Exact-Form Regret and Conservative Correlated Equilibria

Ashkan Soleymani Gabriele Farina Patrick Jaillet
Massachusetts Institute of Technology
{ashkanso,gfarina,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 and Legendre mirror descent 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, and for mirror descent it is the geometry induced by the regularizer. 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 shows how these classes depend on the geometry used by the algorithm. 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 connect these relationships to no-regret learning. Overall, this work gives a unified geometric account of what first-order online learning algorithms are no-regret with respect to, beyond fixed comparators. An extended version of this paper is available at https://arxiv.org/abs/2609.09466.

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 $\boldsymbol{x}' \in K$. 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 $\boldsymbol{x}_t$ with fixed alternatives, while a general deviation may replace $\boldsymbol{x}_t$ by a feasible point $\phi(\boldsymbol{x}_t) \in K$ 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

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

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 $\phi\colon K \to K$. Writing $\boldsymbol{D}_\phi(\boldsymbol{x}) \coloneqq \boldsymbol{x} - \phi(\boldsymbol{x})$, we show that online gradient descent has $O(\sqrt{T})$ regret against every sufficiently regular feasible deviation whose displacement is exact, $\boldsymbol{D}_\phi = \nabla\Psi$, and obtain a uniform bound over normalized families with common smoothness and potential-oscillation parameters. Thus the comparator is the actual endpoint $\phi(\boldsymbol{x}_t)$, 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 gives explicit finite-time bounds for feasible endomorphisms.

The same structure appears in proximal deviations. If $F = f + \iota_K$ is weakly convex, the Moreau-envelope identity gives $\boldsymbol{x} - \operatorname{prox}_F(\boldsymbol{x}) = \nabla M_F(\boldsymbol{x})$, 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 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 $\boldsymbol{D}_\phi(\boldsymbol{x})^{\mathsf{T}}\mathrm{d}\nabla R(\boldsymbol{x})$. We develop the theory in the Legendre interior, where the geometry is most transparent, and recover Bregman proximal deviations as a special case. We defer the extension of mirror descent beyond Legendre regularizers, the corresponding theory for follow-the-regularized-leader (FTRL) and the discussion around the resulting equilibria to the full version of this paper [Soleymani et al., 2026].

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 $\mathrm{ConCE} = \mathrm{PCE}$, where PCE denotes proximal correlated equilibrium. 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 $\mathrm{CE} \subsetneq \mathrm{ConCE} = \mathrm{PCE}$.

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, $\boldsymbol{x}_t$ be the decision at time t, and $\langle \boldsymbol{g}_t, \cdot \rangle$ be the

linear loss incurred at time t by the decision maker, regret takes the form

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

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 generalizes external regret by comparing online decisions $\boldsymbol{x}_t$ with their transformations $\phi(\boldsymbol{x}_t)$ for a family of deviations Φ,

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

External regret corresponds to the special case of constant deviations. Motivated partly by the connection between regret and equilibrium notions [Foster and Vohra, 1997, Hart and Mas-Colell, 2000], a long line of work has developed efficient algorithms for increasingly rich deviation classes, often using sophisticated fixed-point machinery [Farina et al., 2022b, Daskalakis et al., 2025]. Recent results, however, show that online gradient descent already minimizes regret against important nonconstant deviations, including symmetric linear endomorphisms and proximal transformations [Ahunbay, 2026, Ahunbay and Bichler, 2025, Cai et al., 2024, 2026]. These findings raise a natural question about the full class of deviations against which first-order algorithms guarantees sublinear regret.

Additional discussion of first-order equilibrium theory, proximal regret, and gradient equilibrium appears in Appendix A.

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 $\operatorname{relint}(K)$, derivatives and differential forms are taken in affine coordinates on $\operatorname{aff}(K)$. The mirror section introduces its own explicit affine-hull convention. Vectors are written in bold, e.g. $\boldsymbol{x}$, $\boldsymbol{g}$, and coordinates use brackets, e.g. $\boldsymbol{x}[j]$. We write $[T] \coloneqq \{1, \dots, T\}$. 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 $\boldsymbol{x} \in K$ by

$$N_K(\boldsymbol{x}) \coloneqq \left\{ \boldsymbol{n} \in \mathbb{R}^d \,:\, \langle \boldsymbol{n}, \boldsymbol{y} - \boldsymbol{x} \rangle \le 0 \text{ for every } \boldsymbol{y} \in K \right\}.$$

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

$$\boldsymbol{D}_\phi(\boldsymbol{x}) \coloneqq \boldsymbol{x} - \phi(\boldsymbol{x}), \qquad \mathrm{Reg}_T(\phi) \coloneqq \sum_{t=1}^{T} \langle \boldsymbol{g}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle$$

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

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

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

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

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

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

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

Throughout the paper, the subscript identifies the deviation class. Thus Φ_{prox} denotes the Euclidean proximal class, $\Phi^{\text{sm}}_{\text{exact}}$ 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,

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

Bounds for a fixed deviation may depend on the parameters of that deviation and therefore do not automatically give 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 $\phi : K \to K$ for which there is a differentiable potential Ψ_ϕ, defined on a neighborhood of K, satisfying*

$$\boldsymbol{x} - \phi(\boldsymbol{x}) = \nabla\Psi_\phi(\boldsymbol{x}), \qquad \operatorname{osc}_K(\Psi_\phi) \leq B,$$

and whose gradient satisfies

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

Proofs not given in the main text are found in the appendices. Appendix B records the convex-analytic preliminaries, and Appendix E contains the remaining proofs in the order of the main text.

4 Online Gradient Descent Minimizes Exact-Form Regret

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

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

We prove a telescoping regret bound for every smooth feasible exact-form deviation ϕ. Under bounded gradients and bounded potential oscillation, this gives $\operatorname{Reg}_T(\phi) = \mathcal{O}(\sqrt{T})$. 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 $\{\boldsymbol{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online gradient descent* (OGD)*. Let $\phi : K \to K$ be an exact-form deviation $\boldsymbol{D}_\phi(\boldsymbol{x}) = \boldsymbol{x} - \phi(\boldsymbol{x}) = \nabla\Psi(\boldsymbol{x})$ where Ψ is L-smooth with respect to the Euclidean norm. Then, for every sequence $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_T$,*[1]

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

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

Corollary 4.2. *Suppose $\|\boldsymbol{g}_t\| \leq G$ and fix $B, L > 0$. Then OGD with $\eta = \sqrt{2B/(3LG^2T)}$ satisfies*

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

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

[1] Similar bounds can be inferred directly by Ahunbay [2026, Theorem 5.2].

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], which was also observed by Ahunbay [2026]. Let $F = f + \iota_K$ be proper, lower semicontinuous, and ρ-weakly convex with $0 \le \rho < 1$, meaning that $F + \frac{\rho}{2}\|\cdot\|^2$ is convex. The indicator ι_K is zero on K and $+\infty$ outside K. Define

$$\operatorname{prox}_F(\boldsymbol{x}) \coloneqq \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathbb{R}^d} \left\{ F(\boldsymbol{y}) + \tfrac{1}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^2 \right\}, \qquad M_F(\boldsymbol{x}) \coloneqq \min_{\boldsymbol{y}\in\mathbb{R}^d} \left\{ F(\boldsymbol{y}) + \tfrac{1}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^2 \right\}.$$

Because F includes ι_K, the minimization is over K.

$$\boldsymbol{x} - \operatorname{prox}_F(\boldsymbol{x}) = \nabla M_F(\boldsymbol{x}). \tag{2}$$

Theorem 5.1. *Let $F \coloneqq f + \iota_K$ be proper, lower semicontinuous, and ρ-weakly convex with $0 \le \rho < 1$. Define $\phi_F(\boldsymbol{x}) \coloneqq \operatorname{prox}_F(\boldsymbol{x})$. Then $\phi_F : K \to K$ is feasible and Euclidean exact-form, with potential M_F, i.e., $\boldsymbol{x} - \phi_F(\boldsymbol{x}) = \nabla M_F(\boldsymbol{x})$, 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, $\|\boldsymbol{g}_t\| \le G$, $\operatorname{osc}_K(M_F) \le B_F$, and $B_F, G > 0$, then choosing $\eta = \sqrt{2B_F/(3L_F G^2 T)}$ gives*

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

If $B_F = 0$ or $G = 0$, the corresponding regret bound is trivial and the bound follows without optimizing the step size.

