

Rate-Preserving Reductions for Blackwell Approachability

Christoph Dann

Yishay Mansour

Mehryar Mohri

Jon Schneider

Balasubramanian Sivan

Google Research

CDANN@CDANN.NET

MANSOUR.YISHAY@GMAIL.COM

MOHRI@GOOGLE.COM

JSCHNEI@GOOGLE.COM

BALUSIVAN@GOOGLE.COM

Editors: Nika Haghtalab and Ankur Moitra

Abstract

Abernethy et al. (2011) showed that Blackwell approachability and no-regret learning are equivalent, in the sense that any algorithm that solves a specific Blackwell approachability instance can be converted to a sublinear regret algorithm for a specific no-regret learning instance, and vice versa. In this paper, we study a more fine-grained form of such reductions, and ask when this translation between problems preserves not only a sublinear rate of convergence, but also preserves the optimal rate of convergence. That is, in which cases does it suffice to find the optimal regret bound for a no-regret learning instance in order to find the optimal rate of convergence for a corresponding approachability instance?

We show that the reduction of Abernethy et al. (2011) (and of other subsequent work) does not preserve rates: their reduction may reduce a d -dimensional approachability instance \mathcal{I}_1 with optimal convergence rate R_1 to a no-regret learning instance \mathcal{I}_2 with optimal regret-per-round of R_2 , with R_2/R_1 arbitrarily large (in particular, it is possible that $R_1 = 0$ and $R_2 > 0$). On the other hand, we show that it is possible to tightly reduce any approachability instance to an instance of a generalized form of regret minimization we call *improper ϕ -regret minimization* (a variant of the ϕ -regret minimization of Gordon et al. (2008) where the transformation functions may map actions outside of the action set).

Finally, we characterize when linear transformations suffice to reduce improper ϕ -regret minimization problems to standard classes of regret minimization problems (such as external regret minimization and proper ϕ -regret minimization) in a rate preserving manner. We prove that some improper ϕ -regret minimization instances cannot be reduced to either subclass of instance in this way, suggesting that approachability can capture some problems that cannot be easily phrased in the standard language of online learning.

Keywords: Blackwell approachability, regret minimization, swap regret

1. Introduction

Blackwell’s Approachability Theorem is a fundamental result in game theory with far-reaching applications in machine learning, economics, and optimization. It provides a framework for analyzing repeated games where the payoff is a vector rather than a single scalar value. In essence, the theorem allows players to determine whether a specific set of payoff vectors can be “approached” on average over time, even if achieving them individually in a single round is impossible (with more sophisticated variants of this theorem characterizing the rate at which this set can be approached). This concept has been instrumental in developing robust algorithms for online learning, strategic decision-making in economic models, and solving complex optimization problems where the objective is multi-dimensional.

A perhaps even more prevalent paradigm in the area of online learning is that of regret minimization. In the most fundamental form of the problem (external regret minimization), a learner wants to take a sequence of actions that performs at least as well as the optimal static action they could have taken in hindsight. One of the central results of the field of online learning is that there exist learning algorithms which achieve regret that is sublinear in the time horizon T . Understanding how to optimize these regret bounds is a topic of active research, and optimal regret bounds have been established for a variety of problems in this area.

[Abernethy et al. \(2011\)](#) showed that these two problems – Blackwell approachability and regret minimization – are very closely linked, and in fact are “equivalent” in the following sense: given an instance of Blackwell approachability (a vector-valued payoff function and a set to approach), it is possible to reduce it to an instance of regret minimization (specifically, an instance of online linear optimization) such that any sublinear regret algorithm for the regret minimization instance can be used to solve the corresponding Blackwell approachability instance, giving an algorithmic proof of convergence. Conversely, any regret minimization instance can itself be viewed as an instance of Blackwell approachability, so any algorithmic approach to generic Blackwell approachability instances can be applied to solve regret minimization.

In this paper, we explore whether such an equivalence holds at a more fine-grained level, when we care about the specific convergence rates for Blackwell approachability and the specific regret bounds for regret minimization. In other words, imagine we have a problem that we can cast as a specific instance of Blackwell approachability, and we want to understand the optimal convergence rates possible for this instance (this is a common desideratum for many of the aforementioned applications of approachability). Is it sufficient to just understand the optimal regret bounds for the corresponding instance of regret minimization under the reduction of [Abernethy et al. \(2011\)](#)?

1.1. Our results

We begin by answering this question in the negative: there is no direct correspondence between the optimal rate achievable in an approachability instance \mathcal{I} and the optimal regret bound achievable in the corresponding regret minimization instance \mathcal{I}' . More precisely, the algorithmic construction of [Abernethy et al. \(2011\)](#) shows that any upper bound on the regret of \mathcal{I}' translates to an upper bound on the optimal convergence rate of \mathcal{I} . We prove that this translation is lossy, and that there is no analogous translation between lower bounds. In particular, in Theorem 3 we exhibit instances of approachability \mathcal{I} that are *perfectly approachable* (i.e., where the learner can guarantee that the average payoff exactly lies within the set S) but where the corresponding regret-minimization instance \mathcal{I}' has a non-trivial regret lower bound. These same examples provide similar gaps in rates for all subsequent work on approachability that we are aware of, almost all of which share the same core idea as in [Abernethy et al. \(2011\)](#) (see Appendix D for a detailed discussion).

On the other hand, we complement this result by showing that Blackwell approachability can be tightly reduced (in a rate-preserving manner) to a novel variant of regret minimization that we call *improper ϕ -regret minimization* (Theorem 4). To explain what this means, it is helpful to briefly define regret minimization and its relevant variants (we defer formal definitions to Section 2).

We start with the basic setting of online linear optimization. In this problem, a learner must select an action p_t from a convex action set $\mathcal{P} \subseteq \mathbb{R}^d$ each round t for a total of T rounds. At the same time, an adversary selects a loss ℓ_t from a convex loss set $\mathcal{L} \subseteq \mathbb{R}^d$. The learner receives loss $\langle p_t, \ell_t \rangle$ during that round, and would like to minimize their total *regret* over all rounds. In its most

standard form (external regret), this is just the largest gap between their total loss and the total loss of the best fixed action, and can be written as

$$\text{Reg}(\mathbf{p}, \ell) = \max_{p^* \in \mathcal{P}} \left(\sum_{t=1}^T \langle p_t, \ell_t \rangle - \sum_{t=1}^T \langle p^*, \ell_t \rangle \right).$$

Some applications call for obtaining low regret not just with respect to the best fixed action in hindsight, but additionally with respect to transformations of the sequence of actions played by the learner. The most well-known such notion of regret is probably that of *swap regret*, which famously has the property that sublinear swap regret learning algorithms converge to correlated equilibria when used to play normal-form games (Foster and Vohra, 1997; Blum and Mansour, 2007). However, swap regret is just one of a large class of such regret metrics that are succinctly captured by the notion of (linear) ϕ -regret introduced in (Gordon et al., 2008). In ϕ -regret minimization, we are given a collection Φ of linear “transformation” functions ϕ sending \mathcal{P} to \mathcal{P} . The ϕ -regret for this class Φ is given by:

$$\text{Reg}_\Phi(\mathbf{p}, \ell) = \max_{\phi \in \Phi} \left(\sum_{t=1}^T \langle p_t, \ell_t \rangle - \sum_{t=1}^T \langle \phi(p_t), \ell_t \rangle \right).$$

Note that by setting Φ to be the set of all constant functions on \mathcal{P} , we immediately recover the original external regret definition. The case of swap regret corresponds to the setting where \mathcal{P} is the d -simplex, and Φ contains all row-stochastic linear maps. Interestingly, in (Gordon et al., 2008), the authors show that any ϕ -regret minimization problem has a corresponding sublinear regret learning algorithm (this can also be seen from Blackwell approachability).

However, the constraint that every linear function ϕ in Φ maps \mathcal{P} to itself turns out not to be strictly necessary – there exist classes Φ of linear functions that map points in \mathcal{P} to arbitrary points in \mathbb{R}^d (including possibly outside of \mathcal{P}) for which it is still possible to obtain sublinear ϕ -regret. We call any such instance an *improper ϕ -regret minimization* instance (and contrast it with the previous constrained definition by calling those instances *proper ϕ -regret minimization*). In Theorem 4, we show that for any Blackwell approachability instance \mathcal{I} , there exists an improper ϕ -regret minimization instance \mathcal{I}' , such that if you have an algorithm that solves \mathcal{I}' with at most R ϕ -regret, you can use it to get a convergence rate of R/T in the approachability instance \mathcal{I} , and vice versa.

The introduction of improper ϕ -regret raises a natural question: is improper ϕ -regret minimization truly a more general setting than proper ϕ -regret minimization? Or can we reduce – in a rate-preserving manner – any improper ϕ -regret minimization problem to a proper ϕ -regret minimization problem (or even further, to an external regret minimization problem, as the original reduction of Abernethy et al. (2011) partially accomplishes).

Since arbitrary reductions between instances of regret-minimization can be ill-behaved, we study the above questions in the setting of reductions that are entirely specified by linear transformations, which we call *linear equivalences*. It turns out that there are interesting classes of regret-minimization problems that superficially appear to be improper ϕ -regret minimization instances, but that can be shown to be linearly equivalent to proper ϕ -regret minimization or external regret minimization problems. One interesting such class is the class of *weighted regret minimization* problems, where the regret corresponding to the transformation function ϕ is weighted by a positive scalar w_ϕ (this class is discussed in Section 4.1 as a motivating example).

Nonetheless, we show that these three classes of regret minimization problems – external regret, proper ϕ -regret, and improper ϕ -regret – are all distinct (with each class strictly contained in the

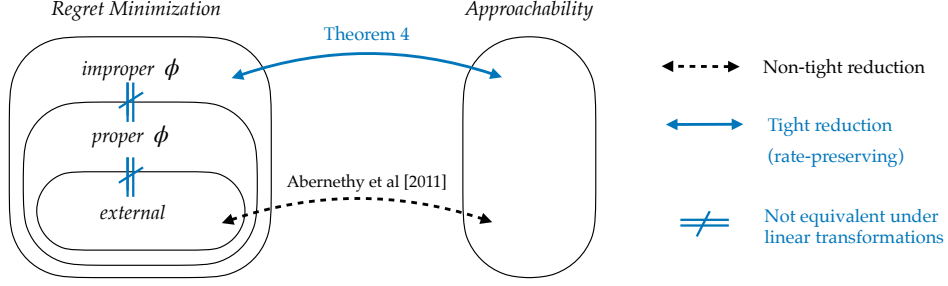


Figure 1: Overview of problem classes and reductions. [Abernethy et al. \(2011\)](#) give a non-tight reduction between approachability and external regret minimization. We give in [Theorem 4](#) a tight (rate-preserving) reduction between approachability and the class of improper ϕ -regret minimization problems. Further, we characterize when an improper ϕ -regret instance is tightly reducible, under linear equivalence, to more well-studied classes of regret minimization problems, like proper ϕ -regret minimization ([Theorem 14](#)) and external regret minimization ([Theorem 5](#)).

next). Specifically, we prove the following results (see [Figure 1](#) for an illustration of the problem classes and reductions):

- We give a clean mathematical characterization of when an improper ϕ -regret instance is linearly equivalent to some external regret instance ([Theorem 5](#)). This characterization puts rather strong constraints on the set of functions Φ (in particular, any two functions in Φ must differ by a constant), and rules out the possibility of (for example) swap regret minimization being linearly equivalent to an external regret minimization instance.
- We provide examples of improper ϕ -regret instances that are provably *not* linearly equivalent to any proper ϕ -regret instance ([Section F.1](#)). This provides evidence that the language of approachability is strictly more powerful than the language of standard regret minimization. These counterexamples also seem mathematically quite rich, connected to concepts like finding linear subspaces of non-invertible matrices.
- Finally, we provide an algorithmic characterization of when an improper ϕ -regret instance is linearly equivalent to a proper ϕ -regret instance, reducing the problem to checking whenever a certain convex cone of d -by- d matrices contains any invertible matrices ([Section F.2](#)). In the case where the action set \mathcal{P} and Φ are specified as the convex hull of N pure actions and M transformation functions respectively, this check can be performed by a randomized algorithm in time polynomial in N , M , and d (and also provides an effective algorithm for producing such a reduction).

1.2. Illustrative examples of improper ϕ -regret and reductions

To illustrate the idea of linear equivalences between different regret minimization problems, in this section we present three different examples of ϕ -regret minimization problems. All examples will have action set $\mathcal{P} = \Delta_3$ (the simplex over 3 actions) and loss set $\mathcal{L} = [-1, 1]^d$. The set of

$$\begin{array}{c|c|c}
 \begin{array}{l} \phi_1(p) = (1, -p_2, -2p_3) \\ \phi_2(p) = (0, 2 - p_2, -2p_3) \\ \phi_3(p) = (0, -p_2, 3 - 2p_3) \end{array} & \begin{array}{l} \phi_1(p) = (1, 0, 0) \\ \phi_2(p) = (-p_1, 2 - p_2, -p_3) \\ \phi_3(p) = (-2p_1, -2p_2, 3 - 2p_3) \end{array} & \begin{array}{l} \phi_1(p) = (p_1 - p_2, p_1 + p_2, p_3) \\ \phi_2(p) = (p_1 + p_3, p_2, p_3 - p_1) \\ \phi_3(p) = (p_1, p_2 - p_3, p_2 + p_3). \end{array} \\
 \text{(a)} & \text{(b)} & \text{(c)}
 \end{array}$$

 Figure 2: Examples of improper ϕ -regret minimization

transformations Φ will always be of the form $\text{conv}(\{\phi_1, \phi_2, \phi_3\})$ but with these specific functions changing between examples as shown in Figure 2.

To begin, note that all these instances are truly *improper* ϕ -regret problems; it is easy to find a $p = (p_1, p_2, p_3) \in \Delta_3$ for which some $\phi_i(p)$ lies outside Δ_3 . Nonetheless, Blackwell approachability can be used to show that it is possible to achieve sublinear ϕ -regret for all three of these instances.

These three instances have different reducibility properties: instance (a) is linearly equivalent to an external regret problem, instance (b) is linearly equivalent to a proper ϕ -regret problem (but not to any external regret problem), and instance (c) cannot be reduced to a proper ϕ -regret problem.

To give an explicit example of what such a linear reduction entails, for instance (a), it turns out that for any $\ell \in \mathcal{L}$ and $p \in \mathcal{P}$ we have the equality $\langle p - \phi_i(p), \ell \rangle = \langle p - e_i, \ell' \rangle$, where e_i is the i th basis vector and $\ell' = \text{diag}(1, 1/2, 1/3)\ell$. Therefore, this strange improper ϕ -regret minimization problem is really an ordinary external regret minimization problem (on the modified loss set $\mathcal{L}' = \text{diag}(1, 1/2, 1/3)\mathcal{L}$) in disguise. One can also check that the different $\phi_i(p)$ differ by a constant, as required by our characterization.

