Attainability in Repeated Games with Vector Payoffs*

Dario Bauso†, Ehud Lehrer‡, Eilon Solan§, and Xavier Venel ¶

May 24, 2013

Abstract

We introduce the concept of attainable sets of payoffs in two-player repeated games with vector payoffs. A set of payoff vectors is called *attainable* by a player if there is a finite horizon T such that the player can guarantee that after time T the distance between the set and the cumulative payoff is arbitrarily small, regardless of the strategy Player 2 is using. We provide a necessary and sufficient condition for the attainability of a convex set, using the concept of B-sets. We then particularize the condition to the case in which the set is a singleton, and provide some equivalent conditions. We finally characterize when all vectors are attainable.

Keywords: Attainability, continuous time, repeated games, vector-payoffs, dynamic games, approachability.

JEL classification: C73, C72

*This research was supported in part by the Google Inter-university center for Electronic Markets and Auctions. Lehrer acknowledges the support of the Israel Science Foundation, Grant #538/11. Solan acknowledges the support of the Israel Science Foundation, Grant #212/09. Venel acknowledges the support of the Israel Science Foundation, Grant #1517/11.

†Dipartimento di Ingegneria Chimica, Gestionale, Informatica, Meccanica, Università di Palermo, viale delle Scienze 90128, Palermo, Italy. e-mail: dario.bauso@unipa.it

‡The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel and INSEAD, Boulevard de Constance, 77305 Fontainebleau, France. e-mail: lehrer@post.tau.ac.il

§The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. e-mail: eilons@post.tau.ac.il

¶The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. e-mail: xavier.venel@gmail.com

1 Introduction

In various dynamic situations the stage-payoff is multidimensional, and the goal of the decision maker is to drive the total vector-payoff as close as possible to a given target set. One such example is dynamic network models, which include a variety of logistic applications such as production, distribution and transportation networks. In the literature on dynamic network flow control [4, 5, 11, 12, 15], the supplier tries to meet a multidimensional demand. His goal is to ensure that the difference between the *total* demand and the *total* supply converges with time to a desirable target. One can model such a situation as a two-player repeated game, where Player 1 is the decision maker and Player 2 represents the adversarial market that controls demand. In the distribution network scenario, for instance, the supplier has a desirable multidimensional inventory level that he would like to maintain, despite erratic behavior of the demand side. Having to deal with an adversarial opponent requires the supplier to cope with the worst possible scenario. This motivates our main objective: to find conditions that characterize when a specific target set can be attained under any possible demand pattern exhibited by the market.

A second example is the Capital Adequacy Ratio. The third Basel Accord states that (a) the bank's Common Equity Tier 1 must be at least 4.5% of its risk-weighted assets at all times, (b) the bank's Tier 1 Capital must be at least 6.0% of its risk-weighted assets at all times, and (c) the total capital, that is, Tier 1 Capital plus Tier 2 Capital, must be at least 8.0% of the bank's risk weighted assets at all times. To accommodate this example in our setup, consider the following 3-dimensional vector. The first coordinate stands for the per-period difference between the bank's Common Equity Tier 1 and 4.5% of its risk-weighted assets; the second coordinate stands for the per-period difference between the bank's Tier 1 Capital and 6.0% of its risk-weighted assets; and the third stands for the per-period difference between the total capital and 8.0% of the bank's risk weighted assets. According to the Capital Adequacy Ratio the coordinates of this vector should be nonnegative. Here, Player 1 represents the bank's managers who control its assets, and Player 2 represents market behavior, which is unpredictable and thought of as adversarial. Thus, the goal of Player 1 is to design a strategy that would drive the 3-dimensional total payoff to the target set – the nonnegative orthant. To ensure that they fulfill the requirements of the Basel Accord, banks try to hold a capital buffer on top of the regulatory minimum, and they periodically adjust their assets to be at the top of the buffer [13, 19].

To model such situations we study two-player repeated games with vector-payoffs in continuous time. We say that a set A in the payoff space is *attainable* by Player 1 if there is a time T such that for every level of proximity, $\varepsilon > 0$, Player 1 has a strategy guaranteeing that against every possible strategy for Player 2, the distance between A and the cumulative payoff up to any time t greater than T is smaller than ε. If a set A is attainable by a supplier, then in order to ensure that the inventory level would converge to A, he can plan his actions based on historical inventory levels and market data.

The definition of attainability is close in spirit to the concept of approachable sets [9], which refers to the average stage-payoff rather than to the cumulative one. While a set A is attainable by Player 1 if he can ensure that the cumulative payoff converges to it, it is approachable by him if he can ensure that the average payoff converges to the set.

To illustrate the difference between these two notions, suppose that the target set of a supplier consists of one point, say x. If this set is attainable by him, it implies that the long-run inventory level is stable around x. On the other hand, if it is approachable, it merely guarantees that the average inventory level converges to x. This may happen also when the actual level itself does not converge to x, and even when any fixed running average does not converge to x. This observation suggests that although the notions of attainability and approachability are close to each other, the flavor of the results and their proofs are completely different.

One of our main results characterizes attainable convex sets. It uses the concept of B-sets (see [9]). It states that a convex set Y is attainable by a player if and only if there exist two B-sets C and C' for that player (or, alternatively, two approachable sets) and a nonnegative real number α such that $\alpha C + \text{Cone}(C') \subseteq Y$. The idea behind this result lies on two main properties that the cumulative payoff along any possible trajectory of the game must have. The first is that the cumulative payoff must reach the set Y within a certain time T, independently of the strategy of the other player. The second property is that it has to remain close to Y at any time after T. The first property is responsible for the existence of C and its role while the second property for that of C'.

In the case where Y is compact then $\text{Cone}(C')$ is necessarily compact, which may happen only if C' consists only of $\vec{0}$. Our characterization entails that $\{\vec{0}\}$ must then be approachable. This observation enables us to provide necessary and sufficient conditions for a singleton (i.e., a set containing a single vector) to be attainable by Player 1. We use it to show that *every* singleton is attainable by Player 1 if and only if the value of all scalar games obtained from the vector-payoff game by projecting the payoff function on any direction $\lambda \neq \vec{0}$ is positive.

The results presented here apply to games played in continuous time and where players are allowed to use a special type of behavior strategies. These strategies are characterized by an increasing sequence of positive real numbers that divide the time span $[0, \infty)$ into subintervals. The play of a player in each interval depends on the play of the other player *before* this interval starts and is independent on the other player's play during this interval. This is equivalent to saying that before the game starts, a player sets an alarm clock to ring at certain pre-specified times, and whenever the clock rings, the player looks at the historical play path and determines how to play until the next time the clock rings. In the literature of differential games this type of strategies is called nonanticipating strategies with delay. We later discuss the interpretation of this type of strategies.

There is a literature on decision problems related to dynamic multiinventory in continuous time (see for instance, the continuous-time control strategy in [11]). The control literature up to this point refers to one-person (the controller) decision problems with uncertainty. To the best of our knowledge, this paper is the first that takes a strategic approach to these problems.

The paper is organized as follows. In Section 2 we provide a motivating example. In Section 3 we introduce the model and main definitions. In Section 4 we present our results, and Section 5 is devoted to discussing a few aspects related to the definition of attainability and to the type of strategies that we are using. Proofs are relegated to Section 6.

2 A motivating example

This section details one motivation of our study: distribution networks. Consider a distributer of a certain product who has two warehouses A and B in different regions. Every month the distributer can order products from factories to each of the warehouses, and he can transport products between the two warehouses, while vendors order products from the warehouses. This situation is described graphically in Figure 2.

In Figure 2, f_A and f_B are the number of products that are sent from factories to the two warehouses A and B, f_T is the number of products that are transported from warehouse A to warehouse B, and w_A and w_B are the number of products sent from the two warehouses to vendors. Negative flows are interpreted as flows in the opposite direction; e.g., if vendors return products to warehouse A (resp. to warehouse B), then w_A (resp. w_B) is negative. If products are transported from warehouse B to warehouse A, then f_T is negative. We analyze this situation in continuous time. The change of stock in the two warehouses is given by the 2-dimensional vector

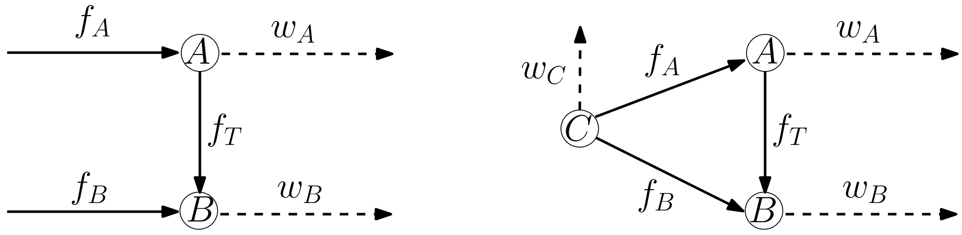

(a) Three distribution flows f_A, f_T, f_B and two vendors requests w_A, w_B.

(b) Factory manager can sell directly to vendors: node C represents factory.

Figure 1: Distribution network with warehouses A and B.

$$u(a_1^t, a_2^t) = \underbrace{\begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & 1 \end{pmatrix}}_{F} \underbrace{\begin{pmatrix} f_A^t \\ f_T^t \\ f_B^t \end{pmatrix}}_{a_1^t} - \underbrace{\begin{pmatrix} w_A^t \\ w_B^t \end{pmatrix}}_{a_2^t},$$

where $a_1^t = (f_A^t, f_B^t, f_T^t)$ is the decision variable of the distributer, and $a_2^t = (w_A^t, w_B^t)$ is the uncontrolled market demand at time t.