The Moreau-envelope identity recovers fixed comparators, projections to convex subsets, and interpolation deviations. Symmetric affine swaps are also exact-form, since their displacements are gradients of quadratic potentials. Proposition E.4 records these examples and their potentials.

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. Every weakly convex proximal map with parameter $0 \le \rho < 1$ is monotone, whereas feasible exact-form deviations can be nonmonotone.

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

The proof is given in Appendix E.3. This strict inclusion does not carry over to the unrestricted equilibrium sets when the approximation error is zero. The following scaling result is the mechanism behind this equivalence.

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

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

coincides on K with $\operatorname{prox}_{F_\alpha}(\boldsymbol{x})$ 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]. Its proof is given in Appendix C. Write Φ_{prox} for the constrained weakly convex proximal maps and $\Phi^{\mathrm{sm}}_{\mathrm{exact}}$ for the feasible exact-form maps whose potentials are $C^{1,1}_{\mathrm{loc}}$ near K. For any class $\mathcal{A}$ of feasible maps with a common domain, define its *identity-radial closure* by

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

Corollary 6.3. *On compact convex K, the Euclidean identity-radial closure satisfies* $\mathrm{Rad}_I(\Phi_{\mathrm{prox}}) = \Phi^{\mathrm{sm}}_{\mathrm{exact}}$.

Thus $\Phi_{\mathrm{prox}} \subsetneq \Phi^{\mathrm{sm}}_{\mathrm{exact}}$, but every smooth exact-form deviation admits a sufficiently small positive interpolation that is proximal. The equilibrium consequence is developed in Section 9.

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 $\boldsymbol{D} : K \to \mathbb{R}^d$ be continuous and let $\gamma : [0,1] \to K$ be a closed piecewise C^1 loop. The circulation of $\boldsymbol{D}$ along γ is*

$$\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^{\mathsf{T}} \mathrm{d}\boldsymbol{x} \coloneqq \int_0^1 \langle \boldsymbol{D}(\gamma(s)), \gamma'(s) \rangle \, \mathrm{d}s.$$

Theorem 7.2. *Let $\boldsymbol{D} : K \to \mathbb{R}^d$ be continuous. Suppose there exists a closed piecewise C^1 loop $\gamma : [0,1] \to K$ such that, after possibly reversing orientation,*

$$-\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^{\mathsf{T}} \mathrm{d}\boldsymbol{x} \coloneqq c_0 > 0.$$

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

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

for a constant $c > 0$ depending only on $\boldsymbol{D}$, 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 proof in Appendix E.4 discretizes the loop and repeatedly drives the iterates around it. The lower bound uses a sufficiently small constant step size, whereas Corollary 4.2 uses $\eta \asymp T^{-1/2}$. Remark E.8 gives an anytime geometric-epoch version.

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 the dual point $\boldsymbol{\theta} = \nabla R(\boldsymbol{x})$, for a distance-generating function R. At an interior iterate, where the normal-cone term vanishes, mirror descent is the dual update $\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \boldsymbol{g}_t$. The same telescoping proof works whenever the primal displacement $\boldsymbol{D}_\phi(\boldsymbol{x})$ is a gradient as a function of the dual coordinate $\boldsymbol{\theta}$.

We work in the Legendre interior. Mirror descent beyond the Legendre setting is discussed in the extended version of this paper [Soleymani et al., 2026]. Throughout 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. We use the standard Legendre regime. Let $R : E \to (-\infty, +\infty]$ be a proper closed Legendre function satisfying

$$\overline{\mathrm{dom}\, R} = K, \qquad \mathcal{D} \coloneqq \mathrm{int}_E(\mathrm{dom}\, R) = \mathrm{relint}(K).$$

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

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

Legendre duality gives mutually inverse C^1 diffeomorphisms $\nabla R : \mathcal{D} \to \Theta$ and $\nabla R^* : \Theta \to \mathcal{D}$. 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 $\boldsymbol{\theta} \coloneqq \nabla R(\boldsymbol{x})$.

We use the Bregman divergence generated by R, defined as

$$D_R(\boldsymbol{y}, \boldsymbol{x}) \coloneqq R(\boldsymbol{y}) - R(\boldsymbol{x}) - \langle \nabla R(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle .$$

With this notation, online mirror descent takes the form

$$\boldsymbol{x}_{t+1} \coloneqq \operatorname*{arg\,min}_{\boldsymbol{x} \in K} \{ \eta \langle \boldsymbol{g}_t, \boldsymbol{x} \rangle + D_R(\boldsymbol{x}, \boldsymbol{x}_t) \} . \tag{3}$$

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

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

Equivalently, the one-form $\boldsymbol{D}_\phi(\boldsymbol{x})^\mathsf{T} \mathrm{d}\nabla R(\boldsymbol{x})$ is exact.

We write $\Phi^R_{\mathrm{exact}}(\mathcal{D})$ for the class of mirror-exact deviations in the Legendre interior.

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

A primal-coordinate characterization is given in Proposition E.9.

8.1 Examples of Mirror-Exact Deviations

To construct mirror-exact deviations, choose a scalar potential $\Psi(\boldsymbol{\theta})$ in dual space and set $\phi(\boldsymbol{x}) = \boldsymbol{x} - \nabla_{\boldsymbol{\theta}} \Psi(\nabla R(\boldsymbol{x}))$, checking that the resulting map is feasible.

In Appendix E.6, 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 $\boldsymbol{a} \in \mathbb{R}^d$,

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

as well as geometric interpolation deviations anchored at $\boldsymbol{u} \in \operatorname{relint}(\Delta_d)$ and mirror-exact deviations that are not Euclidean-exact. Their potentials and feasibility conditions are given in Appendix E.6.

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

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

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

$$D_\Psi(\boldsymbol{\theta}', \boldsymbol{\theta}) \leq L D_{R^*}(\boldsymbol{\theta}', \boldsymbol{\theta}) \tag{4}$$

for every such pair $\boldsymbol{\theta}', \boldsymbol{\theta} \in \Theta$. When $K = E$ and $R(\boldsymbol{x}) = \frac{1}{2}\|\boldsymbol{x}\|^2$, 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 8.3. *Assume the standing Legendre-domain conditions and that R is m-strongly convex on $\mathcal{D}$. Let $\phi : \mathcal{D} \to K$ be feasible and R-mirror-exact with dual potential Ψ. Let $\{\boldsymbol{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online mirror descent (3), initialized in $\mathcal{D}$, with iterates remaining in $\mathcal{D}$. Write*

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

Assume that (4) 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

$$\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle \leq \frac{\Psi(\boldsymbol{\theta}_1) - \Psi(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{L\eta}{m} \sum_{t=1}^{T} \|\boldsymbol{g}_t\|^2 . \tag{5}$$

Consequently, suppose $\|\boldsymbol{g}_t\| \leq G$ 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 $\operatorname{osc}_\Theta(\Psi) \leq B$ is a sufficient trajectory-independent assumption that avoids any dependence of the oscillation bound on the chosen step size.

The relative normal cone vanishes along the interior trajectory, so the proof works directly with the additive dual update. Appendix E.7 gives the telescoping argument. It is enough for (4) to hold along the realized trajectory.

Nonzero mirror circulation gives the corresponding lower bound. For a continuous field $\boldsymbol{D} : \mathcal{D} \to E$ and a closed piecewise C^1 loop $\gamma \subset \mathcal{D}$, this circulation is

$$\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^\mathsf{T} \mathrm{d}\nabla R(\boldsymbol{x}).$$

If it is nonzero, then for every gradient bound $G > 0$ and every sufficiently small constant step size, mirror descent can be forced to incur $cT - O(1/\eta)$ regret against $\boldsymbol{D}$. For a feasible deviation, take $\boldsymbol{D} = \boldsymbol{D}_\phi$. Theorem E.12 states and proves this Legendre consequence of Theorem 7.2.