Similarly, for instance (b), one can check that although the ϕ_i are improper, we can write $\langle p - \phi_i(p), \ell \rangle = \langle p - \phi'_i(p), \ell' \rangle$ for a set of proper ϕ'_i and where $\ell' = \ell/3$. On the other hand, there is no way to reduce this instance to an external regret instance, since different ϕ_i do not differ by a constant function. Interestingly, for this instance we can also write $\langle p - \phi_i(p), \ell \rangle = w_i \cdot \langle p - e_i, \ell \rangle$ (for $(w_1, w_2, w_3) = (1, 2, 3)$), showing that this instance captures a version of *weighted external regret* (where regret w.r.t action i is multiplied by w_i). In particular, this demonstrates the (somewhat surprising fact) that weighted external regret cannot be written exactly as an external regret minimization problem.

Finally, instance (c) is an improper ϕ -regret minimization problem that provably cannot be reduced to either of the smaller classes. Interestingly the construction of this ϕ comes from the fact that the skew-symmetric matrices of odd dimension form a linear subspace of singular matrices (each $p - \phi_i(p)$ is a skew-symmetric linear transformation).

Prior Work We discuss additional prior work in Appendix A.

2. Model and Preliminaries

Given two convex sets $\mathcal{C}_1 \subseteq \mathbb{R}^{d_1}$ and $\mathcal{C}_2 \subseteq \mathbb{R}^{d_2}$, we define their tensor product $\mathcal{C}_1 \otimes \mathcal{C}_2$ to be the subset of $\mathbb{R}^{d_1} \otimes \mathbb{R}^{d_2} \simeq \mathbb{R}^{d_1 d_2}$ equal to the convex hull of all vectors of the form $c_1 \otimes c_2 = c_1 c_2^\top$ for $c_1 \in \mathcal{C}_1$ and $c_2 \in \mathcal{C}_2$. We write $\text{Lin}(\mathcal{C}_1, \mathcal{C}_2)$ to denote the collection of linear maps that send every point in \mathcal{C}_1 to a point in \mathcal{C}_2 , and $\text{Aff}(\mathcal{C}_1, \mathcal{C}_2)$ to denote the collection of affine maps that send every point in \mathcal{C}_1 to a point in \mathcal{C}_2 . Note that if \mathcal{C}_1 belongs to an affine subspace of \mathbb{R}^{d_1} (that does not contain the origin), then any affine function on \mathcal{C}_1 can be equivalently written as a linear function and so $\text{Aff}(\mathcal{C}_1, \mathcal{C}_2) \simeq \text{Lin}(\mathcal{C}_1, \mathcal{C}_2)$. On the other hand, if this is not the case, we can always augment \mathcal{C}_1 (by

adding a $(d + 1)$ st dummy coordinate) and have that $\text{Aff}(\mathcal{C}_1, \mathcal{C}_2) \simeq \text{Lin}(\mathcal{C}'_1, \mathcal{C}_2)$. For any convex set \mathcal{K} , we let $\text{cone}(\mathcal{K}) = \{x \mid \exists \alpha \geq 0 \text{ s.t. } \alpha x \in \mathcal{K}\}$ be the conical hull of \mathcal{K} .

2.1. Approachability

An *approachability problem*¹ is defined by three bounded convex subsets of Euclidean space: an action set \mathcal{P} , a loss set \mathcal{L} , and a “constraint” set of bi-affine functions $\mathcal{U} \subset \text{Aff}(\mathcal{P} \otimes \mathcal{L}, \mathbb{R})$. Traditionally, \mathcal{U} is provided in the form of a multidimensional bilinear function $\mathbf{u}: \mathcal{P} \times \mathcal{L} \rightarrow \mathbb{R}^n$ (which would correspond to the \mathcal{U} given by the convex hull of the n components of \mathbf{u}), but it is easier and more general to work with \mathcal{U} directly. For example, this lets us easily consider approachability problems with infinite-dimensional \mathbf{u} .

For a given time horizon T , an *approachability algorithm* \mathcal{A} is a sequence of functions $\{\mathcal{A}_t\}_{t \in [T]}$, where each \mathcal{A}_t maps a prefix of losses $\ell_1, \dots, \ell_{t-1} \in \mathcal{L}$ to an action $\mathcal{A}_t(\ell_1, \dots, \ell_{t-1}) = p_t$. For a given sequence of T actions \mathbf{p} and T losses ℓ , we define the *approachability loss* $\text{AppLoss}(\mathbf{p}, \ell)$ as

$$\text{AppLoss}(\mathbf{p}, \ell) = \max_{u \in \mathcal{U}} \sum_{t=1}^T u(p_t, \ell_t). \quad (1)$$

The objective of an approachability problem is to construct approachability algorithms which guarantee low AppLoss . For a given algorithm \mathcal{A} and time horizon T , let $\text{AppLoss}_T(\mathcal{A})$ be the worst-case approachability loss of \mathcal{A} on any sequence of T losses, that is, $\text{AppLoss}_T(\mathcal{A}) = \max_{\ell \in \mathcal{L}^T} \text{AppLoss}(\mathcal{A}(\ell), \ell)$. We will omit the subscript T when it is clear from context.

We say an approachability instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ is *approachable* if² for each $\ell \in \mathcal{L}$, there exists a $p \in \mathcal{P}$ such that $u(p, \ell) \leq 0$ for all $u \in \mathcal{U}$. Blackwell’s theorem (Blackwell, 1956) provides the following dichotomy:

1. If an approachability instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ is **not** approachable, then for any algorithm \mathcal{A} , $\text{AppLoss}_T(\mathcal{A}) = \Omega(T)$.
2. If an approachability instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ is approachable, then there exists an algorithm \mathcal{A} where $\text{AppLoss}_T(\mathcal{A}) = O(\sqrt{T})$.

For an approachable instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$, define

$$\text{Rate}(\mathcal{P}, \mathcal{L}, \mathcal{U}) = \lim_{T \rightarrow \infty} \inf_{\mathcal{A}} \frac{\text{AppLoss}_T(\mathcal{A})}{\sqrt{T}}. \quad (2)$$

In words, the optimal algorithm for the instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ achieves a worst-case approachability loss of $\text{Rate}(\mathcal{P}, \mathcal{L}, \mathcal{U})\sqrt{T} + o(\sqrt{T})$.

1. It is straightforward to show that any approachability problem can be written in this form (for completeness, we show this in Appendix B).
 2. This is the “response-satisfiable” definition of approachability. It is known that approachable sets can be equivalently defined via the “half-space-satisfiable” definition of approachability, which asks that for each $u \in \mathcal{U}$, there exists a $p \in \mathcal{P}$ such that $u(p, \ell) \leq 0$ for all $\ell \in \mathcal{L}$. We provide a short proof of this fact in Appendix G (Theorem 13).

2.2. Regret minimization

Much like approachability problems, we specify a *regret minimization problem* by a triple of bounded convex sets: an action set $\mathcal{P} \subseteq \mathbb{R}^d$, a loss set $\mathcal{L} \subseteq \mathbb{R}^d$, and a benchmark set $\Phi \subseteq \text{Aff}(\mathcal{P}, \mathbb{R}^d)$ of affine functions from \mathcal{P} to \mathbb{R}^d . Intuitively, each function $\phi \in \Phi$ represents a “swap” or “transformation” function that the learner is competing against: if the learner outputs a sequence of actions p_1, p_2, \dots, p_T , they want to have low regret compared to the sequence of actions $\phi(p_1), \phi(p_2), \dots, \phi(p_T)$.

For a given regret minimization instance $(\mathcal{P}, \mathcal{L}, \Phi)$, the regret of a sequence of T actions \mathbf{p} and T losses ℓ is given by

$$\text{Reg}(\mathbf{p}, \ell) = \max_{\phi \in \Phi} \left(\sum_{t=1}^T \langle p_t, \ell_t \rangle - \sum_{t=1}^T \langle \phi(p_t), \ell_t \rangle \right). \quad (3)$$

As with approachability, our goal is to design a learning algorithm \mathcal{A} for this problem which achieves low worst-case regret. In fact, it is fairly straightforward to write the above regret minimization problem as a specific instance of approachability: let \mathcal{U} be the set of bilinear functions of the form

$$u_\phi(p, \ell) = \langle p, \ell \rangle - \langle \phi(p), \ell \rangle, \quad (4)$$

where ϕ ranges over all $\phi \in \Phi$, then it is easy to see that the expression for $\text{Reg}(\mathbf{p}, \ell)$ in (3) is equivalent to the expression for $\text{AppLoss}(\mathbf{p}, \ell)$ in (1). We will likewise write $\text{Reg}_T(\mathcal{A})$ to refer to the worst-case regret of algorithm \mathcal{A} , and $\text{Rate}(\mathcal{P}, \mathcal{L}, \Phi)$ to represent the optimal regret rate for this regret minimization instance.

If each $\phi \in \Phi$ contains a fixed point in \mathcal{P} , then the corresponding approachability problem is approachable and $\text{Rate}(\mathcal{P}, \mathcal{L}, \Phi) < \infty$. Furthermore, if the loss set \mathcal{L} contains a vector in every direction ($\text{cone}(\mathcal{L}) = \mathbb{R}^d$), then this condition exactly characterizes regret minimization instances that permit sublinear regret algorithms (this follows easily from Blackwell’s theorem).

Theorem 1 *Consider a regret minimization instance $(\mathcal{P}, \mathcal{L}, \Phi)$. If each $\phi \in \Phi$ has a fixed point $p_\phi \in \mathcal{P}$ (i.e., $\phi(p_\phi) = p_\phi$), then $\text{Rate}(\mathcal{P}, \mathcal{L}, \Phi) < \infty$. Conversely, if $\text{cone}(\mathcal{L}) = \mathbb{R}^d$ and $\text{Rate}(\mathcal{P}, \mathcal{L}, \Phi) < \infty$, then every $\phi \in \Phi$ must have a fixed point in \mathcal{P} .³*

We will categorize regret minimization instances into three classes based on the properties of their set of benchmarks Φ , which we list in increasing order of generality:

- **External regret minimization.** If each $\phi \in \Phi$ has the property that ϕ is a constant function over \mathcal{P} (i.e., there exists a $p_\phi \in \mathcal{P}$ s.t. $\phi(p) = p_\phi$ for all $p \in \mathcal{P}$), then we say that this instance of regret minimization is an *external regret minimization* instance. Note that the classic online learning setting of *online linear optimization* fits into this class (the benchmark set Φ is the set of all constant functions on \mathcal{P}).
- **Proper ϕ -regret minimization.** If each $\phi \in \Phi$ has the property that $\phi(p) \in \mathcal{P}$ whenever $p \in \mathcal{P}$ (i.e., each ϕ maps \mathcal{P} into itself), then we say that this instance of regret minimization is a *proper ϕ -regret minimization* instance. Note that by Brouwer’s fixed-point theorem, each ϕ must contain

3. For the sake of brevity, we defer most proofs of theorems in the main body to Appendix G.

a fixed point in \mathcal{P} , and therefore each proper ϕ -regret minimization problem has a sublinear-regret algorithm.

This class captures the well-known cases of swap regret and internal regret, along with the notion of linear ϕ -regret studied in (Gordon et al., 2008). It also contains the previous class of external regret minimization instances.

- **Improper ϕ -regret minimization.** As long as each $\phi \in \Phi$ has a fixed point in \mathcal{P} (a $p \in \mathcal{P}$ s.t. $\phi(p) = p$), we say that this instance of regret minimization is an *improper ϕ -regret instance*.

Note here that the benchmark functions ϕ are allowed to send points in \mathcal{P} to points outside of \mathcal{P} (that the learner is not even allowed to play). Nonetheless, by Theorem 1, we know that $\text{Rate}(\mathcal{P}, \mathcal{L}, \Phi) < \infty$ for any such improper ϕ -regret instance.

2.3. Reductions between learning problems

Ideally, we would design efficient learning algorithms that provably match the optimal rate for a given approachability or regret minimization instance. Unfortunately, doing this directly seems very challenging – even for the special case of online linear optimization, it is unclear how to construct efficient learning algorithms with near-optimal regret bounds.

Instead, we will settle for being able to reduce an arbitrary approachability instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ to an instance of a simpler learning problem. In particular, all of our simpler learning problems we consider can also be written as approachability problems, so our main goal is to understand when a specific approachability instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ is “reducible” to a second specific instance $(\mathcal{P}', \mathcal{L}', \mathcal{U}')$.

We define our notion of reducibility as follows. We say there is a *tight reduction* between instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \mathcal{U})$ and instance $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \mathcal{U}')$ if both: i. given any algorithm \mathcal{A} for \mathcal{I} , we can construct an algorithm \mathcal{A}' for \mathcal{I}' with $\text{AppLoss}(\mathcal{A}') = \text{AppLoss}(\mathcal{A})$, and ii. given any algorithm \mathcal{A}' for \mathcal{I}' , we can construct an algorithm \mathcal{A} for \mathcal{I} with $\text{AppLoss}(\mathcal{A}) = \text{AppLoss}(\mathcal{A}')$. We say there is a *c-approximate reduction* between the two instances if, for sufficiently large T , we instead have $\text{AppLoss}(\mathcal{A}') \leq c \text{AppLoss}(\mathcal{A})$ and $\text{AppLoss}(\mathcal{A}) \leq c \text{AppLoss}_T(\mathcal{A}')$ in the two constraints respectively. Finally, we say that there is a *weak reduction from* instance \mathcal{I} *to* instance \mathcal{I}' if we only have one direction of the reduction: given an algorithm \mathcal{A}' for \mathcal{I}' , we can efficiently construct an algorithm \mathcal{A} for \mathcal{I} with $\text{AppLoss}(\mathcal{A}) \leq \text{AppLoss}(\mathcal{A}')$. Note that a tight reduction implies that $\text{Rate}(\mathcal{I}') = \text{Rate}(\mathcal{I})$, a *c-approximate reduction* implies that $\text{Rate}(\mathcal{I})/c \leq \text{Rate}(\mathcal{I}') \leq c \text{Rate}(\mathcal{I})$, and a *weak reduction from* \mathcal{I} *to* \mathcal{I}' only implies that $\text{Rate}(\mathcal{I}) \leq \text{Rate}(\mathcal{I}')$.

Since any regret minimization instance can be written as an approachability instance, we can use the same notion of reducibility when talking about reductions among instances of regret minimization. The latter sections of this paper will be concerned with understanding when a specific instance of approachability is equivalent to an instance of a specific class of regret minimization problems. There we will use a more restrictive notion of “linear equivalence”, where all relevant constructions must be provided by affine transformations; we introduce this in Section 4.

2.4. The classical reduction from approachability to external regret minimization

Finally, we present the classical reduction from approachability to external regret minimization. This is the same reduction in (Abernethy et al., 2011) (and subsequent works), adapted to our notation.

Theorem 2 [Abernethy et al. (2011)] *Given any instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \mathcal{U})$ of approachability, there exists a weak reduction from \mathcal{I} to an instance $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$ of external regret minimization.*

The main idea behind the reduction of Theorem 2 is to run a “dual” regret minimization algorithm over the set \mathcal{U} . That is, if the learner chooses actions $u_t \in \mathcal{U}$ every round so to minimize the external regret $\max_{u \in \mathcal{U}} \sum_{t=1}^T u(p_t, \ell_t) - \sum_{t=1}^T u_t(p_t, \ell_t)$, and then chooses p_t so that $u_t(p_t, \ell) \leq 0$ for all $\ell \in \mathcal{L}$ (possible since the instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ is approachable), this external regret upper bounds the approachability loss $\text{AppLoss}(\mathbf{p}, \ell) = \max_{u \in \mathcal{U}} \sum_{t=1}^T u(p_t, \ell_t)$. In particular, our reduced instance \mathcal{I}' has $\mathcal{P}' = \mathcal{U}$ and $\mathcal{L}' = -(\mathcal{P} \otimes \mathcal{L})$.

