

Regret-optimal Strategies for Playing Repeated Games with Discounted Losses

Vijay Kamble *
Dept. of EECS
UC Berkeley

Patrick Loiseau †
EURECOM

Jean Walrand ‡
Dept. of EECS
UC Berkeley

Abstract

The regret-minimization paradigm has emerged as a powerful technique for designing algorithms for online decision-making in adversarial environments. But so far, designing exact minmax-optimal algorithms for minimizing the worst-case regret has proven to be a difficult task in general, with only a few known results in specific settings. In this paper, we present a novel set-valued dynamic programming approach for designing such exact regret-optimal policies for playing repeated games with discounted losses.

Our approach first draws the connection between regret minimization, and determining minimal achievable guarantees in repeated games with vector-valued losses. We then characterize the set of these minimal guarantees as the fixed point of a dynamic programming operator defined on the space of Pareto frontiers of convex and compact sets. This approach simultaneously results in the characterization of the optimal strategies that achieve these minimal guarantees, and hence of regret-optimal strategies in the original repeated game. As an illustration of our approach, we design a simple near-optimal strategy for prediction using expert advice for the case of 2 experts and discounted losses. To the best of our knowledge, this is the first such algorithm for this setting.

1 Introduction

In many scenarios involving repeated decision-making in uncertain environments, one desires robust performance guarantees. The well-known regret minimization paradigm captures this objective by formulating the problem as a repeated game between the decision-maker and the environment, which is modeled as an adversary. The goal of the decision-maker in this game is to design an adaptive and possibly randomized algorithm for choosing actions (henceforth, algorithm/strategy/policy) so as to minimize *regret*, which is defined as the difference between the expected loss incurred by the algorithm, and the loss incurred by the best fixed action that could have been chosen in hindsight against the sequence of actions chosen by the adversary. Since its inception by Hannan in the 1950s [Han57], this formulation of an online decision-making problem under uncertainty has been widely popular both in theory and practice, having been successfully applied to problems like managing investment portfolios, weather forecasting, online spam detection etc., and is now fundamental to the field of online learning.

* vijaykamble.iitkgp@gmail.com

† patrick.loiseau@eurecom.fr

‡ walrand@berkeley.edu

In this setting, several *no-regret* algorithms are now known, which ensure that the time-averaged regret asymptotically vanishes in the long run, not just in expectation, but with probability 1, irrespective of the sequence of actions chosen by the adversary. But in many realistic cases, e.g., when decision-making horizons are finite, the minimal expected regret one can achieve in the worst case over all the sequence of actions chosen by the adversary (henceforth, just optimal regret) is non-zero. In such cases, standard no-regret algorithms can perform quite poorly compared to the optimum and except for a very few special cases, the exact optimal regret and strategies are unknown.

The main contribution of this paper is the proposal of a systematic set-valued dynamic programming approach for designing regret-optimal strategies in repeated games. Such an approach is most fruitful, and hence is best illustrated, in the case of repeated games with discounted losses, which is thus the focus of the paper. In these games, the loss criterion is the weighted sum of per-stage losses, with the loss at stage $t \geq 1$ weighted by β^{t-1} , where $\beta \in (0, 1)$ is called the discount factor. Such discounting is natural in practice, where minimizing current losses is more important than the ones in the future.¹ In this case, since the losses incurred in the initial stages have a non-vanishing contribution to the cumulative loss as the number of stages increases, the optimal long-run average discounted regret for any fixed $\beta \in (0, 1)$ is non-zero (see [CBL03], Thm 2.7).² Several known no-regret algorithms guarantee an average regret of $O(\sqrt{1-\beta})$ asymptotically in this setting, but the performance of these algorithms can be far from optimal for a fixed β . Our approach on the other hand gives a procedure to compute ϵ -regret-optimal strategies in these games for any $\epsilon > 0$. These strategies are extremely simple to implement and require a finite memory (which grows as ϵ decreases). For instance, using our approach, we are able to design a provably near-optimal algorithm for the problem of prediction using expert advice with discounted losses, for the case of 2 experts. To the best of our knowledge, no such algorithm was known before.

Our solution begins with a standard approach in regret minimization, of transforming the repeated game into a repeated game with vector losses. In this game, the number of vector components is the number of actions available to the decision-maker, where each component keeps track of the additional loss incurred relative to the loss incurred if the corresponding action was always chosen in the past. The goal of regret minimization in the original game now translates to the goal of *simultaneously* minimizing the worst-case expected losses on all the components in this vector-valued game. In fact there is a tradeoff here: a better guarantee on one component implies a worse guarantee on some other, and we consider the more ambitious objective of characterizing the entire Pareto-frontier of the minimal losses that can be simultaneously guaranteed across the different components. Our main technical contribution is an effective characterization of this set as the fixed point of a set-valued dynamic programming operator, which simultaneously also characterizes the strategies that achieve the different points on it. This characterization then allows us to design an iterative scheme to approximate this set and compute approximately optimal strategies for arbitrarily low approximation error.

The paper is organized as follows. In Section 2, we formally introduce the notion of regret optimality in repeated games, and demonstrate the transformation of the regret minimization problem into the problem of computing simultaneous guarantees on losses in a vector-valued repeated game. In Section 3 we introduce the set-valued dynamic programming approach for computing these guarantees, and give a characterization of both the optimal regret and the optimal policies. In

¹For instance, such a discounting can capture the existence of a risk-free interest rate on your investments.

²By average discounted regret we mean that the weights are normalized by multiplying with $(1-\beta)$, which ensures that the sum of the weights over an infinite horizon is 1.

Section 4, we present an approximate value iteration procedure to approximate the optimal regret and compute an ϵ -optimal strategy in the case where the decision-maker has only 2 available actions. Section 5 illustrates our approach on the problem of prediction using expert advice with 2 experts. We compare the performance of a near-optimal policy that we compute with that of two algorithms: one is the classical exponentially weighted average forecaster, and another one is the optimal algorithm for prediction using expert advice with 2 experts and with geometrically distributed stopping times, characterized recently by [GPS14]. We summarize our contribution and discuss future directions in Section 6. The proofs of all our results can be found in the appendix.

1.1 Related Literature

The first study of regret minimization in repeated games dates back to the pioneering work of Hannan [Han57], who introduced the notion of regret optimality in repeated games and proposed the earliest known no-regret algorithm. Since then, numerous other such algorithms have been proposed, particularly for the problem of prediction using expert advice, see [LW94, Vov90, CBFH⁺97, FS99], one particularly well-known being the multiplicative weights update algorithm, also known as Hedge. Other settings with limited feedback have been considered, most notably the multi-armed bandit setting [ACBFS02, BCB12]. Stronger notions of regret such as internal regret, have also been studied [FV97, CBL03, BM05, SL05].

Similarly, regret minimization with non-uniformly weighted losses, of which the discounted loss is a special case, has been considered before. While the average regret goes to zero if the weights satisfy a non-summability condition, lower bounds exist ([CBL03], Thm 2.7) that show that the optimal regret is bounded away from 0 if the weights are summable, which is the case with discounting. Natural extensions of no-regret algorithms incur a regret of $O(\sqrt{1-\beta})$ in this case, for instance see [CBL03], Thm 2.8 and [Per14], Prop. 2.22.³

The results on exact regret minimization are few. In an early work, [Cov66] gave the optimal algorithm for the problem of prediction using expert advice over any finite horizon T , for the case of 2 experts, and where the losses are $\{0, 1\}$. [GPS14] recently extended the result to the case of 3 experts for both the finite horizon and geometrically distributed random horizon problems.⁴ [AWY08] considered a related problem, where a gambler places bets from a finite budget repeatedly on a fixed menu of events, the outcomes of which are adversarially chosen from $\{0, 1\}$ (you win or you lose), and characterized the minmax optimal strategies for the gambler and the adversary. [LS13] considered a similar repeated decision-making problem where an adversary is restricted to pick loss vectors (i.e., a loss for each action of the decision-maker in a stage) from a set of basis vectors, and characterized the minmax optimal strategy for the decision-maker under both, a fixed and an unknown horizon. Most of the approaches in these works are specific to their settings, and exploit the assumptions on the structure of the loss vectors.^{5,6} But in general if the loss vectors are

³[CZ10] derive better bounds for the case where future losses are given a higher weight than current ones.

⁴Although a geometric time horizon model seems to be related to the infinite horizon model with discounted losses, the two problem formulations define regret differently, and thus lead to different optimal regrets. We discuss this in Section 5.1.

⁵Many of these works rely on a particular nice property of these settings, which is that the optimal strategy of the adversary is a controlled random walk that makes any algorithm incur the same regret. If the losses are simple, for instance if they are the basis vectors, then this random walk can be exactly analyzed to compute the optimal regret. Knowing the optimal regret then simplifies the computation of the optimal strategy of the decision-maker.

⁶All of these are examples where games with finite action spaces (called matrix games) are repeated, which is the setting that we are concerned with. But there are many works that consider exact minmax optimality in repeated games with continuous action spaces, with specific types of loss functions, see [KMB14, BKM⁺15, KMBA15] and

arbitrary, none of these approaches can be extended and indeed it is recognized that characterizing the optimal regret and algorithm is difficult, cf. [LS13]. On the other hand, using our approach, we are able to characterize the optimal regret and algorithm for any choice of bounded loss vectors.

Similar to our approach, the idea of characterizing the Pareto Frontier of all achievable regrets with respect to different actions has been explored by [Koo13], again in the specific context of prediction with expert advice with 2 experts and $\{0, 1\}$ losses, and for the finite time horizon problem. Our focus on the infinite horizon setting with discounted losses entails the development and use of a more involved machinery of contractive set-valued dynamic programming operators, and at the end it gives results that are much *cleaner* than those in the finite horizon setting. Set-valued dynamic programs have been used in other contexts in the theory of dynamic games, for instance to compute the set of equilibrium payoffs in non-zero sum repeated games with imperfect monitoring [APS86, APS90]. For the use of dynamic programming in zero-sum dynamic games one can refer to the classic paper by Shapley [Sha53] on Stochastic games. For a general theory of dynamic programming, see [Ber05].

An important step in our approach is a reduction of the problem to a vector-valued repeated game. The study of vector-valued repeated games was pioneered by Blackwell [Bla56a]. He gave sufficient conditions for a set to be *approachable* by a player, which means that there exists a strategy for a player that ensures that the average loss approaches this set regardless of the adversary's actions. Moreover, he explicitly defined an adaptive randomized strategy that ensures this. Later he also showed that this theory can be used to obtain no-regret strategies as formulated by Hannan [Han57], using the transformation of the repeated game into a vector-valued game that we described earlier [Bla56b]. This theory was subsequently extended in various ways [Vie92, Leh03], and stronger connections with regret minimization and other learning problems like calibration were shown [ABH11, Per14]. But as far as we know, there has been no prior work on the discounted loss criterion. In this paper, as a by-product of our analysis, we successfully bridge this important gap in the theory of vector-valued repeated games. As a result, this theory bridges significant gaps in other decision-making problems where Blackwell's approachability theory has found applications.⁷

2 Model

Let G be a two player game with m actions $A = \{1, \dots, m\}$ for player 1, who is assumed to be the minimizer and who we will call Alice (the decision-maker), and n actions $B = \{1, \dots, n\}$ for player 2, who is the adversary and who we will call Bob. For each pair of actions $a \in A$ and $b \in B$, the corresponding loss for Alice is $l(a, b) \in \mathbb{R}$. The losses for different pairs of actions are known to Alice. The game G is played repeatedly for T stages $t = 1, 2, \dots, T$. In each stage, both Alice and Bob simultaneously pick their actions $a_t \in A$ and $b_t \in B$ and Alice incurs the corresponding loss $l(a_t, b_t)$. The loss of the repeated game is defined to be the total discounted loss given by $\sum_{t=1}^T \beta^{t-1} l(a_t, b_t)$, where $\beta \in (0, 1)$. We define the total discounted regret of Alice as:

$$\sum_{t=1}^T \beta^{t-1} l(a_t, b_t) - \min_{a \in A} \sum_{t=1}^T \beta^{t-1} l(a, b_t), \quad (1)$$

which is the difference between her actual discounted loss, and the loss corresponding to the single best action that could have been chosen against the sequence of actions chosen by Bob in hindsight.

references therein.