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]. 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 $\boldsymbol{x} \in \mathcal{D}$, define

$$\operatorname{prox}_F^R(\boldsymbol{x}) := \operatorname*{arg\,min}_{\boldsymbol{y} \in E} \{F(\boldsymbol{y}) + D_R(\boldsymbol{y}, \boldsymbol{x})\}, \quad M_F^R(\boldsymbol{x}) := \min_{\boldsymbol{y} \in E} \{F(\boldsymbol{y}) + D_R(\boldsymbol{y}, \boldsymbol{x})\}.$$

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

Theorem 8.4. *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 $\phi_F^R : \mathcal{D} \to K$ defined by $\phi_F^R(\boldsymbol{x}) := \operatorname{prox}_F^R(\boldsymbol{x})$ is feasible and R-mirror-exact. Equivalently,*

$$\big(\boldsymbol{x} - \operatorname{prox}_F^R(\boldsymbol{x})\big)^\mathsf{T} \mathrm{d}\nabla R(\boldsymbol{x}) = \mathrm{d}M_F^R(\boldsymbol{x}). \tag{6}$$

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

We can now directly invoke Theorem 8.3 for this class of deviations in the following proposition.

Proposition 8.5. *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 $\{\boldsymbol{x}_t\}_{t=1}^{T+1} \subset \mathcal{D}$ be generated by online mirror descent (3). Then these iterates control Bregman proximal regret with the bound*

$$\sum_{t=1}^{T} \big\langle \boldsymbol{g}_t, \boldsymbol{x}_t - \operatorname{prox}_F^R(\boldsymbol{x}_t) \big\rangle \leq \frac{\Psi_F(\boldsymbol{\theta}_1) - \Psi_F(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{\eta}{m} \sum_{t=1}^{T} \|\boldsymbol{g}_t\|^2 , \tag{7}$$

where $\Psi_F(\boldsymbol{\theta}_t) = M_F^R(\boldsymbol{x}_t)$ 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 (4), and hence Theorem 8.3 applies.

9 Conservative Correlated Equilibria in Convex Games

The regret guarantees developed above translate directly into equilibrium constraints in convex games. 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 $\ell_i(\boldsymbol{x}_i, \boldsymbol{x}_{-i})$ convex and differentiable in its own action. Let $K := \prod_{i=1}^N K_i$. Assume also that each loss $\ell_i : K \to \mathbb{R}$ 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

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

A distribution is a correlated equilibrium (CE) if $\Gamma_i(\mu, \phi_i) \le 0$ for every player and every feasible Borel-measurable deviation, and a coarse correlated equilibrium (CCE) if the same condition holds for every constant deviation. It is a proximal correlated equilibrium (PCE) if the condition holds for all constrained weakly convex proximal maps [Cai et al., 2026]. We use the unrestricted union over all map-dependent $0 \le \rho < 1$, so CE $\subseteq$ PCE $\subseteq$ CCE.

Definition 9.1 (ConCE for a Specified Deviation Family)**.** *Fix for every player a family Φ_i of feasible Euclidean exact-form deviations. A distribution μ is an ε-conservative correlated equilibrium relative to $(\Phi_i)_i$ if*

$$\sup_{\phi_i\in\Phi_i} \Gamma_i(\mu, \phi_i) \le \varepsilon \qquad \textit{for every player } i.$$

For quantitative conclusions, we use the normalized family $\Phi_{i,\mathrm{exact}}(B_i, L_i)$ from Definition 3.2. When the approximation error is zero, ConCE tests every smooth feasible exact-form deviation, with a $C^{1,1}_{\mathrm{loc}}$ potential on a neighborhood of K_i, without common bounds on these potentials.

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

Theorem 9.2. *Let $\bar\mu_T := T^{-1}\sum_{t=1}^T \delta_{\boldsymbol{x}_t}$ and define $\boldsymbol{g}_{i,t} := \nabla_{\boldsymbol{x}_i}\ell_i(\boldsymbol{x}_{i,t}, \boldsymbol{x}_{-i,t})$. If every player i guarantees*

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

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 $\Phi_{i,\mathrm{exact}}(B_i, L_i)$ with $\|\boldsymbol{g}_{i,t}\| \le G_i$ gives $\varepsilon_T \le \max_i G_i\sqrt{6L_iB_i/T}$.

The proof uses only convexity of the losses and the linearized regret bound. It also applies to mirror-exact deviations along interior play. For a family with common relative-smoothness and trajectory-oscillation bounds, Theorem 8.3 gives the corresponding equilibrium error $2G_i\sqrt{L_iB_i/(m_iT)}$ for player i. These normalized approximate statements are separate from the unrestricted zero-tolerance equalities below.

9.1 Exact and Proximal Equilibrium Constraints

Passing from a deviation class to its identity-radial closure does not change its equilibrium constraints when the approximation error is zero.

Proposition 9.3. *Let $\mathcal{A}_i$ be any family of feasible Borel-measurable deviations for player i. A distribution satisfies $\Gamma_i(\mu, \phi_i) \le 0$ for every $\phi_i \in \mathcal{A}_i$ and every player if and only if it satisfies the same inequalities for every $\phi_i \in \mathrm{Rad}_I(\mathcal{A}_i)$.*

Indeed, convexity gives $\Gamma_i(\mu, (1-\alpha)I + \alpha\phi_i) \ge \alpha\Gamma_i(\mu, \phi_i)$ for every $\alpha \in (0, 1]$. Together with Corollary 6.3, this gives the following.

Theorem 9.4. *Under the unrestricted conventions above, every convex game with compact convex action sets and jointly continuous losses satisfies*

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

This equality relies on the unrestricted deviation classes. For normalized ε-equilibria, the admissible interpolation scale may depend on the deviation, so the difference between the map classes can still affect quantitative guarantees.

9.2 Finite Support and Boundary-Compatible Interpolation

On a finite support, only finitely many recommended actions and deviation values matter. Any such finite collection can be realized, after a sufficiently small common interpolation with the identity, by a single proximal map (Theorem D.2). Consequently, the three equilibrium notions cannot be separated on a finite support.

Corollary 9.5. *For every finitely supported distribution in a convex game,*

$$\mu \in \mathrm{CE} \quad \Longleftrightarrow \quad \mu \in \mathrm{ConCE} \quad \Longleftrightarrow \quad \mu \in \mathrm{PCE}.$$

In particular the three notions coincide in every finite game.

For finite games, the statement uses the standard mixed extension and distributions supported on pure-action profiles. Appendix D gives the interpolation theorem and its proof.

9.3 Full Correlated Equilibrium on Intervals

On an interval, every smooth scalar displacement has an antiderivative. Under joint continuity of the losses, smooth feasible deviations also approximate all measurable CE deviations.

Theorem 9.6 (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 $\ell_i : K \to \mathbb{R}$ is continuous on $K := \prod_{j=1}^{N} K_j$. Then*

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

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

In particular, infinite support alone does not separate these notions. Under continuous losses, such a separation requires at least one player's action set to have affine dimension at least two.

9.4 Strict Separation from CE on Continuous Support

With two-dimensional individual action spaces, the following separation is possible.

Theorem 9.7. *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*

$$\mu \in \mathrm{ConCE} = \mathrm{PCE} \qquad \textit{but} \qquad \mu \notin \mathrm{CE}.$$

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

The construction uses the diagonal law (X, X) with X uniform on a disk. Its marginals are absolutely continuous, while the joint law is singular in the product action space. Appendix E.10 gives the construction and all equilibrium proofs.

10 Conclusion

This paper develops a geometric theory of the action-dependent comparators controlled by first-order methods. In Euclidean geometry, exact feasible displacement fields give explicit regret bounds for projected gradient descent, whereas nonzero circulation can be exploited to force linear regret. In the Legendre setting, mirror descent uses the same principle in dual coordinates, and Bregman proximal deviations are mirror-exact. At the equilibrium level, strict inclusion of deviation maps does not imply strict inclusion of zero-tolerance equilibrium sets. Unrestricted ConCE and PCE coincide, and they coincide with CE on finite supports and on intervals, while two-dimensional continuous supports allow strict separation.

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.

Ashkan Soleymani, Gabriele Farina, and Patrick Jaillet. Exact-form regret for gradient descent, mirror descent and follow-the-regularized-leader. *arXiv preprint arXiv:2609.09466*, 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.

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.

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.

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.

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.

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.

A Additional Related Work

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 give the corresponding exactness picture for projected gradient descent. We focus instead on Φ-regret against the actual endpoints $\phi(\boldsymbol{x}_t)$ 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.