3. Reducing approachability to (improper) regret minimization

3.1. The classical reduction is not tightly rate-preserving

Theorem 2 proves that the classical reduction of Abernethy et al. (2011) is a weak reduction from any approachability problem to an external regret minimization problem. It is natural to wonder whether this reduction is in fact tight, or if not, whether the gap in rates between the original approachability problem and the eventual regret problem is small. The following counterexample proves that this is not the case.

Theorem 3 *There exists an approachability instance \mathcal{I} such that $\text{Rate}(\mathcal{I}) = 0$ but $\text{Rate}(\mathcal{I}') > 0$, where \mathcal{I}' is the regret minimization instance obtained by applying the reduction of Theorem 2 to \mathcal{I} . In particular, there is no finite $c > 0$ for which the reduction in Theorem 2 is a c -approximate reduction for all approachability instances, and the reduction is not tight.*

At a high level, the counter-example in Theorem 3 works as follows. We design an approachability instance where the learner every round must choose between playing two different regret minimization instances. One of these regret minimization instances is trivial (having only one arm), but the other is a standard hard instance with d arms. Because of the trivial instance, this problem is perfectly approachable ($\text{Rate}(\mathcal{I}) = 0$). However, when we perform the reduction to external regret minimization in Theorem 2, we must choose a sequence of u_t that incurs low regret against *any* adversarial sequence of p_t and ℓ_t – even sequences of p_t that we would never pick.

We emphasize here that although we have attributed this reduction to Abernethy et al. (2011), essentially all later works designing algorithms for approachability or ϕ -regret minimization build off the same reduction to a regret minimization problem over the set \mathcal{U} , and all (to our knowledge) lead to sub-optimal rate algorithms for this counter-example. We discuss this in more detail in Appendix D.

3.2. A tight reduction from approachability to improper ϕ -regret minimization

In contrast to Theorem 3, it turns out that it is possible to tightly reduce approachability to improper ϕ -regret.

Theorem 4 *For any approachability instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \mathcal{U})$, there exists a tight reduction from \mathcal{I} to an improper ϕ -regret minimization instance $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$.*

The key observation is that if we “augment” the action space of the learner so as to also include the benchmark $u \in \mathcal{U}$ they are competing against each round (in particular, by setting $\mathcal{P}' = \mathcal{P} \otimes \mathcal{U}$), we can directly interpret $\text{AppLoss}(\mathbf{p}, \ell)$ as a specific ϕ -regret over this space.

Theorem 4 prompts us to question when we can further reduce an improper ϕ -regret minimization instance to a proper one. We study this question through the lens of linear equivalence in the next section.

4. Tight linear reductions for regret minimization

Theorem 4 implies that approachability is equivalent to improper ϕ -regret minimization. Ideally, we would be able to, in turn, tightly reduce this instance of improper ϕ -regret minimization to an instance of the better studied problem of external regret minimization (as the classical reduction attempts to do), or at least to an instance of proper ϕ -regret minimization.

In this section, we examine the question of “which instances of regret minimization are tightly reducible to each other?”, and specifically “when is an improper ϕ -regret instance tightly reducible to a proper ϕ -regret (or even an external regret) instance?”. However, the space of general reductions is hard to work with directly⁴ – for this reason, we will restrict ourselves to a subclass of these reductions that we call *linear equivalences*. We will formally define linear equivalences in the next section. For now, we will begin by motivating them via an application where a non-trivial tight reduction is possible: solving the problem of *weighted regret minimization*.

4.1. Warm-up: weighted regret minimization

In the weighted regret minimization problem, we have an action set \mathcal{P} , a loss set \mathcal{L} , and a collection $\Phi = \{\phi_1(p), \phi_2(p), \dots, \phi_N(p)\}$ of proper affine benchmark functions (i.e., each ϕ_i maps \mathcal{P} to itself). Each ϕ_i has an accompanying scalar weight $w_i > 0$, representing the importance we place on regret with respect to ϕ_i . As with the general regret minimization problem and approachability problem, the learner runs a learning algorithm to decide their action at time t . The goal of the learner is to minimize the expression:

$$\text{WReg}(\mathbf{p}, \ell) = \max_{i \in [N]} \left(w_i \cdot \left[\sum_{t=1}^T \langle p, \ell \rangle - \sum_{t=1}^T \langle \phi_i(p), \ell \rangle \right] \right). \quad (5)$$

In other words, $\text{WReg}(\mathbf{p}, \ell)$ simply equals the corresponding ϕ -regret where the regret with respect to ϕ_i is scaled by w_i .

One way to approach a specific weighted regret minimization problem \mathcal{I} is to simply ignore the weights and run an algorithm for the corresponding proper ϕ -regret minimization problem. This would work, in the sense that a sublinear regret algorithm for the unweighted problem would also obtain sublinear regret for the weighted regret minimization problem. But completely ignoring the weights is, of course, lossy: a regret guarantee of R for the unweighted problem only translates to a regret guarantee of $(\max w_i)R$ for the weighted problem. Conversely, if we have a learning algorithm with a regret guarantee of R for the weighted problem, running it on the unweighted

4. Not only is it hard to work with directly, it quickly becomes meaningless without imposing any constraints on the reductions themselves. Without any such constraints, essentially any two regret-minimization instances \mathcal{I} and \mathcal{I}' with $\text{Rate}(\mathcal{I})$ and $\text{Rate}(\mathcal{I}')$ are equivalent. This is because given an algorithm \mathcal{A} for \mathcal{I} , we can just compute $\text{AppLoss}(\mathcal{A})$ and construct an arbitrary algorithm \mathcal{A}' with $\text{AppLoss}(\mathcal{A}') = \text{AppLoss}(\mathcal{A})$.

instance only translates to a regret guarantee of $R/(\min w_i)$. In the language of our different types of reductions, this naive translation of algorithms is only an ω -approximate reduction, where $\omega = (\max w_i)/(\min w_i)$.

Despite this, it is possible to exactly write the weighted regret minimization problem as a regret minimization problem of the form described in Section 2.2. We simply need to observe that:

$$w_i (\langle p, \ell \rangle - \langle \phi_i(p), \ell \rangle) = \langle p, \ell \rangle - \langle w_i \phi_i(p) - (w_i - 1)p, \ell \rangle. \quad (6)$$

Therefore, if we let $\tilde{\phi}_i(p) = w_i \phi_i(p) - (w_i - 1)p$ and $\tilde{\Phi} = \text{conv}(\{\tilde{\phi}_i\}_{i=1}^N)$, then the regret minimization instance $(\mathcal{P}, \mathcal{L}, \tilde{\Phi})$ is exactly equivalent to our weighted regret minimization problem. However, this regret minimization problem is in general an *improper* ϕ -regret minimization problem (even though all the original ϕ_i are proper). In fact, this is the case even when the benchmarks ϕ_i are all constant and correspond to a weighted external regret minimization problem (e.g., if $\mathcal{P} = \Delta_3$, $\phi_1(p) = (1, 0, 0)$, and $w_1 = 2$, then $\tilde{\phi}_1(p) = (2 - p_1, -p_2, -p_3)$).

It turns out that in this case, there is a tight reduction between this improper ϕ -regret minimization problem and a proper ϕ -regret minimization problem. In particular, if we let $W = \max w_i$, then note that $\langle p - \tilde{\phi}_i(p), \ell \rangle = \langle p - \phi'_i(p), W\ell \rangle$ for ϕ'_i defined via

$$\phi'_i(p) = \frac{w_i}{W} \phi_i(p) + \left(1 - \frac{w_i}{W}\right) p.$$

Moreover, each transformation $\phi'_i(p)$ is proper (since it is now a convex combination of two proper transformations). This means that the improper ϕ -regret minimization problem we originally faced reduces to a proper ϕ -regret minimization problem under the linear transformation of scaling the losses by W (we prove this formally in Section G.6).

Motivated by this, we introduce the following notion of a *linear equivalence* between two regret minimization problems. We say two approachability instances $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ and $\mathcal{I}' = (\mathcal{P}, \mathcal{L}', \Phi')$ are *linearly equivalent*⁵ if there exists an invertible d -by- d linear transformation S with the property that: i. $\mathcal{L}' = (S^T)^{-1} \mathcal{L}$, and ii. each $\phi \in \Phi$ corresponds bijectively to a $\phi' \in \Phi'$ satisfying $\phi'(p) = p + S(\phi(p) - p)$. In particular, one can check that $\langle \phi'(p) - p, (S^T)^{-1} \ell \rangle = \langle \phi(p) - p, \ell \rangle$, which implies that the Φ -regret a learner faces by playing the sequence of actions \mathbf{p} against the sequence of losses ℓ is the same as the Φ' -regret a learner faces by playing \mathbf{p} against the sequence of losses $(S^T)^{-1} \ell$ (for example, the reduction above corresponds to the transformation $S = (1/W) \cdot \text{Id}$).

4.2. Reducing improper ϕ -regret to external regret

In this section we will provide a complete characterization of when (minimal) regret minimization instances are linearly equivalent to external regret minimization instances (recall that these are instances where all of the functions $\phi \in \Phi$ are constant over \mathcal{P}).

Given a set of points U , let $\text{span}_{\text{Aff}}(U)$ denote the *affine span* of U , that is, the affine subspace formed by all vectors of the form $\sum_{i=1}^k \lambda_i v_i$ for any $k > 0$, $v_i \in U$, and $\lambda_i \in \mathbb{R}$ such that $\sum_i \lambda_i = 1$ (note that the λ_i need not be non-negative). We prove the following characterization.

Theorem 5 *Let $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ be a regret minimization instance. Then \mathcal{I} is linearly equivalent to an external regret minimization instance if and only if the following conditions are met:*

5. In Appendix E, we present a considerably more general definition of linear equivalences (that e.g. doesn't require both instances to share the same set \mathcal{P}) and show that it suffices to understand the simplified version we present here under some reasonable assumptions on the sets \mathcal{P} and \mathcal{L} .

1. For all $\phi \in \Phi$, each ϕ has a single fixed point p in $\text{span}_{\text{Aff}}(\mathcal{P})$.
2. For all $\phi_1, \phi_2 \in \Phi$, $\phi_1(p) - \phi_2(p)$ is a constant for all $p \in \mathcal{P}$.

Note that by the assumption that the instance \mathcal{I} is valid, every $\phi \in \Phi$ must contain a fixed point in \mathcal{P} ; condition 1 above rules out the existence of any other fixed points in the entire affine subspace containing \mathcal{P} . We consider a few illustrative examples of Theorem 5 below:

- **Classic external regret:** In the classic external-regret setting of learning with experts, $\mathcal{P} = \Delta_d$, $\mathcal{L} = [0, 1]^d$, and the set Φ contains all constant functions $\phi : \mathcal{P} \rightarrow \mathcal{P}$. If ϕ is the constant map $\phi(p) = p_\phi$ (for some fixed $p_\phi \in \mathcal{P}$), then ϕ has the unique fixed point p_ϕ . It is also clear that the difference of any two constant functions is constant.
- **Improper regret that can be reduced to external regret:** Consider the setting where $\mathcal{P} = \Delta_2$, $\mathcal{L} = [0, 1]^2$, and Φ contains all linear functions of the form $\phi_\alpha(x, y) = (x, (2 - \alpha)y - \alpha x)$ for all $\alpha \in [0, 1]$. Note that this is an improper ϕ -regret minimization instance since ϕ_α often sends points in the simplex to points outside the simplex (e.g., $\phi_{1/2}(2/3, 1/3) = (2/3, 1/6)$). Nonetheless, ϕ_α satisfies the constraints of Theorem 5: $(1 - \alpha, \alpha)$ is a fixed point of ϕ_α , and $\phi_{\alpha_1}(p, 1 - p) - \phi_{\alpha_2}(p, 1 - p) = (0, \alpha_2 - \alpha_1)$ for any $(p, 1 - p) \in \Delta_2$. And indeed, one can check that under the invertible transformation matrix $S = [[0, 1], [1, -1]]$, $\phi_\alpha(x, y)$ gets transformed (via (18)) to $\phi'_\alpha(x, y) = ((1 - \alpha)(x + y), \alpha(x + y))$ (which takes on the constant value $(1 - \alpha, \alpha)$ for $(x, y) \in \Delta_2$).
- **Swap regret:** Swap regret minimization is a proper ϕ -regret minimization where $\mathcal{P} = \Delta_d$, $\mathcal{L} = [0, 1]^d$, and Φ is the convex hull of all $\phi_\pi(p, \ell) = \sum_{i=1}^d p_i(\ell_i - \ell_{\pi(i)})$ where π ranges over all “swap functions” from $[d]$ to $[d]$. It can easily be checked that these functions do not differ by a constant (for $d > 1$)⁶ and therefore by Theorem 5 there is no linear equivalence between swap regret and external regret.
- **Weighted external regret:** Consider the weighted regret minimization setting of Section 4.1 where each original function $\phi_i(p)$ is a constant function (so this captures a weighted version of external regret minimization). The expression (6) describes how to express this in terms of an improper ϕ -regret minimization problem with functions $\tilde{\phi}_i(p)$. Looking at the expressions for $\tilde{\phi}_i(p)$, we can see that a weighted external regret minimization instance is linearly equivalent to an ordinary external regret instance if and only if all the weights w_i are equal.

4.3. Reducing improper ϕ -regret to proper ϕ -regret

We discuss the problem of when improper ϕ -regret instances are linearly equivalent to proper ϕ -regret instances in Appendix F. We show that there exist some improper ϕ -regret instances that are not linearly equivalent to any proper ϕ -regret instance, and give an efficient algorithm for deciding when this is the case.

6. The case $d = 2$ is interesting. When $d = 2$ it is possible to bound swap regret by at most twice external regret, and so sublinear external regret implies sublinear swap regret. But there is still no linear equivalence between swap regret and external regret in this case, just as there is none between weighted external regret and external regret in the next example.

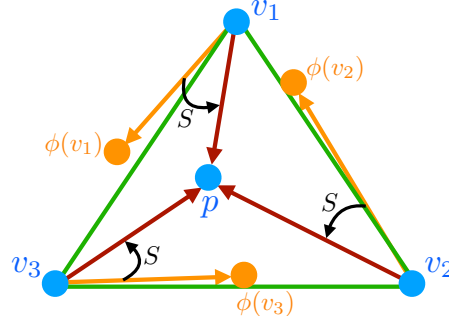


Figure 3: Illustration of the definition of S in the proof of Lemma 15. v_1, v_2, v_3 represent the vertices of the simplex in \mathbb{R}^3 and p is the fixed point of ϕ . The vector $\phi(v_k)$, $k \in [3]$, may not be in the simplex. The linear function S maps vector $(\phi - I)(v_k)$ into $p - v_k$.

Acknowledgments

YM is supported in part by funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 882396), by the Israel Science Foundation, the Yandex Initiative for Machine Learning at Tel Aviv University and a grant from the Tel Aviv University Center for AI and Data Science (TAD).

