k and dominates every j th affine function. Thus $H(\mathbf{z}^k) = a_k$ and $\mathbf{s}^k \in \partial H(\mathbf{z}^k)$. Finally, every closed walk decomposes into simple cycles, so checking nontrivial simple cycles is sufficient. $\square$

Theorem 10.6. *Let $K \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Let $\mathbf{x}^1, \dots, \mathbf{x}^m \in K$ be distinct and choose arbitrary targets $\mathbf{y}^1, \dots, \mathbf{y}^m \in K$. Fix $\rho \in (0, 1)$. There is $\alpha_0 > 0$ such that for every $\alpha \in (0, \alpha_0]$ one can find a proper lower-semicontinuous $F_\alpha = f_\alpha + \iota_K$ that is ρ -weakly convex and satisfies*

$$\text{prox}_{F_\alpha}(\mathbf{x}^j) = (1 - \alpha)\mathbf{x}^j + \alpha\mathbf{y}^j, \quad j = 1, \dots, m.$$

The proximal minimizer is unique.

Proof. Write $\mathbf{d}^j := \mathbf{x}^j - \mathbf{y}^j$ and define

$$\mathbf{z}_\alpha^j := \mathbf{x}^j - \alpha\mathbf{d}^j, \quad \mathbf{s}_\alpha^j := \mathbf{x}^j - (1 - \rho)\mathbf{z}_\alpha^j.$$

Convexity of K gives $\mathbf{z}_\alpha^j \in K$ for $\alpha \in [0, 1]$. At $\alpha = 0$, $\mathbf{z}_0^j = \mathbf{x}^j$ and $\mathbf{s}_0^j = \rho\mathbf{x}^j$. For every nontrivial simple cycle C ,

$$\sum_{(j,k) \in C} \langle \mathbf{s}_0^j, \mathbf{z}_0^k - \mathbf{z}_0^j \rangle = -\frac{\rho}{2} \sum_{(j,k) \in C} \|\mathbf{x}^k - \mathbf{x}^j\|^2 < 0.$$

There are finitely many simple cycles and their cycle sums vary continuously with α . Thus, for all sufficiently small common $\alpha > 0$, the data $(\mathbf{z}_\alpha^j, \mathbf{s}_\alpha^j)$ are cyclically monotone. By Lemma C.1, choose a proper closed convex H_α with $\mathbf{s}_\alpha^j \in \partial H_\alpha(\mathbf{z}_\alpha^j)$.

Set

$$f_\alpha(\mathbf{z}) := H_\alpha(\mathbf{z}) - \frac{\rho}{2} \|\mathbf{z}\|^2, \quad F_\alpha := f_\alpha + \iota_K.$$

Then F_α is ρ -weakly convex. Moreover, using the zero vector from $N_K(\mathbf{z}_\alpha^j)$,

$$0 = \mathbf{s}_\alpha^j - \rho\mathbf{z}_\alpha^j + \mathbf{z}_\alpha^j - \mathbf{x}^j \in \partial \left(F_\alpha(\mathbf{z}) + \frac{1}{2} \|\mathbf{z} - \mathbf{x}^j\|^2 \right)_{\mathbf{z}=\mathbf{z}_\alpha^j}.$$

The proximal objective is $(1 - \rho)$ -strongly convex, so this minimizer is unique and equals $\mathbf{z}_\alpha^j = (1 - \alpha)\mathbf{x}^j + \alpha\mathbf{y}^j$. $\square$

D Other Proofs and Details

D.1 Proofs for Section 4

First, we show that the feasibility of exact-form deviations controls the normal vectors that are introduced due to the projection step in projected online gradient descent (OGD).

Lemma D.1. *Let $\phi : K \rightarrow K$, let $\mathbf{D}_\phi(\mathbf{x}) := \mathbf{x} - \phi(\mathbf{x})$, and let $\mathbf{n} \in N_K(\mathbf{x})$. Then*

$$\langle \mathbf{n}, \mathbf{D}_\phi(\mathbf{x}) \rangle \geq 0.$$

Proof. Since $\phi(\mathbf{x}) \in K$ and $\mathbf{n} \in N_K(\mathbf{x})$,

$$\langle \mathbf{n}, \phi(\mathbf{x}) - \mathbf{x} \rangle \leq 0.$$

Multiplying by -1 gives the statement. $\square$

Theorem 4.1. Let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online gradient descent (OGD). Let $\phi : K \rightarrow K$ be an exact-form deviation $\mathbf{D}_\phi(\mathbf{x}) = \mathbf{x} - \phi(\mathbf{x}) = \nabla \Psi(\mathbf{x})$ where Ψ is L -smooth with respect to the Euclidean norm. Then, for every sequence $\mathbf{g}_1, \dots, \mathbf{g}_T$,

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{\Psi(\mathbf{x}_1) - \Psi(\mathbf{x}_{T+1})}{\eta} + \frac{3L\eta}{2} \sum_{t=1}^T \|\mathbf{g}_t\|^2. \quad (1)$$

Proof. Fix a round t and define

$$\mathbf{s}_t := \mathbf{x}_t - \mathbf{x}_{t+1}, \quad \mathbf{n}_t := \mathbf{x}_t - \eta \mathbf{g}_t - \mathbf{x}_{t+1}.$$

By Lemma A.1, $\mathbf{n}_t \in N_K(\mathbf{x}_{t+1})$. Also

$$\eta \mathbf{g}_t = \mathbf{s}_t - \mathbf{n}_t.$$

Thus,

$$\eta \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle = \langle \mathbf{s}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle - \langle \mathbf{n}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle. \quad (19)$$

We bound the two terms separately. First, L -smoothness of Ψ gives

$$\begin{aligned} \Psi(\mathbf{x}_{t+1}) &\leq \Psi(\mathbf{x}_t) + \langle \nabla \Psi(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle + \frac{L}{2} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &= \Psi(\mathbf{x}_t) - \langle \mathbf{D}_\phi(\mathbf{x}_t), \mathbf{s}_t \rangle + \frac{L}{2} \|\mathbf{s}_t\|^2. \end{aligned}$$

Therefore

$$\langle \mathbf{s}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \Psi(\mathbf{x}_t) - \Psi(\mathbf{x}_{t+1}) + \frac{L}{2} \|\mathbf{s}_t\|^2. \quad (20)$$

Second, by Lemma D.1 applied at $\mathbf{x}_{t+1}$,

$$\langle \mathbf{n}_t, \mathbf{D}_\phi(\mathbf{x}_{t+1}) \rangle \geq 0.$$

Thus

$$\begin{aligned} -\langle \mathbf{n}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle &= -\langle \mathbf{n}_t, \mathbf{D}_\phi(\mathbf{x}_{t+1}) \rangle - \langle \mathbf{n}_t, \mathbf{D}_\phi(\mathbf{x}_t) - \mathbf{D}_\phi(\mathbf{x}_{t+1}) \rangle \\ &\leq \|\mathbf{n}_t\| \|\mathbf{D}_\phi(\mathbf{x}_t) - \mathbf{D}_\phi(\mathbf{x}_{t+1})\| \\ &\leq L \|\mathbf{n}_t\| \|\mathbf{s}_t\|. \end{aligned}$$

Projection is nonexpansive and $\mathbf{x}_t = \Pi_K(\mathbf{x}_t)$, so

$$\|\mathbf{s}_t\| = \|\Pi_K(\mathbf{x}_t) - \Pi_K(\mathbf{x}_t - \eta \mathbf{g}_t)\| \leq \eta \|\mathbf{g}_t\|. \quad (21)$$

Moreover,

$$\|\mathbf{n}_t\| = \text{dist}(\mathbf{x}_t - \eta \mathbf{g}_t, K) \leq \|\mathbf{x}_t - \eta \mathbf{g}_t - \mathbf{x}_t\| = \eta \|\mathbf{g}_t\|.$$

Therefore

$$-\langle \mathbf{n}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq L\eta^2 \|\mathbf{g}_t\|^2. \quad (22)$$

Combining (19), (20), and (22), and using $\|\mathbf{s}_t\| \leq \eta \|\mathbf{g}_t\|$, yields

$$\eta \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \Psi(\mathbf{x}_t) - \Psi(\mathbf{x}_{t+1}) + \frac{3L\eta^2}{2} \|\mathbf{g}_t\|^2.$$

Dividing by η and summing over $t \in [T]$ telescopes the potential terms and proves (1). $\square$

Corollary 4.3. Suppose $\|\mathbf{g}_t\| \leq G$ and fix $B, L > 0$. Then OGD with $\eta = \sqrt{2B/(3LG^2T)}$ satisfies

$$\sup_{\phi \in \Phi_{\text{exact}}(B, L)} \text{Reg}_T(\phi) \leq G\sqrt{6LBT}.$$

Thus the algorithm controls the whole normalized class $\Phi_{\text{exact}}(B, L)$ with one learning rate.

Proof. For every $\phi \in \Phi_{\text{exact}}(B, L)$, the proof of Theorem 4.1 applies the Lipschitz-gradient condition in Definition 3.2. Its telescoping term is bounded by B , uniformly in ϕ , while $\sum_{t=1}^T \|\mathbf{g}_t\|^2 \leq G^2T$. Hence

$$\text{Reg}_T(\phi) \leq \frac{B}{\eta} + \frac{3L\eta G^2T}{2} \quad \text{for every } \phi \in \Phi_{\text{exact}}(B, L).$$

Taking the supremum and substituting the displayed value of η proves the claim. This common choice of η is precisely why the normalization is needed. $\square$

D.2 Proofs for Section 5

Theorem 5.1. Let $F := f + \iota_K$ be proper, lower semicontinuous, and ρ -weakly convex with $0 \leq \rho < 1$. Define $\phi_F(\mathbf{x}) := \text{prox}_F(\mathbf{x})$. Then $\phi_F : K \rightarrow K$ is feasible and Euclidean exact-form, with potential M_F , i.e., $\mathbf{x} - \phi_F(\mathbf{x}) = \nabla M_F(\mathbf{x})$, whose gradient is $L_F = (2 - \rho)/(1 - \rho)$ -Lipschitz. Consequently projected gradient descent controls every such proximal deviation by Theorem 4.1. In particular, if K is compact, $\|\mathbf{g}_t\| \leq G$, $\text{osc}_K(M_F) \leq B_F$, and $B_F, G > 0$, then choosing $\eta = \sqrt{2B_F/(3L_FG^2T)}$ gives

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \text{prox}_F(\mathbf{x}_t) \rangle \leq \frac{B_F}{\eta} + \frac{3L_F\eta G^2T}{2} = \mathcal{O}(\sqrt{T}).$$

If $B_F = 0$ or $G = 0$, the corresponding regret bound is trivial and the displayed optimizing step size need not be used.

Proof. Feasibility holds because $F = f + \iota_K$ forces $\text{prox}_F(\mathbf{x}) \in K$. Exactness is exactly (16). The regret bound follows by applying Theorem 4.1 with $\Psi = M_F$ and the smoothness bound from Lemma A.3. $\square$

Proposition 5.2. The exact-form class contains the following deviations.

1. **External regret.** For $\mathbf{u} \in K$, take $F := \iota_{\{\mathbf{u}\}}$. Then $\text{prox}_F(\mathbf{x}) = \mathbf{u}$ and $\mathbf{x} - \mathbf{u} = \nabla \frac{1}{2} \|\mathbf{x} - \mathbf{u}\|^2$.
2. **Projection to a convex subset.** For a nonempty closed convex set $S \subseteq K$, take $F := \iota_S$. Then $\text{prox}_F(\mathbf{x}) = \Pi_S(\mathbf{x})$.
3. **Projection-based/no-move regret.** For a vector $\mathbf{v}$, take $F(\mathbf{y}) := \langle \mathbf{v}, \mathbf{y} \rangle + \iota_K(\mathbf{y})$. Then $\text{prox}_F(\mathbf{x}) = \Pi_K(\mathbf{x} - \mathbf{v})$.
4. **Interpolation deviations.** For $\mathbf{u} \in K$ and $\alpha \in (0, 1)$, the map $\phi(\mathbf{x}) := (1 - \alpha)\mathbf{x} + \alpha\mathbf{u}$ is the prox map of

$$F(\mathbf{y}) := \frac{\alpha}{2(1 - \alpha)} \|\mathbf{y} - \mathbf{u}\|^2 + \iota_K(\mathbf{y}).$$

The case $\alpha = 1$ is external regret.