The connection to symmetric affine deviations is particularly simple. Consider an affine map $\phi(\boldsymbol{x}) = A\boldsymbol{x}+\boldsymbol{b}$, whose displacement is $\boldsymbol{D}_\phi(\boldsymbol{x}) = (I-A)\boldsymbol{x}-\boldsymbol{b}$. When $I-A$ is symmetric, this displacement is the gradient of a quadratic potential. Thus the symmetric affine deviations studied by Ahunbay and Bichler [2025], Cai et al. [2026] are finite maps with exact displacement fields. Relative to Ahunbay and Bichler [2025], our exact-form class extends the tractable linear–quadratic family to nonlinear feasible deviations, with 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 its regret guarantees and note that the class controlled by the algorithm may be strictly broader than the proximal class. We study this larger finite-map class 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 also develop a finite-map regret bound for mirror descent in the Legendre interior, where the relevant one-form is $\boldsymbol{D}_\phi(\boldsymbol{x})^{\mathsf{T}}\mathrm{d}\nabla R(\boldsymbol{x})$ and Bregman proximal deviations form a special case. In Euclidean geometry, identity interpolation recovers the smooth exact-form class. Finally, we derive the equilibrium consequences of these finite deviation classes. To our knowledge, these 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 $\Psi_{\boldsymbol{v}}(\boldsymbol{x}) = \langle \boldsymbol{v}, \boldsymbol{x}\rangle$, whose exact displacements are the constant vectors $\nabla\Psi_{\boldsymbol{v}} = \boldsymbol{v}$. Our framework replaces these constant directions by nonlinear finite displacements $\boldsymbol{D}_\phi(\boldsymbol{x}) = \nabla\Psi(\boldsymbol{x})$ generated by feasible endomorphisms $\phi\colon K \to K$. Both approaches rely on 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.

A.1 Interval Action Spaces and Binary Games

Two related results help place Theorem 9.6. In finite two-action games, Monnot and Piliouras [2017, Proposition 6.1] show that CE $=$ 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. Our interval result 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.

B Preliminaries

Lemma B.1. *Let $\boldsymbol{u} \in \mathbb{R}^d$ and $\boldsymbol{x}^+ \coloneqq \Pi_K(\boldsymbol{u})$. Then*

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

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

$$\left\langle \boldsymbol{x}^+ - \boldsymbol{u}, \boldsymbol{y} - \boldsymbol{x}^+ \right\rangle \geq 0 \qquad \forall \boldsymbol{y} \in K.$$

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

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

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

is convex. We use the weakly convex subdifferential convention

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

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

$$\langle \boldsymbol{u} - \boldsymbol{v}, \boldsymbol{p} - \boldsymbol{q} \rangle \geq -\rho\|\boldsymbol{p} - \boldsymbol{q}\|^2. \tag{8}$$

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

$$\nabla M_F(\boldsymbol{x}) = \boldsymbol{x} - \text{prox}_F(\boldsymbol{x}). \tag{9}$$

Moreover,

$$\|\text{prox}_F(\boldsymbol{x}) - \text{prox}_F(\boldsymbol{y})\| \leq \frac{1}{1-\rho}\|\boldsymbol{x} - \boldsymbol{y}\|, \tag{10}$$

so ∇M_F is Lipschitz with the crude bound

$$\|\nabla M_F(\boldsymbol{x}) - \nabla M_F(\boldsymbol{y})\| \leq \frac{2-\rho}{1-\rho}\|\boldsymbol{x} - \boldsymbol{y}\|. \tag{11}$$

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

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

The first-order optimality conditions are

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

Applying hypomonotonicity (8) gives

$$\langle (\boldsymbol{x} - \boldsymbol{p}) - (\boldsymbol{y} - \boldsymbol{q}), \boldsymbol{p} - \boldsymbol{q} \rangle \geq -\rho\|\boldsymbol{p} - \boldsymbol{q}\|^2,$$
$$\langle \boldsymbol{x} - \boldsymbol{y}, \boldsymbol{p} - \boldsymbol{q} \rangle - \|\boldsymbol{p} - \boldsymbol{q}\|^2 \geq -\rho\|\boldsymbol{p} - \boldsymbol{q}\|^2.$$

Thus

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

which proves (10).

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

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

This is (11). □

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

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

$$\langle \boldsymbol{x} - \boldsymbol{y}, \operatorname{prox}_F(\boldsymbol{x}) - \operatorname{prox}_F(\boldsymbol{y}) \rangle \ge 0.$$

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

$$\boldsymbol{x} - \operatorname{prox}_F(\boldsymbol{x}) \in \partial F(\operatorname{prox}_F(\boldsymbol{x})), \qquad \boldsymbol{y} - \operatorname{prox}_F(\boldsymbol{y}) \in \partial F(\operatorname{prox}_F(\boldsymbol{y})).$$

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

$$\big\langle \big(\boldsymbol{x} - \operatorname{prox}_F(\boldsymbol{x})\big) - \big(\boldsymbol{y} - \operatorname{prox}_F(\boldsymbol{y})\big), \operatorname{prox}_F(\boldsymbol{x}) - \operatorname{prox}_F(\boldsymbol{y}) \big\rangle \ge -\rho \| \operatorname{prox}_F(\boldsymbol{x}) - \operatorname{prox}_F(\boldsymbol{y}) \|^2.$$

Rearranging gives

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

Thus prox_F is monotone. □

C Scaled Exact Deviations and Proximal Representation

Lemma C.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 $\widetilde{\Psi}$ 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 $\widetilde{\Psi} := \chi\Psi$ 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 $\phi(\boldsymbol{x}) := \boldsymbol{x} - \nabla\Psi(\boldsymbol{x})$ be a feasible exact-form deviation with $\Psi \in C^{1,1}_{\mathrm{loc}}$ on a neighborhood of K. Then there exists $\alpha_0 > 0$ such that, for every $\alpha \in (0, \alpha_0]$, the interpolated map*

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

coincides on K with $\operatorname{prox}_{F_\alpha}(\boldsymbol{x})$ 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 C.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

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

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

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

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

Corollary 6.3. *On compact convex K, the Euclidean identity-radial closure satisfies $\mathrm{Rad}_I(\Phi_{\mathrm{prox}}) = \Phi^{\mathrm{sm}}_{\mathrm{exact}}$.*

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

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

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

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

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

Thus ϕ is smooth exact-form, proving the reverse inclusion. □

D Boundary-Compatible Finite Interpolation

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

$$\sum_{r=1}^{q} \left\langle \boldsymbol{s}^{j_r}, \boldsymbol{z}^{j_{r+1}} - \boldsymbol{z}^{j_r} \right\rangle \le 0 \quad \textit{for every cycle } j_1, \dots, j_q, j_{q+1} = j_1.$$

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 $w(j,k) = \left\langle \boldsymbol{s}^j, \boldsymbol{z}^k - \boldsymbol{z}^j \right\rangle$. 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

$$a_k \ge a_j + w(j,k) = a_j + \left\langle \boldsymbol{s}^j, \boldsymbol{z}^k - \boldsymbol{z}^j \right\rangle \qquad (j,k \in [m]).$$

Define the finite polyhedral function

$$H(\boldsymbol{z}) \coloneqq \max_{j \in [m]} \left\{ a_j + \left\langle \boldsymbol{s}^j, \boldsymbol{z} - \boldsymbol{z}^j \right\rangle \right\}.$$

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

Theorem D.2. *Let $K \subseteq \mathbb{R}^d$ be nonempty, closed, and convex. Let $\boldsymbol{x}^1, \dots, \boldsymbol{x}^m \in K$ be distinct and choose arbitrary targets $\boldsymbol{y}^1, \dots, \boldsymbol{y}^m \in K$. 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*

$$\mathrm{prox}_{F_\alpha}(\boldsymbol{x}^j) = (1-\alpha)\boldsymbol{x}^j + \alpha \boldsymbol{y}^j, \qquad j = 1, \dots, m.$$

The proximal minimizer is unique.