Suppose that the number of products that can be ordered by vendors at each time instance is bounded by 2, and the number of products that can be returned by vendors to each warehouse at every time instance is 3. In other words, w_A^t and w_B^t are in $[-3, 2]$. Suppose also that the amount of product that the distributer can order from or return to the factories and transport between the two warehouses is bounded by 5.

This situation can be described by a two-person game as follows. The distributer (Player 1) has 8 actions

$$(5,5,5), (5,5,-5), (5,-5,5), (5,-5,-5), (-5,5,5), (-5,5,-5), (-5,-5,5), (-5,-5,-5),$$

while the market demand or nature (Player 2) has 4 actions

$$(-3,-3), (-3,2), (2,-3), (2,2).$$

The payoffs correspond to the change of stock in the two warehouses, and are given by the following table:

	(-3,-3)	(-3,2)	(2,-3)	(2,2)
(5,5,5)	(3,13)	(3,8)	(-2,13)	(-2,8)
(5,5,-5)	(3,3)	(3,-2)	(-2,3)	(-2,-2)
(5,-5,5)	(13,3)	(13,-2)	(8,3)	(8,-2)
(5,-5,-5)	(13,-7)	(13,-12)	(8,-7)	(8,-12)
(-5,5,5)	(-7,13)	(-7,8)	(-12,-13)	(-12,8)
(-5,5,-5)	(-7,3)	(-7,-2)	(-12,3)	(-12,-2)
(-5,-5,5)	(3,3)	(3,-2)	(-2,3)	(-2,-2)
(-5,-5,-5)	(3,-7)	(3,-12)	(-2,-7)	(-2,-12)

Figure 2: The strategic-form game corresponding to the situation.

At every time instance the two players choose their actions. Each market behavior translates into a mixed action of Player 2, and each behavior of the distributer corresponds to a mixed action of Player 1. The (2-dimensional) total payoff up to time t is the number of products that are stored in each of the two warehouses. The goal of the distributer is to ensure that the total number of products in each warehouse does not exceed its capacity, that is, that the total payoff should not exceed a certain (2-dimensional) bound.

Figure 2 describes the case where the factory manager can sell directly to vendors, bypassing the distribution to warehouses. This situation can be represented by adding an additional node C modeling the factory, and an edge that represents the market demand. The stock is now a 3-dimensional vector, as we have to take into account the inventory available at the factory, and consequently the change in the stock modifies as shown below:

$$u(a_1^t, a_2^t) = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & 1 \\ -1 & 0 & -1 \end{pmatrix} \begin{pmatrix} f_A^t \\ f_T^t \\ f_B^t \end{pmatrix} - \begin{pmatrix} w_A^t \\ w_B^t \\ w_C^t \end{pmatrix}.$$

A recurrent question in the network flow control literature [4, 5, 11, 12, 15] is about conditions that ensure the existence of a control strategy that drives the excess supply vector to a desired target level in $\mathbb{R}^m$ regardless of the unpredictable realization of the demand. The equivalence between the excess supply and the cumulative payoff in the dynamic game motivates our study. The rest of the paper is devoted to the analysis of conditions under which Player 1 has a strategy ensuring the attainability of a convex set, regardless of the behavior of Player 2.

Situations where the target is to control the total payoff occur also in production and transportation networks. Production networks describe production processes and activities necessary to turn raw materials into intermediate products and eventually into final products. The nodes of the networks represent raw materials and intermediate/final products. The buffer at each single node i models the amount of material or product of type i stored or produced up to the current time, and hyper-arcs describe the materials or products consumed (tail nodes) and produced (head nodes) in each activity or process. Transportation networks model the flow of commodities, information, or traffic; nodes of the networks represent hubs and the buffers at the nodes describe the quantity of commodities present in the hubs. The edges describe transportation routes.

2.1 Related control and optimization literature

We highlight two main streams of related literature, one from the control area and the second from the optimization area. These two bodies of literature have two main elements in common: i) the interest towards robustness, and ii) the presence of a network dynamic flow scenario.

Connections between robust control and noncooperative game theory has a long history (see, e.g., [3]). Robust control is the area of control theory that looks for strategies that "control" the state of a dynamical system, for instance, drive it to a given set, despite the effects of disturbances (see the seminal paper [7]). Among the foundations of robust control we find two main notions that can be related to attainability and are surveyed in [10]. The first notion, *robust global attractiveness*, refers to the property of a set to "attract" the state of the system under a proper control strategy, independently of the effects of the disturbance. The second notion, *robustly controlled invariance*, describes the property of a set to bound the state trajectory under a proper control strategy, independently of the effects of the disturbance. Both notions are widely exploited in a variety of works that contribute to the use of robust control in dynamic network flow models [4, 5, 11, 12, 15].

A second stream of literature can be identified under the name of "robust optimization". This is a relatively recent technique that describes uncertainty via sets and optimizes the worst-case cost over those sets (see, e.g., the introduction to the special issue [6]). The use of robust optimization techniques in dynamic network models is the main focus of [1, 2, 8]. There, a main theme is to "adjust" some of the supplier's decision variables to the uncertain outcome. More specifically, some variables are determined before the outcome is realized while the rest are determined after the outcome is realized. Such a problem formulation is referred to as "Adjustable Robust Counterpart" (ARC) problem, or "two-stage robust optimization with recourse" and as it will be clear later it shares striking similarities with the formulation of attainable strategies presented in the current paper.

This paper focuses on the game theoretic aspects related to attainable sets. A discussion on applications to network flow control problems is introduced in a companion paper [16].

3 Attainability

In the first part of this section we introduce the mathematical model of repeated game in continuous time and elaborate on the type of strategies used by the players. In the remaining part, we provide a formal definition of attainability.

3.1 The model

We study a two-player repeated game with vector payoffs in continuous time Γ. The set of players is $N = \{1, 2\}$, and the finite set of actions of each player i is A_i. The instantaneous payoff is given by a function $u : A_1 \times A_2 \to \mathbb{R}^m$, where m is a natural number. We assume w.l.o.g. that payoffs are bounded by 1, so that $u : A_1 \times A_2 \to [-1, 1]^m$. We extend u to the set of mixed action pairs, $\Delta(A_1) \times \Delta(A_2)$, in a bilinear fashion. The one-shot vector-payoff game (A_1, A_2, u) is denoted by G and we will say that the game in continuous time Γ is *based on* G. For $i \in \{1, 2\}$, $-i$ denotes the opponent of i.

The game Γ is played over the time interval $[0, \infty)$. We assume that the players use nonanticipating behavior strategies with delay, which we define below. Roughly, a nonanticipating behavior strategy with delay divides time into intervals. The behavior of a player in a given interval depends on the behavior of the other player up to the beginning of the interval. In other words, the way a player plays during a given interval of time does not affect the way the opponent plays during that interval. Still, it may affect the other player's play in subsequent intervals.

Formally, denote by $\mathcal{C}_i$ the set of all *controls* of player i, that is, the set of all measurable functions from the time space, $[0, \infty)$, to player i's mixed actions. That is,

$$\mathcal{C}_i := \{a_i : [0, \infty) \to \Delta(A_i),\ a_i \text{ is measurable}\}.$$

Definition 1 *A function $\sigma_i : \mathcal{C}_{-i} \to \mathcal{C}_i$ is a* behavior strategy with delay *(or simply a* strategy*) for player i, if there exists an increasing sequence of real numbers $(\tau_i^k)_{k\in\mathbb{N}}$ such that for every $a_{-i}, a'_{-i} \in \mathcal{C}_{-i}$,*

$$a_{-i}(t) = a'_{-i}(t) \quad \forall t \in [0, \tau_i^k) \quad \Longrightarrow (\sigma_i(a_{-i}))(t) = (\sigma_i(a'_{-i}))(t) \quad \forall k \in \mathbb{N}, \forall t \in [\tau_i^k, \tau_i^{k+1}),$$

where $\tau_i^0 = 0$.

In the sequel we refer to the real numbers $(\tau_i^k)_{k\in\mathbb{N}}$ in Definition 1 as the *updating times related to* σ_i.

Remark 1 *In the literature on differential games a strategy as the one defined above is called a nonanticipating strategy with delay. An equivalent formulation, which may look more transparent to game theorists, is as follows. A strategy for player i is a list* $(\tau_i^k, \sigma_i^k)_{k\in\mathbb{N}}$ *where* $(\tau_i^k)_{k\in\mathbb{N}}$ *is an increasing sequence of real numbers, and for each* $k \in \mathbb{N}$, σ_i^k *is a function that maps play paths (of both players) on the interval* $[0, \tau_i^k)$ *to plays of player i in the interval* $[\tau_i^k, \tau_i^{k+1})$.

Later on when defining or referring to a strategy we use this formulation.

