Absorption Paths and Equilibria in Quitting Games[*]

Galit Ashkenazi-Golan,[†] Ilia Krasikov,[‡] Catherine Rainer,[§] and Eilon Solan[¶]

June 24, 2021

Abstract

We study quitting games and define the concept of *absorption paths*, which is an alternative definition to strategy profiles that accommodates both discrete-time aspects and continuous-time aspects, and is parametrized by the total probability of absorption in past play rather than by time. We then define the concept of *sequentially 0-perfect* absorption paths, which are shown to be limits of ε-equilibrium strategy profiles as ε goes to 0. We establish that any quitting game that does not have simple equilibria (that is, an equilibrium where the game terminates in the first period or one where the game never terminates) has a sequentially 0-perfect absorption path.

We finally identify a class of quitting games that possess sequentially 0-perfect absorption paths.

[†]The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997800, Israel. e-mail: galit.ashkenazi@gmail.com.

[‡]National Research University Higher School of Economics, 20 Myasnitskaya Ulitsa, Moscow 101000, Russia. e-mail: krasikovis.main@gmail.com.

[§]Univ Brest, UMR CNRS 6205, 6, avenue Victor-le-Gorgeu, B.P. 809, 29285 Brest cedex, France. e-mail: Catherine.Rainer@univ-brest.fr.

[¶]The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997800, Israel. e-mail: eilons@post.tau.ac.il.

1 Introduction

Stochastic games were introduced by Shapley (1953) as a dynamic model, where the players' behavior affects the evolution of the state variable. Whether every multiplayer stochastic game admits an ε-equilibrium is one of the most difficult open problems in game theory to date. Mertens and Neyman (1981) proved that the value exists in two-player zero-sum games, Vieille (2000a, 2000b) proved that an ε-equilibrium exists in two-player nonzero-sum games, Solan (1999) extended this result to three-player absorbing games, and Flesch, Schoenmakers, and Vrieze (2008, 2009) proved the existence of an ε-equilibrium when each player controls one component of the state variable.

Solan and Vieille (2001) introduced a new class of stochastic games, called *quitting games*, where each player has two actions, continue and quit, the game terminates once at least one player chooses to quit, and the terminal payoff depends on the set of players who choose to quit at the termination stage. Solan and Vieille (2001) proved that if the payoff function satisfies a certain condition, then an ε-equilibrium exists. Simon (2007, 2012) and Solan and Solan (2020) extended this result to other families of payoff functions. Though the class of quitting games is simple – if the game has not terminated by a given stage, then necessarily all players continued so far – the analysis of these games is intricate, the mathematical tools used to study them are diverse, and include dynamical systems, topological tools, and linear complementarity problems, and the equilibria these games possess may be complex (see, Flesch, Thuijsman, and Vrieze (1997), Solan (2003), and Solan and Vieille (2002)).

The main difficulty in studying ε-equilibria in stochastic games is that the undiscounted payoff is not continuous over the space of strategies, hence one cannot apply a fixed point theorem to prove the existence of an ε-equilibrium. In this paper we provide a new representation for strategy profiles in quitting games, termed *absorption paths* (*AP* for short). This representation allows for both discrete-time aspects and continuous-time aspects in the players' behavior. Moreover, the undiscounted payoff is continuous over the space of absorption paths. In fact, the space of absorption paths is a compactification of the space of absorbing strategy profiles, when such profiles are properly represented.

The representation of strategy profiles via AP's involves parametrizing time ac-

cording to the accumulated probability of absorption. Using a parametrization of time to facilitate analysis of continuous-time models as the limit of discrete-time models was done, e.g., by Vieille (1992) for studying weak approachability in repeated games with vector payoffs and by Sorin and Vigeral (2020) for studying ε-optimal trajectories in discounted zero-sum stochastic games.

We define the concept of *sequentially 0-perfect* AP, denoted *0-AP*, which is the analogue of equilibrium in standard strategy profiles. We then show that when no simple ε-equilibrium exists, limits of ε-equilibria in standard strategy profiles are 0-AP's, and that every 0-AP induces an ε-equilibrium in standard strategy profiles, for every $\varepsilon > 0$. In particular, a quitting game admits an ε-equilibrium for every $\varepsilon > 0$ if and only if it admits a 0-AP.