References

- Jacob D. Abernethy, Peter L. Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of COLT*, volume 19 of *JMLR Proceedings*, pages 27–46, 2011.
- Shipra Agrawal, Constantinos Daskalakis, Vahab S. Mirrokni, and Balasubramanian Sivan. Robust repeated auctions under heterogeneous buyer behavior. In Éva Tardos, Edith Elkind, and Rakesh Vohra, editors, *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, page 171. ACM, 2018.
- Yu Bai, Chi Jin, Song Mei, Ziang Song, and Tiancheng Yu. Efficient phi-regret minimization in extensive-form games via online mirror descent. *Advances in Neural Information Processing Systems*, 35:22313–22325, 2022.
- David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, 2007.
- Mark Braverman, Jieming Mao, Jon Schneider, and S. Matthew Weinberg. Selling to a no-regret buyer. In Éva Tardos, Edith Elkind, and Rakesh Vohra, editors, *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 523–538. ACM, 2018.
- Linda Cai, S. Matthew Weinberg, Evan Wildenhain, and Shirley Zhang. Selling to multiple no-regret buyers. In Jugal Garg, Max Klimm, and Yuqing Kong, editors, *Web and Internet Economics -*

- 19th International Conference, WINE 2023, Shanghai, China, December 4-8, 2023, Proceedings*, volume 14413 of *Lecture Notes in Computer Science*, pages 113–129. Springer, 2023.
- Andrea Celli, Alberto Marchesi, Gabriele Farina, Nicola Gatti, et al. Decentralized no-regret learning algorithms for extensive-form correlated equilibria. In *IJCAI*, pages 4755–4759, 2021.
- Evgenii Chzhen, Christophe Giraud, and Gilles Stoltz. A unified approach to fair online learning via blackwell approachability. In *Proceedings of NeurIPS*, pages 18280–18292, 2021.
- Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. From external to swap regret 2.0: An efficient reduction and oblivious adversary for large action spaces. *arXiv preprint arXiv:2310.19786*, 2023.
- Christoph Dann, Yishay Mansour, Mehryar Mohri, Jon Schneider, and Balasubramanian Sivan. Pseudonorm approachability and applications to regret minimization. In Shipra Agrawal and Francesco Orabona, editors, *International Conference on Algorithmic Learning Theory, February 20-23, 2023, Singapore*, volume 201 of *Proceedings of Machine Learning Research*, pages 471–509. PMLR, 2023.
- A. P. Dawid. The well-calibrated bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982. doi: 10.1080/01621459.1982.10477856. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1982.10477856>.
- Yuan Deng, Jon Schneider, and Balasubramanian Sivan. Strategizing against no-regret learners. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 1577–1585, 2019. URL <http://papers.nips.cc/paper/8436-strategizing-against-no-regret-learners>.
- Dean P. Foster and Sergiu Hart. Smooth calibration, leaky forecasts, finite recall, and nash dynamics. *Games and Economic Behavior*, 109:271–293, 2018. ISSN 0899-8256. doi: <https://doi.org/10.1016/j.geb.2017.12.022>. URL <https://www.sciencedirect.com/science/article/pii/S0899825618300113>.
- Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, 1997.
- Dean P Foster and Rakesh V Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, 29(1-2):7–35, 1999.
- Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of ICML*, volume 307, pages 360–367. ACM, 2008.
- Amy Greenwald and Amir Jafari. A general class of no-regret learning algorithms and game-theoretic equilibria. In *Proceedings of COLT*, volume 2777 of *Lecture Notes in Computer Science*, pages 2–12. Springer, 2003.
- Amy Greenwald, Zheng Li, and Warren Schudy. More efficient internal-regret-minimizing algorithms. In *COLT*, pages 239–250, 2008.

- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- Dileep M. Kalathil, Vivek S. Borkar, and Rahul Jain. Blackwell’s approachability in Stackelberg stochastic games: A learning version. In *Proceedings of CDC*, pages 4467–4472. IEEE, 2014.
- Joon Kwon. Refined approachability algorithms and application to regret minimization with global costs. *The Journal of Machine Learning Research*, 22:200–1, 2021.
- Joon Kwon and Vianney Perchet. Online learning and blackwell approachability with partial monitoring: Optimal convergence rates. In Aarti Singh and Xiaojin (Jerry) Zhu, editors, *Proceedings of AISTATS*, volume 54, pages 604–613. PMLR, 2017.
- Ehud Lehrer. Approachability in infinite dimensional spaces. *Int. J. Game Theory*, 31(2):253–268, 2003.
- Shie Mannor and Vianney Perchet. Approachability, fast and slow. In *Journal of Machine Learning Research: Workshop and Conference Proceedings*, volume 30, pages 1–16, 2013.
- Shie Mannor and Nahum Shimkin. The empirical bayes envelope and regret minimization in competitive markov decision processes. *Math. Oper. Res.*, 28(2):327–345, 2003.
- Shie Mannor, Vianney Perchet, and Gilles Stoltz. Set-valued approachability and online learning with partial monitoring. *J. Mach. Learn. Res.*, 15(1):3247–3295, 2014a.
- Shie Mannor, Vianney Perchet, and Gilles Stoltz. Approachability in unknown games: Online learning meets multi-objective optimization. In *Proceedings of COLT*, volume 35, pages 339–355, 2014b.
- Yishay Mansour, Mehryar Mohri, Jon Schneider, and Balasubramanian Sivan. Strategizing against learners in bayesian games. In Po-Ling Loh and Maxim Raginsky, editors, *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 5221–5252. PMLR, 2022.
- Sobhan Miryoosefi, Kianté Brantley, Hal Daumé III, Miroslav Dudík, and Robert E. Schapire. Reinforcement learning with convex constraints. In *Proceedings of NIPS*, pages 14070–14079, 2019.
- Mehryar Mohri and Scott Yang. Conditional swap regret and conditional correlated equilibrium. In *NIPS*, pages 1314–1322, 2014.
- Mehryar Mohri and Scott Yang. Online learning with transductive regret. In *Proceedings of NIPS*, pages 5214–5224, 2017.
- Binghui Peng and Aviad Rubinstein. Fast swap regret minimization and applications to approximate correlated equilibria. *arXiv preprint arXiv:2310.19647*, 2023.
- Vianney Perchet. *Approchabilité, Calibration et Regret dans les Jeux à Observations Partielles*. PhD thesis, Université Pierre et Marie Curie - Paris VI, 2010.
- Vianney Perchet. Exponential weight approachability, applications to calibration and regret minimization. *Dynamic Games and Applications*, 5(1):136–153, 2015.

- Vianney Perchet and Marc Quincampoix. On a unified framework for approachability with full or partial monitoring. *Mathematics of Operations Research*, 40(3):596–610, 2015.
- Vianney Perchet and Marc Quincampoix. A differential game on Wasserstein space. Application to weak approachability with partial monitoring. *CoRR*, abs/1811.04575, 2018.
- Aviad Rubinstein and Junyao Zhao. Strategizing against no-regret learners in first-price auctions. *CoRR*, abs/2402.08637, 2024.
- Nahum Shimkin. An online convex optimization approach to blackwell’s approachability. *The Journal of Machine Learning Research*, 17(1):4434–4456, 2016.
- Xavier Spinat. A Necessary and Sufficient Condition for Approachability. *Mathematics of Operations Research*, 27(1):31–44, 2002.
- Gilles Stoltz and Gábor Lugosi. Internal regret in on-line portfolio selection. *Machine Learning*, 59(1):125–159, 2005.
- Gilles Stoltz and Gábor Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games Econ. Behav.*, 59(1):187–208, 2007.
- Nicolas Vieille. Weak approachability. *Math. Oper. Res.*, 17(4):781–791, 1992.
- Brian Hu Zhang, Ioannis Anagnostides, Gabriele Farina, and Tuomas Sandholm. Efficient phi-regret minimization with low-degree swap deviations in extensive-form games. *arXiv preprint arXiv:2402.09670*, 2024.

Appendix Contents

1	Introduction	1
1.1	Our results	2
1.2	Illustrative examples of improper ϕ -regret and reductions	4
2	Model and Preliminaries	5
2.1	Approachability	6
2.2	Regret minimization	7
2.3	Reductions between learning problems	8
2.4	The classical reduction from approachability to external regret minimization	8
3	Reducing approachability to (improper) regret minimization	9
3.1	The classical reduction is not tightly rate-preserving	9
3.2	A tight reduction from approachability to improper ϕ -regret minimization	9
4	Tight linear reductions for regret minimization	10
4.1	Warm-up: weighted regret minimization	10
4.2	Reducing improper ϕ -regret to external regret	11
4.3	Reducing improper ϕ -regret to proper ϕ -regret	12
A	Additional Prior Work	18
B	Approaching the negative orthant	19
C	Non-degenerate counterexample to the reduction	20
D	Other reductions for approachability and our counter-example	21
E	General linear equivalences	22
E.1	Simplifying linear equivalences	23
F	Reducing improper ϕ-regret to proper ϕ-regret	26
F.1	Irreducible improper instances	26
F.2	An algorithmic characterization	28
G	Omitted Proofs	29
G.1	Response satisfiability is equivalent to half-space satisfiability	29
G.2	Proof of Theorem 1	30
G.3	Proof of Theorem 2	30
G.4	Proof of Theorem 3	31
G.5	Proof of Theorem 4	32
G.6	Weighted regret reduction	33
G.7	Proof of Theorem 5	33

Appendix A. Additional Prior Work

Blackwell’s approachability, applications and extensions. There is a large body of work studying, extending, and applying Blackwell’s approachability theory to various problems of interest, including regret minimization (Foster and Vohra, 1999), game theory (Hart and Mas-Colell, 2000), reinforcement learning (Mannor and Shimkin, 2003; Kalathil et al., 2014; Miryoosefi et al., 2019), calibration (Dawid, 1982; Foster and Hart, 2018). Approachability and partial monitoring were studied in a series of publications by Perchet (2010); Mannor et al. (2014a,b); Perchet and Quincampoix (2015, 2018); Kwon and Perchet (2017). More recently, approachability has also been used in the analysis of fairness in machine learning (Chzhen et al., 2021). The notion of approachability has been extended in several studies. These include Vieille (1992) studying weak approachability in finite dimensional spaces, Spinat (2002) providing necessary and sufficient approachability conditions for arbitrary sets (not just convex sets), Lehrer (2003) extending approachability theory to infinite-dimensional spaces.

Improved regret guarantees through approachability for other norms. As mentioned earlier, Abernethy et al. (2011) showed how to use Blackwell’s approachability to solve a general class of regret minimization problems. Nevertheless, their reduction could lead to suboptimal regret guarantees: e.g., a \sqrt{TK} regret for an external regret minimization problem (with T steps and K actions), where $\sqrt{T \log(K)}$ is the optimal regret. It has been observed by many papers (Perchet, 2015; Shimkin, 2016; Kwon, 2021; Dann et al., 2023) since then that this is due to the choice of the ℓ_2 norm for measuring the distance between the average payoff and the target set in approachability. With approachability algorithms for the more suitable ℓ_∞ norm, one can recover the optimal regret guarantees for many problems. Dann et al. (2023) address the time and space complexity of such algorithms from being prohibitively high: i.e., from not depending polynomially on the dimension of the space of vector payoffs.

Swap regret and repeated games. The study of swap-regret and its generalizations has seen renewed interest in recent years due to its interesting connections to learning agents playing repeated games against strategic agents. While there is a large body of work on strategic agents playing against each other in repeated games, and also learning agents playing against each other in repeated games, the outcome of strategic agents playing against learners has remained largely unexplored until recently. The work of Deng, Schneider, and Sivan (2019) initiates the study of optimizer-learner interactions and show that a learner playing a no-swap-regret learning algorithm in a repeated game will not let an optimizer’s reward exceed the Stackelberg value of the game, where the latter itself is always obtainable by an optimizer playing against any no-regret learner. Mansour, Mohri, Schneider, and Sivan (2022) show that a learner playing a no-swap-regret algorithm is not just sufficient, but also necessary to ensure that an optimizer gets no more than the Stackelberg value of the game. They further study the class of Bayesian games and give a sufficient condition in the form of no-polytope-swap-regret for the optimizer to not exceed the Bayesian Stackelberg value of the game. This condition of no-polytope-swap-regret was shown to be necessary to cap the optimizer utility at the Bayesian Stackelberg value by Rubinstein and Zhao (2024). Special classes of these include (Braverman et al., 2018) studying the specific 2-player Bayesian game of an auction between a single seller and single buyer, Agrawal, Daskalakis, Mirrokni, and Sivan (2018) studying a similar setting but also other buyer behaviors beyond learning, and Cai, Weinberg, Wildenhain, and Zhang (2023) extending these to a single seller and multiple buyers.

Φ -regret. Several prominent notions of regret are special instances of Φ -regret (Greenwald and Jafari, 2003). These include standard external regret, internal regret (Foster and Vohra, 1997) (see also (Stoltz and Lugosi, 2005, 2007; Greenwald et al., 2008)), and swap regret (Blum and Mansour, 2007; Peng and Rubinstein, 2023; Dagan et al., 2023). The notions of conditional swap regret (Mohri and Yang, 2014) or transductive regret (Mohri and Yang, 2017) are also related to Φ -regret. Recently, Φ -regret minimization has found applications in designing learning algorithms for extensive-form games that converge to certain classes of correlated equilibria (Celli et al., 2021; Bai et al., 2022; Zhang et al., 2024).

Appendix B. Approaching the negative orthant

In this appendix we show that any approachability problem can equivalently be written in terms of approaching the negative orthant.

Traditionally (as in Blackwell (1956)), instead of specifying a constraint set \mathcal{U} , an approachability problem specifies a vector-valued bilinear payoff function $\mathbf{u}: \mathcal{P} \times \mathcal{L} \rightarrow \mathbb{R}^k$ along with a set $\mathcal{S} \subseteq \mathbb{R}^k$ that the learner would like to approach. In particular, they would like to minimize the average distance

$$\text{AppDist}(\mathbf{p}, \ell) = \text{dist}_\nu \left(\frac{1}{T} \sum_{t=1}^T \mathbf{u}(p_t, \ell_t), \mathcal{S} \right),$$

where $\text{dist}_\nu(x, \mathcal{S})$ is the minimum distance between x and the set \mathcal{S} under some norm ν . We prove the following result:

Theorem 6 *For any norm ν , set \mathcal{S} , and payoff \mathbf{u} , there exists a corresponding convex set $\mathcal{U} \subseteq \text{Aff}(\mathcal{P} \otimes \mathcal{L}, \mathbb{R})$ such that*

$$\text{AppDist}(\mathbf{p}, \ell) = \text{AppLoss}(\mathbf{p}, \ell).$$

where here AppLoss is understood to mean the approachability loss (1) with respect to the set of constraints \mathcal{U} .