Remark 2 *Models with continuous time are typically used as a tractable version of discrete-time models, where the gap between two consecutive stages is small. This is the case here as well. Suppose that time is discrete, and that the time interval between any two successive decisions is extremely small. Suppose moreover, that observing opponent's actions is time consuming and possibly costly. Thus, players cannot observe each other's actions at every time. Rather, they observe their opponent's actions relatively rarely compared to the frequency in which actions are taken. Our continuous-time model captures this aspect: although time is continuous, players observe historical play and decide how to play in the next interval in discrete times. Neither updating nor new decision is taken place between two updating times.*

Every pair of strategies $\sigma = (\sigma_1, \sigma_2)$ uniquely determines a play path $(a^\sigma(t))_{t\in\mathbb{R}_+}$. The cumulative payoff-vector up to time T associated with the pair of strategies σ is given by

$$\gamma^T(\sigma) = \int_0^T u(a^\sigma(t))\mathrm{d}t \in \mathbb{R}^m. \tag{1}$$

Sometimes we denote it by γ_G^T when we wish to emphasize that the payoff is in the game based on G. Note that since the payoffs are bounded by 1, the integral in (1) is well-defined.

3.2 Attainability: the definition

The subject matter of this paper is the concept of attainable sets. A set of vectors is attainable by a player if he can guarantee that the distance between the set and the cumulative payoff converges to 0, regardless of the strategy of the opponent. We provide a definition here and two alternative ones are discussed later in Section 5.

Definition 2 *A nonempty closed set $Y \subseteq \mathbb{R}^m$ is* attainable *by Player 1 if there is $T > 0$ such that for every $\varepsilon > 0$ there is a strategy σ_1 of Player 1 such that*[1]

$$d(\gamma^t(\sigma_1, \sigma_2), Y) \leq \varepsilon, \quad \forall t \geq T, \forall \sigma_2.$$

A set Y is attainable if there is a finite horizon T such that Player 1 can ensure, against any possible strategy of Player 2, that the cumulative payoff up to any time $t \geq T$ is within ε from Y. Note that the time T is uniform across all levels of precision. That is, in order for Y to be attainable by Player 1, Player 1 must be able to guarantee that the cumulative payoff at any time longer than T would be within any ε from Y. However, different ε's might require different strategies employed by Player 1. It might therefore happen that although Y is attainable by Player 1, the cumulative payoff would never touch Y itself. We say that the strategy σ_1 in Definition 2 *attains the set Y up to ε.*

When Y contains a single vector and it is attainable by Player 1, we say that the vector x is attainable by him. Denote by W the set of attainable vectors.

Since the notion of attainability is related to that of approachability, we recall the definition of the later (in a continuous-time framework). Denote the mean vector-payoff between times 0 and T by

$$\overline{\gamma}^T(\sigma_1, \sigma_2) = \frac{1}{T}\gamma^T(\sigma_1, \sigma_2).$$

Definition 3 *A nonempty closed set $Y \subseteq \mathbb{R}^m$ is* approachable by Player 1 *if for every $\epsilon > 0$ there exist $T > 0$ and a strategy σ_1 of Player 1 such that*

$$d\left(\overline{\gamma}^t(\sigma_1, \sigma_2), Y\right) \leq \epsilon, \quad \forall t \geq T, \forall \sigma_2.$$

The original definition of approachable sets [9] was given in discrete time repeated games. A set is approachable in the discrete time model if and only if it is approachable in the continuous time one (see, [17]).

The definitions of attainability and approachability are close in spirit. There is, however, a significant difference between the two concepts. A set is approachable if the average payoff converges to it, while a set is attainable if the cumulative payoff converges to the set. In other words, approachability refers to the convergence of the *average* payoff, while attainability to the convergence of the *cumulative* payoff.

[1] The distance referred to throughout the paper is the Euclidean distance and the norm is the L_2-norm $\|.\|_2$. The distance between a point x and a set A is, therefore, $d(x, A) = \min_{y \in Y} \|x - y\|_2$.

4 Results

This section presents the main results. The first, Theorem 1, characterizes closed and convex attainable sets. Using this result, we derive a characterization of attainable compact convex sets and of attainable vectors. Finally, the last result provides a stronger condition that ensures that any vector $x \in \mathbb{R}^m$ is attainable.

For every nonempty closed set $Y \subseteq \mathbb{R}^m$ and every $z \in \mathbb{R}^m$ we denote by $\Pi_Y(z) := \{y \in Y \mid d(z,y) = d(z,Y)\}$ the set of points in Y closest to z. When the set Y is convex, $\Pi_Y(z)$ contains a single point. Our main theorem characterizes closed convex attainable sets. To state the result we borrow from [9] the concept of B-set.

Definition 4 *A nonempty closed set $C \subseteq \mathbb{R}^m$ is a B-*set for Player 1 *if for every $z \in \mathbb{R}^m$ there exists $c \in \Pi_C(z)$ and $x \in \Delta(A_1)$ such that*

$$\langle u(x,a_2) - c, z - c\rangle \leq 0, \quad \forall a_2 \in A_2.$$

If a nonempty closed set contains a B-set, then it is approachable [9]. Conversely, every approachable set contains a B-set [14, 18]. Therefore, a set is approachable if and only if it contains a B-set.

We show that a closed convex set is attainable if and only if it contains a certain sum of two B-sets. For every set $Y \subseteq \mathbb{R}^m$ we denote the cone spanned by Y by $\mathrm{Cone(Y)} = \{\alpha\mathrm{y}, \alpha \in \mathbb{R}^+, \mathrm{y} \in \mathrm{Y}\}$.

Theorem 1 *A closed convex set $Y \subseteq \mathbb{R}^m$ is attainable by Player 1 if and only if there exist $\alpha > 0$ and two B-sets for that player C and C' such that*

$$\alpha C + \mathrm{Cone}(C') \subseteq Y. \tag{2}$$

The idea behind the theorem is that any trajectory attaining a set Y consists of two parts. A first sub-trajectory reaches Y in a fixed finite time. This is represented by the αC. A second sub-trajectory stays close to Y. When Y is unbounded, keeping the trajectory within Y is equivalent to keeping the direction in which the trajectory progresses within a proper range. This is represented by the second term $\mathrm{Cone}(C')$. In the special case where Y is compact (i.e., also bounded) Player 1 can ensure that the trajectory will remain in Y only if he can keep the trajectory close to $\vec{0}$, so that the set $\mathrm{Cone}(\{\vec{0}\})$ is attainable. The inclusion in (2) is justified by the observtion that any superset of an attainable set is attainable too.

We deduce now several corollaries. In the first we focus on compact convex sets. Whenever Y is compact, the characterization of Theorem 1 can be simplified. Indeed, since the only compact cone is $\{\vec{0}\}$ whenever Y is a compact convex attainable sets we must have $C = \{\vec{0}\}$. By Theorem 1 it follows that there exists $\alpha > 0$ and a B-set C such that $C \subseteq \frac{1}{\alpha}Y$. By setting $\delta = \frac{1}{\alpha}$, we infer that δY is approachable. Since every approachable set contains a B-set, this yields the following result.

Corollary 1 *A compact convex set $Y \subseteq \mathbb{R}^m$ is attainable by Player 1 if and only if*

B1 *The vector $\vec{0} \in \mathbb{R}^m$ is approachable by Player 1, and*

B2 *There exists a scalar $\delta > 0$ such that δY is approachable by Player 1.*

The following example, borrowed from [9], shows that the result introduced above does not hold when Y is not convex.

Example 1 *Consider the following payoff repeated game with 2-dimensional payoffs:*

	L	R
T	$(0,0)$	$(0,0)$
B	$(1,0)$	$(1,1)$

Figure 3: The payoff function in Example 3.

Define $Y := \{(1/2, t),\ t \le 1/4\} \cup \{(1,t),\ t \ge 1/4\}$. It was shown in Blackwell [9] that the set Y is not approachable by Player 1. One can verify that none of its dilatations is approachable by Player 1. Nevertheless, the set Y is attainable by Player 1. Indeed the following behavior strategy attains it for Player 1:

$$\sigma_1(t) = \begin{cases} B & t \in [0, \frac{1}{2}), \\ T & t \in [\frac{1}{2}, 1), \gamma_1^{1/2} \le \frac{1}{4}, \\ B & t \in [\frac{1}{2}, 1), \gamma_1^{1/2} > \frac{1}{4}, \\ T & t \ge 1. \end{cases}$$

In the particular case where $Y = \{x\}$ is a singleton, we can again be more precise. We seperate the attainability of $\vec{0}$ and the attainability of $x \neq \vec{0}$. To state the next result we need the following notations. Let $\lambda \in \mathbb{R}^m$. Denote[2] by $\langle \lambda, G \rangle$ the zero-sum one-shot game whose set of players and their action sets are as in the game G, and the payoff that Player

[2]The inner product is defined by $\langle x, y \rangle := \sum_{i=1}^m x_i y_i$ for every $x, y \in \mathbb{R}^m$.

2 pays to Player 1 is $\langle \lambda, u(a_1, a_2) \rangle$ for every $(a_1, a_2) \in A_1 \times A_2$. As a zero-sum one-shot game, the game $\langle \lambda, G \rangle$ has a value, denoted v_λ.

For every mixed action $p \in \Delta(A_1)$ denote

$$D_1(p) = \{u(p, q) \colon q \in \Delta(A_2)\}.$$

$D_1(p)$ is the set of all payoffs that might be realized when Player 1 plays the mixed action p. If $v_\lambda \geq 0$ (resp. $v_\lambda > 0$), then there is a mixed action $p \in \Delta(A_1)$ such that $D_1(p)$ is a subset of the closed half space $\{x \in \mathbb{R}^m \colon \langle \lambda, x \rangle \geq 0\}$ (resp. half space $\{x \in \mathbb{R}^m \colon \langle \lambda, x \rangle > 0\}$). Thus $D_1(p)$ and λ are in the same half-space, or, equivalently, $D_1(p)$ and $-\lambda$ are in two different half-spaces.