⁷A notable example is the analysis of zero-sum repeated games with incomplete information [Zam92, AM95, Sor02]

An adaptive randomized strategy π_A for Alice specifies for each stage t , a mapping from the set of observations till stage t , i.e., $H_t = (a_1, b_1, \dots, a_{t-1}, b_{t-1})$, to a probability distribution on the action set A , denoted by $\Delta(A)$. Let Π_A be the set of all such policies of Alice.

The adversary Bob is assumed to choose a deterministic oblivious strategy, i.e., his choice is simply a sequence of actions $\pi_B = (b_1, b_2, b_3, \dots, b^T)$ chosen before the start of the game. Let Π_B be the set of all such sequences.⁸ We would like to compute the worst case or minmax *expected* discounted regret which is defined as:

$$\min_{\pi_A \in \Pi_A} \max_{\pi_B \in \Pi_B} \left(E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} l(a_t, b_t) \right] - \min_{a \in A} \sum_{t=1}^T \beta^{t-1} l(a, b_t) \right), \quad (2)$$

and the strategy for Alice that guarantees this value. Here the expectation is over the randomness in Alice's strategy. Here one can see that there is no loss of generality in assuming that the adversary is deterministic. Indeed if Π_B is allowed to be the set of possible randomizations over T -length sequences of Bob's actions, the optimal policy of Bob in the problem

$$\max_{\pi_B \in \Pi_B} E_{\pi_A, \pi_B} \left[\sum_{t=1}^T \beta^{t-1} l(a_t, b_t) - \min_{a \in A} \sum_{t=1}^T \beta^{t-1} l(a, b_t) \right]$$

is a deterministic sequence.

We can now equivalently write (2) as:

$$\min_{\pi_A \in \Pi_A} \max_{\pi_B \in \Pi_B} \max_{a \in A} E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} (l(a_t, b_t) - l(a, b_t)) \right]. \quad (3)$$

In order to address this objective, it is convenient to define a vector-valued game \mathbb{G} , in which, for a pair of actions $a \in A$ and $b \in B$, the vector of losses is $\mathbf{r}(a, b)$ with m components (recall that $|A| = m$), where

$$r_k(a, b) = l(a, b) - l(k, b) \quad (4)$$

for $k = 1, \dots, m$. $r_k(a, b)$ is the single-stage additional loss that Alice bears by choosing action a instead of action k , when Bob chooses b : the so called single-stage regret with respect to action k . For a choice of strategies $\pi_A \in \Pi_A$ and $\pi_B \in \Pi_B$ of the two players, the expected loss on component k in this vector-valued repeated game over horizon T is given by

$$R_k^T(\pi_A, \pi_B) = E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} r_k(a_t, b_t) \right], \quad (5)$$

where the expectation is over the randomness in Alice's strategy. Now observe that by playing a fixed policy $\pi_A \in \Pi_A$, irrespective of the strategy chosen by Bob, Alice guarantees that the total expected loss on component k is no more than $\max_{\pi_B^k \in \Pi_B} R_k^T(\pi_A, \pi_B^k)$. Suppose that we determine the set of all *simultaneous guarantees* that correspond to *all* the strategies $\pi_A \in \Pi_A$, defined as:

$$\mathcal{W}^T \triangleq \left\{ \left(\max_{\pi_B^k \in \Pi_B} R_k^T(\pi_A, \pi_B^k) \right)_{k=1, \dots, m} : \pi_A \in \Pi_A \right\}. \quad (6)$$

⁸Having an oblivious adversary is a standard assumption in regret-minimization literature [CBL03] and in fact it is known that in this case, an oblivious adversary is as powerful as a non-oblivious (adaptive) adversary.

Then it is clear that

$$\min_{\pi_A \in \Pi_A} \max_{\pi_B \in \Pi_B} \max_{a \in A} E_{\pi_A} \left[\sum_{t=1}^T \beta^{t-1} (l(a_t, b_t) - l(a, b_t)) \right] = \min_{\mathbf{x} \in \mathcal{W}^T} \max_k x_k.$$

In fact, we are only interested in finding the *minimal* points in the set \mathcal{W}^T , i.e., its *Lower Pareto frontier*, which is defined as the set

$$\Lambda(\mathcal{W}^T) \triangleq \{\mathbf{x} \in \mathcal{W}^T : \forall \mathbf{x}' \in \mathcal{W}^T \setminus \{\mathbf{x}\}, \exists k \text{ s.t. } x_k < x'_k\}, \quad (7)$$

since all other points are strictly sub-optimal. Let this set be denoted as \mathcal{V}^T . Our goal in this paper is to characterize and compute the set \mathcal{V}^∞ that can be achieved in the infinite horizon game and compute policies for Alice in Π^A that guarantee different points in it.

3 Set-valued dynamic programming

In the remainder of the paper, we will consider a general model for a vector-valued repeated game \mathbb{G} with m actions for Alice and n actions for Bob, where m and n are arbitrary but finite, and we assume that the number of components in the vector of losses $\mathbf{r}(a, b)$ is K , where K may be different from m . For ease of exposition, we will consider the case where $K = 2$, but the analysis holds for any finite K .

Our goal is to compute the set of minimal guarantees \mathcal{V}^∞ that Alice can achieve in the infinite horizon vector-valued discounted repeated game and characterize the policies that achieve it. Our results will hold for any vector-valued repeated game, not just the one that arises from the regret minimization formulation discussed before.

In our analysis, we will assume that Bob is not restricted to playing oblivious deterministic strategies, but can play any adaptive randomized strategy. An adaptive randomized strategy π_B for Bob specifies for each stage t , a mapping from the set of observations till stage t , i.e., $H_t = (a_1, b_1, \dots, a_{t-1}, b_{t-1})$, to a probability distribution on the action set B , denoted by $\Delta(B)$. There is no loss in doing so, since we will later show in Section 3.6 that the optimal set \mathcal{V}^∞ and the optimal policies for Alice remain the same irrespective of which of the two strategy spaces are assumed for Bob. Although the analysis is simpler if we restrict Bob to using only deterministic oblivious strategies, the approach we take allows us to demonstrate that our results hold at a higher level of generality in the broad context of characterizing minimal guarantees in repeated games with vector losses.

3.1 Overview of results

Our main contributions are as follows:

1. We show that the set \mathcal{V}^∞ of *minimal* losses that Alice can simultaneously guarantee in an infinitely repeated vector-valued zero-sum repeated game with discounted losses is the fixed point of a set-valued dynamic programming operator defined on the space of lower Pareto frontiers of closed convex sets with an appropriately defined metric. We then show that the optimal policies that guarantee different points in this set are of the following form. \mathcal{V}^∞ can be parametrized so that each point corresponds to a parameter value, which can be thought of as

an “information state”. Each state is associated with an immediate optimal randomized action and a transition rule that depends on the observed action of the adversary. In order to attain a point in \mathcal{V}^∞ , Alice starts with the corresponding state, plays the associated randomized action, transitions into another state depending on Bob’s observed action as dictated by the rule, plays the randomized action associated with the new state and so on. In particular, the strategy does not depend on Alice’s past actions and it depends on Bob’s past actions only through this state that the minimizing player keeps track of.

2. For the case where Alice has only 2 actions, we give a value-iteration based procedure to approximate \mathcal{V}^∞ and to compute an approximately optimal policy that only uses a coarse finite quantization of the parameter space. This strategy can be simply implemented by a randomized finite-state automaton. Any desired diminishing approximation error can be attained by choosing the appropriate quantization granularity and number of iterations. Our procedure in principle can be extended to an arbitrary number of actions. We finally illustrate our theory and the approximation procedure on a simple model of prediction with expert advice with 2 experts.

3.2 Overview of the approach

We now present an informal description of our approach. Let \mathbb{G}^T denote the T - stage repeated game and let \mathbb{G}^∞ denote the infinitely repeated game. Let $\mathcal{V}^0 = \{(0, 0)\}$. We can show that one can obtain the set \mathcal{V}^{T+1} from the set \mathcal{V}^T , by decomposing Alice’s strategy in \mathbb{G}^{T+1} into a strategy for the 1st stage, and a continuation strategy for the remainder of the game from stage 2 onwards, as a function of the action chosen by both the players in the 1st stage. The induction results from the fact that the minimal guarantees that she can guarantee from stage 2 onwards are exactly the set \mathcal{V}^T . Suppose that at the start of \mathbb{G}^{T+1} , Alice fixes the following plan for the entire game: she will play a mixed strategy $\alpha \in \Delta(A)$ in stage 1. Then depending on her realized action a and Bob’s action b , from stage 2 onwards she will play a continuation strategy that achieves the guarantee $\mathbf{R}(a, b) \in \mathcal{V}^T$ (she will choose one such point $\mathbf{R}(a, b)$ for every $a \in A$ and $b \in B$). Note that it is strictly sub-optimal for Alice to choose any points outside \mathcal{V}^T from stage 2 onwards. Now this plan for the entire game \mathbb{G}^{T+1} gives Alice the following expected simultaneous guarantees on the two components:

$$\left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_1(a, b) + \beta R_1(a, b)] \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_2(a, b) + \beta R_2(a, b)] \right\} \right).$$

By varying the choice of α and the map $\mathbf{R}(a, b)$ we can obtain the set of all the simultaneous guarantees that Alice can achieve in the $(T + 1)$ -stage game. The Lower Pareto frontier of this set is exactly \mathcal{V}^{T+1} . Thus there is an operator Φ , such that

$$\mathcal{V}^{T+1} = \Phi(\mathcal{V}^T)$$

for any $T \geq 0$, where \mathcal{V}^0 is defined to be the singleton set $\{(0, 0)\}$. In what follows, we will show that this operator is a contraction in the space of Lower Pareto frontiers of closed convex sets, with an appropriately defined metric. This space is shown to be complete, and thus the sequence \mathcal{V}^T converges in the metric d to a set \mathcal{V}^* , which is the unique fixed point of this operator Φ . As one would guess, this \mathcal{V}^* is indeed the set \mathcal{V}^∞ of minimal simultaneous guarantees that Alice can achieve in the infinitely repeated game \mathbb{G}^∞ .

The rest of this section formalizes these arguments. We will begin the formal presentation of our results by first defining the space of Pareto frontiers that we will be working with.

3.3 The space of Pareto frontiers in $[0, 1]^2$

Definition 3.1. (a) Let $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$. We say that $\mathbf{u} \preceq \mathbf{v}$ if $u_1 \leq v_1$ and $u_2 \leq v_2$. Also, we say that $\mathbf{u} \prec \mathbf{v}$ if $\mathbf{u} \preceq \mathbf{v}$ and $\mathbf{u} \neq \mathbf{v}$. If $\mathbf{u} \prec \mathbf{v}$, we say that \mathbf{v} is dominated by \mathbf{u} .

(b) A Pareto frontier in $[0, 1]^2$ is a subset \mathcal{V} of $[0, 1]^2$ such that no $\mathbf{v} \in \mathcal{V}$ is dominated by another element of \mathcal{V} .