This relation between ε-equilibrium in standard strategy profiles and 0-AP's is useful because an important research agenda is understanding ε-equilibrium in quitting games, and 0-AP's are much simpler to study than ε-equilibria: the sets of AP's and 0-AP's are compact in the weak topology, the payoff function is continuous over the set of AP's, and 0-AP's do not allow for profitable deviations. These properties should be contrasted with the analogous properties for strategies and ε-equilibria: the set of strategies is compact in the product topology, but in this topology the set of ε-equilibria is not compact and the payoff function is not continuous. Moreover, ε-equilibria allow for deviations where profit is low.

Finally, using Viability Theory we identify one class of quitting games where 0-AP's exist, thereby proving that this class admits an ε-equilibrium for every $\varepsilon > 0$.

The paper is organized as follows. The model of quitting games is presented in Section 2, and the equilibrium concept that we study is presented in Section 3. AP's are presented in Section 4, and their application to prove existence of ε-equilibrium in a certain class of quitting games is described in Section 5. Concluding remarks appear in Section 6.

2 The Model

Definition 2.1 *A quitting game is a pair $\Gamma = (I, r)$, where I is a finite set of players and $r : \prod_{i \in I} \{C^i, Q^i\} \to \mathbb{R}^I$ is a payoff function.*

Player i's action set is $A^i := \{C^i, Q^i\}$. These actions are interpreted as continue and quit, respectively. Denote by $A := \prod_{i \in I} A^i$ the set of action profiles. The game is played as follows. At every stage $n \in \mathbb{N}$ each player $i \in I$ chooses an action $a_n^i \in A^i$. If all players continue, the play continues to the next stage; if at least one player

3

quits, the play terminates, and the terminal payoff is $r(a_n)$, where $a_n = (a_n^i)_{i \in I}$. If no player ever quits, the payoff is $r(\vec{C})$, where $\vec{C} := (C^i)_{i \in I}$.

We denote by $A^* := A \setminus \{\vec{C}\}$ the set of all action profiles in which at least one player quits, by $A_1^* := \{(Q^i, C^{-i}), i \in N\}$ the set of all action profiles in which exactly one player quits, where $C^{-i} := (C^j)_{j \neq i}$, and by $A_{\geq 2}^* := A^* \setminus A_1$ the set of all action profiles in which at least two players quit.

A *mixed action profile* is a vector $\xi = (\xi^i)_{i \in I} \in [0,1]^I$, with the interpretation that ξ^i is the probability with which player i quits. The probability of absorption under the mixed action profile ξ is $p(\xi) := 1 - \prod_{i \in I}(1 - \xi^i)$. Extend the absorbing payoff to mixed action profiles that are absorbing with positive probability: for every $\xi \in [0,1]^I$ such that $\xi \neq \vec{0}$, define $r(\xi) := \frac{\sum_{a \in A^*} \xi(a) r(a)}{p(\xi)}$, where

$$\xi(a) := \left(\prod_{\{i:\ a^i = Q^i\}} \xi^i \right) \cdot \left(\prod_{\{i:\ a^i = C^i\}} (1 - \xi^i) \right), \quad \forall a \in A. \tag{1}$$

A (behavior) *strategy* of player i is a function $x^i = (x_n^i)_{n \in \mathbb{N}} : \mathbb{N} \to [0,1]$, with the interpretation that x_n^i is the probability that player i quits at stage n if the game did not terminate before that stage. A *strategy profile* is a vector $x = (x^i)_{i \in I}$ of strategies, one for each player. A strategy profile x is *stationary* if $x_n = x_{n+1}$ for all $n \in \mathbb{N}$; that is, if the players play the same mixed actions repeatedly as long as the game has not terminated.

Denote by $\theta := \min\{n \in \mathbb{N} : a_n \in A^*\}$ the stage of termination; $\theta = \infty$ if all players continue throughout the game. For every strategy profile x, the probability distribution of the random variable (θ, a_θ) is denoted $\mathbf{P}_x$. Denote by $\mathbf{E}_x$ the corresponding expectation operator. A strategy profile x is *absorbing* if $\mathbf{P}_x(\theta < \infty) = 1$.

The *payoff* under strategy profile x is

$$\gamma(x) := \mathbf{E}_x \left[\mathbf{1}_{\{\theta < \infty\}} r(a_\theta) + \mathbf{1}_{\{\theta = \infty\}} r(\vec{C}) \right].$$

Let $\varepsilon \geq 0$. A strategy profile x^* is an ε-*equilibrium* if $\gamma^i(x^*) \geq \gamma^i(x^i, x^{*,-i}) - \varepsilon$ for every player $i \in I$ and every strategy x^i of player i.