5. **Local proximal deviations.** If $\|\mathbf{x} - \text{prox}_F(\mathbf{x})\| \leq \delta$ on K , then the deviation belongs to the local exact-form subclass $\|\nabla M_F(\mathbf{x})\| \leq \delta$.
6. **Symmetric affine swaps.** Let $\phi(\mathbf{x}) := A\mathbf{x} + \mathbf{b}$ map K into K . Let $T := \text{span}(K - K)$ be the direction space of $\text{aff}(K)$, fix any $\mathbf{a} \in K$, and assume $AT \subseteq T$ and the restriction of $I - A$ to T is self-adjoint. Then ϕ is exact-form in the ambient sense of our definition. On $\text{aff}(K)$, a potential is

$$\Psi(\mathbf{a} + \mathbf{h}) := \langle D_\phi(\mathbf{a}), \mathbf{h} \rangle + \frac{1}{2} \langle \mathbf{h}, (I - A)\mathbf{h} \rangle, \quad \mathbf{h} \in T. \quad (3)$$

An ambient extension of this potential is given in the proof.

Proof. Items 1 and 2 are immediate from Theorem 5.1. For item 3, complete the square:

$$\arg \min_{\mathbf{y} \in K} \left\{ \langle \mathbf{v}, \mathbf{y} \rangle + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2 \right\} = \arg \min_{\mathbf{y} \in K} \frac{1}{2} \|\mathbf{y} - (\mathbf{x} - \mathbf{v})\|^2 = \Pi_K(\mathbf{x} - \mathbf{v}).$$

For item 4, the unconstrained first-order condition for

$$\frac{\alpha}{2(1-\alpha)} \|\mathbf{y} - \mathbf{u}\|^2 + \frac{1}{2} \|\mathbf{y} - \mathbf{x}\|^2$$

gives $\mathbf{y} = (1-\alpha)\mathbf{x} + \alpha\mathbf{u}$, which lies in K by convexity. Hence it is also the constrained minimizer. Item 5 is just the Moreau identity. For item 6, write every point in the affine hull as $\mathbf{x} = \mathbf{a} + \mathbf{h}$ with $\mathbf{h} \in T$. Since $\mathbf{a} \in K$ and $\phi(\mathbf{a}) \in K$, the vector $\mathbf{D}_\phi(\mathbf{a}) = \mathbf{a} - \phi(\mathbf{a})$ lies in T . For every tangent direction $\mathbf{r} \in T$,

$$\mathbf{D}_\phi(\mathbf{a} + \mathbf{h} + \mathbf{r}) - \mathbf{D}_\phi(\mathbf{a} + \mathbf{h}) = (I - A)\mathbf{r}.$$

Differentiating (3) in direction $\mathbf{r}$,

$$\begin{aligned} d\Psi(\mathbf{a} + \mathbf{h})[\mathbf{r}] &= \langle \mathbf{D}_\phi(\mathbf{a}), \mathbf{r} \rangle + \langle \mathbf{h}, (I - A)\mathbf{r} \rangle \\ &= \langle \mathbf{D}_\phi(\mathbf{a}) + (I - A)\mathbf{h}, \mathbf{r} \rangle \\ &= \langle \mathbf{D}_\phi(\mathbf{a} + \mathbf{h}), \mathbf{r} \rangle. \end{aligned}$$

The second equality uses self-adjointness on T . Hence the affine-hull gradient of Ψ equals $\mathbf{D}_\phi$. To obtain the ambient potential required by Definition 3.1, let P_T be the orthogonal projection onto T and define, for $\mathbf{z} \in \mathbb{R}^d$,

$$\tilde{\Psi}(\mathbf{z}) := \langle \mathbf{D}_\phi(\mathbf{a}), P_T(\mathbf{z} - \mathbf{a}) \rangle + \frac{1}{2} \langle P_T(\mathbf{z} - \mathbf{a}), (I - A)P_T(\mathbf{z} - \mathbf{a}) \rangle.$$

Self-adjointness on T and $\mathbf{D}_\phi(\mathbf{a}) \in T$ imply $\nabla \tilde{\Psi}(\mathbf{a} + \mathbf{h}) = \mathbf{D}_\phi(\mathbf{a}) + (I - A)\mathbf{h} = \mathbf{D}_\phi(\mathbf{a} + \mathbf{h})$ for every $\mathbf{h} \in T$. Thus $\tilde{\Psi}$ is a smooth ambient extension with the required gradient on K . Feasibility is the assumption $\phi(K) \subseteq K$. $\square$

D.3 Proofs for Section 7

Lemma 7.2. *If $\mathbf{D} = \nabla \Psi$ on a neighborhood of K , then we have $\oint_\gamma \mathbf{D}(\mathbf{x})^\top d\mathbf{x} = 0$ for every closed loop $\gamma \subseteq K$.*

Proof. By the chain rule,

$$\int_0^1 \langle \nabla \Psi(\gamma(s)), \gamma'(s) \rangle ds = \int_0^1 \frac{d}{ds} \Psi(\gamma(s)) ds = \Psi(\gamma(1)) - \Psi(\gamma(0)) = 0.$$

$\square$

Lemma D.2. *Let $\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_m = \mathbf{z}_0$ be a closed polygon in K . Suppose*

$$C := \sum_{i=0}^{m-1} \langle \mathbf{D}(\mathbf{z}_i), \mathbf{z}_i - \mathbf{z}_{i+1} \rangle > 0. \quad (23)$$

Run projected GD with step size η from $\mathbf{x}_1 = \mathbf{z}_0$. For one traversal, choose

$$\mathbf{g}_{i+1} := \frac{\mathbf{z}_i - \mathbf{z}_{i+1}}{\eta}, \quad i = 0, \dots, m-1. \quad (24)$$

Then the iterates follow the polygon exactly and the regret accumulated in one traversal is C/η .

Proof. If $\mathbf{x}_{i+1} = \mathbf{z}_i$, then the unprojected update is

$$\mathbf{x}_{i+1} - \eta \mathbf{g}_{i+1} = \mathbf{z}_i - (\mathbf{z}_i - \mathbf{z}_{i+1}) = \mathbf{z}_{i+1} \in K.$$

Projection leaves it unchanged, so $\mathbf{x}_{i+2} = \mathbf{z}_{i+1}$. This proves the trajectory claim by induction. The regret over the traversal is

$$\sum_{i=0}^{m-1} \langle \mathbf{g}_{i+1}, \mathbf{D}(\mathbf{z}_i) \rangle = \frac{1}{\eta} \sum_{i=0}^{m-1} \langle \mathbf{D}(\mathbf{z}_i), \mathbf{z}_i - \mathbf{z}_{i+1} \rangle = \frac{C}{\eta}.$$

□

Theorem 7.3. *Let $\mathbf{D} : K \rightarrow \mathbb{R}^d$ be continuous. Suppose there exists a closed piecewise C^1 loop $\gamma : [0, 1] \rightarrow K$ such that, after possibly reversing orientation,*

$$-\oint_{\gamma} \mathbf{D}(\mathbf{x})^\top d\mathbf{x} := c_0 > 0.$$

Then for every gradient bound $G > 0$ and every sufficiently small $\eta > 0$, there exists a sequence of losses with $\|\mathbf{g}_t\| \leq G$ such that online gradient descent, initialized at $\gamma(0)$, suffers

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}(\mathbf{x}_t) \rangle \geq cT - O(1/\eta)$$

for a constant $c > 0$ depending only on $\mathbf{D}$, γ , and G . In particular the regret is $\Omega(T)$. If the algorithm starts elsewhere in the same path-connected component of K , the adversary can first steer it to the loop along any fixed polygonal path; this changes only the $O(1/\eta)$ transient term.

Proof. Let L_γ be the length of γ , which is positive because the circulation is nonzero. The quantity

$$-\oint_{\gamma} \mathbf{D}(\mathbf{x})^\top d\mathbf{x}$$

is the limit of left-endpoint Riemann sums

$$\sum_{i=0}^{m-1} \langle \mathbf{D}(\mathbf{z}_i), \mathbf{z}_i - \mathbf{z}_{i+1} \rangle, \quad \mathbf{z}_i = \gamma(s_i),$$

as the mesh of the partition $0 = s_0 < s_1 < \dots < s_m = 1$ goes to zero. Hence, for all sufficiently small η , we may choose a partition with every chord length at most ηG , with

$$\sum_{i=0}^{m-1} \langle \mathbf{D}(\mathbf{z}_i), \mathbf{z}_i - \mathbf{z}_{i+1} \rangle \geq \frac{c_0}{2},$$

and with $m \leq 2L_\gamma/(\eta G)$. For one traversal, define

$$\mathbf{g}_{i+1} := \frac{\mathbf{z}_i - \mathbf{z}_{i+1}}{\eta}, \quad i = 0, \dots, m-1.$$

Then $\|\mathbf{g}_{i+1}\| \leq G$, and Lemma D.2 gives regret at least $c_0/(2\eta)$ over one traversal. Therefore the average regret per step during complete cycles is at least

$$\frac{c_0/(2\eta)}{2L_\gamma/(\eta G)} = \frac{c_0 G}{4L_\gamma}.$$

Repeat the cycle for $\lfloor T/m \rfloor$ complete traversals and set the remaining gradients, if any, to zero. This gives

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}(\mathbf{x}_t) \rangle \geq \frac{c_0 G}{4L_\gamma} T - O(1/\eta),$$

which proves the claim. $\square$

Corollary 7.5. *Let $E_K := \text{span}(K - K)$, let $U := \text{relint}(K)$ be open and simply connected in $\text{aff}(K)$, and let $\mathbf{D} : U \rightarrow E_K$ be C^1 . In any orthonormal affine coordinate system on E_K , $\partial_j \mathbf{D}_i = \partial_i \mathbf{D}_j$ throughout U if and only if $\mathbf{D} = \nabla_{E_K} \Psi$ on U for a C^2 potential Ψ . If a skew derivative is nonzero somewhere, online gradient descent can be forced to incur linear regret against $\mathbf{D}$ along a sufficiently small feasible loop. When the field is curl-free, assume that Ψ admits an ambient extension $\tilde{\Psi}$, defined on a neighborhood of K , that satisfies the smoothness and bounded-oscillation assumptions of Theorem 4.1 and Corollary 4.2. If $\tilde{\phi}(\mathbf{x}) := \mathbf{x} - \nabla \tilde{\Psi}(\mathbf{x})$ lies in K for every $\mathbf{x} \in K$, then projected GD has $O(\sqrt{T})$ regret against $\tilde{\phi}$.*

Proof. The equivalence between vanishing skew derivatives and existence of a potential is the Poincaré lemma for C^1 one-forms on the simply connected open set U , expressed in orthonormal affine coordinates on E_K . Nonzero curl gives, by Stokes' theorem, a sufficiently small loop contained in U with nonzero circulation. Applying the cyclic construction from the proof of Theorem 7.3 to this loop gives the claimed linear regret. Under the additional extension, feasibility, smoothness, and oscillation assumptions, the positive statement follows from Theorem 4.1 and Corollary 4.2. $\square$

D.4 Proofs for Section 8

Proposition 8.3. *A field $\mathbf{D} : \mathcal{D} \rightarrow E$ is R -mirror-exact if and only if there exists a differentiable scalar potential $V : \mathcal{D} \rightarrow \mathbb{R}$ such that*

$$\nabla V(\mathbf{x}) = \nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x}). \quad (5)$$

Equivalently, $\mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x}) = dV(\mathbf{x})$.

Proof. First suppose $\mathbf{D}(\mathbf{x}) = \nabla_\theta \Psi(\nabla R(\mathbf{x}))$. Define

$$V(\mathbf{x}) := \Psi(\nabla R(\mathbf{x})).$$

By the chain rule,

$$\nabla V(\mathbf{x}) = \nabla^2 R(\mathbf{x}) \nabla_\theta \Psi(\nabla R(\mathbf{x})) = \nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x}).$$

Conversely, suppose (5) holds. For $\theta \in \Theta$, write $\mathbf{x} := (\nabla R)^{-1}(\theta)$ and define

$$\Psi(\theta) := V((\nabla R)^{-1}(\theta)).$$

Another application of the chain rule gives

$$\nabla_\theta \Psi(\theta) = (\nabla^2 R(\mathbf{x}))^{-1} \nabla V(\mathbf{x}) = \mathbf{D}(\mathbf{x}),$$

which is the desired dual-gradient representation. $\square$

D.5 Examples of Mirror-exact Deviations for Section 8.1

Throughout this subsection, R^* is the Fenchel conjugate defined in the standing Legendre setup, so that

$$\nabla R^*(\theta) = (\nabla R)^{-1}(\theta)$$

whenever the inverse is defined. Each displayed map is a valid deviation once it is pointwise feasible, i.e. once it maps $\mathcal{D}$ into K .