(c) The Lower Pareto frontier (or simply Pareto frontier) of $\mathcal{S} \subset [0, 1]^2$, denoted by $\Lambda(\mathcal{S})$, is the set of elements of \mathcal{S} that do not dominate any another element of \mathcal{S} .

The Pareto frontier of a set may be empty, as is certainly the case when the set is open. But one can show that the Pareto frontier of a compact set is always non-empty. Since compactness is equivalent to a set being closed and bounded in Euclidean spaces, any closed subset of $[0, 1]^2$ has a non-empty Pareto frontier. We define the following space of Pareto frontiers:

Definition 3.2. \mathcal{F} is the space of Pareto frontiers of closed and convex subsets of $[0, 1]^2$.

We will now define a metric on this space. We first define the *upset* of a set, illustrated in Figure 1.

Definition 3.3. Let \mathcal{A} be a subset of $[0, 1]^2$. The upset $up(\mathcal{A})$ of \mathcal{A} is defined as $up(\mathcal{A}) = \{\mathbf{x} \in [0, 1]^2 \mid x_1 \geq y_1 \text{ and } x_2 \geq y_2, \text{ for some } \mathbf{y} \in \mathcal{A}\}$. Equivalently, $up(\mathcal{A}) = \{\mathbf{x} \in [0, 1]^2 \mid \mathbf{x} = \mathbf{y} + \mathbf{v}, \text{ for some } \mathbf{y} \in \mathcal{A} \text{ and } \mathbf{v} \succeq \mathbf{0}\}$.

It is immediate that the upset of the Pareto frontier of a closed convex set in $[0, 1]^2$ is closed and convex. We recall the definition of *Hausdorff distance* induced by the ∞ -norm.




Figure 1: A Pareto frontier \mathcal{V} and its upset $up(\mathcal{V})$.

Definition 3.4. Let \mathcal{A} and \mathcal{B} be two subsets of \mathbb{R}^2 . The Hausdorff distance $h(\mathcal{A}, \mathcal{B})$ between the two sets is defined as

$$h(\mathcal{A}, \mathcal{B}) = \max\left\{\sup_{\mathbf{x} \in \mathcal{A}} \inf_{\mathbf{y} \in \mathcal{B}} \|\mathbf{x} - \mathbf{y}\|_\infty, \sup_{\mathbf{y} \in \mathcal{B}} \inf_{\mathbf{x} \in \mathcal{A}} \|\mathbf{x} - \mathbf{y}\|_\infty\right\}.$$

We now define the distance between two Pareto frontiers in \mathcal{F} as the Hausdorff distance between their upsets.

Definition 3.5. For two pareto frontiers \mathcal{U} and \mathcal{V} in \mathcal{F} , we define the distance $d(\mathcal{U}, \mathcal{V})$ between them as $d(\mathcal{U}, \mathcal{V}) = h(\text{up}(\mathcal{U}), \text{up}(\mathcal{V}))$.

We can then show that \mathcal{F} is complete in the metric d . This follows from the completeness of the Hausdorff metric for closed convex subsets of $[0, 1]^2$.

Proposition 3.1. Let $(\mathcal{V}_n)_{n \in \mathbb{N}}$ be a sequence in \mathcal{F} . Suppose that $\sup_{m, k > n} d(\mathcal{V}_m, \mathcal{V}_k) \rightarrow 0$. Then there exists a unique $\mathcal{V} \in \mathcal{F}$ such that $d(\mathcal{V}_n, \mathcal{V}) \rightarrow 0$.

3.4 Dynamic programming operator and the existence of a fixed point

By scaling and shifting the losses, we assume without loss of generality that $r_k(a, b) \in [0, 1 - \beta]$ for all (a, b, k) . Accordingly, the total discounted rewards of the game take values in $[0, 1]$ irrespective of the time horizon. Now, for a closed set $\mathcal{S} \subseteq [0, 1]^2$, define the following operator Ψ that maps \mathcal{S} to a subset of \mathbb{R}^2 :

$$\Psi(\mathcal{S}) = \left\{ \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_1(a, b) + \beta R_1(a, b)] \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_2(a, b) + \beta R_2(a, b)] \right\} \right) : \alpha \in \Delta(A), \mathbf{R}(a, b) \in \mathcal{S} \forall a \in A, b \in B \right\}. \quad (8)$$

This operator can be interpreted as follows. Assuming that \mathcal{S} is the set of pairs of expected guarantees on losses that Alice can ensure in \mathbb{G}^T , $\Psi(\mathcal{S})$ is the set of pairs of expected guarantees that she can ensure in \mathbb{G}^{T+1} . We will show that if \mathcal{S} is closed then $\Psi(\mathcal{S})$ is closed as well. But if \mathcal{S} is convex then $\Psi(\mathcal{S})$ is not necessarily convex. Nevertheless we can show that the Pareto frontier of $\Psi(\mathcal{S})$ is the Pareto frontier of some convex and compact set.

Lemma 3.1. Let $\mathcal{S} \subseteq [0, 1]^2$ be a closed set. Then $\Psi(\mathcal{S}) \subseteq [0, 1]^2$ is closed. If in addition, \mathcal{S} is convex, then:

1. Any point \mathbf{u} in $\Lambda(\Psi(\mathcal{S}))$ is of the form:

$$\mathbf{u} = \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} \right).$$

where $\mathbf{Q}(b) \in \Lambda(\mathcal{S})$.

2. $\Lambda(\Psi(\mathcal{S})) \in \mathcal{F}$.

In the context of our discussion in Section 3.2, the first point of the Lemma implies that in any optimal plan for Alice in \mathbb{G}^{T+1} , the continuation strategy from stage 2 onwards need not depend on her own action in stage 1 (but it depends on Bob's action in the first stage). Define the following dynamic programming operator Φ on \mathcal{F} .

Definition 3.6. (Dynamic programming operator) For $\mathcal{V} \in \mathcal{F}$, we define $\Phi(\mathcal{V}) = \Lambda(\Psi(\mathcal{V}))$.

Now since \mathcal{V} is the Lower Pareto frontier of some closed convex subset of \mathbb{R}^2 , say \mathcal{S} , and since $\Lambda(\Psi(\mathcal{V})) = \Lambda(\Psi(\mathcal{S}))$, from Lemma 3.1, we know that $\Phi(\mathcal{V}) \in \mathcal{F}$ whenever $\mathcal{V} \in \mathcal{F}$. Next, we claim that Φ is a contraction in the metric d .

Lemma 3.2.

$$d(\Phi(\mathcal{U}), \Phi(\mathcal{V})) \leq \beta d(\mathcal{U}, \mathcal{V}). \quad (9)$$

Finally we show that the dynamic programming operator has a unique fixed point and starting from a Pareto frontier in \mathcal{F} , the sequence of frontiers obtained by a repeated application of this operator converges to this point.

Theorem 3.1. *Let $\mathcal{V} \in \mathcal{F}$. Then the sequence $(\mathcal{A}_n = \Phi^n(\mathcal{V}))_{n \in \mathbb{N}}$ converges in the metric d to a Pareto frontier $\mathcal{V}^* \in \mathcal{F}$, which is the unique fixed point of the operator Φ , i.e., the unique solution of $\Phi(\mathcal{V}) = \mathcal{V}$.*

We can then show that \mathcal{V}^* is indeed the optimal set \mathcal{V}^∞ that we are looking for.

Theorem 3.2. $\mathcal{V}^\infty = \mathcal{V}^*$.

3.5 Optimal policies: Existence and Structure

For a Pareto frontier $\mathcal{V} \in \mathcal{F}$, one can define a one-to-one function from some parameter set \mathcal{P} to \mathcal{V} . Such a function parameterizes the Pareto frontier. For instance, consider the function $\mathbf{F} : [-1, 1] \times \mathcal{F} \rightarrow \mathbb{R}^2$, where one defines

$$\mathbf{F}(p, \mathcal{V}) = \arg \min_{(x,y)} x + y \quad (10)$$

$$\text{s.t. } x \geq u_1, y \geq u_2, \mathbf{u} \in \mathcal{V}, y = x + p.$$

$\mathbf{F}(p, \mathcal{V})$ is essentially the component-wise smallest point of intersection of the line $y = x + p$ with the upset of \mathcal{V} in $[0, 2]^2$ (see Figure 2).




Figure 2: Parameterization of \mathcal{V} .

Then for a given \mathcal{V} , the function $\mathbf{F}(\cdot, \mathcal{V}) : [-1, 1] \rightarrow \mathcal{V}$ is indeed one-to-one since any point on \mathcal{V} intersects exactly one of the lines. We can now express the dynamic programming operator in the form of such a parametrization. Assume that \mathcal{V}^* is such that $\mathcal{V}^* = \Phi(\mathcal{V}^*)$. Then for $p \in \mathcal{P}$, one can choose $\alpha(p) \in \Delta(A)$ and $q(b, p) \in \mathcal{P}$ for each $b \in B$ such that for $k \in \{1, 2\}$,

$$F_k(p, \mathcal{V}^*) = \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a(p) r_k(a, b) + \beta F_k(q(b, p), \mathcal{V}^*) \right\}. \quad (11)$$

Then we have the following result.

Theorem 3.3. *For any $p_1 \in \mathcal{P}$, the upper bound $\mathbf{F}(p_1, \mathcal{V}^*)$ in \mathcal{V}^* is guaranteed on losses by Alice in the infinite horizon game by first choosing action $a_1 \in A$ with probability $\alpha_{a_1}(p_1)$. Then if Bob chooses an action $b_1 \in B$, the optimal guarantees to choose from the second step onwards are then $\beta\mathbf{F}(p_2, \mathcal{V}^*)$ in $\beta\mathcal{V}^*$, where $p_2 = q(b_1, p_1)$, which can be guaranteed by Alice by choosing action $a_2 \in A$ with probability $\alpha_{a_2}(p_2)$, and so on.*

This implies that \mathcal{P} can be thought of as a *state space* for the policy. Each state is associated with an immediate optimal randomized action and a transition rule that depends on the observed action of Bob. In order to attain a point in \mathcal{V}^* , Alice starts with the corresponding state, plays the associated randomized action, transitions into another state depending on Bob's observed action as dictated by the rule, plays the randomized action associated with the new state and so on. *In particular, the policy does not depend on the past actions of Alice and it depends on the past actions of Bob only through this information state that Alice keeps track of.* Since Alice's optimal policy itself does not depend on her own past actions, Bob's optimal response does not depend on them either. Hence one can see that Bob has an *oblivious* best response to any optimal policy of Alice. This fact is important for the following discussion.