It is easy to check that every two-player quitting game admits an ε-equilibrium, for every $\varepsilon > 0$. Solan (1999) extended this result to three-player quitting games, see also Flesch, Thuijsman, and Vrieze (1997). Whether every quitting game admits an ε-equilibrium for every $\varepsilon > 0$ is an open problem.

4

3 Sequential ε-Perfectness

3.1 ε-Perfectness in Strategic-Form Games

Let $G = (I, (A^i)_{i \in I}, r)$ be a strategic-form game with set of players I, set of actions A^i for each player $i \in I$, and payoff function $r : A \to \mathbb{R}^I$, where $A = \prod_{i \in I} A^i$.

In an ε-equilibrium, no player can profit more than ε by deviating. This does not rule out the possibility that a player plays with small probability an action that generates her a low payoff. This deficiency is taken care of by the following concept, borrowed from Solan and Vieille (2001), which requires that a player does not play with positive probability actions that generates her a low payoff.

Definition 3.1 *Let $G = (I, (A^i)_{i \in I}, r)$ be a strategic-form game, let $i \in I$, and let $\xi \in \prod_{i \in I} \Delta(A^i)$ be a mixed action profile. Player i is ε-perfect at ξ in G if the following conditions hold for every action $a^i \in A^i$:*

$$r^i(a^i, \xi^{-i}) \leq r^i(\xi) + \varepsilon, \tag{2}$$

$$\xi^i(a^i) > 0 \quad \Longrightarrow \quad r^i(a^i, \xi^{-i}) \geq r^i(\xi) - \varepsilon. \tag{3}$$

Eq. (2) means that player i cannot gain more than ε by unilaterally altering her action; Eq. (3) requires that player i cannot lose more than ε, no matter which one of the actions to which she assigns positive probability is played. Player i is 0-perfect if ξ^i is a best response to ξ^{-i}.

Standard continuity arguments yield that if player i is ε_k-perfect at a mixed action profile ξ_k in the game $G_k = (I, (A^i)_{i \in I}, r_k)$, if $(\xi_k)_{k \in \mathbb{N}}$ converges to a limit ξ, if $(\varepsilon_k)_{k \in \mathbb{N}}$ converges to 0, and if r is a payoff function that satisfies $r^i = \lim_{k \to \infty} r_k^i$, then player i is 0-perfect at ξ in $G = (I, (A^i)_{i \in I}, r)$.

3.2 Sequentially ε-Perfect Players in Quitting Games

In this section we extend the concept of ε-perfect players to quitting games. Consider a quitting game $\Gamma = (I, r)$. For every vector $y \in \mathbb{R}^I$ let $G_\Gamma(y)$ be the one shot game with set of players I, set of actions $A^i = \{Q^i, C^i\}$ for each player $i \in I$, and payoff function r_Γ defined by

$$r_\Gamma(y; a) := \begin{cases} r(a) & a \neq \vec{C}, \\ y & a = \vec{C}. \end{cases}$$

The game $G_\Gamma(y)$ represents one stage of the game Γ, when the continuation payoff is y. A strategy profile in $G_\Gamma(y)$ is a vector $\xi \in [0,1]^I$, with the interpretation that ξ^i is the probability that player i chooses the action Q^i, for each $i \in I$.

We now define the concept of sequential ε-perfectness in quitting games. For every $n \in \mathbb{N}$ denote by $\gamma_n(x)$ the expected payoff under x, conditional that the game did not terminate in the first $n-1$ stages.[1]

$$\gamma_n(x) := \mathbf{E}_x[\mathbf{1}_{\{\theta < \infty\}} r(a_\theta) + \mathbf{1}_{\{\theta = \infty\}} r(\vec{C}) \mid \theta \geq n].$$

Definition 3.2 *Let Γ be a quitting game and let $i \in I$ be a player. Player i is sequentially ε-perfect at the strategy profile x in Γ if for every $n \in \mathbb{N}$, player i is ε-perfect at the mixed action profile x_n in the strategic-form game $G_\Gamma(\gamma_{n+1}(x))$.*

Player i is sequentially 0-perfect at the strategy profile x in Γ if x^i is a best response to x^{-i} in every sub-game.