Proof. Write $\boldsymbol{d}^j \coloneqq \boldsymbol{x}^j - \boldsymbol{y}^j$ and define

$$\boldsymbol{z}^j_\alpha \coloneqq \boldsymbol{x}^j - \alpha \boldsymbol{d}^j, \qquad \boldsymbol{s}^j_\alpha \coloneqq \boldsymbol{x}^j - (1-\rho)\boldsymbol{z}^j_\alpha.$$

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

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

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

Set

$$f_\alpha(\boldsymbol{z}) \coloneqq H_\alpha(\boldsymbol{z}) - \frac{\rho}{2} \left\| \boldsymbol{z} \right\|^2, \qquad F_\alpha \coloneqq f_\alpha + \iota_K.$$

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

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

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

E Other Proofs and Details

E.1 Proofs for Section 4

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

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

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

Remark E.2. *The proof of Theorem 4.1 and Corollary E.1 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.*

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 E.3. *Let $\phi : K \to K$, let $\boldsymbol{D}_\phi(\boldsymbol{x}) \coloneqq \boldsymbol{x} - \phi(\boldsymbol{x})$, and let $\boldsymbol{n} \in N_K(\boldsymbol{x})$. Then*

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

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

$$\langle \boldsymbol{n}, \phi(\boldsymbol{x}) - \boldsymbol{x} \rangle \le 0.$$

Multiplying by -1 gives the statement. □

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

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

Proof. Fix a round t and define

$$\boldsymbol{s}_t \coloneqq \boldsymbol{x}_t - \boldsymbol{x}_{t+1}, \qquad \boldsymbol{n}_t \coloneqq \boldsymbol{x}_t - \eta \boldsymbol{g}_t - \boldsymbol{x}_{t+1}.$$

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

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

Thus,

$$\eta \langle \boldsymbol{g}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle = \langle \boldsymbol{s}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle - \langle \boldsymbol{n}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle . \tag{12}$$

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

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

Therefore

$$\langle \boldsymbol{s}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle \le \Psi(\boldsymbol{x}_t) - \Psi(\boldsymbol{x}_{t+1}) + \frac{L}{2} \|\boldsymbol{s}_t\|^2 . \tag{13}$$

Second, by Lemma E.3 applied at $\boldsymbol{x}_{t+1}$,

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

Thus

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

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

$$\|\boldsymbol{s}_t\| = \|\Pi_K(\boldsymbol{x}_t) - \Pi_K(\boldsymbol{x}_t - \eta \boldsymbol{g}_t)\| \le \eta \|\boldsymbol{g}_t\| . \tag{14}$$

Moreover,

$$\|\boldsymbol{n}_t\| = \operatorname{dist}(\boldsymbol{x}_t - \eta \boldsymbol{g}_t, K) \le \|\boldsymbol{x}_t - \eta \boldsymbol{g}_t - \boldsymbol{x}_t\| = \eta \|\boldsymbol{g}_t\| .$$

Therefore

$$-\langle \boldsymbol{n}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle \le L\eta^2 \|\boldsymbol{g}_t\|^2 . \tag{15}$$

Combining (12), (13), and (15), and using $\|\boldsymbol{s}_t\| \le \eta \|\boldsymbol{g}_t\|$, gives

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

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

Corollary 4.2. *Suppose $\|\boldsymbol{g}_t\| \le G$ and fix $B, L > 0$. Then OGD with $\eta = \sqrt{2B/(3LG^2T)}$ satisfies*

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

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 $\sum_{t=1}^T \|\boldsymbol{g}_t\|^2 \le G^2T$. Hence

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

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. □

E.2 Proofs for Section 5

Theorem 5.1. *Let $F := f + \iota_K$ be proper, lower semicontinuous, and ρ-weakly convex with $0 \le \rho < 1$. Define $\phi_F(\boldsymbol{x}) := \operatorname{prox}_F(\boldsymbol{x})$. Then $\phi_F : K \to K$ is feasible and Euclidean exact-form, with potential M_F, i.e., $\boldsymbol{x} - \phi_F(\boldsymbol{x}) = \nabla M_F(\boldsymbol{x})$, 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, $\|\boldsymbol{g}_t\| \le G$, $\operatorname{osc}_K(M_F) \le B_F$, and $B_F, G > 0$, then choosing $\eta = \sqrt{2B_F/(3L_F G^2 T)}$ gives*

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

If $B_F = 0$ or $G = 0$, the corresponding regret bound is trivial and the bound follows without optimizing the step size.

Proof. Feasibility holds because $F = f + \iota_K$ forces $\operatorname{prox}_F(\boldsymbol{x}) \in K$. Exactness is exactly (9). The regret bound follows by applying Theorem 4.1 with $\Psi = M_F$ and the smoothness bound from Lemma B.3. □

Proposition E.4. *The exact-form class contains the following deviations.*

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

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

   The case $\alpha = 1$ is external regret.
5. ***Local proximal deviations.*** *If $\|\boldsymbol{x} - \operatorname{prox}_F(\boldsymbol{x})\| \le \delta$ on K, then the deviation belongs to the local exact-form subclass $\|\nabla M_F(\boldsymbol{x})\| \le \delta$.*
6. ***Symmetric affine swaps.*** *Let $\phi(\boldsymbol{x}) := A\boldsymbol{x} + \boldsymbol{b}$ map K into K. Let $T := \operatorname{span}(K - K)$ be the direction space of $\operatorname{aff}(K)$, fix any $\boldsymbol{a} \in K$, 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 $\operatorname{aff}(K)$, a potential is*

   $$\Psi(\boldsymbol{a} + \boldsymbol{h}) := \langle \boldsymbol{D}_\phi(\boldsymbol{a}), \boldsymbol{h} \rangle + \frac{1}{2} \langle \boldsymbol{h}, (I - A)\boldsymbol{h} \rangle, \qquad \boldsymbol{h} \in T. \tag{16}$$

   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:

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

For item 4, the unconstrained first-order condition for

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

gives $\boldsymbol{y} = (1-\alpha)\boldsymbol{x} + \alpha\boldsymbol{u}$, 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 $\boldsymbol{x} = \boldsymbol{a} + \boldsymbol{h}$ with

$\boldsymbol{h} \in T$. Since $\boldsymbol{a} \in K$ and $\phi(\boldsymbol{a}) \in K$, the vector $\boldsymbol{D}_\phi(\boldsymbol{a}) = \boldsymbol{a} - \phi(\boldsymbol{a})$ lies in T. For every tangent direction $\boldsymbol{r} \in T$,

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

Differentiating (16) in direction $\boldsymbol{r}$,

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

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

$$\widetilde{\Psi}(\boldsymbol{z}) \coloneqq \langle \boldsymbol{D}_\phi(\boldsymbol{a}), P_T(\boldsymbol{z} - \boldsymbol{a}) \rangle + \frac{1}{2} \langle P_T(\boldsymbol{z} - \boldsymbol{a}), (I - A)P_T(\boldsymbol{z} - \boldsymbol{a}) \rangle .$$

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

E.3 Strict Map-Level Separation

Proof of Proposition 6.1. Let $k \geq 2$ be an even integer and let $a \in (0, 1]$ satisfy $a\pi k/2 > 1$. Define

$$\boldsymbol{D}_\phi(x) \coloneqq ax(1 - x)\sin(2\pi kx), \qquad \phi(x) \coloneqq x - \boldsymbol{D}_\phi(x).$$

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

$$\Psi(x) \coloneqq \int_0^x as(1 - s)\sin(2\pi ks)\,\mathrm{d}s,$$

so ϕ is feasible exact-form.