3.6 Oblivious adversary

Recall that when we defined the dynamic programming operator, we assumed that Bob is allowed to choose any adaptive non-oblivious strategy. We will now argue that even if we restrict Bob to an oblivious deterministic policy, the set of minimal losses that Alice can guarantee is still \mathcal{V}^* , the fixed point of our operator. First recall that Bob's best response to Alice's optimal strategy that achieves different points on the frontier \mathcal{V}^* is deterministic and oblivious (offline). This is because Alice does not take into account her own actions to determine the information state transitions. In fact, if Alice is restricted to use strategies that do not depend on her own actions chosen in the past, then the best response to such policy is always an offline deterministic policy, and hence the minimal achievable frontier if Bob is restricted to use offline deterministic policies is \mathcal{V}^* . So all that we need to verify is that Alice does not gain by using policies that depend on her own past actions, when Bob is restricted to using only offline deterministic strategies. To see this, suppose that \mathcal{V}^T is the set of guarantees that Alice by using general randomized adaptive strategies in \mathbb{G}^T , assuming that Bob is restricted to using deterministic offline policies. Then in \mathbb{G}^{T+1} , Alice's policy is a choice of a distribution over her actions, α , and then a mapping from the realized actions (a, b) to some continuation (randomized adaptive) policy $\pi(a, b) \in \Pi_A$. But since Bob's responses that maximize the losses on the different components cannot depend on the realization of Alice's action a , and can only depend on α , his best responses from time 2 onwards would effectively be against the strategy $\pi'(b)$ of Alice that chooses the policy $\pi(a, b)$ with probability α_a for each realized action a . This policy better guarantee a point in \mathcal{V}^T since all other policies are strictly-suboptimal. Thus the guarantees that Alice can achieve in \mathbb{G}^{T+1} is given by the set:

$$\mathcal{V}^T = \Lambda \left(\left\{ \max_{b \in B} \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b), \max_{b \in B} \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right. \right. \\ \left. \left. \mid \alpha \in \Delta(A), \mathbf{Q}(b) \in \mathcal{V}_{t+1} \forall b \in B \right\} \right).$$

But this is exactly the dynamic programming operator in Definition 3.6. Hence we can conclude that \mathcal{V}^* is indeed the set of minimal guarantees, even if Bob is restricted to using deterministic offline policies.

4 Approximation for $K = 2$

Finding the minmax regret and the optimal policy that achieves it requires us to compute \mathcal{V}^* and $\{(\boldsymbol{\alpha}(p), q(b, p)) : p \in \mathcal{P}\}$ that satisfies (11). Except in simple examples, it is not possible to do so analytically. Hence we now propose a computational procedure to approximate the optimal Pareto frontier in \mathbb{R}^2 and devise approximately optimal policies. In order to do so, we need to define an approximation of any Pareto frontier. Consider the following approximation scheme for a Pareto frontier $\mathcal{V} \in \mathcal{F}$. For a fixed positive integer N , define the approximation operator to be

$$\Gamma_N(\mathcal{V}) = \Lambda\left(\text{ch}\left(\left\{\mathbf{F}(p, \mathcal{V}) : p \in \left\{0, \pm\frac{1}{N}, \pm\frac{2}{N}, \dots, \pm\frac{N-1}{N}, \pm 1\right\}\right\}\right)\right). \quad (12)$$

where $\mathbf{F}(p, \mathcal{V})$ was defined in (10). Here ch denotes the convex hull of a set. Thus $\Gamma_N(\mathcal{V})$ is a Lower Pareto frontier of a convex polygon with at most $2N + 1$ vertices.

Now suppose that \mathcal{V} is the Pareto frontier of a convex and compact set. Then we know that $\Phi(\mathcal{V})$ is also the Pareto frontier of a convex and compact set, and we can express the compound operator $\Gamma_N \circ \Phi(\mathcal{V})$ via a set of explicit optimization problems as in (10), that only take \mathcal{V} as input:

$$\begin{aligned} \mathbf{F}(p, \Phi(\mathcal{V})) &= \arg \min_{(x, y)} x + y & (13) \\ \text{s.t. } x &\geq \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \quad \forall b \in B, \quad y \geq \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \quad \forall b \in B, \\ &y = x + p, \quad \boldsymbol{\alpha} \in \Delta(A), \quad \mathbf{Q}(b) \in \mathcal{V} \quad \forall b \in B \end{aligned}$$

If $\mathcal{V} \in \mathcal{F}$ is a Lower Pareto frontier of a convex polygon, then this is a linear programming problem, and further $\Gamma_N \circ \Phi(\mathcal{V})$ is a Lower Pareto frontier of a convex polygon with at most $2N + 1$ vertices. We then we have the following result:

Proposition 4.1. *Let $\mathcal{G}^0 = \{(0, 0)\}$ and let $\mathcal{G}_n = (\Gamma_N \circ \Phi)^n(\mathcal{G}^0)$. Then*

$$d(\mathcal{V}^*, \mathcal{G}_n) \leq \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right) + \beta^n. \quad (14)$$

And thus

$$\limsup_{n \rightarrow \infty} d(\mathcal{V}^*, \mathcal{G}_n) = O\left(\frac{1}{N(1 - \beta)}\right).$$

Hence for any ϵ , there is a pair (N, n) such that $d(\mathcal{V}^*, \mathcal{G}_n) \leq \epsilon$. This result implies an iterative procedure for approximating \mathcal{V}^* by successively applying the compound operator $\Gamma_N \circ \Phi$ to \mathcal{V}^0 , by solving the above linear program at each step. Since \mathcal{G}^n is a Lower Pareto frontier of a convex polygon with at most $2N + 1$ vertices for each n , the size of the linear program remains the same throughout.

4.1 Extracting an approximately optimal policy

From \mathcal{G}^n , one can also extract an approximately optimal strategy π^n in the infinite horizon game. Suppose $\alpha^*(p)$ and $\mathbf{Q}^*(b, p)$ for $b \in B$ are the optimal values that solve the program (13) to compute $\mathbf{F}(p, \Phi(\mathcal{G}^n))$ for different $p \in \{0, \pm\frac{1}{N}, \pm\frac{2}{N}, \dots, \pm\frac{N-1}{N}, \pm 1\}$. Then these define an approximately optimal policy in the following class:

Definition 4.1. A $(2N+1)$ -mode policy π is a mapping from each $p \in \{0, \pm\frac{1}{N}, \pm\frac{2}{N}, \dots, \pm\frac{N-1}{N}, \pm 1\}$ to the pair

1. $\alpha(p) \in \Delta(A)$, and
2. $\left(q(b, p), q'(b, p), \kappa(b, p) \right)$, where for all $b \in B$, $q(b, p), q'(b, p) \in \{0, \pm\frac{1}{N}, \pm\frac{2}{N}, \dots, \pm 1\}$ such that $|q(b, p) - q'(b, p)| = \frac{1}{N}$, and $\kappa(b, p) \in [0, 1]$.

The interpretation is as follows. One starts with some initial *mode*, i.e., a value of $p \in \{0, \pm\frac{1}{N}, \dots, \pm 1\}$. Then at any step, if the current mode is p , then Alice first chooses action $a \in A$ with probability $\alpha_a(p)$. Then if Bob plays action $b \in B$, Alice considers the new mode to be $q(b, p)$ with probability $\kappa(b, p)$, and $q'(b, p)$ with probability $1 - \kappa(b, p)$ and plays accordingly thereafter.

Now, $\alpha^*(p)$ defines $\alpha(p)$ in the policy, and $(q(b, p), q'(b, p), \kappa(b, p))$ are defined such that they satisfy

$$\mathbf{Q}^*(b, p) = \kappa(b, p)\mathbf{F}(q(b, p), \mathcal{G}^n) + (1 - \kappa(b, p))\mathbf{F}(q'(b, p), \mathcal{G}^n).$$

Let \mathcal{V}^{π^n} be the corresponding Pareto frontier that is attained by the policy π_n (each point on this frontier is guaranteed by choosing different possible initial randomizations over the $2N + 1$ modes). We have the following result.

Proposition 4.2.

$$d(\mathcal{V}^{\pi^n}, \mathcal{V}^*) \leq \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right) + 2\beta^n + \frac{1}{N} \left(\frac{2 - \beta^n - \beta^{n+1}}{(1 - \beta)^2} \right). \quad (15)$$

And thus

$$\limsup_{n \rightarrow \infty} d(\mathcal{V}^{\pi^n}, \mathcal{V}^*) = O\left(\frac{1}{N(1 - \beta)^2} \right).$$

Remark: The procedure to approximate the frontier and extract an approximately optimal policy illustrates that our characterization of the minmax optimal policy via the fixed point of a dynamic programming operator opens up the possibility of using several dynamic programming based approximation procedures. Here we have not tried to determine a computation procedure that achieves the optimal error-complexity tradeoff. For fixed (N, n) , in order to approximate the optimal frontier, the procedure needs to solve nN linear programs, each with $O(N)$ variables and constraints, to give the corresponding error bound in the theorem. One can split the error into two terms: the first term is the quantization error which is bounded by $\frac{1}{N(1-\beta)}$ and the iteration error which is bounded by β^n . The second term is relatively benign but the first term requires $N = \frac{1}{(1-\beta)\epsilon}$, which grows rapidly when β is close to 1. For finding an approximately optimal policy, the scaling is like $\frac{1}{(1-\beta)^2\epsilon}$, which grows even faster. Nevertheless, all of this computation can be done offline, and needs to be performed only once. The resulting approximately optimal policy is very simple to implement, and requires a bounded memory ($O(\log \frac{1}{\epsilon(1-\beta)^2})$ bits).

Table 1: Possible loss scenarios

Expert 1	1	0	1	0
Expert 2	0	1	1	0

Table 2: Single-stage regret w.r.t. Expert 1 & 2

Expert 1	(0,1)	(0,-1)	(0,0)	(0,0)
Expert 2	(-1,0)	(1,0)	(0,0)	(0,0)

5 Example: Combining expert advice

Consider the following model of combining expert advice. There are two experts who give Alice recommendation for a decision-making task: say predicting whether a particular stock will rise the next day or fall. Each day Alice decides to act on the recommendation made by one of the experts. The experts' recommendations may be correct or wrong, and if Alice acts on an incorrect recommendation, she bears a loss of 1. Otherwise she does not incur any loss. This model can be represented by Table 1. The rows correspond to the choice made by Alice and the columns correspond to the four different possibilities: either one of the experts gives the correct recommendation and the other doesn't, or both give the correct recommendation or both give the incorrect one. The matrix of single-stage regrets is given in Table 2.

In Figure 3, the computed approximately optimal Pareto frontiers for a range of values of β are shown with the corresponding (theoretical) approximation errors as given by Theorem 4.1.




Figure 3: Approximations of $(1 - \beta)V^*$ for different β values with corresponding errors




Figure 4: Average regret incurred by the 77-mode approximately optimal policy, GPS, and Hedge against 5 arbitrarily chosen strategies of the adversary, for a discount factor $\beta = 0.8$.

5.1 Comparison with other algorithms

In this section, we compare the performance of an approximately optimal algorithm derived from our approach to two known algorithms for the problem of prediction with expert advice with 2 experts.

We will assume a discount factor of 0.8. One algorithm is the well known the exponentially weighted average forecaster, which belongs to the Hedge or “Multiplicative Weights” class of algorithms. In this algorithm, if $L_t(i)$ is the cumulative loss of expert i till time t , then the probability of choosing expert i at time $t + 1$ is

$$p_i(t + 1) \propto \exp(-\eta_t L_t(i)),$$

where $\eta_t = \sqrt{8 \log K / t}$, where K is the number of experts, which in this case is 2. If one uses the discounted cumulative losses in this algorithm, then it achieves an upper bound of $O(\sqrt{1 - \beta})$ on the average discounted regret [CBL03].

The second algorithm is the optimal algorithm given by [GPS14] for the problem with a geometrically distributed time horizon, where at each stage, the game ends with a probability $(1 - \beta)$ and continues with probability β . This is essentially the same as our model of discounted losses, where the discount factor is interpreted as the probability of continuation at each stage. But the difference in that formulation and our formulation is in the definition of regret. In their formulation, the loss of the decision-maker is compared to the expected loss of the best expert in the *realized* time horizon (where the expectation is over the randomness in the time horizon), i.e., to $E_T[\min_{i=1,2} \sum_{t=1}^T l_t(i)]$, where $l_t(i)$ is the loss of expert i at time t . On the other hand, in our formulation, the loss of the decision-maker is compared to the expert with the lowest expected loss, i.e., to $\min_{i=1,2} E_T[\sum_{t=1}^T l_t(i)] = \min_{i=1,2} (1 - \beta) \sum_{t=1}^{\infty} \beta^{t-1} l_t(i)$. Naturally, the optimal regret in their formulation is at least as high as the optimal regret in our formulation, and in fact it turns out to be strictly higher. For example, for $\beta = 0.8$, the optimal regret in their formulation is $1/6 \approx 0.166$, while in our formulation, the regret in this case is 0.136, within an error of up to 0.05 (see Figure 3). Note that in the context where discounting captures the time value of money, our formulation of regret is the natural one; it is difficult to attribute meaning to the notion of regret defined in [GPS14] in this context.