Corollary 2 *The following three properties are equivalent.*

- **C1** *The vector $\vec{0} \in \mathbb{R}^m$ is attainable by Player 1*
- **C2** *The vector $\vec{0} \in \mathbb{R}^m$ is approachable by Player 1.*
- **C3** *For every $\lambda \in \mathbb{R}^m$, $v_\lambda \geq 0$.*

The equivalence between $C1$ and $C2$ is an immediate consequence of Corollary 1. Based on that $\vec{0}$ is approachable if and only if it is a B-set, the equivalence between $C2$ and $C3$ follows from [9].

The following result characterizes when a given vector $x \neq \vec{0}$ is attainable. For every $y \in \mathbb{R}^m$ denote by $(G - y)$ the two-player one-shot game that is identical to G except for its payoff function. The payoff function of $(G-y)$ is $(u-y)$, where $(u-y)(a_1, a_2) = u(a_1, a_2) - y$ for every $a_1 \in A_1$ and $a_2 \in A_2$.

Corollary 3 *Let $\vec{0} \neq x \in \mathbb{R}^m$. The vector x is attainable by Player 1 if and only if*

- **D1** *The vector $\vec{0} \in \mathbb{R}^m$ is attainable by Player 1*

and either one of the following conditions holds:

- **D2** *There is $\delta > 0$ such that the vector $\vec{0} \in \mathbb{R}^m$ is attainable by Player 1 in the game based on $(G - \delta x)$.*
- **D3** *There is $\delta > 0$ such that, for every $\lambda \in \mathbb{R}^m$, $v_\lambda \geq \delta \langle x, \lambda \rangle$.*
- **D4** *There is $\delta_0 > 0$ such that for every $q \in \Delta(A_2)$ there is $p \in \Delta(A_1)$ and $\delta > \delta_0$ satisfyingze $u(p, q) = \delta x$.*

Conditions **D2** and **D3** are reformulations of **B2**. Condition **D4** needs additional work in order to be proven. The proofs are differed to the last section.

Corollary 3 implies that whenever any vector x is attainable, so is the vector $\vec{0}$. Since attainability is concerned with the cumulative payoff, once a target level is (almost) reached, this level should be maintained in the long run. This means that once a neighborhood of a target level x is reached, from that point in time and on $\vec{0}$ ought to be attained. This is the reason why $\vec{0}$ is attainable when any vector x is attainable, and why $\vec{0}$ plays a major role in the theory of attainability. However, Condition **B1** alone is not sufficient for the attainability of other vectors other then $\vec{0}$ itself.

The previous result naturally leads to a sufficient condition for a vector to be attainable.

Proposition 1 *Let $x \in \mathbb{R}^m$ such that*

E1 $v_\lambda \geq 0$ *for every* $\lambda \in \mathbb{R}^m \setminus \{\vec{0}\}$*, and*

E2 *For every* $\lambda \in \mathbb{R}^m \setminus \{\vec{0}\}$*, if* $\langle \lambda, x \rangle \geq 0$ *then* $v_\lambda > 0$*.*

Then x is attainable.

We deduce the following theorem which deals with the case where all the vectors are attainable.

Theorem 2 *The following statements are equivalent:*

F1 *Every vector $x \in \mathbb{R}^m$ is attainable by Player 1;*

F2 $v_\lambda > 0$ *for every* $\lambda \in \mathbb{R}^m \setminus \{\vec{0}\}$*.*

The fact that Condition **F2** implies Condition **F1** is a consequence of Proposition 1. Indeed given that **F2** is true, then for every $x \in \mathbb{R}^m$, Condition **E2** is satisfied and thus every $x \in \mathbb{R}^m$ is attainable.

The converse implication can be obtained by focusing on Condition **D3** in Corollary 3. Assume that Condition **F1** holds. For every $\lambda \in \mathbb{R}^m \setminus \{\vec{0}\}$, the vector $x = \lambda$ is attainable and satisfies $\langle x, \lambda \rangle > 0$. Therefore, Condition **D3** implies that $v_\lambda > 0$, and Condition **F2** holds as well.

Remark 3 *If Condition **F2** is satisfied, then for every open half space H of $\mathbb{R}^m$ there is a mixed action $p \in \Delta(A_1)$ such that $D_1(p) \subseteq H$. Standard continuity and compactness arguments imply that in this case there is $\delta_1 > 0$ such that for every half space H there is $p \in \Delta(A_1)$ satisfying $d(D_1(p), H) \geq \delta_1$. Stated differently, there is $\delta_2 > 0$ such that for every vector λ whose ℓ_1-norm is 1, $\langle \lambda, u(p,q) \rangle > \delta_2$ for every $q \in \Delta(A_2)$.*

Note the difference between Condition **C3** of Corollary 2 and Condition **F2** of Theorem 2. In the former, the value of the scalar-payoff game with payoffs $\langle \lambda, u(p,q) \rangle$ is nonnegative for every direction $\lambda \in \mathbb{R}^m \setminus \{\vec{0}\}$, while in the latter it is strictly positive. The former guarantees attainability of the vector $\vec{0}$, while the latter guarantees that every vector is attainable.

5 Discussion

The model and the results described above give rise to a number of additional questions. (a) What are the analogous results in discrete time repeated game to the ones we obtained? (b) Are there different notions of attainability that do not impose a uniform time of convergence? (c) What happens if the updating times are not predetermined and can be selected as a function of the information available up to the updating time?

We next elaborate on these questions and highlight a few open problems left for future research.

5.1 Continuous time versus discrete time.

The characterization presented in Theorem 1 depends crucially on the continuous time setting. The following example shows that it is invalid when time is discrete.

Example 2 Consider a game in *discrete* time where payoffs are one-dimensional and each player has two actions. Payoffs are given by the following matrix:

	L	R
U	$-2-1$	$-2+1$
B	$2-1$	$2+1$

$=$

	L	R
U	-3	-1
B	1	3

Figure 3: The payoff function in Example 2.

The payoffs in this game are the sum of two numbers, one determined by Player 1 (-2 if he plays U, 2 if he plays B), and the other by Player 2 (-1 if she plays L, 1 if she plays R).

Condition **C3** is satisfied, and therefore 0 is attainable by Player 1. The following strategy guarantees that the cumulative payoff is within $9 \cdot 2\eta$ from 0 at any $t > 2$, where $\eta > 0$ is given; the details of the proof can be found in the proof of Theorem 1. Divide the time line into countably many blocks, where the length of the k-th block is $\frac{\eta}{k}$. In the k-th block Player 1 plays U if the cumulative payoff at the beginning of the block is positive, and he plays B otherwise.

We show that 0 is not attainable by Player 1 in the game in discrete time. When time is discrete, a behavior strategy for a player is a function that assigns a mixed action to each past history. For every $\ell \in \mathbb{N}$, let p^ℓ be the mixed action played by Player 1 at stage ℓ. Note that p^ℓ depends on past play. Let σ_2 be the strategy that at each stage ℓ plays L if $p^\ell(U) \geq \frac{1}{2}$, and R otherwise. The stage payoff is then at least 2 whenever Player 2 plays R, and at most -2 whenever Player 2 plays L. In particular, if the total payoff up to stage ℓ is in the interval $[-\frac{1}{2}, \frac{1}{2}]$, then the payoff up to stage $\ell + 1$ lies outside this interval. Thus, the cumulative payoff does not converge to 0. ■

Example 2 suggests that the characterization of the set of attainable vectors in games in discrete time is more challenging than the characterization in continuous time.

5.2 Alternative definitions of attainability

We here provide two alternative definitions of the concept of attainability, which we term asymptotic attainability and weak asymptotic attainability. We then explore some relations between the three definitions.

For every set $Y \subseteq \mathbb{R}^m$ we denote by $B(Y, \varepsilon)$ the set of all points whose distance from at least one point in Y is less than ε, that is,

$$B(Y, \varepsilon) := \{x \in \mathbb{R}^m : d(x, Y) < \varepsilon\}.$$

When Y contains a single point x, we write $B(x, \varepsilon)$ instead of $B(\{x\}, \varepsilon)$.

Definition 5 *(i) The set $Y \subseteq \mathbb{R}^m$ is* asymptotically attainable *by Player 1 if there is a strategy σ_1 for Player 1 such that for every strategy σ_2 of Player 2,*

$$\lim_{T \to \infty} d(\gamma^T(\sigma_1, \sigma_2), Y) = 0. \tag{3}$$

(ii) The set Y is weakly asymptotically attainable *by Player 1, if the set $B(Y, \varepsilon)$ is asymptotically attainable by Player 1 for every $\varepsilon > 0$.*

Asymptotic attainability requires that a set is asymptotically reached by the cumulative payoff without putting any bound on the time it takes to reach the set. Attainability, on the other hand, requires that a set is approximately reached in a bounded time, independent of the degree of approximation. Weak asymptotic attainability relaxes both time boundedness and the level of the approximation precision. A set Y is weakly asymptotically attainable if any neighborhood $B(Y, \varepsilon)$ of Y can be asymptotically attained, without having a universal bound on the time at which this neighborhood is reached.