1. Fixed comparators: external regret is always mirror-exact. Fix $\mathbf{u} \in K$ and set

$$\phi_{\mathbf{u}}(\mathbf{x}) := \mathbf{u}, \quad \mathbf{D}_{\mathbf{u}}(\mathbf{x}) := \mathbf{x} - \mathbf{u}.$$

Take the dual potential

$$\Psi_{\mathbf{u}}(\boldsymbol{\theta}) := R^*(\boldsymbol{\theta}) - \langle \mathbf{u}, \boldsymbol{\theta} \rangle.$$

Then, with $\boldsymbol{\theta} = \nabla R(\mathbf{x})$,

$$\nabla_{\boldsymbol{\theta}} \Psi_{\mathbf{u}}(\boldsymbol{\theta}) = \nabla R^*(\boldsymbol{\theta}) - \mathbf{u} = \mathbf{x} - \mathbf{u} = \mathbf{D}_{\mathbf{u}}(\mathbf{x}).$$

Thus every fixed-comparator deviation is R -mirror-exact for every mirror map. The identity/no-move deviation $\phi(\mathbf{x}) = \mathbf{x}$ is the special case $\mathbf{D}_{\phi} \equiv 0$, with potential $\Psi \equiv 0$.

2. Dual translations: one mirror step on a linear loss. Let $\mathbf{a} \in E$. Assume that $\boldsymbol{\theta} - \mathbf{a} \in \Theta$ for every $\boldsymbol{\theta} \in \Theta$, and define

$$\phi_{\mathbf{a}}(\mathbf{x}) := \nabla R^*(\nabla R(\mathbf{x}) - \mathbf{a}). \quad (25)$$

This is the unconstrained Bregman proximal map for the linear function $\mathbf{y} \mapsto \langle \mathbf{a}, \mathbf{y} \rangle$. It is mirror-exact with potential

$$\Psi_{\mathbf{a}}(\boldsymbol{\theta}) := R^*(\boldsymbol{\theta}) - R^*(\boldsymbol{\theta} - \mathbf{a}),$$

since

$$\nabla_{\boldsymbol{\theta}} \Psi_{\mathbf{a}}(\boldsymbol{\theta}) = \nabla R^*(\boldsymbol{\theta}) - \nabla R^*(\boldsymbol{\theta} - \mathbf{a}) = \mathbf{x} - \phi_{\mathbf{a}}(\mathbf{x}).$$

So a fixed translation in dual space is an exact-form deviation for mirror descent. This is often the most concrete way to remember the definition.

3. Mirror interpolation toward a comparator. Fix $\mathbf{u} \in \mathcal{D}$, write $\boldsymbol{\theta}_{\mathbf{u}} := \nabla R(\mathbf{u})$, and choose $\lambda \in [0, 1)$. Assume that

$$(1 - \lambda)\boldsymbol{\theta} + \lambda\boldsymbol{\theta}_{\mathbf{u}} \in \Theta \quad \forall \boldsymbol{\theta} \in \Theta.$$

Define

$$\phi_{\lambda, \mathbf{u}}^R(\mathbf{x}) := \nabla R^*((1 - \lambda)\nabla R(\mathbf{x}) + \lambda\nabla R(\mathbf{u})). \quad (26)$$

This moves $\mathbf{x}$ a fraction λ of the way toward $\mathbf{u}$, but the straight line is taken in dual coordinates. The potential

$$\Psi_{\lambda, \mathbf{u}}(\boldsymbol{\theta}) := R^*(\boldsymbol{\theta}) - \frac{1}{1 - \lambda} R^*((1 - \lambda)\boldsymbol{\theta} + \lambda\boldsymbol{\theta}_{\mathbf{u}})$$

works because

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \Psi_{\lambda, \mathbf{u}}(\boldsymbol{\theta}) &= \nabla R^*(\boldsymbol{\theta}) - \nabla R^*((1 - \lambda)\boldsymbol{\theta} + \lambda\boldsymbol{\theta}_{\mathbf{u}}) \\ &= \mathbf{x} - \phi_{\lambda, \mathbf{u}}^R(\mathbf{x}). \end{aligned}$$

When $\lambda = 0$, this is the identity deviation. As $\lambda \uparrow 1$, it approaches the external deviation $\phi_{\mathbf{u}}$, whose potential was given above.

4. Entropy on the simplex: multiplicative deviations. Let $K := \Delta_d$ and work on $\text{relint}(\Delta_d)$. Take the entropic mirror map

$$R(\mathbf{x}) := \sum_{i=1}^d \mathbf{x}[i] \log \mathbf{x}[i].$$

All derivatives are understood on the affine hull of the simplex. Equivalently, one may identify dual vectors that differ by a multiple of $\mathbf{1}$, or fix the gauge $\sum_i \boldsymbol{\theta}[i] = 0$. In these logit coordinates, $\nabla R^*(\boldsymbol{\theta}) = \text{softmax}(\boldsymbol{\theta})$. Hence the dual translation (25) becomes

$$\phi_{\mathbf{a}}(\mathbf{x})[i] := \frac{\mathbf{x}[i] \exp(-\mathbf{a}[i])}{\sum_{j=1}^d \mathbf{x}[j] \exp(-\mathbf{a}[j])}.$$

This is multiplicative-weights reweighting by $\exp(-\mathbf{a})$, and it is pointwise feasible on the simplex. The potential is the log-sum-exp difference

$$\Psi_{\mathbf{a}}(\boldsymbol{\theta}) := \log \sum_{i=1}^d \exp(\boldsymbol{\theta}[i]) - \log \sum_{i=1}^d \exp(\boldsymbol{\theta}[i] - \mathbf{a}[i]),$$

whose gradient is

$$\nabla_{\boldsymbol{\theta}} \Psi_{\mathbf{a}}(\boldsymbol{\theta}) = \text{softmax}(\boldsymbol{\theta}) - \text{softmax}(\boldsymbol{\theta} - \mathbf{a}).$$

Since $\text{softmax}(\nabla R(\mathbf{x})) = \mathbf{x}$, this equals $\mathbf{x} - \phi_{\mathbf{a}}(\mathbf{x})$.

The dual interpolation (26) gives another familiar entropy example. For $\mathbf{u} \in \text{relint}(\Delta_d)$,

$$\phi_{\lambda, \mathbf{u}}^R(\mathbf{x})[i] := \frac{\mathbf{x}[i]^{1-\lambda} \mathbf{u}[i]^{\lambda}}{\sum_{j=1}^d \mathbf{x}[j]^{1-\lambda} \mathbf{u}[j]^{\lambda}}.$$

Thus mirror-exact deviations include geometric interpolation on the simplex, not only Euclidean arithmetic interpolation.

5. Metric-symmetric affine swaps. Let $K := E := \mathbb{R}^d$ and $R_H(\mathbf{x}) := \frac{1}{2} \mathbf{x}^{\top} H \mathbf{x}$, where $H \succ 0$. This is a Legendre function with full domain. Then $\nabla R_H(\mathbf{x}) = H \mathbf{x}$, and Proposition 8.3 says that $\mathbf{D}$ is R_H -mirror-exact exactly when $H \mathbf{D}(\mathbf{x})$ is an ordinary gradient field.

Consider an affine deviation

$$\phi(\mathbf{x}) := A \mathbf{x} + \mathbf{b}, \quad \mathbf{D}(\mathbf{x}) := \mathbf{x} - \phi(\mathbf{x}) = (I - A) \mathbf{x} - \mathbf{b}.$$

If

$$H(I - A) = (I - A)^{\top} H, \tag{27}$$

then $\mathbf{D}$ is R_H -mirror-exact; feasibility is automatic because $K = \mathbb{R}^d$. Indeed,

$$V(\mathbf{x}) := \frac{1}{2} \mathbf{x}^{\top} H (I - A) \mathbf{x} - \langle H \mathbf{b}, \mathbf{x} \rangle$$

satisfies $\nabla V(\mathbf{x}) = H \mathbf{D}(\mathbf{x})$. For $H = I$, condition (27) says that $I - A$ is symmetric, recovering the symmetric affine swaps in Euclidean geometry. For general H , the correct symmetry is symmetry in the mirror metric.

6. A mirror-exact field that is not Euclidean-exact. The previous example gives a feasible separation on an unconstrained action space. Let

$$H := \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \quad K := \mathbb{R}^2,$$

use $R_H(\mathbf{x}) := \frac{1}{2}\mathbf{x}^\top H\mathbf{x}$, and for $\varepsilon \geq 0$ define

$$\mathbf{D}(\mathbf{x})[1] := \varepsilon(\mathbf{x}[1] + \mathbf{x}[2]), \quad \mathbf{D}(\mathbf{x})[2] := \frac{\varepsilon}{2}(\mathbf{x}[1] + \mathbf{x}[2]), \quad \phi(\mathbf{x}) := \mathbf{x} - \mathbf{D}(\mathbf{x}).$$

The map is feasible because it maps $\mathbb{R}^2$ to itself. Writing $S := \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, it is R_H -mirror-exact because

$$H\mathbf{D}(\mathbf{x}) = \varepsilon S\mathbf{x} = \nabla \left(\frac{\varepsilon}{2}(\mathbf{x}[1] + \mathbf{x}[2])^2 \right).$$

But $\mathbf{D}$ is not Euclidean-exact when $\varepsilon > 0$, since

$$\partial_2 \mathbf{D}(\mathbf{x})[1] = \varepsilon, \quad \partial_1 \mathbf{D}(\mathbf{x})[2] = \frac{\varepsilon}{2}.$$

So mirror-exact regret is not just Euclidean exact-form regret written in different notation. The mirror map changes which deviations telescope.

D.6 Proofs for Section 8.2

Theorem 8.4. *Assume the standing Legendre-domain conditions and that R is m -strongly convex on $\mathcal{D}$. Let $\phi : \mathcal{D} \rightarrow K$ be feasible and R -mirror-exact with dual potential Ψ . Let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online mirror descent (4), initialized in $\mathcal{D}$, with iterates remaining in $\mathcal{D}$. Write*

$$\boldsymbol{\theta}_t := \nabla R(\mathbf{x}_t), \quad \mathbf{D}_\phi(\mathbf{x}_t) = \nabla_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta}_t).$$

Assume that (6) holds for every consecutive pair $(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t)$, $t \in [T]$. A sufficient trajectory-independent condition is that it hold for every ordered pair in $\Theta \times \Theta$. Then

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \frac{\Psi(\boldsymbol{\theta}_1) - \Psi(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{L\eta}{m} \sum_{t=1}^T \|\mathbf{g}_t\|^2. \quad (7)$$

Consequently, suppose $\|\mathbf{g}_t\| \leq G$ and that, for the horizon-dependent choices of constant step size considered below, there is a constant B , independent of the horizon, such that

$$\max_{1 \leq t \leq T+1} \Psi(\boldsymbol{\theta}_t) - \min_{1 \leq t \leq T+1} \Psi(\boldsymbol{\theta}_t) \leq B,$$

Then choosing $\eta \asymp 1/\sqrt{T}$ gives $O(\sqrt{T})$ regret against ϕ . More precisely, when $B, L, G > 0$, the choice $\eta = \sqrt{mB/(LG^2T)}$ gives the bound $2G\sqrt{LBT/m}$. The global condition $\text{osc}_\Theta(\Psi) \leq B$ is a sufficient trajectory-independent assumption that avoids any dependence of the oscillation bound on the chosen step size.