Remark 3.3 *In the strategic-form game $G_\Gamma(\gamma_{n+1}(x))$, when the other players play x_n^{-i}, the payoff of player i when she plays x_n^i (resp. Q^i, C^i) is $\gamma_n^i(x)$ (resp. $r^i(Q^i, x^{-i})$, $(1 - p(C^i, x_n^{-i}))\gamma_{n+1}(x) + p(C^i, x_n^{-i})r^i(C^i, x_n^{-i}))$. Therefore, if player i is ε-perfect at x_n in $G_\Gamma(\gamma_{n+1}(x))$, then in particular $r^i(Q^i, x_n^{-i}) \leq \gamma_n^i(x) + \varepsilon$, and, if $x_n^i(Q^i) > 0$ then $r^i(Q^i, x_n^{-i}) \geq \gamma_n^i(x) - \varepsilon$.*

The following two results relate ε-equilibria to sequential ε-perfectness in quitting games.

Theorem 3.4 (Simon, 2007, Theorem 3 + Solan and Vieille, 2001, Proposition 2.13) *A quitting game Γ admits an ε-equilibrium for every $\varepsilon > 0$, if and only if at least one of the following statements holds.*

(S.1) For every $\varepsilon > 0$ sufficiently small the game admits a stationary ε-equilibrium.

(S.2) For every $\varepsilon > 0$ sufficiently small the game admits an ε-equilibrium x that has the following structure: there is a player $i \in I$ who quits with probability 1 at the first stage; from the second stage and on, all players punish player i with a payoff ε-close to her min-max level.[2]

(S.3) For every $\varepsilon > 0$ sufficiently small there is an absorbing strategy profile x such that all players $i \in I$ are sequentially ε-perfect at x.

---

[1]Since a strategy x^i is a function from $\mathbb{N}$ to $[0,1]$, the conditional probability distribution $\mathbf{P}_x(\cdot \mid \theta \geq n)$ is uniquely defined even when $\mathbf{P}_x(\theta < n) = 1$.

Theorem 3.5 (Solan and Vieille, 2001, Propositions 2.4 and 2.13) *Let $\varepsilon >$ 0 be sufficiently small. Every absorbing strategy profile x at which all players are sequentially ε-perfect is an $\varepsilon^{1/6}$-equilibrium.*

4 An Alternative Representation of Strategy Profiles

A strategy profile $x = (x_n)_{n\in\mathbb{N}}$ is parameterized by time: x_n^i is the probability that player i quits at stage n if the game did not terminate before that stage. As is well known, the space of strategies is compact in the product topology. There are two issues with this topology:

- The payoff is not continuous in this topology. Indeed, if for every $k \in \mathbb{N}$, $x(k)$ is the stationary strategy profile in which in every stage each player quits with probability $\frac{1}{k}$, then the sequence $(x(k))_{k\in\mathbb{N}}$ converges to the strategy profile x under which all players always continue. While under the strategy profile $x(k)$ absorption occurs with probability 1 and $\lim_{k\to\infty} \gamma(x(k)) = \frac{1}{|I|}\sum_{i\in I} r(Q^i, C^{-i})$, under the strategy profile x the game is never absorbed and $\gamma(x) = r(\vec{C})$.

- It may not be possible to generate the limit behavior of a sequence of strategy profiles by a strategy profile. For example, when $(x(k))_{k\in\mathbb{N}}$ are the strategy profiles that are defined in the first bullet, we have $\lim_{k\to\infty} \mathbf{P}_{x(k)}[a_\theta = (Q^i, C^{-i}) \mid \theta = n] = \frac{1}{|I|}$ for every $n \in \mathbb{N}$ and $i \in I$, yet there is no strategy profile x that satisfies $\mathbf{P}_x[a_\theta = (Q^i, C^{-i}) \mid \theta = n] = \frac{1}{|I|}$ for every $n \in \mathbb{N}$. Indeed, under such a strategy profile $x = (x^i)_{i\in I}$, for every $n \in \mathbb{N}$ we have $x_n^i > 0$ for each $i \in I$, and then $\sum_{i\in I} \mathbf{P}_x[a_\theta = (Q^i, C^{-i}) \mid \theta = n] < 1$ as soon as $|I| > 1$.