Any attainable set is also weakly asymptotically attainable and any asymptotically attainable set is weakly asymptotically attainable as well. In addition, observe that the set of asymptotically attainable vectors and the set of weakly asymptotically attainable vectors are convex cones. The definition implies that the set of weakly attainable vectors is also closed.

Using Corollary 3 , we now show that attainability of a vector does not imply its asymptotic attainability. This implies in particular that these two concepts are not identical.

Example 3 We provide an example where the vector $\vec{0}$ is attainable but not asymptotically attainable. Consider the following game where payoffs are 2-dimensional, each player has 2 actions, and the payoffs are scalar and given by:

	L	R
U	1	0
B	0	-1

Figure 4: The payoff function in Example 3.

In this game $v_\lambda = 0$ for every $\lambda \in \mathbb{R}$. Thus, for every $\lambda \in \mathbb{R}^2$ one has $v_\lambda \geq 0$, and therefore corollary 3 implies that the vector $\vec{0}$ is attainable by Player 1. We argue that $\vec{0}$ is not asymptotically attainable by Player 1. Assume that Player 1 implements a strategy σ_1. In an initial time interval the strategy σ_1 plays one of the rows with a positive probability. Consider the strategy σ_2 of Player 2 that plays constantly a column that generates a nonzero vector in that initial interval. For instance, if σ_1 plays the action U with positive probability in the initial time interval, then σ_2 play the action L always. The initial period produces a nonzero payoff and this payoff is not diminishing to zero because Player 2 keeps playing the same column forever. This example shows that $\vec{0}$ is attainable by Player 1 but not asymptotically attainable by him.

We point out that the argument mentioned above shows in fact that $\vec{0}$ is not attainable in the corresponding game in discrete time as well.

The following example shows that a weakly attainable vector need not be attainable.

Example 4 Consider a two-player game where payoffs are 2-dimensional, Player 1 has 3 actions, Player 2 has 2 actions, and the payoff function is given by the left-hand side matrix in Figure 5.

	L	R
U	$(1,1)$	$(0,1)$
M	$(0,0)$	$(1,1)$
B	$(0,0)$	$(0,0)$

The game G

Figure 5: The payoff function of the game G in Example 4.

The vector $(0,0)$ is attainable by Player 1, using the strategy that always plays B. The vector $x := (1,1)$ is weakly asymptotically attainable according to Definition 5. Indeed, given $\varepsilon > 0$ consider the strategy σ_1^ε, with updating times $(\tau_1^k)_{k\in\mathbb{N}}$ defined by $\tau_1^k = k\varepsilon$ for $k \in \mathbb{N}$, that is defined as follows.

- If the total payoff up to time τ_1^k is not in the set $B((1,1),\varepsilon)$, during the time interval $[\tau_1^k, \tau_1^{k+1})$ play the mixed action $[\varepsilon(U), (1-\varepsilon)(M)]$.

- If the total payoff up to time τ_1^k is in the set $B((1,1),\varepsilon)$, during the time interval $[\tau_1^k, \tau_1^{k+1})$ play the action B.

For every $t \geq \frac{1}{\varepsilon}$ one has $d(\gamma^t(\sigma_1^\varepsilon, \sigma_2), (1,1)) < \varepsilon$, so that the vector x is indeed weakly asymptotically attainable by Player 1.

The vector x, however, is not attainable by Player 1 (according to Definition 2). To show this claim we use Corollary 3 and prove that Condition **D4** does not hold for x. Indeed, fix $\delta_0 > 0$, and set $q := [(1-\frac{\delta_0}{2})(L), \frac{\delta_0}{2}(R)]$. Let $p \in \Delta(A_1)$ be arbitrary. If $u(p,q) = \delta x = (\delta,\delta)$ for $\delta > 0$, then necessarily $p_U = 0$. One can verify that $u(p,q)$ cannot be equal to δx for $\delta > \delta_0$, and therefore Condition **D4** does not hold for x.

Remark 4 *The proof of Theorem 2 shows that every vector $x \in \mathbb{R}^m$ is attainable by Player 1 if and only if every vector $x \in \mathbb{R}^m$ is asymptotically attainable by Player 1. Example 3 shows that attainability does not imply asymptotic attainability. We are unable to tell whether or not asymptotic attainability implies attainability.*

5.3 Alternative strategies in continuous time.

The strategies we use here are nonanticipating strategies with delay. In these strategies the times $(\tau_i^k)_{k\in\mathbb{N}}$ at which a player observes past play are independent of the play of the other player. One could consider a broader class of strategies in which $(\tau_i^k)_{k\in\mathbb{N}}$ are stopping times.

In other words, τ_i^{k+1} is a time that depends on (that is, it is measurable with respect to) the information available to player i at time τ_i^k, for each $k \in \mathbb{N}$. In this type of strategies, the updating times $(\tau_i^k)_{k\in\mathbb{N}}$, are not predetermined real numbers, as in Definition 1. Our results remain valid even if Player 2 is allowed to use a strategy from this broader class of strategies.

5.4 Additional open problems

The results above refer to attainability of a convex set, and did not discuss attainability, asymptotic attainability, or weak asymptotic attainability of nonconvex sets. Characterizing when a set of payoffs is attainable (according to these three definitions) remains open. We also leave attainability in discrete time and attainability when payoffs are discounted for future investigations.

6 Proofs

6.1 Proof of Theorem 1

The aim of this section is to prove the characterization of attainable closed convex sets. We first provide the outline of the proof.

6.1.1 Outline

Continuous approachability and discrete approachability are equivalent [17]. This justifies the use of the notion of B-sets in the study of games in continuous time.

Given $\alpha > 0$, a closed convex set $Y \subseteq \mathbb{R}^m$, and two B-sets $C, C' \subseteq \mathbb{R}^m$, two B-sets such that

$$\alpha C + \mathrm{Cone}(C') \subset Y,$$

we check that, by convexity of Y, we can replace the B-sets with their convex hull which are approachable:

$$\alpha \mathrm{Conv(C)} + \mathrm{Cone(Conv(C'))} \subset \mathrm{Y}.$$

Then we prove that given two approachable convex sets C and C' the set $\alpha C + \mathrm{Cone}(C')$ is attainable. Thus Y is attainable.

To show the converse implication, given a set Y, we define the family of sets $\overline{Y}_t$ as the intersection of $[-1, 1]^m$ and $\frac{1}{t}Y$. We show that the family $\overline{Y}_t$ admits limit values $\overline{Y}_\infty$.

Moreover, for every $t > 0$, we have

$$t\overline{Y}_t + \mathrm{Cone}(\overline{\mathrm{Y}}_\infty) \subset \mathrm{Y}$$

We prove that there exists $T \in \mathbb{R}_+$ such that for all $t \geq T$, the set $\overline{Y}_t$ is approachable. This implies that $\overline{Y}_T$ and $\overline{Y}_\infty$ are approachable and each one contains a B-set [18]. It follows that (2) holds with $C = \overline{Y}_T$ and $C' = \overline{Y}_\infty$.

6.1.2 The condition is sufficient

Let Y be a closed convex set. Suppose that there exists $\alpha > 0$ and two B-sets C and C' such that

$$\alpha C + \mathrm{Cone}(C') \subset Y.$$

We prove that Y is attainable. Since any superset of an approachable set is approachable, the sets $\mathrm{Conv(C)}$ and $\mathrm{Conv(C')}$ are approachable. Since Y is convex, these two sets are subsets of Y and satisfy

$$\alpha\mathrm{Conv(C)} + \mathrm{Cone}(\mathrm{Conv(C')}) = \mathrm{Conv}(\alpha\mathrm{C} + \mathrm{Cone(C')}) \subset \mathrm{Conv(Y)} = \mathrm{Y}.$$

We now prove the following.

Proposition 2 *Let $\alpha > 0$ and C, C' be two closed convex approachable subsets of $\mathbb{R}^m$. Then $\alpha C + \mathrm{Cone}(C')$ is attainable.*

The following lemma claims that the distance between a point and a set is a convex function. For every finite collection $(C_i)_{i=1}^n$ of nonempty subsets of $\mathbb{R}^m$ and every collection $(\lambda_i)_{i=1}^m$ of scalars, denote

$$\sum_{i=1}^n \lambda_i C_i := \{z \in \mathbb{R}^m \mid z = \sum_{i=1}^n \lambda_i c_i,\ \forall c_i \in C_i,\ \forall i = 1, \ldots, n\}.$$

Lemma 1 *Let $n \in \mathbb{N}$, let $(x_i)_{i=1}^n$ be points in $\mathbb{R}^m$, and let $(C_i)_{i=1}^n$ be nonempty closed subsets of $\mathbb{R}^m$. For every collection of positive real numbers $(\lambda_i)_{i=1}^n$ one has*

$$d\left(\sum_{i=1}^n \lambda_i x_i, \sum_{i=1}^n \lambda_i C_i\right) \leq \sum_{i=1}^n \lambda_i d(x_i, C_i).$$

Proof. For every $i \in \{1, 2, \ldots, n\}$ let c_i be a point in C_i that satisfies $d(x_i, c_i) = d(x_i, C_i)$. Note that $c := \sum_{i=1}^n \lambda_i x_i \in \sum_{i=1}^n \lambda_i C_i$.

$$\begin{aligned} d\left(\sum_{i=1}^{n}\lambda_i x_i, \sum_{i=1}^{n}\lambda_i C_i\right) &\leq \left\|\sum_{i=1}^{n}\lambda_i x_i - c\right\| \\ &\leq \sum_{i=1}^{n}\lambda_i \|x_i - c_i\| \\ &= \sum_{i=1}^{n}\lambda_i d(x_i, C_i). \end{aligned}$$

■

When C is convex, $\sum_{i=1}^{n}\lambda_i C = \left(\sum_{i=1}^{n}\lambda_i\right)C$, and therefore we obtain the following corollary.