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

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

Thus ϕ is not monotone. By Lemma B.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. □

E.4 Proofs for Section 7

Lemma E.5. *If $\boldsymbol{D} = \nabla\Psi$ on a neighborhood of K, then we have $\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^\mathsf{T}\mathrm{d}\boldsymbol{x} = 0$ for every closed loop $\gamma \subseteq K$.*

Proof. By the chain rule,

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

□

Lemma E.6. *Let $\boldsymbol{z}_0, \boldsymbol{z}_1, \ldots, \boldsymbol{z}_m = \boldsymbol{z}_0$ be a closed polygon in K. Suppose*

$$C \coloneqq \sum_{i=0}^{m-1} \langle \boldsymbol{D}(\boldsymbol{z}_i), \boldsymbol{z}_i - \boldsymbol{z}_{i+1} \rangle > 0. \tag{17}$$

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

$$\boldsymbol{g}_{i+1} \coloneqq \frac{\boldsymbol{z}_i - \boldsymbol{z}_{i+1}}{\eta}, \qquad i = 0, \ldots, m - 1. \tag{18}$$

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

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

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

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

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

□

Theorem 7.2. *Let $\boldsymbol{D} : K \to \mathbb{R}^d$ be continuous. Suppose there exists a closed piecewise C^1 loop $\gamma : [0,1] \to K$ such that, after possibly reversing orientation,*

$$-\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^\mathsf{T} \mathrm{d}\boldsymbol{x} \coloneqq c_0 > 0.$$

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

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

for a constant $c > 0$ depending only on $\boldsymbol{D}$, 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

$$-\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^\mathsf{T} \mathrm{d}\boldsymbol{x}$$

is the limit of left-endpoint Riemann sums

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

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

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

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

$$\boldsymbol{g}_{i+1} \coloneqq \frac{\boldsymbol{z}_i - \boldsymbol{z}_{i+1}}{\eta}, \qquad i = 0, \dots, m-1.$$

Then $\|\boldsymbol{g}_{i+1}\| \le G$, and Lemma E.6 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

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

which proves the claim. □

Corollary E.7 (Euclidean Curl Criterion)**.** *Let $E_K := \mathrm{span}(K - K)$, let $U := \mathrm{relint}(K)$ be open and simply connected in $\mathrm{aff}(K)$, and let $\boldsymbol{D} : U \to E_K$ be C^1. In any orthonormal affine coordinate system on E_K, $\partial_j \boldsymbol{D}_i = \partial_i \boldsymbol{D}_j$ throughout U if and only if $\boldsymbol{D} = \nabla_{E_K} \Psi$ 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 $\boldsymbol{D}$ along a sufficiently small feasible loop. When the field is curl-free, assume that Ψ admits an ambient extension $\widetilde{\Psi}$, defined on a neighborhood of K, that satisfies the smoothness and bounded-oscillation assumptions of Theorem 4.1 and Corollary E.1. If $\widetilde{\phi}(\boldsymbol{x}) := \boldsymbol{x} - \nabla\widetilde{\Psi}(\boldsymbol{x})$ lies in K for every $\boldsymbol{x} \in K$, then projected GD has $O(\sqrt{T})$ regret against $\widetilde{\phi}$.*

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 7.2 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 E.1. □

Remark E.8. *Theorem 7.2 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.*

E.5 Proofs for Section 8

Proposition E.9. *A field $\boldsymbol{D} : \mathcal{D} \to E$ is R-mirror-exact if and only if there exists a differentiable scalar potential $V : \mathcal{D} \to \mathbb{R}$ such that*

$$\nabla V(\boldsymbol{x}) = \nabla^2 R(\boldsymbol{x}) \boldsymbol{D}(\boldsymbol{x}). \tag{19}$$

Equivalently, $\boldsymbol{D}(\boldsymbol{x})^{\mathsf{T}} \mathrm{d}\nabla R(\boldsymbol{x}) = \mathrm{d}V(\boldsymbol{x})$.

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

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

By the chain rule,

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

Conversely, suppose (19) holds. For $\boldsymbol{\theta} \in \Theta$, write $\boldsymbol{x} := (\nabla R)^{-1}(\boldsymbol{\theta})$ and define

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

Another application of the chain rule gives

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

which is the desired dual-gradient representation. □

E.6 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

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

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.

Fixed Comparators and External Regret Fix $\boldsymbol{u} \in K$ and set

$$\phi_{\boldsymbol{u}}(\boldsymbol{x}) \coloneqq \boldsymbol{u}, \qquad \boldsymbol{D}_{\boldsymbol{u}}(\boldsymbol{x}) \coloneqq \boldsymbol{x} - \boldsymbol{u}.$$

Take the dual potential

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

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

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

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

Dual Translations Let $\boldsymbol{a} \in E$. Assume that $\boldsymbol{\theta} - \boldsymbol{a} \in \Theta$ for every $\boldsymbol{\theta} \in \Theta$, and define

$$\phi_{\boldsymbol{a}}(\boldsymbol{x}) \coloneqq \nabla R^*(\nabla R(\boldsymbol{x}) - \boldsymbol{a}). \tag{20}$$

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

$$\Psi_{\boldsymbol{a}}(\boldsymbol{\theta}) \coloneqq R^*(\boldsymbol{\theta}) - R^*(\boldsymbol{\theta} - \boldsymbol{a}),$$

since

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

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.

Mirror Interpolation Toward a Comparator Fix $\boldsymbol{u} \in \mathcal{D}$, write $\boldsymbol{\theta}_{\boldsymbol{u}} \coloneqq \nabla R(\boldsymbol{u})$, and choose $\lambda \in [0, 1)$. Assume that

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

Define

$$\phi^R_{\lambda, \boldsymbol{u}}(\boldsymbol{x}) \coloneqq \nabla R^*\big((1 - \lambda)\nabla R(\boldsymbol{x}) + \lambda \nabla R(\boldsymbol{u})\big). \tag{21}$$

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

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

works because

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

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

Entropy on the Simplex and Multiplicative Deviations Let $K \coloneqq \Delta_d$ and work on $\operatorname{relint}(\Delta_d)$. Take the entropic mirror map

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

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, $\nabla R^*(\boldsymbol{\theta}) = \operatorname{softmax}(\boldsymbol{\theta})$. Hence the dual translation (20) becomes

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

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

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

whose gradient is

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

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

The dual interpolation (21) gives another familiar entropy example. For $\boldsymbol{u} \in \operatorname{relint}(\Delta_d)$,

$$\phi^{R}_{\lambda,\boldsymbol{u}}(\boldsymbol{x})[i] \coloneqq \frac{\boldsymbol{x}[i]^{1-\lambda}\boldsymbol{u}[i]^{\lambda}}{\sum_{j=1}^{d} \boldsymbol{x}[j]^{1-\lambda}\boldsymbol{u}[j]^{\lambda}}.$$

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

Metric-Symmetric Affine Swaps Let $K \coloneqq E \coloneqq \mathbb{R}^d$ and $R_H(\boldsymbol{x}) \coloneqq \frac{1}{2}\boldsymbol{x}^{\mathsf{T}} H \boldsymbol{x}$, where $H \succ 0$. This is a Legendre function with full domain. Then $\nabla R_H(\boldsymbol{x}) = H\boldsymbol{x}$, and Proposition E.9 says that $\boldsymbol{D}$ is R_H-mirror-exact exactly when $H\boldsymbol{D}(\boldsymbol{x})$ is an ordinary gradient field.

Consider an affine deviation

$$\phi(\boldsymbol{x}) \coloneqq A\boldsymbol{x} + \boldsymbol{b}, \qquad \boldsymbol{D}(\boldsymbol{x}) \coloneqq \boldsymbol{x} - \phi(\boldsymbol{x}) = (I - A)\boldsymbol{x} - \boldsymbol{b}.$$

If

$$H(I - A) = (I - A)^{\mathsf{T}} H, \tag{22}$$

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

$$V(\boldsymbol{x}) \coloneqq \frac{1}{2}\boldsymbol{x}^{\mathsf{T}} H(I - A)\boldsymbol{x} - \langle H\boldsymbol{b}, \boldsymbol{x}\rangle$$

satisfies $\nabla V(\boldsymbol{x}) = H\boldsymbol{D}(\boldsymbol{x})$. For $H = I$, condition (22) 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.

A Mirror-Exact Field That Is Not Euclidean-Exact The previous example gives a feasible separation on an unconstrained action space. Let

$$H \coloneqq \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \qquad K \coloneqq \mathbb{R}^2,$$

use $R_H(\boldsymbol{x}) \coloneqq \frac{1}{2}\boldsymbol{x}^{\mathsf{T}} H \boldsymbol{x}$, and for $\varepsilon \geq 0$ define

$$\boldsymbol{D}(\boldsymbol{x})[1] \coloneqq \varepsilon(\boldsymbol{x}[1] + \boldsymbol{x}[2]), \qquad \boldsymbol{D}(\boldsymbol{x})[2] \coloneqq \frac{\varepsilon}{2}(\boldsymbol{x}[1] + \boldsymbol{x}[2]), \qquad \phi(\boldsymbol{x}) \coloneqq \boldsymbol{x} - \boldsymbol{D}(\boldsymbol{x}).$$

The map is feasible because it maps $\mathbb{R}^2$ to itself. Writing $S \coloneqq \left(\begin{smallmatrix} 1 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, it is R_H-mirror-exact because

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

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

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

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

E.7 Proofs for Section 8.2

Theorem 8.3. *Assume the standing Legendre-domain conditions and that R is m-strongly convex on $\mathcal{D}$. Let $\phi : \mathcal{D} \to K$ be feasible and R-mirror-exact with dual potential Ψ. Let $\{\boldsymbol{x}_t\}_{t=1}^{T+1}$ be the iterates generated by online mirror descent (3), initialized in $\mathcal{D}$, with iterates remaining in $\mathcal{D}$. Write*

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

Assume that (4) 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

$$\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle \le \frac{\Psi(\boldsymbol{\theta}_1) - \Psi(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{L\eta}{m} \sum_{t=1}^{T} \|\boldsymbol{g}_t\|^2 . \tag{5}$$

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

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

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 $\operatorname{osc}_\Theta(\Psi) \le B$ is a sufficient trajectory-independent assumption that avoids any dependence of the oscillation bound on the chosen step size.