Figure 4 compares the expected regret incurred by the two algorithms (labeled as Hedge and GPS), to the expected regret incurred by a *77 - mode* approximately optimal policy, against 5 different arbitrarily chosen strategies of the adversary.⁹ Note that the expected regret incurred by our policy remains close to the near-optimal value of 0.136, while that of the other two algorithms exceed it by a considerable amount. Also note that as expected, the regret incurred by the GPS algorithm never exceeds the optimal regret value of 0.166.

6 Conclusion and future directions

We presented a novel set-valued dynamic programming approach to characterize and approximate the set of minimal guarantees that a player can achieve in a discounted repeated game with vector losses. In particular, this gives us a way to compute approximately regret-minimizing strategies for playing repeated games with discounted losses. We showed that this optimal set is the fixed point of a contractive dynamic programming operator, and it is the Pareto frontier of some convex and closed set. We also established the structure of the optimal strategies that achieve the different points on this set. Finally we gave a value-iteration based procedure to approximately compute this set and also find approximately optimal strategies for the case where there are two actions. Of course, this set-valued dynamic programming approach can also be used to determine exactly

⁹In adversary’s strategy $i = 1, 2, \dots, 5$, the probability that expert 1 incurs a loss (and expert 2 doesn’t) at time t is $[1 - (1/i + 1)]^{1/t}$ if t is odd, and $[1 - (1/i + 1)]^t$ if t is even.

regret-optimal strategies for repeated games with a finite time horizon. In this case, one can show that the optimal policy will depend on an information state and also the stage t .

The extension of this approach to the case of long-run average losses in infinitely repeated games is much less straightforward, despite the fact that average cost dynamic programming for standard dynamic optimization problems like MDPs is quite well understood. Such an extension could improve our understanding of the construction of no-regret algorithms and would fill the last remaining gap in seeing the classical dynamic programming paradigm as a holistic and methodical approach to regret minimization.

References

- [ABH11] Jacob Abernethy, Peter L. Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of COLT*, 2011.
- [ACBFS02] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The non-stochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [AM95] Robert J. Aumann and Michael Maschler. *Repeated Games with Incomplete Information*. MIT Press, 1995.
- [APS86] Dilip Abreu, David Pearce, and Ennio Stacchetti. Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory*, 39(1):251 – 269, 1986.
- [APS90] Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):pp. 1041–1063, 1990.
- [AWY08] Jacob Abernethy, Manfred K Warmuth, and Joel Yellin. Optimal strategies from random walks. 2008.
- [BCB12] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- [Ber05] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 2005.
- [BKM⁺15] Peter L Bartlett, Wouter M Koolen, Alan Malek, Eiji Takimoto, and Manfred K Warmuth. Minimax fixed-design linear regression. In *Proceedings of The 28th Annual Conference on Learning Theory (COLT)*, pages 226–239, 2015.
- [Bla56a] David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific J. Math.*, 6(1):1–8, 1956.
- [Bla56b] David Blackwell. Controlled random walks. In J De Groot and J.C.H Gerretsen, editors, *Proceedings of the International Congress of Mathematicians 1954*, volume 3, pages 336–338, 1956.
- [BM05] Avrim Blum and Yishay Mansour. From external to internal regret. In *Proceedings of COLT*, pages 621–636, 2005.

- [CBFH⁺97] Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth. How to use expert advice. *J. ACM*, 44(3):427–485, 1997.
- [CBL03] Nicol Cesa-Bianchi and Gbor Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51(3):239–261, 2003.
- [Cov66] Thomas M Cover. Behavior of sequential predictors of binary sequences. Technical report, DTIC Document, 1966.
- [CZ10] Alexey Chernov and Fedor Zhdanov. Prediction with expert advice under discounted loss. In *Algorithmic Learning Theory*, pages 255–269. Springer, 2010.
- [FS99] Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1–2):79–103, 1999.
- [FV97] Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1–2):40–55, 1997.
- [GPS14] Nick Gravin, Yuval Peres, and Balasubramanian Sivan. Towards optimal algorithms for prediction with expert advice. *arXiv preprint arXiv:1409.3040*, 2014.
- [Han57] James Hannan. Approximation to Bayes risk in repeated plays. In M. Dresher, A. W. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games*, volume 3, pages 97–139. Princeton University Press, 1957.
- [KMB14] Wouter M Koolen, Alan Malek, and Peter L Bartlett. Efficient minimax strategies for square loss games. In *Advances in Neural Information Processing Systems*, pages 3230–3238, 2014.
- [KMBA15] Wouter M Koolen, Alan Malek, Peter L Bartlett, and Yasin Abbasi. Minimax time series prediction. In *Advances in Neural Information Processing Systems*, pages 2548–2556, 2015.
- [Koo13] Wouter M Koolen. The pareto regret frontier. In *Advances in Neural Information Processing Systems*, pages 863–871, 2013.
- [Leh03] Ehud Lehrer. Approachability in infinite dimensional spaces. *International Journal of Game Theory*, 31(2):253–268, 2003.
- [LS13] Haipeng Luo and Robert E Schapire. Towards minimax online learning with unknown time horizon. *arXiv preprint arXiv:1307.8187*, 2013.
- [LW94] N. Littlestone and M.K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.
- [Per14] Vianney Perchet. Approachability, regret and calibration: Implications and equivalences. *Journal of Dynamics and Games*, 1(2):181–254, 2014.
- [Rud86] Walter Rudin. *Real and Complex Analysis*. McGraw-Hill, 3rd edition, 1986.
- [Sha53] Lloyd S Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095, 1953.

- [SL05] Gilles Stoltz and Gbor Lugosi. Internal regret in on-line portfolio selection. *Machine Learning*, 59(1–2):125–159, 2005.
- [Sor02] Sylvain Sorin. *A First Course on Zero Sum Repeated Games*. Springer, 2002.
- [Vie92] Nicolas Vieille. Weak approachability. *Mathematics of Operations Research*, 17(4):pp. 781–791, 1992.
- [Vov90] Volodimir G. Vovk. Aggregating strategies. In *Proceedings of COLT*, pages 371–386, 1990.
- [Zam92] Shmuel Zamir. Chapter 5 repeated games of incomplete information: Zero-sum. In Robert Aumann and Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1, pages 109–154. Elsevier, 1992.

7 Appendix

Proof of Proposition 3.1. In order to prove the result, we need the following set of results about the Hausdorff distance:

Lemma 7.1. a) h is a metric on the space of closed subsets of \mathbb{R}^2 .

b) Assume that $\{\mathcal{A}_n\}$ is a Cauchy sequence of closed subsets of $[0, 1]^2$. Then there is a unique closed subset \mathcal{A} of $[0, 1]^2$ such that $h(\mathcal{A}_n, \mathcal{A}) \rightarrow 0$. This set \mathcal{A} is defined as follows:

$$\mathcal{A} = \{\mathbf{x} \in [0, 1]^2 \mid \exists \mathbf{x}_n \in \mathcal{A}_n \text{ s.t. } \mathbf{x}_n \rightarrow \mathbf{x}\}.$$

c) If the sets $\{\mathcal{A}_n\}$ in b) are convex, then \mathcal{A} is convex.

d) $h(\text{up}(\mathcal{A}), \text{up}(\mathcal{B})) \leq h(\mathcal{A}, \mathcal{B})$.

Proof. a)-b) This is the well-known completeness property of the Hausdorff metric; see [Rud86].

c) Say that $\mathbf{x}, \mathbf{y} \in \mathcal{A}$. Then $\mathbf{x} = \lim_n \mathbf{x}_n$ and $\mathbf{y} = \lim_n \mathbf{y}_n$ for $\mathbf{x}_n \in \mathcal{A}_n$ and $\mathbf{y}_n \in \mathcal{A}_n$. By convexity of each \mathcal{A}_n , $\bar{\mathbf{z}}_n := \lambda \mathbf{x}_n + (1 - \lambda) \mathbf{y}_n \in \mathcal{A}_n$. But then, $\bar{\mathbf{z}}_n \rightarrow \bar{\mathbf{z}} := \lambda \mathbf{x} + (1 - \lambda) \mathbf{y}$. It follows that $\bar{\mathbf{z}} \in \mathcal{A}$, so that \mathcal{A} is convex.

d) Let $\epsilon := h(\mathcal{A}, \mathcal{B})$. Pick $\mathbf{x} \in \text{up}(\mathcal{A})$. Then $\mathbf{x} = \mathbf{y} + \mathbf{v}$ for some $\mathbf{y} \in \mathcal{A}$ and $v \succeq 0$. There is some $\mathbf{y}' \in \mathcal{B}$ with $\|\mathbf{y} - \mathbf{y}'\|_\infty \leq \epsilon$. Then $\mathbf{x}' = \min\{\mathbf{y}' + v, \mathbf{1}\} \in \text{up}(\mathcal{B})$, where the minimization is component-wise. We claim that $\|\mathbf{x}' - \mathbf{x}\|_\infty \leq \epsilon$.

If $\mathbf{y}' + \mathbf{v} \in [0, 1]^2$, this is clear. Assume $y'_1 + v_1 > 1$. Then,

$$x'_1 = 1 < y'_1 + v_1 \text{ and } x_1 = y_1 + v_1 \leq 1.$$

Thus,

$$0 \leq x'_1 - x_1 < y'_1 + v_1 - y_1 - v_1 = y'_1 - y_1.$$

Hence, $|x'_1 - x_1| \leq |y'_1 - y_1|$. Similarly, $|x'_2 - x_2| \leq |y'_2 - y_2|$. Thus, one has

$$\|\mathbf{x}' - \mathbf{x}\|_\infty \leq \|\mathbf{y}' - \mathbf{y}\|_\infty \leq \epsilon.$$

□

Now we can prove the proposition. Under the given assumptions, $\{up(\mathcal{V})_n, n \geq 1\}$ is Cauchy in the Hausdorff metric, so that, by Lemma 7.1, there is a unique closed convex set such that $h(up(\mathcal{V})_n, \mathcal{A}) \rightarrow 0$. But since $h(up(\mathcal{V})_n, up(\mathcal{A})) \leq h(up(\mathcal{V})_n, \mathcal{A})$ (from Lemma 7.1), we have that $h(up(\mathcal{V})_n, up(\mathcal{A})) \rightarrow 0$ and hence $up(\mathcal{A}) = \mathcal{A}$. Thus the Pareto frontier \mathcal{V} of \mathcal{A} is then such that $d(\mathcal{V}_n, \mathcal{V}) \rightarrow 0$.