Proof. Because $\mathbf{x}_{t+1} \in \mathcal{D} = \text{relint}(K)$ and normal cones are taken in the affine hull of K , the relative normal cone at $\mathbf{x}_{t+1}$ is $\{0\}$. The first-order optimality condition for (4) therefore reduces to the unconstrained dual update

$$\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1} = \eta \mathbf{g}_t.$$

Since $\mathbf{D}_\phi(\mathbf{x}_t) = \nabla_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta}_t)$, relative smoothness and Legendre conjugacy give

$$\begin{aligned} \eta \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle &= \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + D_\Psi(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \\ &\leq \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + L D_{R^*}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t). \end{aligned}$$

By the Legendre conjugacy identity stated with the standing assumptions,

$$D_{R^*}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) = D_R(\mathbf{x}_t, \mathbf{x}_{t+1}).$$

Therefore

$$\eta \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + LD_R(\mathbf{x}_t, \mathbf{x}_{t+1}). \quad (28)$$

Let $\mathbf{s}_t := \mathbf{x}_t - \mathbf{x}_{t+1}$. The symmetrized Bregman identity and the dual update imply

$$D_R(\mathbf{x}_t, \mathbf{x}_{t+1}) + D_R(\mathbf{x}_{t+1}, \mathbf{x}_t) = \langle \boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1}, \mathbf{s}_t \rangle = \eta \langle \mathbf{g}_t, \mathbf{s}_t \rangle.$$

Strong convexity of R therefore yields

$$m \|\mathbf{s}_t\|^2 \leq \eta \langle \mathbf{g}_t, \mathbf{s}_t \rangle \leq \eta \|\mathbf{g}_t\| \|\mathbf{s}_t\|.$$

Hence $\|\mathbf{s}_t\| \leq (\eta/m) \|\mathbf{g}_t\|$, with the conclusion trivial when $\mathbf{s}_t = 0$. Consequently,

$$D_R(\mathbf{x}_t, \mathbf{x}_{t+1}) \leq \eta \langle \mathbf{g}_t, \mathbf{s}_t \rangle \leq \frac{\eta^2}{m} \|\mathbf{g}_t\|^2.$$

Combining this estimate with (28) gives the slightly stronger one-step bound

$$\eta \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + \frac{L\eta^2}{m} \|\mathbf{g}_t\|^2.$$

Dividing by η and summing over t proves (7). The $O(\sqrt{T})$ statement follows by bounding the telescoping term by B/η and optimizing η . $\square$

D.7 Proofs for Section 8.3

Lemma D.3. *For fixed $\mathbf{y} \in \text{dom } R$,*

$$\nabla_{\mathbf{x}} D_R(\mathbf{y}, \mathbf{x}) = \nabla^2 R(\mathbf{x})(\mathbf{x} - \mathbf{y}).$$

Proof. By definition,

$$D_R(\mathbf{y}, \mathbf{x}) := R(\mathbf{y}) - R(\mathbf{x}) - \langle \nabla R(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

Only the last two terms depend on $\mathbf{x}$. Let

$$h(\mathbf{x}) = \langle \nabla R(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

Then

$$\nabla h(\mathbf{x}) = \nabla^2 R(\mathbf{x})(\mathbf{y} - \mathbf{x}) - \nabla R(\mathbf{x}).$$

Thus

$$\begin{aligned} \nabla_{\mathbf{x}} D_R(\mathbf{y}, \mathbf{x}) &= -\nabla R(\mathbf{x}) - \nabla h(\mathbf{x}) \\ &= \nabla^2 R(\mathbf{x})(\mathbf{x} - \mathbf{y}). \end{aligned}$$

$\square$

Lemma D.4 (Bregman-Moreau identity). *Assume the standing Legendre-domain conditions. Let $F := f + \iota_K$ be proper and lower semicontinuous, assume that every Bregman proximal minimum is attained, and suppose that prox_F^R is single-valued and continuous on $\mathcal{D}$. Then M_F^R is differentiable on $\mathcal{D}$ and*

$$\nabla M_F^R(\mathbf{x}) = \nabla^2 R(\mathbf{x})(\mathbf{x} - \text{prox}_F^R(\mathbf{x})). \quad (29)$$

Consequently, if $\boldsymbol{\theta} = \nabla R(\mathbf{x})$ and

$$\Psi_F(\boldsymbol{\theta}) = M_F^R((\nabla R)^{-1}(\boldsymbol{\theta})),$$

then

$$\nabla_{\boldsymbol{\theta}} \Psi_F(\boldsymbol{\theta}) = \mathbf{x} - \text{prox}_F^R(\mathbf{x}), \quad \mathbf{x} = (\nabla R)^{-1}(\boldsymbol{\theta}). \quad (30)$$

Proof. The key point is that we do not differentiate the minimizer. Continuity of prox_F^R makes the minimizing points locally bounded as $\mathbf{x}$ varies, so the standard local form of Danskin's theorem applies even though the ambient minimization domain need not be bounded. Uniqueness then implies

$$\nabla M_F^R(\mathbf{x}) = \nabla_{\mathbf{x}} D_R(\text{prox}_F^R(\mathbf{x}), \mathbf{x}).$$

Using Lemma D.3, we obtain

$$\nabla M_F^R(\mathbf{x}) = \nabla^2 R(\mathbf{x})(\mathbf{x} - \text{prox}_F^R(\mathbf{x})),$$

which proves (29).

For the dual identity, write $\mathbf{x} := (\nabla R)^{-1}(\boldsymbol{\theta})$. Since

$$\frac{\partial \mathbf{x}}{\partial \boldsymbol{\theta}} = (\nabla^2 R(\mathbf{x}))^{-1},$$

the chain rule gives

$$\nabla_{\boldsymbol{\theta}} \Psi_F(\boldsymbol{\theta}) = (\nabla^2 R(\mathbf{x}))^{-1} \nabla M_F^R(\mathbf{x}) = \mathbf{x} - \text{prox}_F^R(\mathbf{x}).$$

□

Theorem 8.6. *Assume the standing Legendre-domain conditions. Let $F := f + \iota_K$ be proper and lower semicontinuous, assume that every Bregman proximal minimum above is attained, and suppose that prox_F^R is single-valued and continuous on $\mathcal{D}$. Then the Bregman proximal deviation $\phi_F^R : \mathcal{D} \rightarrow K$ defined by $\phi_F^R(\mathbf{x}) := \text{prox}_F^R(\mathbf{x})$ is feasible and R -mirror-exact. Equivalently,*

$$(\mathbf{x} - \text{prox}_F^R(\mathbf{x}))^\top \text{d}\nabla R(\mathbf{x}) = \text{d}M_F^R(\mathbf{x}). \quad (8)$$

Thus Bregman proximal regret is a special case of mirror-exact regret.

Proof. Feasibility follows from the indicator ι_K , because every minimizer of the Bregman proximal problem lies in K . By (30) in Lemma D.4, the displacement

$$\mathbf{x} - \phi_F^R(\mathbf{x}) = \mathbf{x} - \text{prox}_F^R(\mathbf{x})$$

equals the dual gradient $\nabla_{\boldsymbol{\theta}} \Psi_F(\nabla R(\mathbf{x}))$. This is precisely R -mirror-exactness. The differential form (8) is the same identity written in primal coordinates using (29) in Lemma D.4. □

Proposition 8.7. *Assume the standing Legendre-domain conditions. Assume, in the extended-valued sense on E , that R is m -strongly convex and that $F := f + \iota_K$ is proper, closed, and ρ -weakly convex, where $0 \leq \rho < m$ and $H := F + R$ is proper. Then the Bregman proximal subproblem has a unique solution on $\mathcal{D}$,*

and its solution map is continuous there. Let $\{\mathbf{x}_t\}_{t=1}^{T+1} \subset \mathcal{D}$ be generated by online mirror descent (4). Then these iterates control Bregman proximal regret with the bound

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \text{prox}_F^R(\mathbf{x}_t) \rangle \leq \frac{\Psi_F(\boldsymbol{\theta}_1) - \Psi_F(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{\eta}{m} \sum_{t=1}^T \|\mathbf{g}_t\|^2, \quad (9)$$

where $\Psi_F(\boldsymbol{\theta}_t) = M_F^R(\mathbf{x}_t)$ for all $t \in [T+1]$.

Proof. Since F is ρ -weakly convex and R is m -strongly convex in the extended-valued sense, the proper closed function $H = F + R$ is $(m - \rho)$ -strongly convex. For $\boldsymbol{\theta} \in \Theta$, the Bregman proximal problem differs by an $\mathbf{x}$ -dependent constant from

$$\min_{\mathbf{y} \in E} \{H(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle\}.$$

It therefore has the unique solution $\nabla H^*(\boldsymbol{\theta})$. Moreover, H^* is differentiable with $1/(m - \rho)$ -Lipschitz gradient, so the proximal solution map is continuous on $\mathcal{D}$.

Let $\mathbf{x} := \nabla R^*(\boldsymbol{\theta})$. By the definition of the Bregman envelope,

$$\begin{aligned} \Psi_F(\boldsymbol{\theta}) &= M_F^R(\mathbf{x}) \\ &= \min_{\mathbf{y}} \{F(\mathbf{y}) + R(\mathbf{y}) - R(\mathbf{x}) - \langle \nabla R(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle\} \\ &= \min_{\mathbf{y}} \{F(\mathbf{y}) + R(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle\} + \langle \boldsymbol{\theta}, \mathbf{x} \rangle - R(\mathbf{x}) \\ &= -(F + R)^*(\boldsymbol{\theta}) + R^*(\boldsymbol{\theta}). \end{aligned}$$

This proves $\Psi_F(\boldsymbol{\theta}) = R^*(\boldsymbol{\theta}) - H^*(\boldsymbol{\theta})$ for the function $H := F + R$.

Because H^* is convex and differentiable on the dual region,

$$D_{H^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) \geq 0.$$

Taking the Bregman divergence of the identity $\Psi_F = R^* - H^*$ gives

$$D_{\Psi_F}(\boldsymbol{\theta}', \boldsymbol{\theta}) = D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) - D_{H^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) \leq D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}),$$

which proves the relative-smoothness with $L = 1$ for the potential function Ψ_F . Finally, Theorem 8.6 identifies the Bregman proximal displacement with $\nabla_{\boldsymbol{\theta}} \Psi_F(\boldsymbol{\theta}_t)$. Applying Theorem 8.4 with $L = 1$ gives (9). $\square$

D.8 Proofs for Section 8.4

Theorem 8.9. Assume R is differentiable on a neighborhood of K , so that $\mathcal{X} = K$, and let $\{\mathbf{x}_t\}_{t=1}^{T+1}$ be generated by mirror descent (4) with a constant step size $\eta > 0$. Let $\phi : K \rightarrow K$ be R -mirror-exact with potential Ψ . Suppose (10) holds and $\mathbf{D}_\phi$ is L_ϕ -Lipschitz on K . Then

$$\begin{aligned} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle &\leq \frac{\Psi(\nabla R(\mathbf{x}_1)) - \Psi(\nabla R(\mathbf{x}_{T+1}))}{\eta} + \frac{L\eta}{m} \sum_{t=1}^T \|\mathbf{g}_t\|_*^2 \\ &\quad + \frac{L_\phi}{m} \sum_{t=1}^T \|\mathbf{g}_t\|_* \left[\eta \|\mathbf{g}_t\|_* + \omega_R\left(\frac{\eta}{m} \|\mathbf{g}_t\|_*\right) \right]. \end{aligned}$$

Proof. Set

$$\boldsymbol{\theta}_t := \nabla R(\mathbf{x}_t), \quad \mathbf{s}_t = \mathbf{x}_t - \mathbf{x}_{t+1}.$$

The first-order optimality condition for the mirror-descent subproblem provides $\mathbf{n}_{t+1} \in N_K(\mathbf{x}_{t+1})$ such that

$$\eta \mathbf{g}_t = \boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1} - \mathbf{n}_{t+1}. \quad (31)$$

The symmetrized Bregman identity gives

$$\begin{aligned} D_R(\mathbf{x}_t, \mathbf{x}_{t+1}) + D_R(\mathbf{x}_{t+1}, \mathbf{x}_t) &= \langle \boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1}, \mathbf{s}_t \rangle \\ &= \eta \langle \mathbf{g}_t, \mathbf{s}_t \rangle + \langle \mathbf{n}_{t+1}, \mathbf{s}_t \rangle. \end{aligned}$$

Since $\mathbf{x}_t \in K$ and $\mathbf{n}_{t+1} \in N_K(\mathbf{x}_{t+1})$,

$$\langle \mathbf{n}_{t+1}, \mathbf{s}_t \rangle \leq 0.$$

Strong convexity of R therefore implies

$$m \|\mathbf{s}_t\|^2 \leq \eta \|\mathbf{g}_t\|_* \|\mathbf{s}_t\|.$$

Consequently,

$$\|\mathbf{s}_t\| \leq \frac{\eta}{m} \|\mathbf{g}_t\|_*, \quad D_R(\mathbf{x}_t, \mathbf{x}_{t+1}) \leq \frac{\eta^2}{m} \|\mathbf{g}_t\|_*^2. \quad (32)$$

The normal decomposition and the definition of ω_R also yield

$$\|\mathbf{n}_{t+1}\|_* \leq \eta \|\mathbf{g}_t\|_* + \omega_R\left(\frac{\eta}{m} \|\mathbf{g}_t\|_*\right). \quad (33)$$

Write $\mathbf{D}_t := \mathbf{D}_\phi(\mathbf{x}_t)$. By (31),

$$\eta \langle \mathbf{g}_t, \mathbf{D}_t \rangle = \langle \boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1}, \mathbf{D}_t \rangle - \langle \mathbf{n}_{t+1}, \mathbf{D}_t \rangle.$$

Mirror exactness and (10) give

$$\begin{aligned} \langle \boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1}, \mathbf{D}_t \rangle &= \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + D_\Psi(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \\ &\leq \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + LD_R(\mathbf{x}_t, \mathbf{x}_{t+1}). \end{aligned}$$