Proof. Because $\boldsymbol{x}_{t+1} \in \mathcal{D} = \operatorname{relint}(K)$ and normal cones are taken in the affine hull of K, the relative normal cone at $\boldsymbol{x}_{t+1}$ is $\{0\}$. The first-order optimality condition for (3) therefore reduces to the unconstrained dual update

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

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

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

By Legendre conjugacy,

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

Therefore

$$\eta \langle \boldsymbol{g}_t, \boldsymbol{D}_\phi(\boldsymbol{x}_t) \rangle \le \Psi(\boldsymbol{\theta}_t) - \Psi(\boldsymbol{\theta}_{t+1}) + L D_R(\boldsymbol{x}_t, \boldsymbol{x}_{t+1}). \tag{23}$$

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

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

Strong convexity of R therefore gives

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

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

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

Combining this estimate with (23) gives the one-step bound

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

Dividing by η and summing over t proves (5). The $O(\sqrt{T})$ statement follows by bounding the telescoping term by B/η and optimizing η. □

E.8 Proofs for Section 8.3

Lemma E.10. *For fixed $\boldsymbol{y} \in \operatorname{dom} R$,*

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

Proof. By definition,

$$D_R(\boldsymbol{y}, \boldsymbol{x}) \coloneqq R(\boldsymbol{y}) - R(\boldsymbol{x}) - \langle \nabla R(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle .$$

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

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

Then

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

Thus

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

□

Lemma E.11 (Bregman-Moreau Identity)**.** *Assume the standing Legendre-domain conditions. Let $F \coloneqq f + \iota_K$ 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*

$$\nabla M_F^R(\boldsymbol{x}) = \nabla^2 R(\boldsymbol{x}) \big(\boldsymbol{x} - \operatorname{prox}_F^R(\boldsymbol{x})\big). \tag{24}$$

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

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

then

$$\nabla_{\boldsymbol{\theta}} \Psi_F(\boldsymbol{\theta}) = \boldsymbol{x} - \operatorname{prox}_F^R(\boldsymbol{x}), \qquad \boldsymbol{x} = (\nabla R)^{-1}(\boldsymbol{\theta}). \tag{25}$$

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

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

Using Lemma E.10, we obtain

$$\nabla M_F^R(\boldsymbol{x}) = \nabla^2 R(\boldsymbol{x}) \big(\boldsymbol{x} - \operatorname{prox}_F^R(\boldsymbol{x})\big),$$

which proves (24).

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

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

the chain rule gives

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

□

Theorem 8.4. *Assume the standing Legendre-domain conditions. Let $F \coloneqq f + \iota_K$ 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 $\phi_F^R : \mathcal{D} \to K$ defined by $\phi_F^R(\boldsymbol{x}) \coloneqq \operatorname{prox}_F^R(\boldsymbol{x})$ is feasible and R-mirror-exact. Equivalently,*

$$\big(\boldsymbol{x} - \operatorname{prox}_F^R(\boldsymbol{x})\big)^{\mathsf{T}} \mathrm{d}\nabla R(\boldsymbol{x}) = \mathrm{d}M_F^R(\boldsymbol{x}). \tag{6}$$

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 (25) in Lemma E.11, the displacement

$$\boldsymbol{x} - \phi_F^R(\boldsymbol{x}) = \boldsymbol{x} - \operatorname{prox}_F^R(\boldsymbol{x})$$

equals the dual gradient $\nabla_{\boldsymbol{\theta}} \Psi_F(\nabla R(\boldsymbol{x}))$. This is precisely R-mirror-exactness. The differential form (6) is the same identity written in primal coordinates using (24) in Lemma E.11. □

Proposition 8.5. *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 $0 \le \rho < m$ 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 $\{\boldsymbol{x}_t\}_{t=1}^{T+1} \subset \mathcal{D}$ be generated by online mirror descent (3). Then these iterates control Bregman proximal regret with the bound*

$$\sum_{t=1}^{T} \left\langle \boldsymbol{g}_t, \boldsymbol{x}_t - \operatorname{prox}_F^R(\boldsymbol{x}_t) \right\rangle \le \frac{\Psi_F(\boldsymbol{\theta}_1) - \Psi_F(\boldsymbol{\theta}_{T+1})}{\eta} + \frac{\eta}{m} \sum_{t=1}^{T} \|\boldsymbol{g}_t\|^2 , \tag{7}$$

where $\Psi_F(\boldsymbol{\theta}_t) = M_F^R(\boldsymbol{x}_t)$ 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 $\boldsymbol{x}$-dependent constant from

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

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

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

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,

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

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

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

which proves the relative-smoothness with $L = 1$ for the potential function Ψ_F. Finally, Theorem 8.4 identifies the Bregman proximal displacement with $\nabla_{\boldsymbol{\theta}} \Psi_F(\boldsymbol{\theta}_t)$. Applying Theorem 8.3 with $L = 1$ gives (7). □

E.9 Mirror Circulation in the Legendre Interior

Let $\boldsymbol{D} : \mathcal{D} \to E$ be continuous, and let $\gamma : [0,1] \to \mathcal{D}$ be a closed piecewise C^1 loop. Its mirror circulation is

$$\oint_\gamma \boldsymbol{D}(\boldsymbol{x})^{\mathsf{T}} \mathrm{d}\nabla R(\boldsymbol{x}) = \int_0^1 \left\langle \boldsymbol{D}(\gamma(s)), \frac{\mathrm{d}}{\mathrm{d}s} \nabla R(\gamma(s)) \right\rangle \mathrm{d}s.$$

Every mirror-exact field has zero mirror circulation by the chain rule. The converse lower bound follows by applying the Euclidean loop argument in dual coordinates.

Theorem E.12. *Assume the standing Legendre-domain conditions. Suppose that, after reversing the orientation of γ if necessary,*

$$-\oint_{\gamma} \boldsymbol{D}(\boldsymbol{x})^{\mathsf{T}} \mathrm{d}\nabla R(\boldsymbol{x}) = c_0 > 0.$$

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

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

for a constant $c > 0$ depending only on $\boldsymbol{D}$, γ, R, and G. All iterates remain in $\mathcal{D}$.

Proof. Set $\widetilde{\boldsymbol{D}}(\boldsymbol{\theta}) = \boldsymbol{D}(\nabla R^*(\boldsymbol{\theta}))$ and $\widetilde{\gamma} = \nabla R \circ \gamma$. The change of coordinates gives

$$\oint_{\widetilde{\gamma}} \widetilde{\boldsymbol{D}}(\boldsymbol{\theta})^{\mathsf{T}} \mathrm{d}\boldsymbol{\theta} = \oint_{\gamma} \boldsymbol{D}(\boldsymbol{x})^{\mathsf{T}} \mathrm{d}\nabla R(\boldsymbol{x}).$$

The dual domain Θ is open and convex, so the polygonal loop construction in the proof of Theorem 7.2 can be carried out inside Θ. For successive dual points $\boldsymbol{\theta}_t, \boldsymbol{\theta}_{t+1}$, choose $\boldsymbol{g}_t = (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t+1})/\eta$. The subdivision ensures $\|\boldsymbol{g}_t\| \le G$. The mirror-descent optimality condition is then satisfied at $\boldsymbol{x}_{t+1} = \nabla R^*(\boldsymbol{\theta}_{t+1}) \in \mathcal{D}$, which is the unique minimizer by strict convexity. The regret is exactly

$$\sum_{t=1}^{T} \langle \boldsymbol{g}_t, \boldsymbol{D}(\boldsymbol{x}_t) \rangle = \sum_{t=1}^{T} \left\langle \boldsymbol{g}_t, \widetilde{\boldsymbol{D}}(\boldsymbol{\theta}_t) \right\rangle.$$

Applying the loop estimate from Theorem 7.2 in these dual coordinates proves the claim. □

E.10 Proofs for Section 9

Theorem 9.2. *Let $\bar{\mu}_T := T^{-1} \sum_{t=1}^{T} \delta_{\boldsymbol{x}_t}$ and define $\boldsymbol{g}_{i,t} := \nabla_{\boldsymbol{x}_i} \ell_i(\boldsymbol{x}_{i,t}, \boldsymbol{x}_{-i,t})$. If every player i guarantees*

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

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 $\Phi_{i,\mathrm{exact}}(B_i, L_i)$ with $\|\boldsymbol{g}_{i,t}\| \le G_i$ gives $\varepsilon_T \le \max_i G_i \sqrt{6 L_i B_i / T}$.