Proof For any r , consider the set of points \mathcal{S}_r within distance r (under ν) of \mathcal{S} . We can write \mathcal{S}_r as $\mathcal{S} + rB$, where B is the unit ball in the norm ν . We can then write \mathcal{S}_r as the intersection of the extremal halfspaces in all directions, where the extremal halfspace in direction v (for any $v \in \mathbb{S}^{d-1}$ on the unit sphere) can be written in the form:

$$\{x \in \mathbb{R}^k \mid \langle v, x \rangle \leq a_v + b_v r\}$$

for some constants a_v and b_v (with $b_v \geq 0$). It follows that for any $x \in \mathbb{R}^k$, we can write

$$\text{dist}_\nu(x, \mathcal{S}) = \max_{v \in \mathbb{S}^{d-1}} \left(\frac{\langle v, x \rangle - a_v}{b_v} \right).$$

If we choose our set \mathcal{U} to contain all functions u_v of the form

$$u_v(p, \ell) = \frac{\langle v, \mathbf{u}(p, \ell) \rangle - a_v}{b_v},$$

(for all $v \in \mathbb{S}^{d-1}$), then it follows that $\text{AppDist}(\mathbf{p}, \ell) = \text{AppLoss}(\mathbf{p}, \ell)$. ■

Appendix C. Non-degenerate counterexample to the reduction

One may object that the approachability instance \mathcal{I} presented in Theorem 3 with $\text{Rate}(\mathcal{I}) = 0$ is somewhat degenerate, as there is a single action by the learner that perfectly approaches the negative orthant. In this appendix we show that this is easily addressed; we can use a similar construction to obtain non-degenerate approachability instances where $\text{Rate}(\mathcal{I}')/\text{Rate}(\mathcal{I})$ is arbitrarily large. The main idea is to embed a non-trivial (but very easy) regret-minimization problem on the “unused” coordinates of the instance.

Theorem 7 *For any $c > 0$, we can construct an approachability instance \mathcal{I} such that $\text{Rate}(\mathcal{I}') > c \cdot \text{Rate}(\mathcal{I})$, where \mathcal{I}' is the regret minimization instance obtained by applying the reduction of Theorem 2 to \mathcal{I} . In particular, there is no finite $c > 0$ for which the reduction in Theorem 2 is a c -approximate reduction for all approachability instances, and the reduction is not tight.*

Proof Consider the following approachability instance \mathcal{I} . Fix any $\varepsilon > 0$, $d_1, d_2 > 1$ and let $d = d_1 + d_2$. We will let $\mathcal{P} = \Delta_d$, $\mathcal{L} = [0, 1]^d$, and \mathcal{U} to be the convex hull of the d bilinear functions u_i , where for $i \in [d_1]$,

$$u_i(p, \ell) = \varepsilon \sum_{j=1}^{d_1} p_j (\ell_j - \ell_i),$$

and for $i \in [d_2]$,

$$u_{d_1+i}(p, \ell) = \sum_{j=d_1+1}^{d_1+d_2} p_j (\ell_j - \ell_{d_1+i}).$$

Here the u_i constraint for $i \in [d_1]$ can be thought of as the regret of moving all probability mass on p_1 through p_{d_1} to p_i (but weighted by ε); similarly, the u_{d_1+i} constraint can be thought of as the regret of moving all probability mass on p_{d_1+1} through p_d to p_{d_1+i} .

We will first show that $\text{Rate}(\mathcal{I}) = O(\varepsilon \sqrt{\log d_1})$. To see this, note that if we ignore the last d_2 coordinates (always assigning weight 0 to them) and run a standard online-learning algorithm over the simplex Δ_{d_1} (e.g., Hedge) it is possible to construct an approachability algorithm \mathcal{A} with $\text{AppLoss}_T(\mathcal{A}) = O(\sqrt{T \log d_1})$.

However, the instance of external regret minimization we reduce to will have a worse rate. Let $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$ be the regret minimization instance formed by applying the reduction of Theorem 2. This instance has $\mathcal{P}' = \mathcal{U}$, $\mathcal{L}' = -(\mathcal{P} \otimes \mathcal{L})$, and \mathcal{U}' the set of all constant functions on \mathcal{P}' . We will restrict the loss set further, and insist that the only losses ℓ'_t are of the form $\ell'_t = -(U_{d_2} \otimes \ell_t)$, where $U_{d_2} = (0, 0, \dots, 0, 1/d_2, 1/d_2, \dots, 1/d_2) \in \mathcal{P}$ is the uniform distribution over the last d_2 coordinates, and ℓ_t is chosen from the subset $\mathcal{L}_2 \subseteq \mathcal{L}$ contains all elements of \mathcal{L} whose first d_1 coordinates equal 0 (so $\mathcal{L}_2 \cong [0, 1]^{d_2}$). Since this restricts the adversary, it only makes the regret minimization problem easier (and the rate smaller).

The regret of a pair of sequences \mathbf{p}' and ℓ' for this new problem can be written as

$$\text{Reg}(\mathbf{p}', \ell') = \max_{x^* \in \mathcal{P}'} \left(\sum_{t=1}^T \langle p'_t, \ell'_t \rangle - \sum_{t=1}^T \langle x^*, \ell'_t \rangle \right). \quad (7)$$

To simplify this further, note that the set $\mathcal{P}' = \mathcal{U}$ is given by the convex hull of the d bilinear functions u_i , so we can write each element of \mathcal{P}' uniquely as a convex combination of this d functions. For a $p'_t \in \mathcal{P}'$, we will write p'_i to be the coefficient of u_i in this convex combination (in this way, we identify \mathcal{P}' with the simplex Δ_d).

Now, for any $p' \in \mathcal{P}'$ and $\ell' \in \mathcal{L}'$ (of the above restricted form), we can write

$$\langle p', \ell' \rangle = - \sum_{i=d_1+1}^{d_1+d_2} p'_i \left(\frac{1}{d_2} \sum_{j=d_1+1}^{d_1+d_2} (\ell_j - \ell_i) \right) = \langle \pi_2(p'), \pi_2(\ell) \rangle - \frac{1}{d_2} \left(\sum_{i=1}^{d_2} \pi_2(p')_i \right) \left(\sum_{i=1}^{d_2} \pi_2(\ell)_i \right), \quad (8)$$

where $\pi_2 : \mathbb{R}^d \rightarrow \mathbb{R}^{d_2}$ is the projection map onto the last d_2 coordinates. We can simplify this even further by introducing the map $\bar{\pi} : \Delta_d \rightarrow \Delta_{d_2}$ defined via

$$\bar{\pi}(p')_i = p'_{d_1+i} + \frac{1}{d_2} \left(1 - \sum_{j=1}^{d_2} p'_{d_1+j} \right). \quad (9)$$

This allows us to rewrite (8) as

$$\langle p', \ell' \rangle = \langle \bar{\pi}(p'), \pi_2(\ell) \rangle - \sum_{j=d_1+1}^{d_2} \ell_j. \quad (10)$$

Substituting this in turn into (7), we have that

$$\text{Reg}(\mathbf{p}', \ell') = \max_{x^* \in \mathcal{P}'} \left(\sum_{t=1}^T \langle \bar{\pi}(p'_t), \pi_2(\ell_t) \rangle - \sum_{t=1}^T \langle \bar{\pi}(x^*), \pi_2(\ell_t) \rangle \right). \quad (11)$$

Now, note that the adversary can choose ℓ_t so that $\pi_2(\ell_t)$ takes on any value in $\mathcal{L}_2 = [0, 1]^{d_2}$. Similarly, $\bar{\pi}(p')$ can take on any value in $\mathcal{P}_2 = \Delta_{d_2}$ as p' ranges over \mathcal{P}' . Therefore, this problem is at least as hard as the online linear optimization problem with action set $\mathcal{P}_2 = \Delta_{d_2}$ and loss set $\mathcal{L}_2 = [0, 1]^{d_2}$. But this is exactly the online learning with experts problem (with d_2 experts), which has a regret lower bound of $\Omega(\sqrt{T \log d_2})$. It follows that $\text{Rate}(\mathcal{I}') \geq \Omega(\sqrt{\log d_2})$, and therefore that $\text{Rate}(\mathcal{I}') / \text{Rate}(\mathcal{I}) > \Omega(\sqrt{(\log d_2) / (\log d_1) / \varepsilon})$. This quantity can be made arbitrarily large (for a fixed d_1 and d_2) as we decrease ε . \blacksquare

Appendix D. Other reductions for approachability and our counter-example

There has been a substantial line of research on the algorithmic problems of Blackwell approachability and ϕ -regret minimization since the reduction of [Abernethy et al. \(2011\)](#). As many of these works also concern themselves with getting better rates for instances of Blackwell approachability, it is natural to wonder if these algorithms somehow side-step our counterexample in Section 3 and produce a rate-preserving reduction to an external regret minimization question.

In this appendix, we survey the main approaches we are aware of (in particular, ones that result in improved rates for general classes of approachability problems), and argue that all of them are susceptible to the counter-examples discussed in Section 3.

- [Shimkin \(2016\)](#) performs a similar reduction from the classical definition of approachability to orthant-approachability that we describe in Appendix B, and then notes that for certain classes of approachability problems (e.g., trying to approach a positively curved set in ℓ_2 distance), the resulting online linear optimization problem can be reparametrized as an online convex optimization problem, allowing for $O(\log T)$ rates by running follow-the-leader. In the language of Appendix B, this means the \mathcal{U} can be parametrized in a way so that the function that sends $u \in \mathcal{U}$ to $u(p, \ell) \in \mathbb{R}$ for some fixed choice of $(p, \ell) \in \mathcal{P} \times \mathcal{L}$ is a strongly convex function. The set \mathcal{U} in our counter-example does not have this property.
- [Kwon \(2021\)](#) studies approachability to sets in norms other than the ℓ_2 norm (obtaining better rates than those naively obtained by bounding the ratio between norms). They perform a similar reduction to an online linear optimization problem to the reduction we present in Appendix B.
- Similar to [Kwon \(2021\)](#), [Dann et al. \(2023\)](#) study the generalization of approachability to arbitrary “pseudo-norms” (norms defined by a not-necessarily-symmetric convex body), showing that these distances can better characterize certain forms of ϕ -regret. They likewise reduce this general form of approachability to an ℓ_∞ -variant of approachability in essentially the same way as we present in Appendix B (indeed, we model our reduction in Appendix B on their reduction)
- [Mannor and Perchet \(2013\)](#) characterize approachability instances where it is possible to obtain $O(1)$ rates for $\text{AppLoss}_T(\mathcal{A})$ instead of $O(\sqrt{T})$ rates (e.g., perfectly approachable instances, where there is an action $p^* \in \mathcal{P}$ which guarantee $u(p^*, \ell) \leq 0$ for all $u \in \mathcal{U}$). Actually, the instance \mathcal{I} we present in Theorem 3 has $\text{Rate}(\mathcal{I}) = 0$ and is therefore one of the perfectly approachable instances of [Mannor and Perchet \(2013\)](#); however, the non-degenerate instance \mathcal{I}' we present in Theorem 7 (with arbitrarily large gaps between rates) has $\text{Rate}(\mathcal{I}') > 0$, and none of the positive results in [Mannor and Perchet \(2013\)](#) apply.
- [Gordon et al. \(2008\)](#) introduced the problem of ϕ -regret minimization and the first algorithms for minimizing it. Their algorithm involves a reduction to online linear optimization over the set Φ in exactly the same way as the reduction we present in Section 3.
- Recently, [Peng and Rubinstein \(2023\)](#) and [Dagan et al. \(2023\)](#) have presented new algorithms for minimizing swap regret that obtain much better dependence on the total number of actions (at the cost of a poorer $T/\text{poly}(\log T)$ dependence on T). These algorithms are the only algorithms we are aware of for a non-external-regret minimization problem that cannot be viewed through the reduction we present in Section 3 (i.e., via providing a learning algorithm for the resulting OLO problem). However, the increased dependence on T makes these algorithms inapplicable in the settings we examine.

Appendix E. General linear equivalences

In this appendix, we provide a more general definition of linear equivalence between two approachability instances, capturing the idea that a linear equivalence is a tight reduction between two regret minimization instances where the reduction is given entirely by (perhaps invertible) linear transformations.

Formally, let $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ be a regret minimization instance with $\dim(\mathcal{P}) = \dim(\mathcal{L}) = d$. We say that there is a *linear equivalence* between this instance and the instance $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$ if there exists a bijective correspondence between $\phi \in \Phi$ and $\phi' \in \Phi'$ such that the following conditions hold:

- There exist affine maps $S_{\mathcal{P}', \mathcal{P}} \in \text{Aff}(\mathcal{P}', \mathcal{P})$ and $S_{\mathcal{L}, \mathcal{L}'} \in \text{Aff}(\mathcal{L}, \mathcal{L}')$, such that for any $\ell \in \mathcal{L}$, $p' \in \mathcal{P}'$, $\ell' = S_{\mathcal{L}, \mathcal{L}'}\ell$, and $p = S_{\mathcal{P}', \mathcal{P}}p'$, we have:

$$\langle p - \phi(p), \ell \rangle = \langle p' - \phi'(p'), \ell' \rangle. \quad (12)$$

- There exist affine maps $S_{\mathcal{P}, \mathcal{P}'} \in \text{Aff}(\mathcal{P}, \mathcal{P}')$ and $S_{\mathcal{L}', \mathcal{L}} \in \text{Aff}(\mathcal{L}', \mathcal{L})$ such that for any $\ell' \in \mathcal{L}'$, $p \in \mathcal{P}$, $\ell = S_{\mathcal{L}', \mathcal{L}}\ell'$ and $p' = S_{\mathcal{P}, \mathcal{P}'}p$, we have:

$$\langle p - \phi(p), \ell \rangle = \langle p' - \phi'(p'), \ell' \rangle. \quad (13)$$

Note that the first point above has the following implication. If you have a learning algorithm \mathcal{A}' for \mathcal{I}' , you can use it to construct a learning algorithm \mathcal{A} for \mathcal{I} by following the procedure:

1. At the beginning of round t , \mathcal{A} asks \mathcal{A}' for the p'_t it will play.
2. \mathcal{A} then plays $p_t = S_{\mathcal{P}', \mathcal{P}}p'_t$.
3. \mathcal{A} observes their loss vector ℓ_t .
4. \mathcal{A} passes along the loss vector $\ell'_t = S_{\mathcal{L}, \mathcal{L}'}\ell_t$ to \mathcal{A}' .

The guarantee of equation (12) implies that $\text{Reg}_T(\mathcal{A}) \leq \text{Reg}_T(\mathcal{A}')$. Likewise, the guarantee of equation (13) (and its surrounding bullet) implies that given an algorithm \mathcal{A} for \mathcal{I} , we can efficiently construct an algorithm \mathcal{A}' for \mathcal{I}' such that $\text{Reg}_T(\mathcal{A}') \leq \text{Reg}_T(\mathcal{A})$. Together, this shows that a linear equivalence between \mathcal{I} and \mathcal{I}' implies a tight reduction between \mathcal{I} and \mathcal{I}' .

E.1. Simplifying linear equivalences

We introduce a handful of simplifications to the definition of linear equivalences which will make them easier to characterize and line up with the definition we work with in the main paper. In particular, note that our linear equivalence in Theorem 14 (for weighted regret) involved only a single invertible linear transformation. We will ultimately show that, under some mild assumptions, it suffices to consider linear equivalences specified by only a single invertible linear transformation.