It remains to control the normal term. Feasibility of ϕ gives $\phi(\mathbf{x}_{t+1}) \in K$, hence

$$\langle \mathbf{n}_{t+1}, \mathbf{D}_\phi(\mathbf{x}_{t+1}) \rangle \geq 0.$$

Therefore

$$\begin{aligned} -\langle \mathbf{n}_{t+1}, \mathbf{D}_t \rangle &= -\langle \mathbf{n}_{t+1}, \mathbf{D}_\phi(\mathbf{x}_{t+1}) \rangle - \langle \mathbf{n}_{t+1}, \mathbf{D}_\phi(\mathbf{x}_t) - \mathbf{D}_\phi(\mathbf{x}_{t+1}) \rangle \\ &\leq L_\phi \|\mathbf{n}_{t+1}\|_* \|\mathbf{s}_t\|. \end{aligned}$$

Substituting (32) and (33) gives

$$\begin{aligned} \eta \langle \mathbf{g}_t, \mathbf{D}_t \rangle &\leq \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + \frac{L\eta^2}{m} \|\mathbf{g}_t\|_*^2 \\ &\quad + \frac{L_\phi \eta}{m} \|\mathbf{g}_t\|_* \left[\eta \|\mathbf{g}_t\|_* + \omega_R\left(\frac{\eta}{m} \|\mathbf{g}_t\|_*\right) \right]. \end{aligned}$$

Divide by η and sum over t . The potential differences telescope. □

Corollary 8.10. Assume the hypotheses of Theorem 8.9, $\|\mathbf{g}_t\|_* \leq G$, and $\text{osc}_{\nabla R(K)}(\Psi) \leq B$. For $\eta = T^{-1/2}$,

$$\frac{1}{T} \sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \frac{B}{\sqrt{T}} + \frac{(L + L_\phi)G^2}{m\sqrt{T}} + \frac{L_\phi G}{m} \omega_R\left(\frac{G}{m\sqrt{T}}\right).$$

Consequently, mirror descent has $o(T)$ regret against ϕ . If ∇R is M -Lipschitz on K , then

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}_\phi(\mathbf{x}_t) \rangle \leq \frac{B}{\eta} + C_{R,\phi} \eta G^2 T, \quad C_{R,\phi} := \frac{L + L_\phi}{m} + \frac{L_\phi M}{m^2}.$$

The optimized bound is $2G\sqrt{BC_{R,\phi}T}$ whenever $BC_{R,\phi} > 0$.

Proof. Apply Theorem 8.9 and bound the telescoping term by B . Uniform continuity of ∇R on K gives

$$\omega_R\left(\frac{G}{m\sqrt{T}}\right) \rightarrow 0,$$

so the average regret converges to zero. If $\omega_R(\delta) \leq M\delta$, substitution gives the stated value of $C_{R,\phi}$. The minimum of $B/\eta + C_{R,\phi}\eta G^2 T$ is $2G\sqrt{BC_{R,\phi}T}$ when $BC_{R,\phi} > 0$. $\square$

Theorem 8.11. Suppose that, after reversing the orientation of γ if necessary,

$$-\oint_{\gamma} \mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x}) := c_0 > 0.$$

For every $G > 0$ and every sufficiently small constant step size $\eta > 0$, there is a sequence with $\|\mathbf{g}_t\|_* \leq G$ such that standard constrained mirror descent, initialized at $\gamma(0)$, satisfies

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}(\mathbf{x}_t) \rangle \geq cT - O(1/\eta)$$

for a constant $c > 0$ depending only on $\mathbf{D}$, γ , R , and G .

Proof. Let

$$\boldsymbol{\theta}(s) = \nabla R(\gamma(s)).$$

The curve $\boldsymbol{\theta}$ is closed and rectifiable. Its length is positive because the circulation is nonzero. The negative circulation is the limit of the left Riemann sums

$$\sum_{j=0}^{q-1} \langle \mathbf{D}(\mathbf{z}_j), \boldsymbol{\theta}_j - \boldsymbol{\theta}_{j+1} \rangle, \quad \mathbf{z}_j = \gamma(s_j), \quad \boldsymbol{\theta}_j = \boldsymbol{\theta}(s_j).$$

For every sufficiently small η , choose a partition $0 = s_0 < \dots < s_q = 1$ for which

$$\|\boldsymbol{\theta}_j - \boldsymbol{\theta}_{j+1}\|_* \leq \eta G, \quad q \leq \frac{C_\gamma}{\eta},$$

and

$$\sum_{j=0}^{q-1} \langle \mathbf{D}(\mathbf{z}_j), \boldsymbol{\theta}_j - \boldsymbol{\theta}_{j+1} \rangle \geq \frac{c_0}{2}. \tag{34}$$

Such partitions exist because $\boldsymbol{\theta}$ is piecewise C^1 and $\mathbf{D} \circ \gamma$ is continuous.

For one traversal, initialize at $\mathbf{z}_0$ and set

$$\mathbf{g}_{j+1} := \frac{\boldsymbol{\theta}_j - \boldsymbol{\theta}_{j+1}}{\eta}, \quad j = 0, \dots, q-1.$$

The dual norm of every gradient is at most G . The mirror-descent objective at this step is

$$\mathbf{y} \mapsto \eta \langle \mathbf{g}_{j+1}, \mathbf{y} \rangle + D_R(\mathbf{y}, \mathbf{z}_j).$$

Its gradient at $\mathbf{z}_{j+1}$ is

$$\eta \mathbf{g}_{j+1} + \nabla R(\mathbf{z}_{j+1}) - \nabla R(\mathbf{z}_j) = 0.$$

The objective is m -strongly convex on K , so $\mathbf{z}_{j+1}$ is its unique minimizer. The constrained mirror-descent iterates therefore follow the loop exactly, including at boundary points.

By (34), the regret during one traversal is at least

$$\sum_{j=0}^{q-1} \langle \mathbf{g}_{j+1}, \mathbf{D}(\mathbf{z}_j) \rangle \geq \frac{c_0}{2\eta}.$$

Since $q \leq C_\gamma/\eta$, complete traversals have average regret at least $c_0/(2C_\gamma)$. Repeating the traversal and assigning zero gradients to the remaining rounds gives

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{D}(\mathbf{x}_t) \rangle \geq \frac{c_0}{2C_\gamma} T - O(1/\eta).$$

□

Corollary 8.12. *Let $U := \text{relint}(K)$. The set U is open and convex in $\text{aff}(K)$, and hence simply connected there. Suppose that $R \in C^2(U)$ and $\nabla^2 R(\mathbf{x})$ is positive definite for every $\mathbf{x} \in U$, and the coefficient field $\mathbf{x} \mapsto \nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x})$ is C^1 . All derivatives and coordinate components below are taken in affine coordinates on $\text{aff}(K)$. Then the following conditions are equivalent.*

(i) *The one-form $\mathbf{D}(\mathbf{x})^\top d\nabla R(\mathbf{x})$ is exact on U .*

(ii) *For every coordinate pair i, j ,*

$$\partial_j [\nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x})]_i = \partial_i [\nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x})]_j \quad (\mathbf{x} \in U).$$

(iii) *The mirror circulation vanishes on every closed piecewise C^1 loop in U .*

If these conditions fail, standard mirror descent can be forced to incur linear regret along a sufficiently small loop contained in U . Suppose instead that they hold, that $\mathbf{D}$ extends to the displacement $\mathbf{D}_\phi$ of a feasible map on K , and that the resulting potential and displacement satisfy the regularity hypotheses of Theorem 8.9. Then mirror descent has sublinear regret. Under the Lipschitz-gradient assumptions in Corollary 8.10, the regret is $O(\sqrt{T})$.

Proof. In affine coordinates on $U = \text{relint}(K)$, the coefficient field of the one-form is

$$\boldsymbol{\beta}(\mathbf{x}) = \nabla^2 R(\mathbf{x}) \mathbf{D}(\mathbf{x}).$$

The Poincaré lemma gives the equivalence between exactness of $\boldsymbol{\beta}(\mathbf{x})^\top d\mathbf{x}$ and symmetry of the cross derivatives of $\boldsymbol{\beta}$. Exact forms have zero integral around every closed loop. Conversely, if every closed-loop integral vanishes, the line integral from a fixed base point is path independent and defines a potential.

If a cross derivative is nonzero at some point, continuity and Stokes' theorem give a sufficiently small rectangular loop in U with nonzero mirror circulation. The linear lower bound follows from Theorem 8.11.

Suppose instead that the form is exact. Strong convexity makes ∇R injective on U , and positive definiteness of $\nabla^2 R$ makes it a local diffeomorphism. Hence $\nabla R : U \rightarrow \nabla R(U)$ is a diffeomorphism onto an open set. If V is a potential for the one-form, define

$$\Psi(\boldsymbol{\theta}) := V((\nabla R)^{-1}(\boldsymbol{\theta})).$$

The chain rule gives

$$\nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\boldsymbol{x})) = \boldsymbol{D}(\boldsymbol{x}).$$

The extension, feasibility, and regularity assumptions in the corollary then allow Theorem 8.9 to be applied on K . $\square$

D.9 Proofs for Section 8.5

Proposition 8.14. *Assume $\mathcal{X} = K$. For every admissible generator F in Definition 8.13, the function $\Psi_F(\boldsymbol{\theta}) := R_K^*(\boldsymbol{\theta}) - (F + R)^*(\boldsymbol{\theta})$ is differentiable and satisfies*

$$\boldsymbol{x} - \text{prox}_F^R(\boldsymbol{x}) = \nabla_{\boldsymbol{\theta}} \Psi_F(\nabla R(\boldsymbol{x})) \quad (\boldsymbol{x} \in K).$$

Consequently, $\Phi_{\text{prox}}^R(K) \subseteq \Phi_{\text{rad}}^R(K) \subseteq \Phi_{\text{exact}}^R(K)$, where $\Phi_{\text{exact}}^R(K)$ is the class from Definition 8.8. More precisely, if $\phi_{\alpha} = \text{prox}_F^R$, then ϕ is mirror-exact with potential Ψ_F/α .

Proof. Since R_K is proper, closed, and m -strongly convex, and $F + R$ is proper, closed, and $(m - \rho)$ -strongly convex, their conjugates are differentiable. For every $\boldsymbol{x} \in K$,

$$\nabla R(\boldsymbol{x}) \in \partial R_K(\boldsymbol{x}),$$

because $0 \in N_K(\boldsymbol{x})$. Fenchel duality therefore gives

$$\nabla R_K^*(\nabla R(\boldsymbol{x})) = \boldsymbol{x}.$$

If $\boldsymbol{p} = \text{prox}_F^R(\boldsymbol{x})$, the proximal optimality condition gives

$$\nabla R(\boldsymbol{x}) \in \partial(F + R)(\boldsymbol{p}),$$

and hence

$$\nabla(F + R)^*(\nabla R(\boldsymbol{x})) = \boldsymbol{p}.$$

Subtracting the two identities yields

$$\nabla_{\boldsymbol{\theta}} \Psi_F(\nabla R(\boldsymbol{x})) = \boldsymbol{x} - \text{prox}_F^R(\boldsymbol{x}),$$

so every map in $\Phi_{\text{prox}}^R(K)$ is mirror-exact.