Corollary 4 *Let $n \in \mathbb{N}$, let $(x_i)_{i=1}^{n}$ be points in $\mathbb{R}^m$, and let C be a nonempty closed and convex subset of $\mathbb{R}^m$. For every collection of positive real numbers $(\lambda_i)_{i=1}^{n}$ one has*

$$d\left(\sum_{i=1}^{n}\lambda_i x_i, \left(\sum_{i=1}^{n}\lambda_i\right)C\right) \leq \sum_{i=1}^{n}\lambda_i d(x_i, C).$$

For every strategy σ_i of Player i let σ_i^{β} be the strategy σ_i accelerated by a factor β. That is, $(\sigma_i^{\beta}(a_{-i}))(t) := (\sigma_i(\widehat{a}_{-i}))(\beta t)$, where $\widehat{a}_{-i}(t) = a_{-i}(\beta t)$.

The following result, which holds since time is continuous, states that if the strategy σ_1 of Player 1 guarantees that the distance between the average payoff up to time t and a given set C is less than ε, then the accelerated strategy σ_1^{β} ensures that the distance between the average payoff up to time $\frac{t}{\beta}$ and the set $\frac{C}{\beta}$ is at most $\frac{\epsilon}{\beta}$.

Lemma 2 *Let $C \subseteq \mathbb{R}^m$, $T > 0$, and $\varepsilon > 0$. If the strategy σ_1 satisfies*

$$d\left(\overline{\gamma}^t(\sigma_1, \sigma_2), C\right) \leq \epsilon, \quad \forall t \geq T,\ \forall \sigma_2,$$

then the strategy σ_1^{β} satisfies

$$d\left(\overline{\gamma}^t(\sigma_1^{\beta}, \sigma_2), \frac{C}{\beta}\right) \leq \frac{\epsilon}{\beta}, \quad \forall t \geq \frac{T}{\beta},\ \forall \sigma_2.$$

Proof. For every strategy σ_2 of Player 2 one has

$$\begin{aligned} \gamma^t(\sigma_1^{\beta}, \sigma_2) &= \int_0^t u(a^s(\sigma_1^{\beta}, \sigma_2))\mathrm{d}s \\ &= \frac{1}{\beta}\int_0^{\beta t} u(a^s(\sigma_1, \sigma_2^{1/\beta}))\mathrm{d}s \qquad (4) \\ &= \frac{1}{\beta}\gamma^{\beta t}(\sigma_1, \sigma_2^{1/\beta}). \end{aligned}$$

We deduce that

$$d\left(\overline{\gamma}^t(\sigma_1^\beta,\sigma_2),\frac{x}{\beta}\right)\leq\frac{\varepsilon}{\beta},\quad\forall t\geq\frac{T}{\beta},\forall\sigma_2,$$

as desired. ■

A corollary of this result, which has its own interest but will not be used here, is that the set of attainable vectors is a convex set.

A second corollary of Lemma 2 is that if C is a set that is approachable by Player 1, then for every $t>0$ he can ensure that the total payoff up to time t is arbitrarily close to tC.

Corollary 5 *Let C be a set that is approachable by Player 1. For every $\delta>0$ and every $s\in(0,1)$ there exists a strategy σ_1^* such that*

$$d\left(\gamma^s(\sigma_1^*,\sigma_2),sC\right)\leq\delta,\quad\forall\sigma_2.$$

Proof. Because the set C is approachable by Player 1, there is a constant K, a strategy σ_1 and $T>0$ such that

$$d\left(\overline{\gamma}^t(\sigma_1,\sigma_2),C\right)\leq\frac{K}{\sqrt{t}},\quad\forall\sigma_2,\forall t\geq T.$$

W.l.o.g. we can assume that $T\geq\left(\frac{K}{\delta}\right)^2$. It follows that

$$d\left(\gamma^t(\sigma_1,\sigma_2),tC\right)\leq t\frac{K}{\sqrt{t}},\quad\forall\sigma_2,\forall t\geq T,$$

and by Lemma 2 we have for every $\beta>0$,

$$d\left(\gamma^{\beta t}(\sigma_1^\beta,\sigma_2),\beta tC\right)\leq\beta t\frac{K}{\sqrt{t}},\quad\forall\sigma_2,\forall t\geq T.$$

For every $s\in(0,1)$ substitute $t:=T$ and $\beta:=\frac{s}{T}$, to obtain

$$d\left(\gamma^s(\sigma_1^\beta,\sigma_2),sC\right)\leq s\frac{K}{\sqrt{T}},\quad\forall\sigma_2.$$

The result follows since $s\leq 1$ and $T\geq\left(\frac{K}{\delta}\right)^2$. ■

We now strengthen Corollary 5 to prove that if a set C is approachable by Player 1, then he can guarantee that the total payoff remains close to the cone generated by it, for every $s\geq 0$.

Lemma 3 *Let C be a closed and convex set that is approachable by Player 1. For every $\varepsilon>0$ there exists a strategy σ_1^* such that*

$$d\left(\gamma^s(\sigma_1^*,\sigma_2),sC\right)\leq\epsilon,\quad\forall\sigma_2,\forall s\geq 0.$$

Proof. The strategy σ_1 is given by concatenating strategies that satisfy Corollary 5, with properly chosen s's and ε's.

Let us start by fixing $\varepsilon > 0$. For each $k \in \mathbb{N}$ let σ_1^k be a strategy that satisfies Corollary 5 with $s = \frac{\varepsilon}{k}$ and $\delta = \frac{\varepsilon}{2^k}$. Set $t_0 := 0$ and

$$t_{k+1} := t_k + \frac{\varepsilon}{k}, \quad \forall k \geq 0.$$

Let σ_1 be the strategy of Player 1 that for each $k \geq 0$, at time t_k forgets past play and follows the strategy σ_1^k until time t_{k+1}.

By construction and by Corollary 4, for each $k \geq 0$ we have

$$d(\gamma^{t_k}(\sigma_1, \sigma_2), t_k C) \leq \sum_{j=1}^{k} \frac{\varepsilon}{2^k} \leq \frac{\varepsilon}{2}.$$

Since $t_{k+1} - t_k \leq \frac{\varepsilon}{2}$, and since payoffs are bounded by 1, the triangle inequality implies that

$$d(\gamma^t(\sigma_1, \sigma_2), tC) \leq \varepsilon, \quad \forall t \geq 0,$$

as desired. ∎

We are now ready to complete the proof of Proposition 2.

Proof of Proposition 2. Let $\epsilon > 0$, we build a strategy σ_1^* such that

$$d\left(\gamma^t(\sigma_1^*, \sigma_2), \alpha C + \mathrm{Cone(C')}\right) \leq 2\epsilon, \quad \forall t \geq \alpha, \forall \sigma_2.$$

We define $\sigma_1(C)$ and $\sigma_1(C')$ – two strategies given by the previous lemma applied respectively to C and C'. We define a strategy σ^* of Player 1: follow $\sigma_1(C)$ until time α and then follow $\sigma_1(C')$.

Let $t \geq \alpha$ and σ_2 a strategy of Player 2. By construction, we have

$$\begin{aligned} d\left(\gamma^t(\sigma_1^*, \sigma_2), \alpha C + (t-\alpha)(C')\right) &\leq d\left(\gamma^\alpha(\sigma_1^*(C), \sigma_2), \alpha C\right) + d\left(\gamma^{t-\alpha}(\sigma_1^*(C'), \sigma_2'), (t-\alpha)(C')\right) \\ &\leq 2\epsilon, \end{aligned}$$

where σ_2' being the continuation strategy of Player 2 after time α. The set $\alpha C + \mathrm{Cone(C')}$ is attainable. ∎

6.1.3 The condition is necessary

In this section we prove that if a closed and convex set Y is attainable by Player 1, then there exists $\alpha > 0$ and two B-sets C and C' such that

$$\alpha C + \mathrm{Cone}(C') \subset Y.$$

For every $t \geq 0$ define

$$Y_t := Y \cap [-t, t]^m$$

and

$$\overline{Y}_t := \frac{1}{t} Y_t \subseteq [-1, 1]^m.$$

The Hausdorff metric is a metric over closed subsets of $\mathbb{R}^m$, and defined as follows:

$$d^H(X, Y) := \sup_{x \in X} \inf_{y \in Y} \|x - y\|.$$

It is well known that the set of closed subsets of a compact set is compact in this metric. This implies that the sequence $(\overline{Y}_t)_{t>0}$ has an accumulation point, with respect to the Hausdorff metric.

Lemma 4 *Let Y be a nonempty closed convex set and $\overline{Y}_\infty$ be an accumulation point of the sequence $(\overline{Y}_t)_{t>0}$. If $\vec{0} \in Y$ then*

$$\mathrm{Cone}(\overline{Y}_\infty) \subseteq Y.$$

Proof. For each $t \geq 0$ the set $\overline{Y}_t$ is a compact subset of $[-1, 1]^m$. The first claim follows by observing that the collection of compact subsets of $[-1, 1]^m$ is itself compact in the Hausdorff metric.