First, we will assume (possibly by augmenting the sets \mathcal{P} , \mathcal{L} , and \mathcal{L}' with a fixed extra coordinate) that it is possible to express any affine map over any of these sets as a direct linear transformation, and henceforth restrict our attention to purely linear maps (for both ϕ and $S_{*,*}$). We will also let M_ϕ denote the linear transformation $\text{Id} - \phi$; note that this lets us rewrite the regret term $\langle p, \ell \rangle - \langle \phi(p), \ell \rangle$ in the form $\langle M_\phi p, \ell \rangle$. This also lets us write equations (12) and (13) in the more compact forms

$$\langle M_\phi S_{\mathcal{P}', \mathcal{P}} p', \ell \rangle = \langle M_{\phi'} p', S_{\mathcal{L}, \mathcal{L}'} \ell' \rangle, \quad (14)$$

and

$$\langle M_\phi p, S_{\mathcal{L}', \mathcal{L}} \ell' \rangle = \langle M_{\phi'} S_{\mathcal{P}, \mathcal{P}'} p, \ell' \rangle. \quad (15)$$

We now say that a regret minimization instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ is *minimal* if the following conditions hold:

- There does not exist a convex subset \mathcal{P}' of \mathcal{P} such that $\text{Rate}(\mathcal{P}', \mathcal{L}, \Phi) = \text{Rate}(\mathcal{P}, \mathcal{L}, \Phi)$.
- There does not exist a convex subset \mathcal{L}' of \mathcal{L} such that $\text{Rate}(\mathcal{P}, \mathcal{L}', \Phi) = \text{Rate}(\mathcal{P}, \mathcal{L}, \Phi)$.

Intuitively, the first constraint captures the property that every (extremal) action in \mathcal{P} should be useful for the learner; it should be impossible to achieve the same regret bound by only playing a subset of the actions. Similarly, the second constraint captures an analogous property for the adversary – it should be the case that every extremal loss in \mathcal{L} is useful for constructing an optimal lower bound.

One advantage of working with minimal regret minimization instances is that reductions between minimal regret minimization instances are specified by invertible linear transformations.

Lemma 8 *Let \mathcal{I} and \mathcal{I}' be two minimal regret minimization instances. If there is a linear equivalence between \mathcal{I} and \mathcal{I}' , then the maps $S_{\mathcal{P}', \mathcal{P}}$, $S_{\mathcal{P}, \mathcal{P}'}$, $S_{\mathcal{L}, \mathcal{L}'}$ and $S_{\mathcal{L}', \mathcal{L}}$ must all be bijective linear transformations between the two sets.*

Proof We will first show that $S_{\mathcal{L}, \mathcal{L}'}$ must be surjective. Consider the regret minimization instance $\mathcal{I}'_L = (\mathcal{P}', S_{\mathcal{L}, \mathcal{L}'}\mathcal{L}, \Phi')$, whose loss set $S_{\mathcal{L}, \mathcal{L}'}\mathcal{L}$ contains all losses of the form $S_{\mathcal{L}, \mathcal{L}'}\ell$ for $\ell \in \mathcal{L}$. Note that the reduction from \mathcal{I}' to \mathcal{I} described at the end of Section E only passes loss vectors in $S_{\mathcal{L}, \mathcal{L}'}\mathcal{L}$ to algorithm \mathcal{A}' , and therefore shows that $\text{Rate}(\mathcal{I}) \leq \text{Rate}(\mathcal{I}'_L)$. Since $S_{\mathcal{L}, \mathcal{L}'}\mathcal{L} \subseteq \mathcal{L}'$, we in turn have that $\text{Rate}(\mathcal{I}'_L) \leq \text{Rate}(\mathcal{I}')$. But finally, since there exists a linear equivalence between \mathcal{I} and \mathcal{I}' , we have $\text{Rate}(\mathcal{I}) = \text{Rate}(\mathcal{I}')$, and therefore that $\text{Rate}(\mathcal{I}'_L) = \text{Rate}(\mathcal{I}')$.

Now, if $S_{\mathcal{L}, \mathcal{L}'}\mathcal{L}$ were a strict subset of \mathcal{L}' , this would violate the assumption that \mathcal{I}' is minimal. It follows that we must have $S_{\mathcal{L}, \mathcal{L}'}\mathcal{L} = \mathcal{L}'$, and therefore that $S_{\mathcal{L}, \mathcal{L}'}$ is surjective. By symmetry, the linear transformation $S_{\mathcal{L}', \mathcal{L}}$ is surjective onto \mathcal{L} . These two facts imply that the dimensions of the convex sets \mathcal{L} and \mathcal{L}' must be equal, and the two linear transformations $S_{\mathcal{L}, \mathcal{L}'}$ and $S_{\mathcal{L}', \mathcal{L}}$ must in fact be bijective linear transformations between \mathcal{L} and \mathcal{L}' .

Similarly, to show that $S_{\mathcal{P}', \mathcal{P}}$ is surjective, consider the regret minimization instance $\mathcal{I}_P = (S_{\mathcal{P}', \mathcal{P}}\mathcal{P}', \mathcal{L}, \Phi)$. Since in the reduction from \mathcal{I}' to \mathcal{I} , \mathcal{A} plays only actions in $S_{\mathcal{P}', \mathcal{P}}\mathcal{P}'$, this reduction actually shows that $\text{Rate}(\mathcal{I}_P) \leq \text{Rate}(\mathcal{I}')$. But since $S_{\mathcal{P}', \mathcal{P}}\mathcal{P}' \subseteq \mathcal{P}$, we also have $\text{Rate}(\mathcal{I}) \leq \text{Rate}(\mathcal{I}_P)$. Finally, since $\text{Rate}(\mathcal{I}) = \text{Rate}(\mathcal{I}')$ (by the linear equivalence), all three of these rates must be equal. This means that if $S_{\mathcal{P}', \mathcal{P}}\mathcal{P}'$ were a strict subset of \mathcal{P} , this would violate the minimality of \mathcal{P} . Therefore $S_{\mathcal{P}', \mathcal{P}}$ is surjective onto \mathcal{P} , $S_{\mathcal{P}, \mathcal{P}'}$ is likewise surjective onto \mathcal{P}' , and both transformations must be bijective linear transformations. \blacksquare

When the transformation matrices are guaranteed to be bijective transformations between the sets (as in Lemma 8), we can specify linear equivalences more succinctly. In particular, we can always without loss of generality take $S_{\mathcal{P}', \mathcal{P}} = S_{\mathcal{P}, \mathcal{P}'}^{-1}$ and $S_{\mathcal{L}, \mathcal{L}'} = S_{\mathcal{L}', \mathcal{L}}^{-1}$. The following Lemma shows that for the sake of classifying different regret minimization problems, we can always take $\mathcal{P}' = \mathcal{P}$ and $S_{\mathcal{P}, \mathcal{P}'} = \text{Id}$ (as we did in our reduction in Section 4.1).

Lemma 9 *Let $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ and $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$ be two linearly equivalent minimal regret minimization instances, related via $\mathcal{P}' = S_{\mathcal{P}, \mathcal{P}'}\mathcal{P}$ and $\mathcal{L}' = S_{\mathcal{L}, \mathcal{L}'}\mathcal{L}$ for invertible linear transformations $S_{\mathcal{P}, \mathcal{P}'}$ and $S_{\mathcal{L}, \mathcal{L}'}$. Then \mathcal{I} is also linearly equivalent to the regret minimization instance*

$\mathcal{I}'' = (\mathcal{P}'', \mathcal{L}'', \Phi'')$ where:

$$\begin{aligned}\mathcal{P}'' &= \mathcal{P} \\ \mathcal{L}'' &= S_{\mathcal{P}, \mathcal{P}'}^T S_{\mathcal{L}, \mathcal{L}'} \mathcal{L} \\ M_{\phi''} &= S_{\mathcal{P}, \mathcal{P}'}^{-1} M_{\phi'} S_{\mathcal{P}, \mathcal{P}'}.\end{aligned}$$

(here in the last line, ϕ'' is the element of Φ'' corresponding to ϕ' in Φ' and ϕ in Φ). Moreover, if \mathcal{I}' is a proper ϕ -regret minimization instance, so is \mathcal{I}'' ; similarly, if \mathcal{I}' is an external regret minimization instance, so is \mathcal{I}'' .

Proof We can verify that the necessary conditions for a linear equivalence (equations (14) and (15)) hold for the above equivalence defined between \mathcal{I} and \mathcal{I}'' . Since all transformations are invertible, it suffices to verify the single equation

$$\langle M_{\phi} p, \ell \rangle = \langle M_{\phi''} S_{\mathcal{P}, \mathcal{P}'} p, S_{\mathcal{L}, \mathcal{L}'} \ell \rangle.$$

Substituting in $S_{\mathcal{P}, \mathcal{P}'} = \text{Id}$, $S_{\mathcal{L}, \mathcal{L}'} = S_{\mathcal{P}, \mathcal{P}'}^T S_{\mathcal{L}, \mathcal{L}'}$, and $M_{\phi''} = S_{\mathcal{P}, \mathcal{P}'}^{-1} M_{\phi'} S_{\mathcal{P}, \mathcal{P}'}$ the RHS of the above expression becomes

$$\langle S_{\mathcal{P}, \mathcal{P}'}^{-1} M_{\phi'} S_{\mathcal{P}, \mathcal{P}'} p, S_{\mathcal{P}, \mathcal{P}'}^T S_{\mathcal{L}, \mathcal{L}'} \ell \rangle = \langle M_{\phi'} S_{\mathcal{P}, \mathcal{P}'} p, S_{\mathcal{L}, \mathcal{L}'} \ell \rangle.$$

The RHS of this expression is in turn equal to $\langle M_{\phi} p, \ell \rangle$ by the guarantee of the original linear equivalence, as desired. To check that the property of being a proper ϕ -regret minimization instance / external regret minimization instance is preserved, note that from the definition of $M_{\phi''}$ we can read off:

$$\phi''(p) = S_{\mathcal{P}, \mathcal{P}'}^{-1} \phi'(S_{\mathcal{P}, \mathcal{P}'} p) = S_{\mathcal{P}, \mathcal{P}'}^{-1} \phi'(p').$$

(defining $p' = S_{\mathcal{P}, \mathcal{P}'} p$ in the last line). Now, if ϕ' is proper, $\phi'(p') \in \mathcal{P}'$, and therefore $S_{\mathcal{P}, \mathcal{P}'}^{-1} \phi'(p) \in \mathcal{P}$, and it follows that ϕ'' is proper. Similarly, if $\phi'(p')$ is constant over $p' \in \mathcal{P}'$, then $\phi''(p)$ is also constant over $p \in \mathcal{P}$. \blacksquare

Inspired by Lemma 9, we can now restrict ourselves to reductions from $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ to instances of the form $\mathcal{I}' = (\mathcal{P}, (S^T)^{-1} \mathcal{L}, \Phi')$, where S is an invertible linear transformation, where $S_{\mathcal{P}, \mathcal{P}'} = \text{Id}$, $S_{\mathcal{L}', \mathcal{L}} = S^T$, and $S_{\mathcal{L}, \mathcal{L}'} = (S^T)^{-1}$. For such a reduction, the ϕ' corresponding to ϕ must satisfy

$$\langle M_{\phi} p, S^T \ell' \rangle = \langle M_{\phi'} p, \ell' \rangle, \quad (16)$$

for all $\ell' \in \mathcal{L}'$ and $p \in \mathcal{P}$ (this follows from substituting in our specific transformations into (14)). We can rewrite this in the form:

$$\langle (S M_{\phi} - M_{\phi'}) p, \ell' \rangle = 0$$

for all $p \in \mathcal{P}$ and $\ell' \in \mathcal{L}'$. Since we assume that $\text{span}(\mathcal{P}) = \text{span}(\mathcal{L}') = \mathbb{R}^d$, this implies that

$$M_{\phi'} = S M_{\phi} \quad (17)$$

and in particular that

$$\phi'(p) = p + S(\phi(p) - p). \quad (18)$$

Note that this now coincides with the definition introduced in Section 4.

Appendix F. Reducing improper ϕ -regret to proper ϕ -regret

Theorem 5 shows that the property of being equivalent to an external regret minimization problem is rather stringent: only very structured improper ϕ -regret minimization problems (and hence, approachability instances) can be reduced via linear transformations to such instances.

However, the class of proper ϕ -regret minimization problems is far broader than the class of external regret minimization problems. Notably, every regret minimization problem we have considered thus far either is or is linearly equivalent to a proper ϕ -regret minimization instance. This raises the natural question of whether every improper ϕ -regret minimization instance can be linearly transformed into a proper one.

In this section we show that the answer to this question is no. In Section F.1 we give some examples of “atypical” improper ϕ -regret minimization problems that we prove are not linearly equivalent to any proper ϕ -regret minimization problems.

Given this, we can also ask whether (in the same vein as Theorem 5) we can cleanly characterize the set of regret-minimization problems which are equivalent to proper ϕ -regret minimization. In Section F.2 we give an algorithmic procedure for deciding whether a given regret-minimization instance can be reduced to a proper one in the case where the sets \mathcal{P} , \mathcal{L} , and Φ are all polytopes.

F.1. Irreducible improper instances

In this section we provide two (classes of) examples of improper ϕ -regret minimization problems that cannot be linearly reduced to a proper ϕ -regret minimization problem. These two examples will each illustrate different obstructions to the property of being reducible.

Throughout this section, we will assume that $\mathcal{P} = \Delta_d$ and $\mathcal{L} = [0, 1]^d$, and only vary the set Φ . Before we introduce our first class of examples, we will prove the following lemma, which shows that if the instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ is linearly equivalent to a proper ϕ -regret minimization problem, then all transformations in Φ must share a left eigenvector of eigenvalue 1.

Lemma 10 *Assume \mathcal{P} is contained within an affine subspace of \mathbb{R}^d . If $(\mathcal{P}, \mathcal{L}, \Phi)$ is linearly equivalent to an instance of proper ϕ -regret minimization, then there exists a non-zero $v \in \mathbb{R}^d$ such that $\langle \phi(p), v \rangle = \langle p, v \rangle$ for any $\phi \in \Phi$ and $p \in \mathcal{P}$.*

Proof By (18), this linear equivalence sends $\phi(p)$ to the function $\phi'(p) = p + S(\phi(p) - p)$, for some invertible linear transformation S . Since the target regret minimization instance is proper, $\phi'(p)$ is guaranteed to lie in \mathcal{P} .

Now, since \mathcal{P} is contained within an affine subspace of \mathbb{R}^d , there exists some $w \in \mathbb{R}^d$ and $b \in \mathbb{R}$ such that $\langle p, w \rangle = b$ for all $p \in \mathcal{P}$. This implies that

$$\langle \phi'(p), w \rangle = \langle p, w \rangle + \langle S(\phi(p) - p), w \rangle.$$

Since $\langle \phi'(p), w \rangle = \langle p, w \rangle = b$, this reduces to $\langle S(\phi(p) - p), w \rangle = 0$, which in turn can be written as $\langle \phi(p) - p, S^T w \rangle = 0$. Therefore the vector $S^T w$ satisfies the constraints for v in the lemma statement. \blacksquare

Note that since the simplex Δ_d is contained within the affine hyperplane $\sum_i p_i = 1$, Lemma 10 applies for all examples in this section. Also, we can alternatively think of the statement of Lemma 10 as stating that the matrices M_ϕ must all share a non-trivial left kernel element (an element v such

that $v^T M_\phi = 0$). To construct an irreducible instance, it suffices to find examples of convex sets of matrices that do not all share the same left kernel element.