The inclusion $\Phi_{\text{prox}}^R(K) \subseteq \Phi_{\text{rad}}^R(K)$ follows by taking $\alpha = 1$. If $\phi_{\alpha} = \text{prox}_F^R$, then

$$\boldsymbol{x} - \phi(\boldsymbol{x}) = \frac{1}{\alpha}(\boldsymbol{x} - \phi_{\alpha}(\boldsymbol{x})) = \nabla_{\boldsymbol{\theta}} \left(\frac{\Psi_F}{\alpha} \right) (\nabla R(\boldsymbol{x})).$$

Thus ϕ is mirror-exact with potential Ψ_F/α , which proves $\Phi_{\text{rad}}^R(K) \subseteq \Phi_{\text{exact}}^R(K)$ and completes the proof. $\square$

Proposition 8.15. *Suppose $\phi_\alpha = \text{prox}_{F_\alpha}^R$ on $\mathcal{D}$ for an admissible Bregman proximal generator F_α . Then the dual endpoint potential u_α agrees, up to an additive constant, with the convex conjugate of $F_\alpha + R$. More precisely, there is a constant $c_\alpha \in \mathbb{R}$ such that $u_\alpha = (F_\alpha + R)^* + c_\alpha$, on Θ . Consequently, u_α is convex, and hence we have*

$$D_\Psi(\boldsymbol{\theta}', \boldsymbol{\theta}) \leq \frac{1}{\alpha} D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) \quad (\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta).$$

Proof. Let $H_\alpha := F_\alpha + R$. Since F_α is ρ_α -weakly convex with $\rho_\alpha < m$, the function H_α is strongly convex, and H_α^* is differentiable. For $\boldsymbol{x} = \nabla R^*(\boldsymbol{\theta})$, the Bregman proximal problem differs by a constant from

$$\min_{\boldsymbol{y}} \{H_\alpha(\boldsymbol{y}) - \langle \boldsymbol{\theta}, \boldsymbol{y} \rangle\}.$$

Hence

$$\phi_\alpha(\nabla R^*(\boldsymbol{\theta})) = \nabla H_\alpha^*(\boldsymbol{\theta}).$$

Mirror exactness gives the independent identity

$$\phi_\alpha(\nabla R^*(\boldsymbol{\theta})) = \nabla(R^* - \alpha\Psi)(\boldsymbol{\theta}) = \nabla u_\alpha(\boldsymbol{\theta}).$$

The open convex set Θ is connected, so two differentiable functions with the same gradient differ by a constant. Thus

$$u_\alpha = H_\alpha^* + c_\alpha$$

on Θ . Convexity follows. Finally,

$$D_{u_\alpha}(\boldsymbol{\theta}', \boldsymbol{\theta}) = D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) - \alpha D_\Psi(\boldsymbol{\theta}', \boldsymbol{\theta}) \geq 0,$$

which is the stated relative-curvature inequality. $\square$

Proposition 8.16. *Suppose u_α is convex on Θ , and let*

$$\bar{u}_\alpha := \text{cl}(u_\alpha + \iota_\Theta).$$

Assume $\bar{u}_\alpha$ is proper and that $\phi_\alpha(\mathcal{D}) \subseteq \text{dom } R$. Define the candidate proximal generator

$$F_\alpha(\boldsymbol{y}) := \begin{cases} \bar{u}_\alpha^*(\boldsymbol{y}) - R(\boldsymbol{y}), & \boldsymbol{y} \in \text{dom } R, \\ +\infty, & \boldsymbol{y} \notin \text{dom } R. \end{cases}$$

Then, for every $\boldsymbol{x} \in \mathcal{D}$,

$$\phi_\alpha(\boldsymbol{x}) \in \arg \min_{\boldsymbol{y}} \{F_\alpha(\boldsymbol{y}) + D_R(\boldsymbol{y}, \boldsymbol{x})\}.$$

Under these domain assumptions, convexity of u_α yields a variational Bregman representation of ϕ_α . If, in addition, F_α is proper, lower semicontinuous, and ρ_α -weakly convex for some $\rho_\alpha < m$, the minimizer is unique and $\phi_\alpha = \text{prox}_{F_\alpha}^R$. In that case $\phi \in \Phi_{\text{rad}}^R(\mathcal{D})$.

Proof. A finite differentiable convex function on the open set Θ is continuous there, so

$$\bar{u}_\alpha = \text{cl}(u_\alpha + \iota_\Theta)$$

agrees with u_α on Θ . Fix $\boldsymbol{\theta} \in \Theta$ and set

$$\boldsymbol{z} := \nabla u_\alpha(\boldsymbol{\theta}) = \phi_\alpha(\nabla R^*(\boldsymbol{\theta})).$$

Convexity gives $\mathbf{z} \in \partial \bar{u}_\alpha(\boldsymbol{\theta})$. Fenchel duality then gives

$$\boldsymbol{\theta} \in \partial \bar{u}_\alpha^*(\mathbf{z}).$$

The endpoint-domain hypothesis gives $\mathbf{z} \in \text{dom } R$, and Fenchel equality gives $\bar{u}_\alpha^*(\mathbf{z}) < +\infty$. In particular, $F_\alpha + R$ is proper. Consequently, $\mathbf{z}$ minimizes

$$\mathbf{y} \mapsto \bar{u}_\alpha^*(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle.$$

For $\mathbf{x} = \nabla R^*(\boldsymbol{\theta})$ and $\mathbf{y} \in \text{dom } R$, the objective in the proposition satisfies

$$\begin{aligned} F_\alpha(\mathbf{y}) + D_R(\mathbf{y}, \mathbf{x}) &= \bar{u}_\alpha^*(\mathbf{y}) - R(\mathbf{y}) + R(\mathbf{y}) - R(\mathbf{x}) - \langle \boldsymbol{\theta}, \mathbf{y} - \mathbf{x} \rangle \\ &= \bar{u}_\alpha^*(\mathbf{y}) - \langle \boldsymbol{\theta}, \mathbf{y} \rangle + \langle \boldsymbol{\theta}, \mathbf{x} \rangle - R(\mathbf{x}). \end{aligned}$$

The last two terms do not depend on $\mathbf{y}$. Since $\mathbf{z} \in \text{dom } R$ and the objective is infinite outside $\text{dom } R$, $\mathbf{z} = \phi_\alpha(\mathbf{x})$ is a minimizer. If F_α is ρ_α -weakly convex with $\rho_\alpha < m$, then $F_\alpha + R$ is strongly convex. The minimizer is unique and equals $\text{prox}_{F_\alpha}^R(\mathbf{x})$. $\square$

Corollary 8.17. *For $0 < \alpha \leq 1$, the function $u_\alpha = R^* - \alpha\Psi$ is convex exactly when $\alpha L_R(\Psi) \leq 1$. Define the convexity threshold*

$$\alpha_{\text{cvx}} := \min \left\{ 1, \frac{1}{L_R(\Psi)} \right\},$$

with the conventions $1/0 := +\infty$ and $1/(+\infty) := 0$; a zero threshold means that no positive scale passes the convexity test. Every Bregman proximal interpolation must satisfy $0 < \alpha \leq \alpha_{\text{cvx}}$. Conversely, every such α for which $\bar{u}_\alpha$ is proper and $\phi_\alpha(\mathcal{D}) \subseteq \text{dom } R$ yields the variational representation in Proposition 8.16. It is an admissible Bregman proximal interpolation if the resulting F_α also satisfies the weak-convexity and regularity conditions in that proposition.

Proof. For every ordered pair in Θ ,

$$D_{R^* - \alpha\Psi}(\boldsymbol{\theta}', \boldsymbol{\theta}) = D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) - \alpha D_\Psi(\boldsymbol{\theta}', \boldsymbol{\theta}).$$

A differentiable function on an open convex set is convex if and only if its Bregman divergence is nonnegative on every ordered pair. Hence u_α is convex if and only if

$$D_\Psi(\boldsymbol{\theta}', \boldsymbol{\theta}) \leq \frac{1}{\alpha} D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) \quad \text{for all } \boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta.$$

If $L_R(\Psi) < +\infty$, take a decreasing sequence of valid constants tending to the infimum. Passing to the pointwise limit shows that the infimum is valid. If $L_R(\Psi) = +\infty$, no positive interpolation scale satisfies the inequality. Otherwise, the largest $\alpha \in (0, 1]$ satisfying it is $\min\{1, 1/L_R(\Psi)\}$. Necessity for a Bregman proximal interpolation follows from Proposition 8.15. For the converse, the additional assumptions that $\bar{u}_\alpha$ is proper and $\phi_\alpha(\mathcal{D}) \subseteq \text{dom } R$ permit application of Proposition 8.16; its remaining hypotheses ensure admissibility of the reconstructed generator. $\square$

Proposition 8.18. *Let $1 < p < 2$, let $q := p/(p-1) > 2$, and take the full-domain Legendre regularizer*

$$R(\mathbf{x}) := \frac{1}{2} \|\mathbf{x}\|_p^2 \quad (\mathbf{x} \in \mathbb{R}^2).$$

There is a feasible R -mirror-exact deviation whose dual potential is C^∞ with globally Lipschitz gradient, but which does not belong to $\Phi_{\text{rad}}^R(\mathbb{R}^2)$.

Proof. Let $q := p/(p-1) > 2$. The conjugate regularizer is

$$R^*(\boldsymbol{\theta}) = \frac{1}{2} \|\boldsymbol{\theta}\|_q^2.$$

Set

$$\Psi(\boldsymbol{\theta}) := \frac{1}{2} \boldsymbol{\theta}[2]^2, \quad \phi(\mathbf{x}) := \mathbf{x} - \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\mathbf{x})).$$

The map is feasible because the action space is $\mathbb{R}^2$. It is mirror-exact by construction, and $\nabla \Psi$ is globally 1-Lipschitz.

Fix $\alpha > 0$ and restrict $u_\alpha = R^* - \alpha \Psi$ to the line $\boldsymbol{\theta} = \mathbf{e}_1 + t\mathbf{e}_2$. Then

$$u_\alpha(\mathbf{e}_1 + t\mathbf{e}_2) = \frac{1}{2}(1 + |t|^q)^{2/q} - \frac{\alpha}{2}t^2.$$

The expansion $(1 + s)^{2/q} = 1 + (2/q)s + O(s^2)$ gives

$$u_\alpha(\mathbf{e}_1 + t\mathbf{e}_2) - u_\alpha(\mathbf{e}_1) = \frac{1}{q}|t|^q - \frac{\alpha}{2}t^2 + O(|t|^{2q}).$$

Since $q > 2$, this difference is negative for every sufficiently small nonzero t . The restricted function is even, so convexity would imply

$$u_\alpha(\mathbf{e}_1) \leq \frac{1}{2}u_\alpha(\mathbf{e}_1 + t\mathbf{e}_2) + \frac{1}{2}u_\alpha(\mathbf{e}_1 - t\mathbf{e}_2) = u_\alpha(\mathbf{e}_1 + t\mathbf{e}_2),$$

a contradiction. Thus u_α is not convex for any $\alpha > 0$. By Proposition 8.15, no positive interpolation can be an admissible Bregman proximal map. $\square$

D.10 Proofs for Section 9.1

Proposition 9.1. *Under the standing Legendre assumptions of Section 8, let $\boldsymbol{\theta}_1 := \nabla R(\mathbf{x}_1)$ and use a constant step size η . The mirror-descent iterates satisfy*

$$\mathbf{x}_t = \arg \min_{\mathbf{x} \in E} \left\{ \eta \sum_{s=1}^{t-1} \langle \mathbf{g}_s, \mathbf{x} \rangle + R(\mathbf{x}) - \langle \boldsymbol{\theta}_1, \mathbf{x} \rangle \right\} = \arg \min_{\mathbf{x} \in K} \left\{ \eta \sum_{s=1}^{t-1} \langle \mathbf{g}_s, \mathbf{x} \rangle + D_R(\mathbf{x}, \mathbf{x}_1) \right\}.$$

Thus they are the FTRL iterates for the linearized losses and the tilted regularizer $R - \langle \boldsymbol{\theta}_1, \cdot \rangle$.

Proof. The interior optimality condition for mirror descent is

$$\nabla R(\mathbf{x}_{t+1}) = \nabla R(\mathbf{x}_t) - \eta \mathbf{g}_t.$$

Writing $\boldsymbol{\theta}_t := \nabla R(\mathbf{x}_t)$ and unrolling gives $\boldsymbol{\theta}_t = \boldsymbol{\theta}_1 - \eta \sum_{s < t} \mathbf{g}_s$. Legendre duality therefore gives

$$\mathbf{x}_t = \nabla R^*(\boldsymbol{\theta}_t) = \arg \min_{\mathbf{x} \in E} \{ R(\mathbf{x}) - \langle \boldsymbol{\theta}_t, \mathbf{x} \rangle \}.$$

This is the first FTRL formula. The second follows because $R(\mathbf{x}) - \langle \boldsymbol{\theta}_1, \mathbf{x} \rangle$ differs from $D_R(\mathbf{x}, \mathbf{x}_1)$ by a constant independent of $\mathbf{x}$. $\square$

D.11 Proofs for Section 9.2

Theorem 9.3. Assume that R is proper, closed, (σ, r) -uniformly convex, and that R^* is finite on E . Let the iterates follow (12). If ϕ is R -FTRL-exact with potential Ψ and

$$D_\Psi(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \leq LD_{R^*}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \quad (t \in [T]),$$

then

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{\Psi(\boldsymbol{\theta}_1) - \Psi(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{L\eta^{q-1}}{q\sigma^{q-1}} \sum_{t=1}^T \|\mathbf{g}_t\|_*^q.$$

Proof. Uniform convexity makes $\partial R^*(\boldsymbol{\theta})$ a singleton whenever it is nonempty. Since R^* is finite on E , it has a nonempty subdifferential everywhere and is therefore differentiable on E .