We now turn to the second claim. To show that $\mathrm{Cone}(\overline{Y}_\infty) \subset Y$ we fix a point $z \in \mathrm{Cone}(\overline{Y}_\infty)$ and construct a sequence of points $(x'_n)_{n \in \mathbb{N}}$ in Y that converges to it. Since Y is closed, this will prove that $z \in Y$. Because $z \in \mathrm{Cone}(\overline{Y}_\infty)$, there exist $\alpha > 0$ and $y \in \overline{Y}_\infty$ such that $z = \alpha y$.

Since the sequence $(\overline{Y}_t)_{t>0}$ converges to $\overline{Y}_\infty$ in the Hausdorff metric, for every $n \in \mathbb{N}$ there exists $t_n \geq \alpha$ satisfying

$$d^H(\overline{Y}_{t_n}, \overline{Y}_\infty) \leq \frac{1}{n}.$$

In particular, there is $w_n \in \overline{Y}_{t_n}$ such that $d(w_n, y) \leq \frac{1}{n}$. Setting $x_n := t_n w_n \in Y_{t_n}$ we deduce that $d\left(\frac{1}{t_n} x_n, y\right) \leq \frac{1}{n}$, or equivalently,

$$d\left(\frac{\alpha}{t_n} x_n, \alpha y\right) \leq \frac{\alpha}{n}.$$

Since (i) $\alpha \leq t_n$, (ii) $x_n, \vec{0} \in Y$, and (iii) Y is convex, it follows that $x'_n := \frac{\alpha}{t_n} x_n$ is in Y, and the result follows. ■

Lemma 5 *Let Y be a nonempty closed and convex set, and let $\overline{Y}_\infty$ be an (Hausdorff metric) accumulation point of the sequence $(\overline{Y}_t)_{t>0}$. For every $t \geq 1$ one has*

$$t\overline{Y}_t + \mathrm{Cone}(\overline{Y}_\infty) \subseteq Y.$$

Proof. For every $y \in \mathbb{R}^m$ denote by $Z^y = Y - y$. Then $\lim_{t\to\infty} d(\overline{Z^y}_t, \overline{Y}_t) = 0$, so that $\overline{Y}_\infty$ is an accumulation point of the sequence $(\overline{Z^y}_t)_{t>0}$.

Fix now $y \in Y$. The set Z^y is nonempty, closed, convex, and contains $\vec{0}$, so that by Lemma 4

$$\mathrm{Cone}(\overline{Y}_\infty) = \mathrm{Cone}(\overline{Z^y}_\infty) \subseteq Z^y = Y - y.$$

In particular,

$$y + \mathrm{Cone}(\overline{Y}_\infty) \subseteq Y.$$

The result follows from the fact that this inclusion holds for every $y \in Y$, and because $t\overline{Y}_t \subseteq Y$ for every $t > 0$. ∎

To conclude the proof that the condition is necessary we show that the sets $\overline{Y}_t$ are approachable by Player 1, provided t is large enough. This will imply that the set $\overline{Y}_\infty$, as an accumulation point of approachable sets, is approachable itself.

We start by finding a condition, lightly weaker than that in the definition of approachable sets, which is equivalent to it.

Lemma 6 *A nonempty closed set Y is approachable by Player 1 if and only if for every $\epsilon > 0$ there exists a strategy σ_1 and $T > 0$ such that*

$$d\left(\overline{\gamma}^T(\sigma_1, \sigma_2), Y\right) \leq \epsilon, \quad \forall \sigma_2.$$

Proof. The fact that if Y is approachable by Player 1 then it satisfies the condition in the lemma follows from the definition of approachability. For the converse implication, fix $\varepsilon > 0$, and let σ_1 and T be the strategy of Player 1 and the positive real number that are given by the condition in the lemma. Let σ_1' be the strategy of Player 1 that plays in blocks of length T; at the beginning of each block the strategy forgets past play and starts implementing σ_1 anew. The reader can verify that σ_1' approaches Y. ∎

We are finally ready to prove that the condition in Theorem 1 is necessary. Let Y be a closed and convex set that is attainable by Player 1. Therefore, there exists $T > 0$ such that for every $\epsilon > 0$ there exists a strategy σ_1 satisfying

$$d(\gamma^t(\sigma_1, \sigma_2), Y) \leq \epsilon, \quad \forall t \geq T, \ \forall \sigma_2.$$

By Lemma 6 this implies that the set $\frac{1}{t}Y$ is approachable by Player 1, provided that $t \geq T$. Since payoffs are bounded by 1, it follows that the set $\overline{Y}_t = \frac{1}{t}Y \cap [-1,1]^m$ is also approachable by Player 1, provided that $t \geq T$. Finally, the definition of approachability implies that the set $\overline{Y}_\infty$, as the Hausdorff limit of sets which are approachable by Player 1, is also approachable by Player 1. Since every set that is approachable by Player 1 contains a B-set for that player ([14, 18]), the proof of the necessity of the condition is complete.

6.2 Proof of Corollary 3

By Corollary 1, the set $\{x\}$ is attainable by Player 1 if and only if $\vec{0}$ is approachable by Player 1 and there exists $\delta > 0$ such that the vector δx is approachable by Player 1. We will show that the second property is equivalent to the Conditions **D2** and **D3**. We will then prove that, given that **D1** is satisfied, we can replace any one of these conditions with **D4**.

Part 1: δx is approachable by Player 1 if and only if Condition **D2** holds.

Note that the vector δx is approachable by Player 1 in the game with matrix payoff G if and only if the vector $\vec{0}$ is approachable by him in the game with matrix payoff $G - \delta x$. Since $\vec{0}$ is approachable by a player if and only if it is attainable by him, the result follows.

Part 2: δx is approachable by Player 1 if and only if Condition **D3** holds.

Let us write the B-set condition with respect to the singleton δx. The vector δx is approachable by Player 1 if and only if it is a B-set for him, that is,

$$\forall z \in \mathbb{R}^m,\ \exists x \in \Delta(A_1)\ \forall y \in \Delta(A_2) \langle u(x,y) - \delta x, z - \delta x \rangle \leq 0.$$

Setting $\lambda = \delta x - z$, we obtain

$$\begin{aligned} \forall \lambda \in \mathbb{R}^m,\ \exists x \in \Delta(A_1)\ \forall y \in \Delta(A_2) \langle u(x,y), \lambda \rangle &\geq \langle \delta x, \lambda \rangle, \\ \langle u(x,y), \lambda \rangle &\geq \delta \langle x, \lambda \rangle, \end{aligned}$$

which is equivalent to $v_\lambda \geq \delta \langle x, \lambda \rangle$ for every $\lambda \in \mathbb{R}^m$, which is Condition **D3**

Part 3: If the vector x is attainable by Player 1, then Condition **D4** is satisfied.

Suppose to the contrary that Condition **D4** is not satisfied. That is, for every $\delta_0 > 0$ there is $q \in \Delta(A_2)$ such that for every $p \in \Delta(A_1)$ one has $u(p,q) \neq \delta x$ for every $\delta > \delta_0$. We divide the argument into two cases.

Case A: There is $q \in \Delta(A_2)$ such that $u(p,q) \neq \delta x$ for every $p \in \Delta(A_1)$ and every $\delta > 0$.

We show that by playing constantly q (a strategy that we denote by q^*) Player 2 can prevent Player 1 from attaining x, contradicting the assumption. Let σ_1 be any strategy of Player 1. Denote by p_t the average mixed action played by Player 1 up to time t, that is, $p^t = \frac{1}{t}\int_0^t \sigma_1(s) \mathrm{d}s$. Then, $\gamma^t(\sigma_1, q^*) = tu(p_t, q)$. Thus, $\gamma^t(\sigma_1, q^*)$ is in the cone generated by $R_1(q) := \{u(p,q);\ p \in \Delta(A_1)\}$. This cone is closed and by assumption it does not contain x. Thus, there is a positive distance between x and this cone, implying that $\gamma^t(\sigma_1, q^*)$ cannot get arbitrarily close to x. This contradicts the fact that the vector x is attainable.

Case B: For every $q \in \Delta(A_2)$ there is $p \in \Delta(A_1)$ and $\delta > 0$ such that $u(p,q) = \delta x$, but the δ's are *not* bounded away from zero.

In this case, for every $\delta > 0$, there is $q_\delta \in \Delta(A_2)$ such that $\delta \geq \max\{\delta';\ \exists p \text{ such that } u(p, q_\delta) = \delta' x\}$. We show that for every $\delta > 0$, if Player 2 plays constantly q_δ (a strategy that we denote by q^*_δ), then there is $\varepsilon > 0$ such that for every σ_1, $\|\gamma^T(\sigma_1, q^*_\delta) - x\| < \varepsilon$ implies $T > \frac{1}{4\delta}$.