Proof. Fix player i and $\phi_i \in \Phi_i$. By convexity of ℓ_i in its own action,

$$\begin{aligned} &\ell_i(\boldsymbol{x}_{i,t}, \boldsymbol{x}_{-i,t}) - \ell_i(\phi_i(\boldsymbol{x}_{i,t}), \boldsymbol{x}_{-i,t}) \\ &\quad \le \langle \nabla_{\boldsymbol{x}_i} \ell_i(\boldsymbol{x}_{i,t}, \boldsymbol{x}_{-i,t}), \boldsymbol{x}_{i,t} - \phi_i(\boldsymbol{x}_{i,t}) \rangle. \end{aligned}$$

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 4.2 player by player. □

Proposition 9.3. *Let $\mathcal{A}_i$ be any family of feasible Borel-measurable deviations for player i. A distribution satisfies $\Gamma_i(\mu, \phi_i) \le 0$ for every $\phi_i \in \mathcal{A}_i$ and every player if and only if it satisfies the same inequalities for every $\phi_i \in \mathrm{Rad}_I(\mathcal{A}_i)$.*

Proof. Since $\mathcal{A}_i \subseteq \mathrm{Rad}_I(\mathcal{A}_i)$ by taking $\alpha = 1$, one implication is immediate. For the converse, suppose that μ satisfies every constraint in $\mathcal{A}_i$, and let $\phi_i \in \mathrm{Rad}_I(\mathcal{A}_i)$. 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

$$\begin{aligned} &\ell_i(\boldsymbol{x}_i, \boldsymbol{x}_{-i}) - \ell_i(\phi_{i,\alpha}(\boldsymbol{x}_i), \boldsymbol{x}_{-i}) \\ &\quad \ge \alpha \left(\ell_i(\boldsymbol{x}_i, \boldsymbol{x}_{-i}) - \ell_i(\phi_i(\boldsymbol{x}_i), \boldsymbol{x}_{-i}) \right). \end{aligned}$$

The joint continuity and compactness assumptions ensure that both loss differences are bounded and measurable. Thus both gains are integrable. Taking expectations gives $\Gamma_i(\mu, \phi_{i,\alpha}) \ge \alpha \Gamma_i(\mu, \phi_i)$. The left-hand side is nonpositive, so $\Gamma_i(\mu, \phi_i) \le 0$. □

Theorem 9.4. *Under the unrestricted conventions above, every convex game with compact convex action sets and jointly continuous losses satisfies*

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

Proof. By Corollary 6.3, the smooth exact-form class is the identity-radial closure of the unrestricted proximal class. Apply Proposition 9.3 player by player. □

Corollary 9.5. *For every finitely supported distribution in a convex game,*

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

In particular the three notions coincide in every finite game.

Proof. Only $\mathrm{PCE} \subseteq \mathrm{CE}$ needs proof. If a general deviation τ_i has positive gain, list the finitely many recommendations $\boldsymbol{x}^1, \dots, \boldsymbol{x}^m$ to player i and set $\boldsymbol{y}^j := \tau_i(\boldsymbol{x}^j)$. By Theorem D.2, for sufficiently small common $\alpha > 0$ there is a proximal map satisfying

$$\phi_{i,\alpha}(\boldsymbol{x}^j) = (1-\alpha)\boldsymbol{x}^j + \alpha\tau_i(\boldsymbol{x}^j)$$

at every supported recommendation. Convexity gives $\Gamma_i(\mu, \phi_{i,\alpha}) \geq \alpha\Gamma_i(\mu, \tau_i) > 0$, so $\mu \notin \mathrm{PCE}$. □

Theorem 9.6. *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 $\ell_i : K \to \mathbb{R}$ is continuous on $K := \prod_{j=1}^N K_j$. Then*

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

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

Proof. The inclusion $\mathrm{CE} \subseteq \mathrm{ConCE}$ follows directly from the deviation classes, and $\mathrm{ConCE} = \mathrm{PCE}$ follows from Theorem 9.4. It remains to prove $\mathrm{ConCE} \subseteq \mathrm{CE}$.

Suppose $\mu \notin \mathrm{CE}$. Then for some player i, there is a Borel-measurable deviation $\tau_i : K_i \to K_i$ 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 $h_n : K_i \to K_i$ 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 $\nu_i\big(|\phi_{i,n} - \tau_i| > \frac{1}{n}\big) \leq \nu_i(K_i \setminus A_n) < \frac{1}{n}$. So $\phi_{i,n}(X_i) \to \tau_i(X_i)$ in probability under μ.

Continuity on compact K makes ℓ_i bounded and uniformly continuous. So we get $\ell_i(\phi_{i,n}(X_i), X_{-i}) \longrightarrow \ell_i(\tau_i(X_i), X_{-i})$ in probability. The variables are uniformly bounded so the convergence is in $L^1(\mu)$. This shows that $\Gamma_i(\mu, \phi_{i,n}) \longrightarrow \Gamma_i(\mu, \tau_i) > 0$. 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

$$\Psi_{i,n}(x) = \int_{a_i}^{x} \big(s - \phi_{i,n}(s)\big)\,\mathrm{d}s$$

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 $\mu \notin \mathrm{ConCE}$. This proves $\mathrm{ConCE} \subseteq \mathrm{CE}$ and completes the equality. □

Theorem 9.7. *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*

$$\mu \in \mathrm{ConCE} = \mathrm{PCE} \qquad \textit{but} \qquad \mu \notin \mathrm{CE}.$$

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

Proof. Let $K \subseteq \mathbb{R}^2$ be the unit disk, and let

$$J \coloneqq \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

be the counterclockwise 90° rotation matrix. Player 1 has loss $\ell_1(\boldsymbol{a}, \boldsymbol{s}) = \langle J\boldsymbol{s}, \boldsymbol{a} \rangle$, 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

$$\Gamma_1(\mu, \phi) = \frac{1}{\pi} \int_K \langle J\boldsymbol{x}, \nabla\Psi(\boldsymbol{x}) \rangle \, \mathrm{d}\boldsymbol{x} = 0,$$

because $\mathrm{div}(J\boldsymbol{x}) = 0$ and $J\boldsymbol{x}$ has zero normal flux. Thus $\mu \in \mathrm{ConCE} = \mathrm{PCE}$. The feasible measurable map $\tau(\boldsymbol{x}) \coloneqq -J\boldsymbol{x}/\|\boldsymbol{x}\|$ for $\boldsymbol{x} \neq 0$, with $\tau(0) \coloneqq 0$, has gain $\mathbb{E}\|X\| = 2/3$, so $\mu \notin \mathrm{CE}$. □

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Complete proofs are provided.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Concretely in the main text.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: Complete proofs and clear assumptions are provided.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: [N/A]

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [N/A]

Justification: [N/A]

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [N/A]

   Justification: [N/A]

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [N/A]

   Justification: [N/A]

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: [N/A]

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: This is primarily a theoretical study.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

    Justification: This is primarily a theoretical study.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [No]

    Justification: This is primarily a theoretical study.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [No]

    Justification: This is primarily a theoretical study.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: [N/A]

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This is primarily a theoretical study.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This is primarily a theoretical study.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: [N/A]

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.