FTRL-exactness and the additive dual update give

$$\begin{aligned} \eta \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle &= \langle \boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1}, \nabla_{\boldsymbol{\theta}} \Psi(\boldsymbol{\theta}_t) \rangle \\ &= \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + D_\Psi(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t) \\ &\leq \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + LD_{R^*}(\boldsymbol{\theta}_{t+1}, \boldsymbol{\theta}_t). \end{aligned} \tag{35}$$

Fix $\boldsymbol{\theta} \in E$, let $\mathbf{x} := \nabla R^*(\boldsymbol{\theta})$, and let $\mathbf{d} \in E$. Uniform convexity gives, for every $\mathbf{y} \in \text{dom } R$,

$$R(\mathbf{y}) \geq R(\mathbf{x}) + \langle \boldsymbol{\theta}, \mathbf{y} - \mathbf{x} \rangle + \frac{\sigma}{r} \|\mathbf{y} - \mathbf{x}\|^r.$$

Hence

$$\begin{aligned} R^*(\boldsymbol{\theta} + \mathbf{d}) &\leq R^*(\boldsymbol{\theta}) + \langle \mathbf{d}, \mathbf{x} \rangle + \sup_{\mathbf{z} \in E} \left\{ \langle \mathbf{d}, \mathbf{z} \rangle - \frac{\sigma}{r} \|\mathbf{z}\|^r \right\} \\ &= R^*(\boldsymbol{\theta}) + \langle \mathbf{d}, \nabla R^*(\boldsymbol{\theta}) \rangle + \frac{1}{q\sigma^{q-1}} \|\mathbf{d}\|_*^q. \end{aligned}$$

Thus

$$D_{R^*}(\boldsymbol{\theta} + \mathbf{d}, \boldsymbol{\theta}) \leq \frac{1}{q\sigma^{q-1}} \|\mathbf{d}\|_*^q.$$

Apply this with $\mathbf{d} := \boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_t = -\eta \mathbf{g}_t$ in (35), divide by η , and sum. $\square$

Corollary 9.4. Suppose $\|\mathbf{g}_t\|_* \leq G$ and the oscillation of Ψ along the dual trajectory is at most B . When $B, L, G > 0$, choosing $\eta = \left(\frac{rB\sigma^{q-1}}{LG^qT} \right)^{1/q}$ gives

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq r^{1/r} \frac{GL^{1/q}B^{1/r}}{\sigma^{1/r}} T^{1/q}.$$

For $r = 2$, this is $G\sqrt{2LBT/\sigma}$.

Proof. The theorem gives

$$\frac{B}{\eta} + \frac{LG^qT}{q\sigma^{q-1}}\eta^{q-1}.$$

Differentiation yields the displayed optimizer and bound. $\square$

Theorem 9.5. Suppose $\tilde{D}_\phi$ is continuous and there is a closed piecewise C^1 loop $\gamma \subset \Theta$ with

$$\oint_{\gamma} \langle \tilde{D}_\phi(\boldsymbol{\theta}), d\boldsymbol{\theta} \rangle \neq 0.$$

For every $G > 0$ and every sufficiently small constant $\eta > 0$, bounded gradients $\|\mathbf{g}_t\|_* \leq G$ can force FTRL (12) to incur $cT - O(1/\eta)$ regret against ϕ , for some $c > 0$. The statement holds from every initial dual state in Θ since steering to the loop changes only the transient term.

Proof. Reverse the orientation if necessary so that

$$-\oint_{\gamma} \langle \tilde{D}_\phi(\boldsymbol{\theta}), d\boldsymbol{\theta} \rangle = c_0 > 0.$$

Because Θ is open and convex, join the initial state to a point of the loop by a compact polygonal path in Θ . Discretizing that path into dual increments of norm at most ηG steers FTRL to the loop in $O(1/\eta)$ rounds. Continuity bounds the displacement on the path, so the magnitude of the steering regret is $O(1/\eta)$.

Let L_γ be the length of the loop in the dual norm. For every sufficiently small η , choose a partition $\boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}_s = \boldsymbol{\theta}_0$ with

$$\|\boldsymbol{\theta}_j - \boldsymbol{\theta}_{j+1}\|_* \leq \eta G, \quad s \leq \frac{2L_\gamma}{\eta G},$$

and whose left Riemann sum is at least $c_0/2$. Set

$$\mathbf{g}_{j+1} := \frac{\boldsymbol{\theta}_j - \boldsymbol{\theta}_{j+1}}{\eta}.$$

The FTRL dual state follows the polygon exactly, and one traversal earns at least $c_0/(2\eta)$ regret. Its average regret per round is at least $c_0G/(4L_\gamma)$. Repeating the loop gives

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \geq \frac{c_0G}{4L_\gamma} T - O(1/\eta).$$

□

Corollary 9.6. Assume that $\tilde{D}_\phi$ is C^1 on the open convex set Θ . It is R -FTRL-exact if and only if its Jacobian is symmetric. If a skew derivative is nonzero, Theorem 9.5 gives linear regret along a small dual loop. Under the regularity and normalization assumptions of Theorem 9.3, an exact field has sublinear regret.

Proof. The equivalence between a symmetric Jacobian and a potential follows from the Poincaré lemma on the open convex set Θ . A nonzero skew derivative has nonzero circulation on a sufficiently small rectangle, so Theorem 9.5 applies. The positive statement follows from Corollary 9.4; its horizon-independent bounds give $O(T^{1/q})$ regret, which is sublinear because $q > 1$. □

D.12 Proofs for Section 9.3

Proposition 9.7. Let K be compact and convex, and let R be differentiable and strongly convex on a neighborhood of K . If a feasible map $\phi : K \rightarrow K$ is R_K -FTRL-exact, then it is R -mirror-exact in the sense of Definition 8.8. Consequently,

$$\Phi_{\text{FTRL}}^{R_K}(K) \subseteq \Phi_{\text{exact}}^R(K).$$

Proof. Strong convexity makes R_K^* differentiable. For every $\mathbf{x} \in K$,

$$\nabla R(\mathbf{x}) \in \nabla R(\mathbf{x}) + N_K(\mathbf{x}) = \partial R_K(\mathbf{x}),$$

because $0 \in N_K(\mathbf{x})$. Fenchel duality and uniqueness give

$$\nabla R_K^*(\nabla R(\mathbf{x})) = \mathbf{x}.$$

If ϕ is R_K -FTRL-exact with potential Ψ , substitute $\boldsymbol{\theta} = \nabla R(\mathbf{x})$ in Definition 9.2 to obtain

$$\mathbf{x} - \phi(\mathbf{x}) = \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\mathbf{x})).$$

This is Definition 8.8. □

Theorem 9.8. *The map $\phi_{p,\varepsilon}$ is a smooth feasible deviation. Standard constrained mirror descent with distance generator R_p has $O(\sqrt{T})$ regret against it for bounded gradients and the usual $T^{-1/2}$ step size. In contrast, FTRL with the extended regularizer $R_{p,K}$ can be forced, for every sufficiently small constant step size, to incur $cT - O(1/\eta)$ regret against the same deviation. Hence the inclusion in Proposition 9.7 is strict and $\Phi_{\text{FTRL}}^{R_{p,K}}(K) \subsetneq \Phi_{\text{exact}}^{R_p}(K)$.*

Proof. Let $c := (a+b)/2$, $d := (b-a)/2$, choose $\beta \in (0,1)$, and use the map in (14). The Hessian is

$$\nabla^2 R_p(\mathbf{x}) = (p-1) \text{diag}(\mathbf{x}[1]^{p-2}, \mathbf{x}[2]^{p-2}),$$

and every diagonal entry is at least $m_p := (p-1) \min_{s \in [a,b]} s^{p-2} > 0$. Write $\mathbf{z} := \mathbf{x} - c\mathbf{1}$. For each coordinate i and the other coordinate j ,

$$\begin{aligned} \mathbf{D}_{p,\varepsilon}(\mathbf{x})[i] &= \frac{\varepsilon}{(p-1)\mathbf{x}[i]^{p-2}} (\mathbf{z}[i] + \beta \mathbf{z}[j]), \\ |\mathbf{D}_{p,\varepsilon}(\mathbf{x})[i]| &\leq \frac{\varepsilon(1+\beta)d}{m_p}. \end{aligned}$$

If $\mathbf{D}_{p,\varepsilon}(\mathbf{x})[i] > 0$, then $\mathbf{z}[i] > -\beta \mathbf{z}[j] \geq -\beta d$, and hence $\mathbf{x}[i] - a > (1-\beta)d$. The assumed upper bound on ε gives $\mathbf{D}_{p,\varepsilon}(\mathbf{x})[i] \leq \mathbf{x}[i] - a$. Therefore $\phi_{p,\varepsilon}(\mathbf{x})[i] \geq a$. Its upper bound is immediate because subtracting a positive displacement moves the coordinate downward. If $\mathbf{D}_{p,\varepsilon}(\mathbf{x})[i] < 0$, the same argument gives $b - \mathbf{x}[i] > (1-\beta)d$ and $|\mathbf{D}_{p,\varepsilon}(\mathbf{x})[i]| \leq b - \mathbf{x}[i]$. Thus $\phi_{p,\varepsilon}(K) \subseteq K$.

Define

$$V(\mathbf{x}) := \frac{\varepsilon}{2} (\mathbf{x} - c\mathbf{1})^\top S_\beta (\mathbf{x} - c\mathbf{1}).$$

Then

$$\nabla V(\mathbf{x}) = \varepsilon S_\beta (\mathbf{x} - c\mathbf{1}) = \nabla^2 R_p(\mathbf{x}) \mathbf{D}_{p,\varepsilon}(\mathbf{x}).$$

Hence $\mathbf{D}_{p,\varepsilon}(\mathbf{x})^\top d\nabla R_p(\mathbf{x}) = dV(\mathbf{x})$. Equivalently, on the positive dual orthant,

$$\Psi(\boldsymbol{\theta}) := \frac{\varepsilon}{2} \left(\begin{pmatrix} \boldsymbol{\theta}[1]^{1/(p-1)} \\ \boldsymbol{\theta}[2]^{1/(p-1)} \end{pmatrix} - c\mathbf{1} \right)^\top S_\beta \left(\begin{pmatrix} \boldsymbol{\theta}[1]^{1/(p-1)} \\ \boldsymbol{\theta}[2]^{1/(p-1)} \end{pmatrix} - c\mathbf{1} \right)$$

satisfies $\mathbf{D}_{p,\varepsilon}(\mathbf{x}) = \nabla_{\boldsymbol{\theta}} \Psi(\nabla R_p(\mathbf{x}))$. Because the box is compact and bounded away from the coordinate hyperplanes, R_p , $\mathbf{D}_{p,\varepsilon}$, and Ψ are smooth on neighborhoods of the relevant primal and dual sets. The Hessian of R_p^* is positive definite there. Compactness therefore gives finite constants for the relative curvature, displacement Lipschitz modulus, mirror-gradient Lipschitz modulus, and potential oscillation required by Corollary 8.10. Standard mirror descent consequently has $O(\sqrt{T})$ regret against the map under bounded gradients.

For FTRL, the minimization defining $\nabla R_{p,K}^*$ is separable, and its i th coordinate is

$$\nabla R_{p,K}^*(\boldsymbol{\theta})[i] = \begin{cases} a, & \boldsymbol{\theta}[i] \leq a^{p-1}, \\ \boldsymbol{\theta}[i]^{1/(p-1)}, & a^{p-1} < \boldsymbol{\theta}[i] < b^{p-1}, \\ b, & \boldsymbol{\theta}[i] \geq b^{p-1}. \end{cases}$$

Choose numbers $u < v$ with $b^{p-1} < u < v$ and $s < t$ with $a^{p-1} < s < t < b^{p-1}$. On the rectangle $[u, v] \times [s, t]$, the first primal coordinate is b and the second is $\boldsymbol{\theta}[2]^{1/(p-1)}$. Consequently the first component of the dual displacement is strictly increasing in $\boldsymbol{\theta}[2]$:

$$\frac{\partial \tilde{D}_{p,\varepsilon}[1]}{\partial \boldsymbol{\theta}[2]} = \frac{\varepsilon \beta}{(p-1)^2} b^{2-p} \boldsymbol{\theta}[2]^{(2-p)/(p-1)} > 0,$$

whereas $\partial \tilde{D}_{p,\varepsilon}[2]/\partial \boldsymbol{\theta}[1] = 0$. Equivalently, the circulation around the positively oriented rectangle is

$$(v - u) \left[\tilde{D}_{p,\varepsilon}[1](u, s) - \tilde{D}_{p,\varepsilon}[1](u, t) \right] < 0.$$

The dual displacement is continuous, so Theorem 9.5 produces bounded gradients with $cT - O(1/\eta)$ FTRL regret. This proves both the claimed algorithmic separation and the strictness of Proposition 9.7. $\square$

D.13 Proofs for Section 9.4

Corollary 9.9. *Let $1 < p \leq 2$ and let $q := p/(p-1)$. Then*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{B}{\eta} + \frac{L\eta}{2(p-1)} \sum_{t=1}^T \|\mathbf{g}_t\|_q^2.$$

If $\|\mathbf{g}_t\|_q \leq G$, the optimized bound is $G\sqrt{2LBT/(p-1)}$.