Fix $\delta > 0$. Denote

$$\delta_0 := \max\{\delta' \colon \exists p \text{ such that } u(p, q_\delta) = \delta' x\} < \delta.$$

In particular, $\delta_0 x \in R_1(q_\delta)$, and $\delta' x \notin R_1(q_\delta)$ for every $\delta' > \delta_0$. Let $E := \operatorname{conv}\left(R_1(q_\delta) \cup \{\vec{0}\}\right)$ be the convex hull of $R_1(q_\delta)$ and $\vec{0}$. The set E is convex, compact and it does not contain $\delta' x$ for every $\delta' > \delta_0$. In particular, $2\delta_0 x \notin E$. Thus, there is an open ball $F = B(2\delta_0 x, \eta)$ which is disjoint of E. By the hyperplane separation theorem there is a nonzero vector $\alpha \in \mathbb{R}^m$ such that $\langle e, \alpha\rangle \leq \langle f, \alpha\rangle$ for every $e \in E$ and $f \in F$. Since $\vec{0} \in E$, it follows that $0 = \langle \vec{0}, \alpha\rangle \leq \langle f, \alpha\rangle$ for every $f \in F$.

Without loss of generality assume that $\|\alpha\| = 1$. We claim that $0 < \langle x, \alpha\rangle$. Indeed, if $0 = \langle x, \alpha\rangle$, then every $f \in F$ can be expressed as $f = 2\delta_0 x + v$, where $v = v(f) \in B(\vec{0}, \eta)$. In particular, $0 \leq \langle f, \alpha\rangle = \langle v, \alpha\rangle$. It follows that $\langle v, \alpha\rangle = 0$ for every $v \in B(\vec{0}, \eta)$, which implies that $\alpha = 0$, contradicting the fact that $\|\alpha\| = 1$.

Suppose that $e \in R_1(q_\delta)$ and $T \cdot e \in B(x, \varepsilon)$, with $\varepsilon = \langle x, \alpha\rangle/2$. Then, $T \cdot e = x + z$, where $\|z\| \leq \varepsilon$. Thus, $\langle T \cdot e, \alpha\rangle = \langle x + z, \alpha\rangle$. Since $e \in E$ and $2\delta_0 x \in F$,

$$\langle e, \alpha\rangle \leq \langle 2\delta_0 x, \alpha\rangle \leq \langle 2\delta x, \alpha\rangle.$$

Hence,

$$T = \frac{\langle x + z, \alpha\rangle}{\langle e, \alpha\rangle} \geq \frac{\langle x, \alpha\rangle + \langle z, \alpha\rangle}{2\langle \delta x, \alpha\rangle} \geq \frac{\langle x, \alpha\rangle - \varepsilon}{2\langle \delta x, \alpha\rangle} = \frac{1}{4\delta}. \tag{5}$$

Recall that q^*_δ is the strategy of Player 2 that constantly plays q_δ. To derive a contradiction we will show that the vector x is not attainable; that is, for every T there is $\varepsilon > 0$ such that for every strategy σ_1 of Player 1 there is a strategy σ_2 of Player 2 and $t \leq T$ satisfying $d(\gamma^t(\sigma_1, \sigma_2), x) > \varepsilon$. Fix a strategy σ_1 of Player 1, and suppose that the cumulative payoff up to time T is within ε from x, that is, $\|\gamma^T(\sigma_1, q^*_\delta) - x\| \leq \varepsilon$. Let $p^T := \frac{1}{T}\int_0^T \sigma_1(s)\mathrm{d}s$ be the average mixed action played by σ_1 until time T. Thus, $Tu(p_T, q_\delta) = x + z$, where $\|z\| \leq \varepsilon$. Letting $e = u(p_T, q_\delta)$ we obtain by Eq. (5) that $T > \frac{1}{4\delta}$. In words, the time it takes to reach $B(x, \varepsilon)$ is at least $\frac{1}{4\delta}$. This shows that there is no uniform bound on the time at which the total payoff gets close to x. Thus, x is not attainable, which contradicts the assumption.

Part 4: If Condition **D4** and Condition **D1** are satisfied, then Condition **D3** is satisfied and x is attainable.

We will show that $v_\lambda \geq \delta_0\langle x, \lambda\rangle$ for every $\lambda \in \mathbb{R}^m$. If $\langle x, \lambda\rangle \leq 0$, then by Condition **D1**

$$v_\lambda \geq 0 \geq \delta_0\langle x, \lambda\rangle,$$

as required. If $\langle x, \lambda\rangle > 0$ then Condition **D3** implies that

$$v_\lambda = \inf_{q\in\Delta(A_2)} \sup_{p\in\Delta(A_1)} \langle u(p,q), \lambda\rangle \geq \langle \delta_0 x, \lambda\rangle = \delta_0\langle x, \lambda\rangle,$$

and the proof is complete.

6.3 Proof of Proposition 1

To prove that Conditions **E1** and **E2** are sufficient conditions, we prove that they imply Conditions **D1** and **D3**. By Corollary 2, Condition **E1** implies Condition **D1**. We now show that Condition **D3** holds as well.

It is sufficient to prove that Condition **D3** holds for every λ in the unit ball. The set $\mathcal{S}^{\geq} := \{\lambda \in \mathbb{R}^m \colon \|\lambda\| = 1, \langle x, \lambda\rangle \geq 0\}$ is compact. Since the function $\lambda \to v_\lambda$ is continuous, Condition **E2** implies that there exists $\epsilon > 0$ such that $v_\lambda \geq \epsilon$ for every $\lambda \in \mathcal{S}^{\geq}$. Let $\delta > 0$ such that $\delta\|x\| < \epsilon$. By Cauchy–Schwartz inequality,

$$v_\lambda \geq \langle \delta x, \lambda\rangle = \delta\langle x, \lambda\rangle, \quad \forall \lambda \in \mathcal{S}^{\geq}.$$

If $\langle \lambda, x\rangle < 0$ then Condition **E1** implies that

$$v_\lambda \geq 0 \geq \delta\langle x, \lambda\rangle,$$

and the proof is complete.

References

[1] E. Adida and G. Perakis, "A robust optimization approach to dynamic pricing and inventory control with no backorders", *Mathematical Programming*, Ser. B, vol. 107, 2006, pp. 97–129.

[2] A. Atamturk and M. Zhang, "Two–stage robust network flow and design under demand uncertainty", *Operations Research*, vol. 55, no. 4, 2007, pp. 662–673.

[3] T. Başar and P. Bernhard, *H-infinity optimal control and related minimax design problems: a dynamic game approach*, Birkhäuser, Boston, MA, 1991.

[4] D. Bauso, F. Blanchini, and R. Pesenti, "Optimization of long-run average-flow cost in networks with time-varying unknown demand", *IEEE Transactions on Automatic Control*, vol. 55, no. 1, 2010, pp. 20–31.

[5] D. Bauso, F. Blanchini, and R. Pesenti, "Robust control policies for multi-inventory systems with average flow constraints", *Automatica*, Special Issue on Optimal Control Applications to Management Sciences, vol. 42, no. 8, 2006, pp. 1255–1266.

[6] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, "Foreword: special issue on robust optimization", *Mathematical Programming*, Ser. B, vol. 107, 2006, pp. 1–3.

[7] D. P. Bertsekas and I. B. Rhodes, "On the minimax reachability of target set and target tubes", *Automatica*, vol. 7, 1971, pp. 233–247.

[8] D. Bertsimas and A. Thiele, "A Robust Optimization Approach to Inventory Theory", *Operations Research*, vol. 54, no. 1, 2006, pp. 150–168.

[9] D. Blackwell, "An analog of the minimax theorem for vector payoffs", *Pacific J. Math.*, vol. 6, no. 1, 1956, pp. 1–8.

[10] F. Blanchini, "Set invariance in control", *Automatica*, vol. 35, no. 11, 1999, pp. 1747–1768.

[11] F. Blanchini, S. Miani, and W. Ukovich, "Control of production-distribution systems with unknown inputs and system failures", *IEEE Transactions on Automatic Control*, vol. 45, no. 6, 2000, pp. 1072–1081.

[12] F. Blanchini, F. Rinaldi, and W. Ukovich, "Least inventory control of multi-storage systems with non-stochastic unknown input", *IEEE Transactions on Robotics and Automation*, vol. 13, 1997, pp. 633–645.

[13] F. Heid, D. Porath, and S. Stoltz, "Does capital regulation matter for bank behavior? evidence for german savings banks", Deutsche Bundesbank, Discussion Paper 03/04, 2004.

[14] T.-F. Hou, "Approachability in a two-person game", *The Annals of Mathematical Statistics*, vol. 42, no. 2, pp. 735–744.

[15] E. Khmelnitsky and M. Tzur, "Parallelism of continuous- and discrete-time production planning problems", *IIE Transactions*, vol. 36, 2004, pp. 611–628.

[16] E. Lehrer, E. Solan, and D. Bauso, "Repeated games over networks with vector payoffs: the notion of attainability", in Proc. of the Int. Conf. on NETwork Games, COntrol and OPtimization (NetGCooP 2011), 12-14 Oct. 2011, Paris.

[17] S. Soulaimani, M. Quincampoix, and S. Sorin, "Repeated games and qualitative differential games: approachability and comparison of strategies", *SIAM Journal on Control and Optimization*, vol. 48, no.4, 2009, pp. 2461–2479.

[18] X. Spinat, "A necessary and sufficient condition for approachability", *Mathematics of Operations Research*, vol. 27, no. 1, 2002, pp. 31–44.

[19] P. van Roy, "The impact of the 1988 Basel accord on banks capital ratios and credit risk-taking: an international study", 2005, EFMA 2004 Basel Meetings available at SSRN: http://ssrn.com/abstract=497264.