Constructing such a set is made more difficult by the fact that each $\phi \in \Phi$ must have a fixed point in \mathcal{P} , which translates to the fact that every matrix M_ϕ must have a non-trivial right kernel element p_ϕ whose entries are all non-negative (any such element can be scaled to lie in $\mathcal{P} = \Delta_d$). We would also like the different M_ϕ to not all share the same right kernel element (because in that case each $\phi \in \Phi$ would have the same fixed point, and the instance would not be minimal).

We can construct such a set by noticing that the set of skew-symmetric matrices of odd dimension is a linear subspace of the set of matrices. In particular, if for some $a, b, c \in \mathbb{R}_{\geq 0}$ we consider the skew-symmetric matrix

$$M_{\phi_{a,b,c}} = \begin{pmatrix} 0 & a & -b \\ -a & 0 & c \\ b & -c & 0 \end{pmatrix},$$

then the vector $v = (c, b, a)$ belongs to both the left-kernel and right-kernel of $M_{\phi_{a,b,c}}$. In particular, this means that if we take $\Phi = \text{conv}(\phi_{1,0,0}, \phi_{0,1,0}, \phi_{0,0,1})$ where:

$$\begin{aligned} \phi_{1,0,0}(p_1, p_2, p_3) &= (p_1 - p_2, p_1 + p_2, p_3) \\ \phi_{0,1,0}(p_1, p_2, p_3) &= (p_1 + p_3, p_2, p_3 - p_1) \\ \phi_{0,0,1}(p_1, p_2, p_3) &= (p_1, p_2 - p_3, p_2 + p_3), \end{aligned}$$

then $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$ is a valid improper ϕ -regret minimization instance which, by Lemma 10, is *not* linearly equivalent to a proper ϕ -regret minimization instance.

One might wonder whether Lemma 10 is the sole obstacle to linear equivalence to proper ϕ -regret minimization. Interestingly, this is not the case. We now give a second Φ that does not violate Lemma 10 (i.e., all M_ϕ share a left-kernel element), but where the instance $(\mathcal{P}, \mathcal{L}, \Phi)$ is still provably not linearly equivalent to proper ϕ -regret minimization. To do so, we will rely on the following obstruction.

Lemma 11 *Consider an improper ϕ -regret minimization instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \Phi)$. If there exists an extreme point x of \mathcal{P} and $\phi_1, \phi_2 \in \Phi$ with the property that*

$$\phi_2(x) - x = -\alpha(\phi_1(x) - x) \neq 0$$

for some real $\alpha > 0$, then \mathcal{I} is not linearly equivalent to a proper ϕ -regret minimization instance.

Proof If such a linear equivalence existed, by (18), we would have that $\phi'_1(x) = x + S(\phi_1(x) - x)$ and $\phi'_2(x) = x + S(\phi_2(x) - x) = x - \alpha S(\phi_1(x) - x)$. But since x is an extreme point of \mathcal{P} , it is impossible for both $x + w$ and $x - \alpha w$ to both lie in \mathcal{P} for a non-zero vector w . ■

We now present a valid improper ϕ -regret minimization instance (found by computer search) where Lemma 11 applies but Lemma 10 does not. This example is parametrized by the two matrices:

$$A = \begin{pmatrix} -6 & 8 & -9 \\ 2 & -1 & -9 \\ 4 & -7 & 18 \end{pmatrix}, B = \begin{pmatrix} 10 & -3 & -7 \\ 6 & -6 & 10 \\ -16 & 9 & -3 \end{pmatrix}.$$

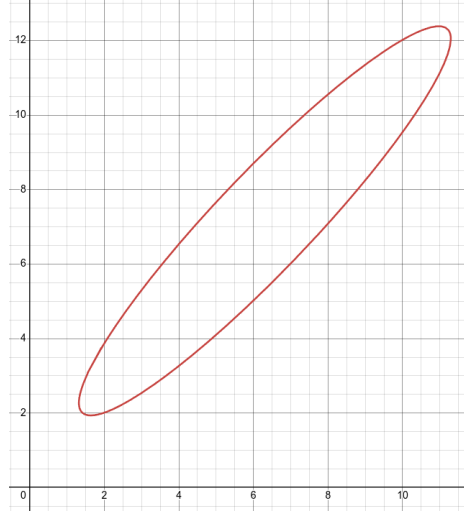


Figure 4: For any t , the matrix $M(t) = A + tB$ has a right kernel element of the form $(x(t), y(t), 1)$, where $x(t) = \frac{72t^2 - 46t + 81}{2(21t^2 + 8t + 5)}$ and $y(t) = \frac{72t^2 - 41t + 36}{21t^2 + 8t + 5}$. The above diagram plots $(x(t), y(t))$ for all $t \in \mathbb{R}$, and shows that this right kernel element always has non-negative entries (and hence has a multiple that lies in Δ_3).

Note that both A and B share the same left-kernel element $(1, 1, 1)$ (so they both map the hyperplane containing \mathcal{P} to itself). We will consider the set of regret functions $\Phi = \{\text{Id} + \gamma_a A + \gamma_b B \mid \gamma_a \in [-1, 1], \gamma_b \in [-1, 1]\}$. One can also verify computationally that all elements of Φ contain a fixed point within the simplex Δ_3 and that this fixed point is not static and changes depending on γ_b/γ_a (see Figure 4).

But now, consider the two functions $\phi_1(p) = p + Ap$ and $\phi_2(p) = p - Ap$ both belonging to Φ . For any $x \in \mathcal{P}$ it is the case that $\phi_1(x) - x = Ax = -(\phi_2(x) - x)$. Choosing x to be the extreme point $(1, 0, 0)$, Ax is non-zero and thus by Lemma 10 this instance is not equivalent to a proper ϕ -regret minimization instance.

F.2. An algorithmic characterization

Finally, we will show how to algorithmically decide whether a given regret minimization instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \mathcal{I})$ is linearly equivalent to a proper ϕ -regret minimization problem via solving an appropriate linear program for the transformation S . For simplicity, we will only handle the case where \mathcal{P} and \mathcal{I} are both polytopes, with \mathcal{P} being the convex hull of the N vertices p_1, p_2, \dots, p_N and \mathcal{I} being the convex hull of the M regret functions $\phi_1, \phi_2, \dots, \phi_M$.

Recall that (by (18)) we have $\phi'_i(p) = p + S(\phi(p) - p)$. Note that ϕ' is a proper ϕ -function if and only if $\phi'_i(p_j)$ lies within \mathcal{P} for all vertices p_j (by convexity, this implies that $\phi'_i(p) \in \mathcal{P}$ for any other $p \in \mathcal{P}$). So there exists a reduction iff there exists an invertible S such that for every $i \in [M]$ and $j \in [N]$, we have the constraint

$$p_j + S(\phi_i(p_j) - p_j) \in \mathcal{P}. \quad (19)$$

We can in turn rephrase (19) as several linear constraints on the entries of the matrix S . In particular, if we introduce the auxiliary variables $\lambda_{i,j,k}$ (for $k \in [N]$) we can rewrite the set of

constraints expressed by (19) as the following linear program (whose variables are the d^2 entries of S and the $\lambda_{i,j,k}$):

$$\begin{aligned} p_j + S(\phi_i(p_j) - p_j) &= \sum_{k=1}^N \lambda_{i,j,k} p_k \text{ for all } i \in [M], j \in [N] \\ \sum_{k=1}^N \lambda_{i,j,k} &= 1 \text{ for all } i \in [M], j \in [N] \\ \lambda_{i,j,k} &\geq 0 \text{ for all } i \in [M], j, k \in [N] \end{aligned}$$

The above constraints define a convex cone of possible values of S which we denote by \mathcal{S} . For example, this cone always contains the solution $S = 0$; we would like to now decide whether it contains any invertible matrices. Luckily, this is straightforward to do in a randomized manner by the following lemma.

Lemma 12 *Given any convex set $\mathcal{S} \subseteq \mathbb{R}^{d \times d}$ of d -by- d matrices, either every matrix in $\text{span}(\mathcal{S})$ is non-invertible, or almost all⁷ matrices in $\text{span}(\mathcal{S})$ are invertible.*

Proof Given any linear subspace of matrices $V \subseteq \mathbb{R}^{d \times d}$, the constraint $\det(M) = 0$ defines an algebraic variety on this space. Any algebraic variety in a Euclidean space either has measure zero or is equal to the entire space. ■

Given this, it suffices to simply test whether a random element in $\text{span}(\mathcal{S})$ is invertible. We can efficiently generate a basis of $\text{span}(\mathcal{S})$ by repeatedly solving the linear program above (finding the extreme points in \mathcal{S} in a direction orthogonal to all basis elements found so far), and once we have such a basis, we can test a random linear combination of the basis elements for invertibility.

This entire procedure takes time polynomial in N , M , and d . It can also likely be extended beyond polytopes to any case where we have a membership oracle for \mathcal{S} (i.e., a procedure that decides whether (19) is satisfied for all $\phi \in \Phi$ and $p \in \mathcal{P}$), modulo numerical issues with testing invertibility. Note also that this procedure not only checks whether a given instance is reducible, but also provides a valid transformation S in the case that it is.

Appendix G. Omitted Proofs

G.1. Response satisfiability is equivalent to half-space satisfiability

Theorem 13 *Let $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ be an approachability instance. The following two conditions on this instance are equivalent:*

- **Response satisfiability:** *For every $\ell \in \mathcal{L}$, there exists a $p \in \mathcal{P}$ such that $u(p, \ell) \leq 0$ for all $u \in \mathcal{U}$.*
- **Half-space satisfiability:** *For every $u \in \mathcal{U}$, there exists a $p \in \mathcal{P}$ such that $u(p, \ell) \leq 0$ for all $\ell \in \mathcal{L}$.*

⁷ All but a measure zero subset in the Euclidean measure of $\text{span}(\mathcal{S})$.

Proof If an instance is response satisfiable, then this means that

$$\max_{u \in \mathcal{U}} \min_{p \in \mathcal{P}} \max_{\ell \in \mathcal{L}} u(p, \ell) \leq 0.$$

Since \mathcal{P} and \mathcal{L} are both convex, applying the minimax theorem to the inner min-max optimization problem, this implies that

$$\max_{u \in \mathcal{U}} \max_{\ell \in \mathcal{L}} \min_{p \in \mathcal{P}} u(p, \ell) \leq 0.$$

Interchanging the order of the two outer maxima, this in turn implies

$$\max_{\ell \in \mathcal{L}} \max_{u \in \mathcal{U}} \min_{p \in \mathcal{P}} u(p, \ell) \leq 0.$$

Finally, applying the minimax theorem to the inner min-max optimization problem once again, we have that

$$\max_{\ell \in \mathcal{L}} \min_{p \in \mathcal{P}} \max_{u \in \mathcal{U}} u(p, \ell) \leq 0,$$

which implies that the instance is half-space satisfiable. The opposite implication can be proved symmetrically. \blacksquare

G.2. Proof of Theorem 1

Proof To prove the forward direction, we will prove the corresponding approachability problem (defined via (4)) is approachable. For this, it suffices to show (via the “half-space satisfiable” definition of approachable) that for any $\phi \in \Phi$, there exists a $p \in \mathcal{P}$ such that $u_\phi(p, \ell) \leq 0$ for all $\ell \in \mathcal{L}$. But since $u_\phi(p, \ell) = \langle p - \phi(p), \ell \rangle$, if we take $p = p_\phi$, $p_\phi - \phi(p_\phi) = 0$ and $u_\phi(p, \ell) = 0$ for all $\ell \in \mathcal{L}$.

Conversely, assume there exists a $\phi \in \Phi$ without a fixed point. Consider the set $V = \{p - \phi(p) \mid p \in \mathcal{P}\}$. V is a convex bounded set (since it is a linear transformation of \mathcal{P} that by assumption does not contain 0). Therefore, there exists a hyperplane separating V from 0 – in particular, there exists a w such that $\langle v, w \rangle > 0$ for any $v \in V$. If we take ℓ proportional to w (possible since $\text{cone}(\mathcal{L}) = \mathbb{R}^d$), then for this ϕ and ℓ , there is no $p \in \mathcal{P}$ such that $u_\phi(p, \ell) \leq 0$. It follows that the corresponding approachability problem is not approachable, and therefore $\text{Rate}(\mathcal{P}, \mathcal{L}, \Phi) = \infty$. \blacksquare

G.3. Proof of Theorem 2

Proof We will define the target instance \mathcal{I}' as follows. We will let the new action set $\mathcal{P}' = \mathcal{U}$ and the new loss set $\mathcal{L}' = -(\mathcal{P} \otimes \mathcal{L})$. Note that both \mathcal{P}' and \mathcal{L}' are convex subsets of $(\dim \mathcal{P})(\dim \mathcal{L})$ -dimensional Euclidean space, and we can define a bilinear pairing between \mathcal{P}' and \mathcal{L}' via $\langle u, -(p \otimes \ell) \rangle = -u(p, \ell)$ for any $u \in \mathcal{P}'$ and $-(p \otimes \ell) \in \mathcal{L}'$. Finally, we will let the new benchmark set \mathcal{U}' consist of all constant functions over \mathcal{P}' (i.e., for each $x \in \mathcal{P}'$, there will exist a $u'_x \in \mathcal{U}'$ such that $u'_x(p') = x$ for all $p' \in \mathcal{P}'$). In other words, \mathcal{I}' is the online linear optimization problem with action set \mathcal{P}' and loss set \mathcal{L}' (and is an external regret minimization problem).

Given a low-regret algorithm \mathcal{B} for solving \mathcal{I}' , we construct the following algorithm \mathcal{A} for solving \mathcal{I} using \mathcal{B} as a black box. In round t :

1. Set $p'_t = \mathcal{B}_t(\ell'_1, \ell'_2, \dots, \ell'_{t-1}) \in \mathcal{P}'$.
2. Play $p_t \in \mathcal{P}$ s.t. $\langle p'_t, p_t \otimes \ell \rangle \leq 0$ for all $\ell \in \mathcal{L}$. Note that such a p_t exists since the instance $(\mathcal{P}, \mathcal{L}, \mathcal{U})$ is approachable – in particular, $\langle p'_t, p_t \otimes \ell \rangle = p'_t(p_t, \ell)$ for some $p'_t \in \mathcal{U}$, and the approachability condition implies that there must exist a p_t where $p'_t(p_t, \ell) \leq 0$ for all $\ell \in \mathcal{L}$.
3. Receive the loss ℓ_t .
4. Set $\ell'_t = -(p_t \otimes \ell_t) \in \mathcal{L}'$.