Proof. The function $\frac{1}{2} \|\cdot\|_p^2$ is $(p-1)$ -strongly convex with respect to $\|\cdot\|_p$ for $1 < p \leq 2$. If $\boldsymbol{\theta} = \nabla(\frac{1}{2} \|\mathbf{x}\|_p^2) + \mathbf{n}$ with $\mathbf{n} \in N_K(\mathbf{x})$, then $\langle \mathbf{n}, \mathbf{y} - \mathbf{x} \rangle \leq 0$ for $\mathbf{y} \in K$, so adding ι_K preserves the same strong-convexity inequality for every selected subgradient. Apply Theorem 9.3 with uniform-convexity exponent $r = 2$, conjugate exponent 2, and $\sigma = p-1$. $\square$

Corollary 9.10. *Let $2 \leq p < \infty$ and let $q := p/(p-1)$. Then*

$$\sum_{t=1}^T \langle \mathbf{g}_t, \mathbf{x}_t - \phi(\mathbf{x}_t) \rangle \leq \frac{B}{\eta} + \frac{L\eta^{q-1}}{q(2^{2-p})^{q-1}} \sum_{t=1}^T \|\mathbf{g}_t\|_q^q.$$

If $\|\mathbf{g}_t\|_q \leq G$, the optimized bound is $p^{1/p} 2^{(p-2)/p} GL^{1/q} B^{1/p} T^{1/q}$, which is $O(T^{1-1/p})$.

Proof. For $A(s) = |s|^{p-2} s$, the scalar inequality

$$(A(a) - A(b))(a - b) \geq 2^{2-p} |a - b|^p$$

holds for all real a, b . Integrating coordinatewise along the segment from $\mathbf{x}$ to $\mathbf{y}$ gives

$$D_{p^{-1}\|\cdot\|_p^p}(\mathbf{y}, \mathbf{x}) \geq \frac{2^{2-p}}{p} \|\mathbf{y} - \mathbf{x}\|_p^p.$$

The same normal-cone argument as above shows that adding ι_K preserves the estimate for every selected subgradient. Thus the regularizer is $(2^{2-p}, p)$ -uniformly convex. Apply Theorem 9.3 and Corollary 9.4. $\square$

D.14 Proofs for Section 10

Theorem 10.3. Let $\bar{\mu}_T := T^{-1} \sum_{t=1}^T \delta_{\mathbf{x}_t}$ and define $\mathbf{g}_{i,t} := \nabla_{\mathbf{x}_i} \ell_i(\mathbf{x}_{i,t}, \mathbf{x}_{-i,t})$. If every player i guarantees

$$\sup_{\phi_i \in \Phi_i} \sum_{t=1}^T \langle \mathbf{g}_{i,t}, \mathbf{x}_{i,t} - \phi_i(\mathbf{x}_{i,t}) \rangle \leq R_{i,T},$$

then $\bar{\mu}_T$ is an ε_T -ConCE relative to $(\Phi_i)_i$, where $\varepsilon_T = \max_i R_{i,T}/T$. In particular, OGD on $\Phi_{i,\text{exact}}(B_i, L_i)$ with $\|\mathbf{g}_{i,t}\| \leq G_i$ gives $\varepsilon_T \leq \max_i G_i \sqrt{6L_i B_i/T}$.

Proof. Fix player i and $\phi_i \in \Phi_i$. By convexity of ℓ_i in its own action,

$$\begin{aligned} \ell_i(\mathbf{x}_{i,t}, \mathbf{x}_{-i,t}) - \ell_i(\phi_i(\mathbf{x}_{i,t}), \mathbf{x}_{-i,t}) \\ \leq \langle \nabla_{\mathbf{x}_i} \ell_i(\mathbf{x}_{i,t}, \mathbf{x}_{-i,t}), \mathbf{x}_{i,t} - \phi_i(\mathbf{x}_{i,t}) \rangle. \end{aligned}$$

Summing, taking the supremum over Φ_i , and dividing by T gives the definition of ε_T -ConCE under $\bar{\mu}_T$. For the final statement apply Corollary 4.3 player by player. $\square$

Proposition 10.1. Let $\mathcal{A}_i$ be any family of feasible Borel-measurable deviations for player i . A distribution satisfies $\Gamma_i(\mu, \phi_i) \leq 0$ for every $\phi_i \in \mathcal{A}_i$ and every player if and only if it satisfies the same inequalities for every $\phi_i \in \text{Rad}_I(\mathcal{A}_i)$.

Proof. Since $\mathcal{A}_i \subseteq \text{Rad}_I(\mathcal{A}_i)$ by taking $\alpha = 1$, one implication is immediate. For the converse, suppose that μ satisfies every constraint in $\mathcal{A}_i$, and let $\phi_i \in \text{Rad}_I(\mathcal{A}_i)$. Choose $\alpha \in (0, 1]$ such that $\phi_{i,\alpha} = (1 - \alpha)I + \alpha\phi_i$ belongs to $\mathcal{A}_i$. Convexity of ℓ_i in its own action gives pointwise

$$\begin{aligned} \ell_i(\mathbf{x}_i, \mathbf{x}_{-i}) - \ell_i(\phi_{i,\alpha}(\mathbf{x}_i), \mathbf{x}_{-i}) \\ \geq \alpha (\ell_i(\mathbf{x}_i, \mathbf{x}_{-i}) - \ell_i(\phi_i(\mathbf{x}_i), \mathbf{x}_{-i})). \end{aligned}$$

The joint continuity and compactness assumptions ensure that both loss differences are bounded and measurable. Thus both gains are integrable. Taking expectations yields $\Gamma_i(\mu, \phi_{i,\alpha}) \geq \alpha \Gamma_i(\mu, \phi_i)$. The left-hand side is nonpositive, so $\Gamma_i(\mu, \phi_i) \leq 0$. $\square$

Theorem 10.4. Under the unrestricted conventions above, every convex game with compact convex action sets and jointly continuous losses satisfies

$$\text{ConCE} = \text{PCE}.$$

Proof. By Corollary 6.3, the smooth exact-form class is the identity-radial closure of the unrestricted proximal class. Apply Proposition 10.1 player by player. $\square$

Corollary 10.5. For any set of regularizers R_i for $i \in [N]$, we have

$$\text{PCE}_{(R_i)_i}^{\text{Breg}} = \text{Eq}_{(R_i)_i}^{\text{rad}}.$$

Proof. For every player, $\Phi_{\text{rad}}^{R_i}(K_i) = \text{Rad}_I(\Phi_{\text{prox}}^{R_i}(K_i))$ by Definition 8.13. Apply Proposition 10.1 player by player. $\square$

Corollary 10.7. For every finitely supported distribution in a convex game,

$$\mu \in \text{CE} \iff \mu \in \text{ConCE} \iff \mu \in \text{PCE}.$$

In particular the three notions coincide in every finite game.

Proof. Only $\text{PCE} \subseteq \text{CE}$ needs proof. If a general deviation τ_i has positive gain, list the finitely many recommendations $\mathbf{x}^1, \dots, \mathbf{x}^m$ to player i and set $\mathbf{y}^j := \tau_i(\mathbf{x}^j)$. By Theorem 10.6, for sufficiently small common $\alpha > 0$ there is a proximal map satisfying

$$\phi_{i,\alpha}(\mathbf{x}^j) = (1 - \alpha)\mathbf{x}^j + \alpha\tau_i(\mathbf{x}^j)$$

at every supported recommendation. Convexity yields $\Gamma_i(\mu, \phi_{i,\alpha}) \geq \alpha\Gamma_i(\mu, \tau_i) > 0$, so $\mu \notin \text{PCE}$. $\square$

Theorem 10.8. *Suppose that every player has a compact interval action set $K_i := [a_i, b_i] \subseteq \mathbb{R}$ and, in addition to the standing assumptions, every loss $\ell_i : K \rightarrow \mathbb{R}$ is continuous on $K := \prod_{j=1}^N K_j$. Then*

$$\text{CE} = \text{ConCE} = \text{PCE}$$

under the unrestricted definitions used in this section when the approximation error is zero.

Proof. The inclusion $\text{CE} \subseteq \text{ConCE}$ follows directly from the deviation classes, and $\text{ConCE} = \text{PCE}$ follows from Theorem 10.4. It remains to prove $\text{ConCE} \subseteq \text{CE}$.

Suppose $\mu \notin \text{CE}$. Then for some player i , there is a Borel-measurable deviation $\tau_i : K_i \rightarrow K_i$ such that $\Gamma_i(\mu, \tau_i) > 0$. Let ν_i be the i -th marginal of μ . Assume that $K_i := [a_i, b_i]$ with $a_i < b_i$ (the case in which K_i is a singleton is trivial).

For every n , by Lusin's theorem, we have a compact set $A_n \subseteq K_i$ with $\nu_i(K_i \setminus A_n) < 1/n$ such that $\tau_i|_{A_n}$ is continuous. From classical extension theorem (bounded form of the Tietze extension theorem), we can get a continuous map $h_n : K_i \rightarrow K_i$ agreeing with τ_i on A_n . After rescaling $u := (x - a_i)/(b_i - a_i)$, let $B_m h_n$ denote the Bernstein polynomial

$$(B_m h_n)(x) = \sum_{k=0}^m h_n\left(a_i + \frac{k}{m}(b_i - a_i)\right) \binom{m}{k} u^k (1-u)^{m-k}.$$

Bernstein approximation is uniform for continuous functions, so choose m_n such that $\|B_{m_n} h_n - h_n\|_\infty < 1/n$ and set $\phi_{i,n} := B_{m_n} h_n$. Bernstein coefficients are nonnegative and sum to one, so we have $\phi_{i,n}(K_i) \subseteq K_i$ and $\nu_i(|\phi_{i,n} - \tau_i| > \frac{1}{n}) \leq \nu_i(K_i \setminus A_n) < \frac{1}{n}$. So $\phi_{i,n}(X_i) \rightarrow \tau_i(X_i)$ in probability under μ .

Continuity on compact K makes ℓ_i bounded and uniformly continuous. So we get $\ell_i(\phi_{i,n}(X_i), X_{-i}) \rightarrow \ell_i(\tau_i(X_i), X_{-i})$ in probability. The variables are uniformly bounded so the convergence is in $L^1(\mu)$. This shows that $\Gamma_i(\mu, \phi_{i,n}) \rightarrow \Gamma_i(\mu, \tau_i) > 0$. In particular, $\phi_{i,n}$ has positive gain for all sufficiently large n .

But every $\phi_{i,n}$ is a smooth exact-form deviation. Indeed, the polynomial potential

$$\Psi_{i,n}(x) = \int_{a_i}^x (s - \phi_{i,n}(s)) \, ds$$

is defined on a neighborhood of K_i and satisfies $\Psi'_{i,n}(x) = x - \phi_{i,n}(x)$.

So a measurable profitable deviation produces a smooth feasible exact-form profitable deviation, so $\mu \notin \text{ConCE}$. This proves $\text{ConCE} \subseteq \text{CE}$ and completes the equality. $\square$

Theorem 10.9. *There is a two-player convex game and a distribution μ with uncountable support—indeed, with marginals absolutely continuous with respect to two-dimensional Lebesgue measure—for which*

$$\mu \in \text{ConCE} = \text{PCE} \quad \text{but} \quad \mu \notin \text{CE}.$$

Consequently $\text{CE} \subsetneq \text{ConCE} = \text{PCE}$ under the definitions above.

Proof. Let $K \subseteq \mathbb{R}^2$ be the unit disk, and let

$$J := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

be the counterclockwise 90° rotation matrix. Player 1 has loss $\ell_1(\mathbf{a}, \mathbf{s}) = \langle J\mathbf{s}, \mathbf{a} \rangle$, player 2 has zero loss, and μ is the law of (X, X) for X uniform on the disk. Thus both marginals are continuous and the joint law has uncountable support, although it is singular and supported on the diagonal of $K \times K$. For every exact map $\phi = I - \nabla\Psi$, the gain is

$$\Gamma_1(\mu, \phi) = \frac{1}{\pi} \int_K \langle J\mathbf{x}, \nabla\Psi(\mathbf{x}) \rangle \, d\mathbf{x} = 0,$$

because $\operatorname{div}(J\mathbf{x}) = 0$ and $J\mathbf{x}$ has zero normal flux. Thus $\mu \in \operatorname{ConCE} = \operatorname{PCE}$. The feasible measurable map $\tau(\mathbf{x}) := -J\mathbf{x}/\|\mathbf{x}\|$ for $\mathbf{x} \neq 0$, with $\tau(0) := 0$, has gain $\mathbb{E}\|X\| = 2/3$, so $\mu \notin \operatorname{CE}$. $\square$