To show uniqueness of \mathcal{V} , assume that there is some $\mathcal{U} \in \mathcal{F}$ such that $d(\mathcal{V}_n, \mathcal{U}) \rightarrow 0$. Then, the closed convex set $up(\mathcal{U})$ is such that $h(up(\mathcal{V})_n, up(\mathcal{U})) \rightarrow 0$. By Lemma 7.1, this implies that $up(\mathcal{U}) = up(\mathcal{V})$, so that $\mathcal{U} = \mathcal{V}$. \square

Proof of Lemma 3.1. In order to prove this lemma, we need a few intermediate results. First, we need the following fact:

Lemma 7.2. *Let \mathcal{V} be the Lower Pareto frontier of a closed convex set. Then \mathcal{V} is closed.*

Proof. Suppose that $\{\mathbf{x}^n\}$ is a sequence of points in \mathcal{V} that converge to some point \mathbf{x} . Then since S is closed, $\mathbf{x} \in S$. We will show that $\mathbf{x} \in \mathcal{V}$. Suppose not. Then there is some $\mathbf{u} \in \mathcal{V}$ such that $\mathbf{u} \preceq \mathbf{x}$. Suppose first that $u_1 < x_1$ and $u_2 < x_2$. Then let $\epsilon = \frac{\min(x_1 - u_1, x_2 - u_2)}{2}$ and consider the ball of radius ϵ around \mathbf{x} , i.e.

$$B_{\mathbf{x}}(\epsilon) = \{\mathbf{y} \in \mathbb{R}^2 : \|\mathbf{y} - \mathbf{x}\|_2 \leq \epsilon\}.$$

Then for any point \mathbf{y} in $B_{\mathbf{x}}(\epsilon)$, we have that $\mathbf{u} \preceq \mathbf{y}$. But since $\{\mathbf{x}^n\}$ converges to \mathbf{x} , there exists some point in the sequence that is in $B_{\mathbf{x}}(\epsilon)$, and \mathbf{u} is dominated by this point, which is a contradiction. Hence either $u_1 = x_1$ or $u_2 = x_2$. Suppose w.l.o.g. that $u_1 < x_1$ and $u_2 = x_2$. See Figure 5. Let $\delta = \frac{x_1 - u_1}{2}$ and consider the ball of radius δ centered at \mathbf{x} , i.e. $B_{\mathbf{x}}(\delta)$. Let \mathbf{x}^n be a point in the sequence such that $\mathbf{x}^n \in B_{\mathbf{x}}(\delta)$. Now $x_1^n > u_1$ and hence it must be that $x_2^n < u_2$.

Now some $\lambda \in (0, 1)$, consider a point $\mathbf{r} = \lambda \mathbf{u} + (1 - \lambda) \mathbf{x}^n$ such that $r_1 = u_1 + \frac{\delta}{2}$. It is possible to pick such a point since $x_1 = u_1 + 2\delta$ and $|x_1^n - x_1| \leq \delta$, and hence $x_1^n > u_1 + \frac{\delta}{2}$ (please see the figure). Now $\mathbf{r} \in S$ since S is convex. Now $r_1 = x_1 - \frac{3\delta}{2} < x_1$ and also $r_2 < u_2 = x_2$ since $\lambda > 0$




Figure 5: Construction in the proof of Lemma 7.2.

and $x_2^n < u_2$. Let $\delta' = \frac{x_2 - r_2}{2}$. Then consider the ball $B_{\mathbf{x}}(\delta')$ centered at \mathbf{x} . Clearly $\mathbf{r} \preceq \mathbf{y}$ for any $\mathbf{y} \in B_{\mathbf{x}}(\delta')$. But since $\{\mathbf{x}^n\}$ converges to \mathbf{x} , there exist some point in the sequence that is in $B_{\mathbf{x}}(\delta')$, and \mathbf{r} is dominated by this point, which is again a contradiction. Thus $\mathbf{x} \in \mathcal{V}$ and hence \mathcal{V} is closed. \square

Remark: The Pareto frontier of a closed set need not be a closed set, as the example in Figure 6 shows.




Figure 6: A closed set S whose Pareto frontier \mathcal{V} is not closed.

Next we define the following notion of convexity of pareto frontiers.

Definition 7.1. A Pareto frontier \mathcal{V} is p -convex if for any $\mathbf{v}, \mathbf{u} \in \mathcal{V}$ and for each $\lambda \in [0, 1]$, there exists a point $\mathbf{r}(\lambda) \in \mathcal{V}$ such that $\mathbf{r}(\lambda) \preceq \lambda \mathbf{v} + (1 - \lambda) \mathbf{u}$.

We then show the following equivalences.

Lemma 7.3. For a Pareto frontier $\mathcal{V} \subset [0, 1]^2$, the following statements are equivalent:

1. \mathcal{V} is p -convex and a closed set.
2. \mathcal{V} is p -convex and the lower Pareto frontier of a closed set $\mathcal{S} \subseteq [0, 1]^2$.
3. \mathcal{V} is the lower Pareto frontier of a closed convex set $H \subseteq [0, 1]^2$.

Proof. 1 is a special case of 2 and hence 1 implies 2. To show that 2 implies 3, if \mathcal{S} is convex then there is nothing to prove. So assume \mathcal{S} is not convex. Then let \mathcal{H} be the convex hull of \mathcal{S} . First, since $[0, 1] \times [0, 1]$ is convex, $\mathcal{H} \subseteq [0, 1] \times [0, 1]$. Then since \mathcal{S} is closed and bounded, it is also compact. Hence \mathcal{H} is the convex hull of a compact set, which is compact and hence closed and bounded. Now we will show that \mathcal{V} is the Lower Pareto frontier of \mathcal{H} . To see this, any $\mathbf{u} \in \mathcal{S}$ is of the form $\mathbf{u} = \lambda \mathbf{x} + (1 - \lambda) \mathbf{y}$ where $\mathbf{x}, \mathbf{y} \in \mathcal{S}$. But then there are points $\mathbf{x}', \mathbf{y}' \in \mathcal{V}$ such that $\mathbf{x}' \preceq \mathbf{x}$ and $\mathbf{y}' \preceq \mathbf{y}$. Thus we have that

$$\lambda \mathbf{x}' + (1 - \lambda) \mathbf{y}' \preceq \lambda \mathbf{x} + (1 - \lambda) \mathbf{y} = \mathbf{u}.$$

But since \mathcal{V} is convex, there exists some $\mathbf{r}(\lambda) \in \mathcal{V}$ such that

$$\mathbf{r}(\lambda) \preceq \lambda \mathbf{x}' + (1 - \lambda) \mathbf{y}' \preceq \lambda \mathbf{x} + (1 - \lambda) \mathbf{y} = \mathbf{u}.$$

Thus the Pareto frontier of \mathcal{H} is a subset of \mathcal{V} , but since $\mathcal{V} \in \mathcal{H}$ and it is a Pareto frontier, \mathcal{V} is the Lower Pareto frontier of \mathcal{H} . Finally Lemma 7.2 shows that 3 implies that \mathcal{V} is closed. To show it is convex, suppose that \mathbf{u} and \mathbf{v} are two points in \mathcal{V} . Since they also belong to H , which is convex, for each $\lambda \in [0, 1]$, $\lambda \mathbf{u} + (1 - \lambda) \mathbf{v} \in \mathcal{S}$ and thus there is some $\mathbf{r}(\lambda) \in \mathcal{V}$ such that $\mathbf{r}(\lambda) \preceq \lambda \mathbf{u} + (1 - \lambda) \mathbf{v}$. Thus \mathcal{V} is convex. \square

We can now finally prove the lemma. Note that $\Psi(\mathcal{S})$ is the image of the continuous function f from the product space $\mathcal{S}^{m \times n} \times \Delta(A)$ to a point in \mathbb{R}^2 , which is a Hausdorff space. Since \mathcal{S} is closed and bounded, it is compact. Also the simplex $\Delta(A)$ is compact. Thus by Tychonoff's theorem, the product space $\mathcal{S}^{m \times n} \times \Delta(A)$ is compact. Hence by the closed map lemma, f is a closed map and hence $\Psi(\mathcal{S})$ is closed.

Now assume that \mathcal{S} is a closed convex set. Then $\Lambda(\mathcal{S})$ exists by Lemma 3.1 and further it is p -convex by Lemma 7.3. Let $\mathcal{U} = \Lambda(\mathcal{S})$. Clearly, $\Lambda(\Psi(\mathcal{S})) = \Lambda(\Psi(\mathcal{U}))$. Recall that any point \mathbf{u} in $\Lambda(\Psi(\mathcal{U}))$ is of the form:

$$\mathbf{u} = \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_1(a, b) + \beta R_1(a, b)] \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a [r_2(a, b) + \beta R_2(a, b)] \right\} \right)$$

for some $\alpha \in \Delta(A)$ and $\mathbf{R}(a, b) \in \mathcal{U}$. But since \mathcal{U} is p -convex, for each $b \in B$, there exists some $\mathbf{Q}(b) \in \mathcal{U}$ such that $\mathbf{Q}(b) \preceq \sum_{a=1}^m \alpha_a \mathbf{R}(a, b)$. Hence statement 1 follows.

Now let

$$\mathbf{u} = \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} \right)$$

and

$$\mathbf{v} = \left(\max_{b \in B} \left\{ \sum_{a \in A} \eta_a r_1(a, b) + \beta R_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \eta_a r_2(a, b) + \beta R_2(b) \right\} \right)$$

be two points in $\Lambda(\Psi(\mathcal{U}))$, where $\alpha, \eta \in \Delta(A)$ and $\mathbf{Q}(b), \mathbf{R}(b) \in \mathcal{V}$.

For a fixed $\lambda \in [0, 1]$, let $\kappa_a = \alpha_a \lambda + \eta_a (1 - \lambda)$. Then

$$\begin{aligned} \lambda \mathbf{u} + (1 - \lambda) \mathbf{v} &= \left(\lambda \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \right\} + (1 - \lambda) \max_{b \in B} \left\{ \sum_{a \in A} \eta_a r_1(a, b) + \beta R_1(b) \right\}, \right. \\ &\quad \left. \lambda \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} + (1 - \lambda) \max_{b \in B} \left\{ \sum_{a \in A} \eta_a r_2(a, b) + \beta R_2(b) \right\} \right) \\ &\succeq \left(\max_{b \in B} \left\{ \sum_{a \in A} \kappa_a r_1(a, b) + \beta \lambda Q_1(b) + (1 - \lambda) R_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \kappa_a r_2(a, b) + \beta \lambda Q_2(b) + (1 - \lambda) R_2(b) \right\} \right) \\ &\succeq \left(\max_{b \in B} \left\{ \sum_{a \in A} \kappa_a r_1(a, b) + \beta L_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \kappa_a r_2(a, b) + \beta L_2(b) \right\} \right) \end{aligned}$$

The first inequality holds since \max is a convex function and the second follows since \mathcal{U} is p -convex, and hence $\mathbf{L}(b) = (L_1(b), L_2(b)) \in \mathcal{U}$ that satisfy the given relation exist. Thus $\Lambda(\Psi(\mathcal{S}))$ is p -convex. And hence from Lemma 7.3, it is the Lower Pareto frontier of a closed convex set in $[0, 1]^2$, i.e., it is in \mathcal{F} . \square

Proof of Lemma 3.2. In order to prove this Lemma, we first define another metric on the space \mathcal{F} that is equivalent to d .

Definition 7.2. For two Pareto frontiers \mathcal{A} and \mathcal{B} of $[0, 1]^2$, we define

$$e(\mathcal{A}, \mathcal{B}) \triangleq$$

$$\inf\{\epsilon \geq 0 : \forall \mathbf{u} \in \mathcal{A}, \exists \mathbf{v} \in \mathcal{B} \text{ s.t. } \mathbf{v} \preceq \mathbf{u} + \epsilon \mathbf{1} \text{ and } \forall \mathbf{v} \in \mathcal{B}, \exists \mathbf{u} \in \mathcal{A} \text{ s.t. } \mathbf{u} \preceq \mathbf{v} + \epsilon \mathbf{1}\}. \quad (16)$$

Here $\mathbf{1} = (1, 1)$.

We can show that the two metrics are equivalent.