In this section we will provide an alternative representation of strategy profiles, that takes care of these two issues by allowing both discrete-time behavior and continuous-time behavior. The representation will be based on a change of parametrization: instead of parameterizing the strategy profile according to time n, we will parameterize it according to the probability of termination t. The parameter t will run from 0 to 1, representing the total probability of absorption. In addition, for each $t \in [0, 1]$ and action profile $a \in A^*$ we will indicate the probability by which the game is absorbed by the action profile a up to that moment in which the total probability of absorption is t.

---

[2] The min-max level of player i is $v^i := \inf_{x^{-i}} \sup_{x^i} \gamma^i(x^i, x^{-i})$.

4.1 A Motivating Example

The following example motivates the representation of strategy profiles by parametrizing time differently. For every $\eta \geq 0$, consider the three-player quitting game Γ_η displayed in Figure 1.

	C^3			Q^3	
	C^2	Q^2		C^2	Q^2
C^1	$-1,-1,-1$	$-1,0,2$		$2,-1,0$	$0,\eta,-1$
Q^1	$0,2,-1$	$\eta,-1,0$		$-1,0,\eta$	$-1,-1,-1$

Figure 1: The three-player game Γ_η.

The game Γ_0 was studied in Flesch, Thuijsman, and Vrieze (1997), who characterized the set of its 0-equilibria. They showed that the following periodic strategy profile x^* with period 3 is an equilibrium, and that all players are 0-perfect at x^*: at stage 1 (resp. 2, 3) Player 1 (resp. 2, 3) quits with probability $\frac{1}{2}$, while the other two players continue.

Solan (2003) studied the game Γ_η for η small, and showed that this game admits no 0-equilibrium. For example, the strategy profile x^* described in the previous paragraph is not a 0-equilibrium when $\eta > 0$, because Player 3 is better off quitting in stage 1 and obtaining $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot \eta = \frac{\eta}{2}$, while her payoff under x^* is 0.

Solan (1999) implies that every three-player quitting game, and in particular the games Γ_η for $\eta > 0$, admit an ε-equilibrium for every $\varepsilon > 0$. In fact, the following strategy profile x^m, which depends on a positive integer m and is periodic with period $3m$, is an ε-equilibrium of Γ_η, provided m is sufficiently large: at stages $1, 2, \ldots, m$ (resp. $m+1, m+2, \ldots, 2m$, resp. $2m+1, 2m+2, \ldots, 3m$) Player 1 (resp. 2, 3) quits with probability ρ, where $(1-\rho)^m = \frac{1}{2}$, while the other two players continue.

The limit of the sequence $(x^m)_{m \in \mathbb{N}}$ as m goes to infinity is the strategy profile under which all players continue in all stages, which is not an ε-equilibrium of Γ_η, provided $\varepsilon < 1$. Another natural limit of the sequence $(x^m)_{m \in \mathbb{N}}$ is a strategy profile in continuous time: First Player 1 quits in continuous time, until the total probability that she quits is $\frac{1}{2}$, then Player 2 quits in continuous time, until the total probability that she quits is $\frac{1}{2}$, then Player 3 quits in continuous time, until the total probability that she quits is $\frac{1}{2}$, and, if the play has not terminated (which happens with probability $\frac{1}{8}$), the players repeat this behavior.

This example motivates a definition of strategy profiles that include both discrete-time aspects and continuous-time aspects, which we provide in the next section.

4.2   Absorption Paths - Definition

For clarity of exposition, before providing a general definition of our new concept, absorption paths, we define absorption paths that originate from absorbing strategy profiles.

Definition 4.1 *For every absorbing strategy profile x and every $n \in \mathbb{N}$ denote $t_n :=$ $\mathbf{P}_x(\theta < n)$. The absorption path (AP for short) defined by x is the function $\pi^x :$ $[0,1] \times A^* \to [0,1]$ given by $\pi_t^x(a) := \mathbf{P}_x(\theta \leq n, a_\theta = a)$ for every $a \in A^*$, $n \in \mathbb{N}$, and $t \in [t_n, t_{n+1})$, and is continuous from the left at $t = 1$.*

Remark 4.2   *1. While a strategy profile is a vector of strategies, each describing the behavior of an individual player, an AP highlights information about the absorbing entries that are played along the game. For example, it clearly indicates through which action profiles the play is more probable to be absorbed early in the game, and which later on. In particular, AP's cannot be defined for each player separately, but only for strategy profiles.*

   2. The function $x \mapsto \pi^x$ that is defined in Definition 4.1 is not one-to-one. Indeed, fix an absorbing strategy profile x and let x' be the strategy profile in which all players continue in the first stage, and from the second stage on they follow x :

$$x_n'^i = \begin{cases} C^i, & \text{if } n = 1, \\ x_{n-1}^i, & \text{if } n > 1. \end{cases}$$

   Then $\pi^{x'} = \pi^x$. In fact, given an absorbing strategy profile x, the addition or elimination of stages in which all players continue is the only way to create an absorbing strategy profile x' such that $\pi^x = \pi^{x'}$.