By the definition of $\text{Reg}(\mathbf{p}', \ell')$, we have

$$\begin{aligned}
 \text{Reg}(\mathbf{p}', \ell') &= \sum_{t=1}^T \langle p'_t, \ell'_t \rangle - \min_{x^* \in \mathcal{P}'} \sum_{t=1}^T \langle x^*, \ell'_t \rangle \\
 &= \sum_{t=1}^T \langle p'_t, -(p_t \otimes \ell_t) \rangle - \min_{u^* \in \mathcal{U}} \sum_{t=1}^T \langle u^*, -(p_t \otimes \ell_t) \rangle \\
 &= \max_{u^* \in \mathcal{U}} \sum_{t=1}^T u^*(p_t, \ell_t) - \sum_{t=1}^T p'_t(p_t, \ell_t) \\
 &= \text{AppLoss}(\mathbf{p}, \ell) - \sum_{t=1}^T p'_t(p_t, \ell_t) \geq \text{AppLoss}(\mathbf{p}, \ell).
 \end{aligned}$$

This very final inequality follows from the fact that $p'_t(p_t, \ell) \leq 0$ for all $\ell \in \mathcal{L}$ via our choice of p_t in step 2. Altogether, this analysis implies that $\text{AppLoss}(\mathcal{A}) \leq \text{Reg}(\mathcal{B})$, and therefore provides a weak reduction from \mathcal{I} to \mathcal{I}' . \blacksquare

G.4. Proof of Theorem 3

Proof Consider the following approachability instance \mathcal{I} . Fix any $d' \geq 2$ and let $d = d' + 1$. We will let $\mathcal{P} = \Delta_d$, $\mathcal{L} = [0, 1]^d$, and \mathcal{U} to be the convex hull of the d' bilinear functions u_i , where for $i \in [d']$,

$$u_i(p, \ell) = \sum_{j=1}^{d'} p_j(\ell_j - \ell_i).$$

Here the u_i constraint for $i \in [d']$ can be thought of as the regret of moving all probability mass on p_1 through $p_{d'}$ to p_i . Note, however, that the learner also has the option of placing probability mass on action $d' + 1$, which has no approachability constraint associated with this. As a result, $\text{Rate}(\mathcal{I}) = 0$, since the learner can perfectly approach the negative orthant by always playing $p_t = e_{d'+1}$.

However, the instance of external regret minimization we reduce to will have a worse rate, as it will require solving a genuine regret minimization problem. Let $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$ be the regret minimization instance formed by applying the reduction of Theorem 2. This instance has $\mathcal{P}' = \mathcal{U}$, $\mathcal{L}' = -(\mathcal{P} \otimes \mathcal{L})$, and \mathcal{U}' the set of all constant functions on \mathcal{P}' . We will restrict the loss set further, and insist that the only losses ℓ'_t are of the form $\ell'_t = -(U_{d'} \otimes \ell_t)$, where $U_{d'} = (1/d', 1/d', \dots, 1/d', 0) \in \mathcal{P}$ is the uniform distribution over the first d' coordinates, and ℓ_t is chosen from the subset $\mathcal{L}_2 \subseteq \mathcal{L}$ containing all elements of \mathcal{L} whose last coordinate equals 0 (so $\mathcal{L}_2 \cong [0, 1]^{d'}$). Since this restricts the adversary, it only makes the regret minimization problem easier (and the rate smaller).

The regret of a pair of sequences \mathbf{p}' and ℓ' for this new problem can be written as

$$\text{Reg}(\mathbf{p}', \ell') = \max_{x^* \in \mathcal{P}'} \left(\sum_{t=1}^T \langle p'_t, \ell'_t \rangle - \sum_{t=1}^T \langle x^*, \ell'_t \rangle \right). \quad (20)$$

To simplify this further, note that the set $\mathcal{P}' = \mathcal{U}$ is given by the convex hull of the d' bilinear functions u_i , so we can write each element of \mathcal{P}' uniquely as a convex combination of these d' functions. For a $p' \in \mathcal{P}'$, we will write p'_i to be the coefficient of u_i in this convex combination (in this way, we identify \mathcal{P}' with the simplex $\Delta_{d'}$).

Now, for any $p' \in \mathcal{P}'$ and $\ell' \in \mathcal{L}'$ (of the above restricted form), we can write

$$\langle p', \ell' \rangle = - \sum_{i=1}^{d'} p'_i \left(\frac{1}{d'} \sum_{j=1}^{d'} (\ell_j - \ell_i) \right) = \langle p', \pi(\ell) \rangle - \frac{1}{d'} \left(\sum_{i=1}^{d'} \pi(\ell)_i \right), \quad (21)$$

where $\pi(\ell)$ is the projection of ℓ onto the first d' coordinates. Substituting this in turn into (20) yields

$$\text{Reg}(\mathbf{p}', \ell') = \max_{x^* \in \mathcal{P}'} \left(\sum_{t=1}^T \langle p'_t, \pi(\ell_t) \rangle - \sum_{t=1}^T \langle x^*, \pi(\ell_t) \rangle \right). \quad (22)$$

Now, note that the adversary can choose ℓ_t so that $\pi(\ell_t)$ takes on any value in $\mathcal{L}_2 = [0, 1]^{d'}$. Similarly, p' can take on any value in $\mathcal{P}_2 = \mathcal{P}' = \Delta_{d'}$. Therefore, this problem is at least as hard as the online linear optimization problem with action set $\mathcal{P}_2 = \Delta_{d'}$ and loss set $\mathcal{L}_2 = [0, 1]^{d'}$. But this is exactly the online learning with experts problem (with d' experts), which has a regret lower bound of $\Omega(\sqrt{T \log d'})$. It follows that $\text{Rate}(\mathcal{I}) \geq \Omega(\sqrt{\log d'}) > 0$. ■

One may object that the approachability instance in the proof of Theorem 3 is somewhat degenerate, since the approachability instance has a clear optimal action for the learner (which guarantees perfect approachability). In Appendix C we show that it is possible to slightly perturb this example in a way that avoids the existence of such an optimal action, while maintaining an arbitrarily large gap in rates between the two instances.

G.5. Proof of Theorem 4

Proof Let $\mathcal{X} = \mathcal{U} \otimes \mathcal{P}$ and $\mathcal{Y} = \mathcal{L}$. Let $B(x, y) : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ be the bilinear form defined via $B(u \otimes p, \ell) = -u(p, \ell)$. Note that we can write

$$\begin{aligned} \text{AppLoss}(\mathbf{p}, \ell) &= \max_{u \in \mathcal{U}} \sum_{t=1}^T u(p_t, \ell_t) = \max_{u \in \mathcal{U}} \sum_{t=1}^T -B(u \otimes p_t, \ell_t) \\ &= \max_{u \in \mathcal{U}} \left(\sum_{t=1}^T B(x_t, \ell_t) - \sum_{t=1}^T B(\phi_u(x_t), \ell_t) \right), \end{aligned}$$

where $x_t \in \mathcal{X}$ is any action whose projection onto \mathcal{P} is equal to p_t , and $\phi_u : \mathcal{X} \rightarrow \mathcal{X}$ is the linear function defined via $\phi_u(u' \otimes p) = (u' + u) \otimes p$ (for any $u' \in \mathcal{U}$); note that this is improper as $(u' + u) \otimes p$ is not guaranteed to be within \mathcal{X} . Now, the bilinear form B corresponds to a matrix M_B such that $B(x, y) = \langle x, M_B y \rangle$. From this it follows that if we set $\mathcal{P}' = \mathcal{X}$, $\mathcal{L}' = M_B \mathcal{Y}$, and $\Phi' = \{\phi_u \mid u \in \mathcal{U}\}$, then $\text{AppLoss}(\mathbf{p}, \ell) = \text{Reg}(\mathbf{p}', \ell')$ and this is a tight reduction between \mathcal{I} and \mathcal{I}' . ■

G.6. Weighted regret reduction

Theorem 14 *There exists a tight reduction from the improper ϕ -regret instance $\mathcal{I} = (\mathcal{P}, \mathcal{L}, \tilde{\Phi})$ to a proper ϕ -regret instance \mathcal{I}' .*

Proof We can construct the proper ϕ -regret minimization problem as follows. Let $W = \max w_i$, and let $\mathcal{P}' = \mathcal{P}$, $\mathcal{L}' = W\mathcal{L}$, and $\Phi' = \text{conv}(\{\phi'_i\}_{i=1}^N)$, where

$$\phi'_i(p) = \frac{w_i}{W}\phi_i(p) + \left(1 - \frac{w_i}{W}\right)p.$$

Note that each ϕ'_i is the convex combination of the two proper functions ϕ_i and the identity, and is therefore also proper. We let $\mathcal{I}' = (\mathcal{P}', \mathcal{L}', \Phi')$ denote this proper ϕ -regret minimization problem.

The reduction from \mathcal{I} to \mathcal{I}' is very simple and boils down to rescaling the losses by W . The necessary observation is just that:

$$\langle p - \tilde{\phi}_i(p), \ell \rangle = \langle p - \phi'_i(p), W\ell \rangle. \quad (23)$$

In particular, given an algorithm \mathcal{A}' for \mathcal{I}' , we can construct an algorithm \mathcal{A} for \mathcal{I} with $\text{Reg}(\mathcal{A}') = \text{Reg}(\mathcal{A})$ as follows:

1. At the beginning of round t , \mathcal{A} asks \mathcal{A}' for the p'_t it will play. \mathcal{A} then plays $p_t = p'_t$.
2. \mathcal{A} observes their loss vector ℓ_t . \mathcal{A} passes along the loss vector $\ell'_t = W\ell_t$ to \mathcal{A}' .

Equation (23) implies that not only do the worst-case regrets $\text{Reg}(\mathcal{A}') = \text{Reg}(\mathcal{A})$ agree for the above pair of algorithms, but that in any execution of the reduction, both algorithms achieve exactly the same regret. A similar reduction (dividing the loss by W instead of multiplying it by W) suffices to show that we can transform an algorithm \mathcal{A} for \mathcal{I} to an algorithm \mathcal{A}' for \mathcal{I}' with the same regret bound. ■

G.7. Proof of Theorem 5

To prove Theorem 5, we will use the following lemma (which proves the above characterization in the case of a single transformation ϕ).

Lemma 15 *Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a linear function, and let $\mathcal{P} \subseteq \mathbb{R}^d$ be a convex set. Then the following two conditions are equivalent:*

- *There exists an invertible linear transformation S such that $\phi'(p) = p + S(\phi(p) - p)$ is constant for all $p \in \mathcal{P}$.*
- *$\phi(p)$ has a single fixed point p_ϕ in $\text{span}_{\text{Aff}}(\mathcal{P})$.*

Proof First, assume there does exist such a linear transformation S . Assume to the contrary that ϕ has two distinct fixed points $p_1, p_2 \in \text{span}_{\text{Aff}}(\mathcal{P})$. Note that $\phi'(p_1) = p_1 + S(\phi(p_1) - p_1) = p_1$ so p_1 is a fixed point of ϕ' ; likewise p_2 is a fixed point of ϕ' . But if ϕ' is constant on \mathcal{P} , it must (as a linear transformation) must also be constant on $\text{span}_{\text{Aff}}(\mathcal{P})$, implying that we must have $\phi'(p_1) = \phi'(p_2)$, a contradiction.

Now, assume $\phi(p)$ contains a single fixed point $p_\phi \in \text{span}_{\text{Aff}}(\mathcal{P})$. If $\dim(\text{span}_{\text{Aff}}(\mathcal{P})) = k$, pick any k points $v_1, v_2, \dots, v_k \in \mathcal{P}$ that together with p_ϕ affinely span \mathcal{P} . In particular, the k vectors $(v_i - p_\phi)$ are linearly independent.

For each $i \in [k]$, let $w_i = \phi(v_i) - v_i$. We claim that the k vectors w_i are also linearly independent. If they were not, there would exist a non-trivial linear combination $\sum_i \lambda_i w_i$ equal to 0. Expanding this out (and using the fact that ϕ is linear), we have

$$\phi\left(\sum_i \lambda_i v_i\right) = \sum_i \lambda_i v_i.$$

Let $L = \sum_i \lambda_i$. Adding $(1 - L)p_\phi$ to both sides (and using the fact that $\phi(p_\phi) = p_\phi$) we have

$$\phi\left((1 - L)p_\phi + \sum_i \lambda_i v_i\right) = (1 - L)p_\phi + \sum_i \lambda_i v_i.$$

But the element $v = (1 - L)p_\phi + \sum_i \lambda_i v_i$ belongs to $\text{span}_{\text{Aff}}(\mathcal{P})$, so (by assumption) the only way it could be a fixed point of ϕ is if $v = p_\phi$. But $v - p_\phi = \sum_i \lambda_i (v_i - p_\phi)$, so this would imply that the original vectors $p - v_i$ were not linearly independent, a contradiction.

We can now directly construct our transformation S . We will choose any invertible S that satisfies

$$S(\phi(v_i) - v_i) = p_\phi - v_i$$

for all $i \in [k]$. Since both sets of vectors $w_i = \phi(v_i) - v_i$ and $p_\phi - v_i$ are linearly independent, we can choose an invertible such S (technically we also need to specify the action of S on vectors not contained in the span of w_i , but for these we can choose an arbitrary invertible transformation between the remaining sets of basis vectors).

We now claim that the resulting $\phi'(p) = p + S(\phi(p) - p)$ is a constant function (and in fact, equals p_ϕ for all $p \in \mathcal{P}$). To see this, write any p in the form $p = p_\phi + \sum_{i=1}^k \lambda_i (v_i - p_\phi)$ for some $\lambda_i \in \mathbb{R}$. Then

$$S(\phi(p) - p) = \sum_{i=1}^k \lambda_i S(\phi(v_i) - v_i) = \sum_{i=1}^k \lambda_i (p_\phi - v_i) = - \sum_{i=1}^k \lambda_i (v_i - p_\phi),$$

and therefore $p + S(\phi(p) - p) = p_\phi$, as desired. ■

We can now complete the proof of Theorem 5.

Proof [Proof of Theorem 5] We first prove that both conditions are necessary. The necessity of the first condition follows immediately from Lemma 15. To see that the second condition is necessary, consider any $\phi_1, \phi_2 \in \Phi$. By (18) we have

$$\begin{aligned}\phi'_1(p) &= p + S(\phi_1(p) - p) \\ \phi'_2(p) &= p + S(\phi_2(p) - p)\end{aligned}$$

and therefore

$$\phi'_1(p) - \phi'_2(p) = S(\phi_1(p) - \phi_2(p)).$$

If S is invertible and $\phi'_1(p) - \phi'_2(p)$ is a constant for all $p \in \mathcal{P}$, then $\phi_1(p) - \phi_2(p)$ must also be a constant for all $p \in \mathcal{P}$.

To show that these two conditions are sufficient, single out a representative $\psi \in \Phi$. By Lemma 15 we can construct an S such that the corresponding ψ' satisfies

$$\psi'(p) = p + S(\psi(p) - p) = p_\psi$$

for all $p \in \mathcal{P}$. Now, since any pair of functions in Φ differ by a constant, for each $\phi \in \Phi$, write $\phi(p) = \psi(p) + \delta_\phi$ where δ_ϕ is some p -independent constant in \mathbb{R}^d . If we let p_ϕ denote the fixed point of ϕ in \mathcal{P} , then since $\phi(p_\phi) = p_\phi$, so we can deduce that $\delta_\phi = p_\phi - \psi(p_\phi)$. Note then that

$$\begin{aligned} \phi'(p) &= p + S(\phi(p) - p) \\ &= p + S(\psi(p) + \delta_\phi - p) \\ &= p + S(\psi(p) - p + p_\phi - \psi(p_\phi)) \\ &= p + (p_\psi - p) - (p_\psi - p_\phi) \\ &= p_\phi, \end{aligned}$$

so each $\phi'(p)$ is constant for all $p \in \mathcal{P}$. ■