Lemma 7.4.

$$e(\mathcal{A}, \mathcal{B}) = d(\mathcal{A}, \mathcal{B}).$$

Proof. Suppose that $e(\mathcal{A}, \mathcal{B}) \leq \epsilon$. Consider a point $\mathbf{x} \in \text{up}(\mathcal{A})$ such that $\mathbf{x} = \mathbf{y} + \mathbf{v}$ where $\mathbf{y} \in \mathcal{A}$ and $\mathbf{v} \succeq 0$. Suppose that there is no $\mathbf{x}' \in \text{up}(\mathcal{B})$ such that $\|\mathbf{x} - \mathbf{x}'\|_\infty \leq \epsilon$. This means that $\text{up}(\mathcal{B})$ is a subset of the region \mathcal{S} shown in the Figure 7. But since $\mathbf{y} = \mathbf{x} - \mathbf{v}$, \mathbf{y} is in region \mathcal{S}' . But for any $\mathbf{u} \in \mathcal{S}$ and $\bar{w} \in \mathcal{S}'$, $\|\mathbf{u} - \bar{w}\|_\infty \geq \epsilon$. This contradicts the fact that for \mathbf{y} there is some $\mathbf{y}' \in \mathcal{B}$, such that $\mathbf{y} + \epsilon \mathbf{1} \succeq \mathbf{y}'$. Thus $d(\mathcal{A}, \mathcal{B}) \leq \epsilon$. Now suppose that $d(\mathcal{A}, \mathcal{B}) \leq \epsilon$. Then for any $\mathbf{x} \in \mathcal{A}$,




Figure 7: Construction for the proof of Lemma 7.4

there is a $\mathbf{x}' \in \text{up}(\mathcal{B})$ such that $\|\mathbf{x} - \mathbf{x}'\|_\infty \leq \epsilon$ where $\mathbf{x}' = \mathbf{y} + \mathbf{v}$ for $\mathbf{y} \in \mathcal{B}$ and $\mathbf{v} \succeq 0$. Thus $\mathbf{x} + \epsilon \mathbf{1} \succeq \mathbf{x}' = \mathbf{y} + \mathbf{v}$. Thus $e(\mathcal{A}, \mathcal{B}) \leq \epsilon$. \square

Now suppose $e(\mathcal{U}, \mathcal{V}) = \epsilon$. Let

$$\left(\max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_1(a, b) + \beta R_1(b) \right\}, \max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_2(a, b) + \beta R_2(b) \right\} \right)$$

be some point in $\Phi(\mathcal{V})$, where $\boldsymbol{\alpha} \in \Delta(\mathcal{A})$. Then for each $\mathbf{R}(b)$, for each b , we can choose $\mathbf{Q}(b) \in \mathcal{U}$ such that $\mathbf{Q}(b) \preceq \mathbf{R}(b) + \epsilon \mathbf{1}$. We then have

$$\begin{aligned} \max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_1(a, b) + \beta Q_1(b) \right\} &= \max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_1(a, b) + \beta R_1(b) + \beta(Q_1(b) - R_1(b)) \right\} \\ &\leq \max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_1(a, b) + \beta R_1(b) + \beta \epsilon \right\} \\ &= \max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_1(a, b) + \beta R_1(b) \right\} + \beta \epsilon. \end{aligned}$$

Similarly, we can show that

$$\max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} \leq \max_{b \in \mathcal{B}} \left\{ \sum_{a \in \mathcal{A}} \alpha_a r_2(a, b) + \beta R_2(b) \right\} + \beta \epsilon. \quad (17)$$

Thus

$$\begin{aligned} & \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} \right) \\ & \preceq \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta R_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta R_2(b) \right\} \right) + \beta \epsilon \mathbf{1}. \end{aligned}$$

But since

$$\left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} \right) \in \Psi(\mathcal{U}),$$

and since $\Phi(\mathcal{U}) = \Lambda(\Psi(\mathcal{U}))$, there exists some $(L_1, L_2) \in \Phi(\mathcal{U})$ such that

$$(L_1, L_2) \preceq \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta Q_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta Q_2(b) \right\} \right).$$

Thus

$$(L_1, L_2) \preceq \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_1(a, b) + \beta R_1(b) \right\}, \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a r_2(a, b) + \beta R_2(b) \right\} \right) + \beta \epsilon \mathbf{1}.$$

We can show the other direction (roles of $\Phi(\mathcal{U})$ and $\Phi(\mathcal{V})$ reversed) similarly and thus we have that

$$e(\Phi(\mathcal{U}), \Phi(\mathcal{V})) \leq \beta \epsilon = \beta e(\mathcal{U}, \mathcal{V}). \quad (18)$$

□

Proof of Theorem 3.1. Since Φ is a contraction in the metric d , the sequence $\{A_n\}$ is Cauchy in \mathcal{F} . Hence by Lemma 7.1, $\{A_n\}$ converges to a Pareto frontier $\mathcal{V}^* \in \mathcal{F}$. The continuity of the operator further implies that $\mathcal{V}^* = \Phi(\mathcal{V}^*)$. To show uniqueness, observe that if there are two fixed points \mathcal{U} and \mathcal{V} , then we have $d(\mathcal{U}, \mathcal{V}) = d(\Phi(\mathcal{U}), \Phi(\mathcal{V})) \leq \beta d(\mathcal{U}, \mathcal{V})$, which implies that $d(\mathcal{U}, \mathcal{V}) = 0$ and hence $\mathcal{U} = \mathcal{V}$. □

Proof of Theorem 3.2. In \mathbb{G}^∞ , fix $T \geq 1$ and consider a truncated game where Alice can guarantee the cumulative losses in $\beta^{T+1} \mathcal{V}^*$ after time $T + 1$. Then the minimal losses that she can guarantee after time T is the set:

$$\begin{aligned} & \Lambda \left(\left\{ \max_{b \in B} \beta^T \sum_{a \in A} \alpha_a r_1(a, b) + \beta^{T+1} Q_1(b), \max_{b \in B} \beta^T \sum_{a \in A} \alpha_a r_2(a, b) + \beta^{T+1} Q_2(b) \right. \right. \\ & \quad \left. \left. | \alpha \in \Delta(A), \mathbf{Q}(b) \in \mathcal{V}^* \forall b \in B \right\} \right). \end{aligned}$$

This set is $\beta^T \mathcal{V}^*$. By induction, this implies that the set of minimal losses that she can guarantee after time 0 is \mathcal{V}^* .

The losses of the truncated game and of the original game differ only after time $T + 1$. Since the losses at each step are bounded by $(1 - \beta)$, the cumulative losses after time $T + 1$ are bounded by $\frac{\beta^{T+1}(1-\beta)}{1-\beta} = \beta^{T+1}$. Consequently, the minimal losses of the original game must be in the set

$$\left\{ \mathbf{u} \in [0, 1]^2 : u_1 \in [x_1 - \beta^{T+1}, x_1 + \beta^{T+1}], u_2 \in [x_2 - \beta^{T+1}, x_2 + \beta^{T+1}], x \in \mathcal{V}^* \right\}.$$

Since $T \geq 1$ is arbitrary, the minimal losses that Alice can guarantee in the original game must be in \mathcal{V}^* . \square

Proof of Theorem 3.3. Assume that Alice can guarantee every pair $\beta^{T+1}\mathbf{u}$ of cumulative losses with $\mathbf{u} \in \mathcal{V}^*$ after time $T + 1$ by choosing some continuation strategy in Π_A . Let $\mathbf{x} = \mathbf{F}(p, \mathcal{V}^*)$. We claim that after time T , Alice can guarantee a loss of no more than $\beta^T \mathbf{x}$ on each component by first choosing $a_T = a$ with probability $\alpha_a(p)$ and then if Bob chooses $b \in B$, choosing a continuation strategy that guarantees her $\mathbf{F}(p', \mathcal{V}^*)$, where $p' = q(b, p)$. Indeed by following this strategy, her expected loss after time T is then

$$\left\{ \beta^T \sum_a \alpha_a(p) r_k(a, b) + \beta^{T+1} F_k(q(b, p), \mathcal{V}^*) \right\} \leq \beta^T F_k(p, \mathcal{V}^*) = \beta^T x_k$$

when the game is G_k . Thus, this strategy for Alice guarantees that her loss after time T is no more than $\beta^T \mathcal{V}^*$. Hence by induction, following the indicated strategy (in the statement of the theorem) for the first T steps and then using the continuation strategy from time $T + 1$ onwards, guarantees that her loss is not more than $\mathbf{F}(p_1, \mathcal{V}^*)$ after time 0. Now, even if Alice plays arbitrarily after time $T + 1$ after following the indicated strategy for the first T steps, she still guarantees that her loss is no more than $\mathbf{F}(p_1, \mathcal{V}^*) + \beta^{T+1}(1, 1)^T$. Since this is true for arbitrarily large values of T , playing the given policy indefinitely guarantees that her loss is no more than $\mathbf{F}(p_1, \mathcal{V}^*)$. \square

Proof of Proposition 4.1. We first need the following lemma about the approximation operator.

Lemma 7.5. *Consider a $\mathcal{V} \in \mathcal{F}$. Then*

$$d(\mathcal{V}, \Gamma_N(\mathcal{V})) \leq \frac{1}{N}.$$

Proof. Any point in $\Gamma_N(\mathcal{V})$ is of the form $\lambda \mathbf{u} + (1 - \lambda) \mathbf{v}$ where $\mathbf{u}, \mathbf{v} \in \mathcal{V}$. By the p -convexity of \mathcal{V} , there is some $\mathbf{r}(\lambda) \in \mathcal{V}$, such that $\mathbf{r}(\lambda) \preceq \lambda \mathbf{u} + (1 - \lambda) \mathbf{v}$. Also clearly for any $\mathbf{u} \in \mathcal{V}$,

$$\min \left\{ \|\mathbf{u} - \mathbf{v}\|_\infty : \mathbf{v} \in \Gamma_N(\mathcal{V}) \right\} \leq \max \left\{ \|\mathbf{F}(p, \mathcal{V}) - \mathbf{F}(p', \mathcal{V})\|_\infty : |p' - p| = \frac{1}{N} \right\} = \frac{1}{N}.$$

\square

Next, we have $\mathcal{G}_n = \Gamma_N \circ \Phi(\mathcal{G}_{n-1})$. Consider another sequence of Pareto frontiers

$$\left(\mathcal{A}_n = \Phi^n(\mathcal{G}_0) \right)_{n \in \mathbb{N}} \tag{19}$$

Then we have

$$\begin{aligned}
d(\mathcal{A}_n, \mathcal{G}_n) &= d(\Phi(\mathcal{A}_{n-1}), \Gamma_N(\Phi(\mathcal{G}_{n-1}))) \\
&\stackrel{(a)}{\leq} d(\Phi(\mathcal{A}_{n-1}), \Phi(\mathcal{G}_{n-1})) + d(\Phi(\mathcal{G}_{n-1}), \Gamma_N(\Phi(\mathcal{G}_{n-1}))) \\
&\stackrel{(b)}{\leq} \beta d(\mathcal{A}_{n-1}, \mathcal{G}_{n-1}) + \frac{1}{N}
\end{aligned} \tag{20}$$

where inequality (a) is the triangle inequality and (b) follows from (18) and Lemma 7.5. Coupled with the fact that $d(\mathcal{A}_0, \mathcal{G}_0) = 0$, we have that

$$\begin{aligned}
d(\mathcal{A}_n, \mathcal{G}_n) &\leq \frac{1}{N} \left(1 + \beta + \beta^2 + \dots + \beta^{n-1} \right) \\
&= \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right)
\end{aligned} \tag{21}$$

Since Φ is a contraction, the sequence $\{\mathcal{A}_n\}$ converges to some pareto frontier \mathcal{V}^* . Suppose that we stop the generation of the sequences $\{\mathcal{A}_n\}$ and $\{\mathcal{G}_n\}$ at some n . Now since $\mathcal{A}_0 = \mathcal{G}_0 = \{(0, 0)\}$, and since the stage payoffs $r_k(a, b) \in [0, 1 - \beta]$, we have that $d(\mathcal{A}_1, \mathcal{A}_0) \leq 1 - \beta$. This implies from the contraction property that $d(\mathcal{V}^*, \mathcal{A}_n) \leq \frac{\beta^n(1-\beta)}{1-\beta} = \beta^n$ and thus by triangle inequality we have

$$d(\mathcal{V}^*, \mathcal{G}_n) \leq \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right) + \beta^n. \tag{22}$$

□

Proof of Proposition 4.2. In order to prove this result, we need a few intermediate definitions and results. First, we need to characterize the losses guaranteed by any $(2N + 1)$ -mode policy. Such a policy π defines the following operator on any function $\mathbf{F} : \{0, \pm \frac{1}{N}, \pm \frac{2}{N}, \dots, \pm \frac{N-1}{N}, \pm 1\} \rightarrow \mathbb{R}^2$:

$$\begin{aligned}
\Delta_N^\pi(\mathbf{F})(p) &= \left(\max_{b \in B} \left\{ \sum_{a \in A} \alpha_a(p) r_1(a, b) + \kappa(b, p) \beta F_1(q(b, p)) + (1 - \kappa(b, p)) \beta F_1(q'(b, p)) \right\}, \right. \\
&\quad \left. \max_{b \in B} \left\{ \sum_{a \in A} \alpha_a(p) r_2(a, b) + \kappa(b, p) \beta F_2(q(b, p)) + (1 - \kappa(b, p)) \beta F_2(q'(b, p)) \right\} \right).
\end{aligned} \tag{23}$$

For a function $\mathbf{F} : \{0, \pm \frac{1}{N}, \pm \frac{2}{N}, \dots, \pm \frac{N-1}{N}, \pm 1\} \rightarrow \mathbb{R}^2$, define the following norm:

$$\|\mathbf{F}\| = \max_{p \in \{0, \pm \frac{1}{N}, \pm \frac{2}{N}, \dots, \pm \frac{N-1}{N}, \pm 1\}} \|\mathbf{F}(p)\|_\infty.$$

We can easily show that Δ_N^π is a contraction in the norm.

Lemma 7.6.

$$\|\Delta_N^\pi(\mathbf{F}) - \Delta_N^\pi(\mathbf{G})\| \leq \beta \|\mathbf{F} - \mathbf{G}\|. \tag{24}$$

We can then show the following result.

Lemma 7.7. Consider a $(2N + 1)$ -mode policy π . Then there is a unique function

$$\mathbf{F}^\pi : \{0, \pm \frac{1}{N}, \pm \frac{2}{N}, \dots, \pm \frac{N-1}{N}, \pm 1\} \rightarrow \mathbb{R}^2$$

such that $\Delta_N^\pi(\mathbf{F}^\pi) = \mathbf{F}^\pi$. Further, The policy π initiated at mode p where $p \in \{0, \pm \frac{1}{N}, \pm \frac{2}{N}, \dots, \pm 1\}$, guarantees the vector of losses $\mathbf{F}^\pi(p)$.

The first part of the result follows from the fact that the operator is a contraction and the completeness of the space of vector-valued functions with a finite domain for the given norm. The second part follows from arguments similar to those in the proof of Theorem 3.3.

Now let

$$\mathcal{V}^{\pi_n} = \Lambda \left(ch(\{\mathbf{F}^{\pi_n}(p) : p \in \{0, \pm \frac{1}{N}, \dots, \pm \frac{N-1}{N}, \pm 1\}\}) \right),$$

where \mathbf{F}^{π_n} is the fixed point of the operator $\Delta_N^{\pi_n}$.

Now define a sequence of functions $\mathbf{F}^n : \{0, \pm \frac{1}{N}, \pm \frac{2}{N}, \dots, \pm \frac{N-1}{N}, \pm 1\} \rightarrow \mathbb{R}^2$ where $\mathbf{F}^n(p) = \mathbf{F}(p, \Phi(\mathcal{G}^{n-1})) = \mathbf{F}(p, \mathcal{G}^n)$. We then have that

$$d(\mathcal{V}^{\pi_n}, \mathcal{V}^*) \leq d(\mathcal{V}^{\pi_n}, \mathcal{G}_n) + d(\mathcal{G}_n, \mathcal{V}^*) \quad (25)$$

$$\leq d(\mathcal{V}^{\pi_n}, \mathcal{G}_n) + \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right) + \beta^n. \quad (26)$$

From the definition of d , it is clear that $d(\mathcal{V}^{\pi_n}, \mathcal{G}_n) \leq \|\mathbf{F}^{\pi_n} - \mathbf{F}^n\|$. Next we have

$$\|\mathbf{F}^{\pi_n} - \mathbf{F}^n\| \leq \|\mathbf{F}^{\pi_n} - \Delta_N^{\pi_n}(\mathbf{F}^n)\| + \|\Delta_N^{\pi_n}(\mathbf{F}^n) - \mathbf{F}^n\| \quad (27)$$

$$\stackrel{(a)}{=} \|\Delta_N^{\pi_n}(\mathbf{F}^{\pi_n}) - \Delta_N^{\pi_n}(\mathbf{F}^n)\| + \|\mathbf{F}^{n+1} - \mathbf{F}^n\| \quad (28)$$

$$\stackrel{(b)}{\leq} \beta \|\mathbf{F}^{\pi_n} - \mathbf{F}^n\| + \|\mathbf{F}^{n+1} - \mathbf{F}^n\|. \quad (29)$$

$$(30)$$

Here (a) holds because $\Delta_N^{\pi_n}(\mathbf{F}^n) = \mathbf{F}^{n+1}$ by the definition of the policy π_n , and because \mathbf{F}^{π_n} is a fixed point of the operator $\Delta_N^{\pi_n}$. (b) holds because $\Delta_N^{\pi_n}$ is a contraction. Thus we have

$$d(\mathcal{V}^{\pi_n}, \mathcal{G}_n) \leq \|\mathbf{F}^{\pi_n} - \mathbf{F}^n\| \leq \frac{\|\mathbf{F}^{n+1} - \mathbf{F}^n\|}{1 - \beta}. \quad (31)$$

And finally we have:

$$d(\mathcal{V}^{\pi_n}, \mathcal{V}^*) \leq \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right) + \beta^n + \frac{\|\mathbf{F}^{n+1} - \mathbf{F}^n\|}{1 - \beta}. \quad (32)$$

To finish up, we need the following result:

Lemma 7.8.

$$\|\mathbf{F}^{n+1} - \mathbf{F}^n\| \leq d(\mathcal{G}_{n+1}, \mathcal{G}_n).$$

Proof. Let $\mathbf{u} = \mathbf{F}^{n+1}(p)$ and $\mathbf{v} = \mathbf{F}^n(p)$ for some p . Now \mathbf{u} is the point of intersection of \mathcal{G}_{n+1} and the line $y = x + p$. \mathbf{v} is the point of intersection of the frontier \mathcal{G}_n and the line $y = x + p$. Now suppose that $\|\mathbf{u} - \mathbf{v}\|_\infty > d(\mathcal{G}_{n+1}, \mathcal{G}_n)$. Then either for \mathbf{u} , there is no $\mathbf{r} \in \mathcal{G}_n$ such that $\mathbf{r} \preceq \mathbf{u} + \mathbf{1}d(\mathcal{G}_{n+1}, \mathcal{G}_n)$ or for \mathbf{v} , there is no $\mathbf{r} \in \mathcal{G}_{n+1}$ such that $\mathbf{r} \preceq \mathbf{v} + \mathbf{1}d(\mathcal{G}_{n+1}, \mathcal{G}_n)$. Either of the two cases contradict the definition of $d(\mathcal{G}_{n+1}, \mathcal{G}_n)$. Thus $\|\mathbf{u} - \mathbf{v}\|_\infty \leq d(\mathcal{G}_{n+1}, \mathcal{G}_n)$. \square

Finally, by the triangle inequality we have

$$d(\mathcal{G}_{n+1}, \mathcal{G}_n) \leq d(\mathcal{A}_{n+1}, \mathcal{A}_n) + d(\mathcal{G}_{n+1}, \mathcal{A}_{n+1}) + d(\mathcal{G}_n, \mathcal{A}_n) \quad (33)$$

$$\leq (1 - \beta)\beta^n + \frac{1}{N} \left(\frac{1 - \beta^{n+1}}{1 - \beta} \right) + \frac{1}{N} \left(\frac{1 - \beta^n}{1 - \beta} \right). \quad (34)$$

Combining with (32) we have the result. □

Table 3: An approximately optimal 77-mode policy for $\beta = 0.8$.

Modes	Pr. of choosing Ex. 1	Transition if Ex. 1 incurs loss	Transition if Ex. 2 incurs loss
38	1	37	38
37	0.9999	36	38
36	0.9998	35	38
35	0.9996	33	38
34	0.9992	32	38
33	0.9987	31	38
32	0.998	30	38
31	0.9971	28	38
30	0.9958	27	38
29	0.9942	26	38
28	0.992	25	38
27	0.9894	23	38
26	0.9862	22	38
25	0.9823	21	38
24	0.9775	20	38
23	0.972	18	38
22	0.9654	17	38
21	0.9577	16	37
20	0.9489	15	35
19	0.9389	13	34
18	0.9276	12	33
17	0.9149	11	32
16	0.9008	10	30
15	0.8852	8	29
14	0.8682	7	28
13	0.8497	6	27
12	0.8296	5	25
11	0.8082	3	24
10	0.7855	2	23
9	0.7612	1	22
8	0.7357	0	20
7	0.709	2	19
6	0.6813	-3	18
5	0.6525	-4	17
4	0.623	-5	15
3	0.5928	-7	14
2	0.5621	-8	13
1	0.5311	-9	12
0	0.5	-10	10

Table 4: An approximately optimal 77-mode policy for $\beta = 0.8$ (continued).

Modes	Pr. of choosing Ex. 1	Transition if Ex. 1 incurs loss	Transition if Ex. 2 incurs loss
-1	0.4687	-12	9
-2	0.4377	-13	8
-3	0.407	-14	7
-4	0.3768	-15	5
-5	0.3473	-17	4
-6	0.3185	-18	3
-7	0.2908	-19	-2
-8	0.2642	-20	0
-9	0.2387	-22	-1
-10	0.2143	-23	-2
-11	0.1916	-24	-3
-12	0.1702	-25	-5
-13	0.1501	-27	-6
-14	0.1316	-28	-7
-15	0.1147	-29	-8
-16	0.0992	-30	-10
-17	0.0851	-32	-11
-18	0.0724	-33	-12
-19	0.0611	-34	-13
-20	0.0511	-35	-15
-21	0.0422	-37	-16
-22	0.0346	-38	-17
-23	0.028	-38	-18
-24	0.0224	-38	-20
-25	0.0177	-38	-21
-26	0.0137	-38	-22
-27	0.0106	-38	-23
-28	0.008	-38	-25
-29	0.0058	-38	-26
-30	0.0042	-38	-27
-31	0.0029	-38	-28
-32	0.002	-38	-30
-33	0.0013	-38	-31
-34	0.0008	-38	-32
-35	0.0004	-38	-33
-36	0.0002	-38	-35
-37	0.0001	-38	-36
-38	0	-38	-37