   3. For simplicity we defined AP's for absorbing strategy profiles. The definition can be adapted to strategy profiles x for which $\mathbf{P}_x(\theta < \infty) < 1$.

For each strategy profile x, the AP π^x is càdlàg. For all càdlàg maps $\pi : [0,1] \times A^* \to [0,1]$ and every $a \in A^*$, set $\pi_{0-}(a) := 0$ and $\pi_{t-}(a) := \lim_{s \nearrow t} \pi_s(a)$ for $t \in (0,1]$. Set also $\widehat{\pi}_t := \sum_{a \in A^*} \pi_t(a)$, $\Delta\pi_t := \pi_t - \pi_{t-}$ and $\Delta\widehat{\pi}_t := \widehat{\pi}_t - \widehat{\pi}_{t-}$ for every $t \in [0,1]$.

The AP defined by x satisfies the following properties. (1) For $t \in [t_n, t_{n+1})$, $\widehat{\pi}_t^x$ is the probability that the game is absorbed before or at period n, that is, $\widehat{\pi}_t^x = \mathbf{P}_x(\theta \leq n) = t_{n+1}$, and so $\widehat{\pi}_t^x \geq t$ for all $t \in [0,1)$. (2) On the interval (t_n, t_{n+1}), $\widehat{\pi}_t^x$ is constant and equals $\mathbf{P}_x(\theta \leq n)$. (3) The ratio $\frac{\Delta\pi_{t_n}^x}{1-t_n}$ equals the probability that the game terminates at period n by action profile a, and is therefore equal to $x_n(a)$.

9

As we have seen in Section 4.1, the set of AP's defined by absorbing strategy profiles is not closed in the weak topology of càdlàg paths. The new concept of the paper, absorption paths, are the elements in the closure of this set of AP's.

Let $\mathbf{F}$ be the set of càdlàg paths $\pi = (\pi_t(a), a \in A^*)_{t \in [0,1]}$ with values in $[0,1]^{A^*}$, such that, for all $a \in A^*$, $t \mapsto \pi_t(a)$ is nondecreasing. We endow $\mathbf{F}$ with the weak topology: a sequence $(\pi^k)_{k \in \mathbb{N}}$ converges to π if $\int_{[0,1]} f(t) d\pi_t^k(a) \to \int_{[0,1]} f(t) d\pi_t(a)$, for every continuous map $f : [0,1] \to \mathbb{R}$ and every $a \in A^*$. In such a case we write $\pi^k \Rightarrow \pi$. Recall that $\pi^k \Rightarrow \pi$ if and only if $\pi_t^k \to \pi_t$ for every $t \in [0,1]$ at which π is continuous. The set $\mathbf{F}$ is sequentially compact in the weak topology.

For each $\pi \in \mathbf{F}$ define

$$
\begin{aligned}
T(\pi) &:= \{t \in [0,1], \widehat{\pi}_t = t\}, \\
S(\pi) &:= \{t \in [0,1], \Delta\pi_t \neq 0\}.
\end{aligned}
$$

$S(\pi)$ is the set of jumps of π, that is, the play is in discrete time, and as we will see, $T(\pi)$ is the set of t's in which the play is in continuous time.

Finally we introduce the right-hand side derivative of $t \mapsto \pi_t$: for every $t \in [0,1)$ set $\dot{\pi}_t := \liminf_{s \searrow t} \frac{\pi_s - \pi_t}{s-t}$. By Lebesgue's Theorem for the differentiability of monotone functions, for every $\pi \in \mathbf{F}$ the liminf is in fact a limit almost everywhere in $[0,1)$.

Definition 4.3 *An element π of F is an* absorption path *(AP) if*

(A.1) for every $t \in [0,1]$, we have $\widehat{\pi}_t \geq t$,

(A.2) on each connected component (t_1, t_2) of $[0,1] \setminus (S(\pi) \cup T(\pi))$, $\widehat{\pi}$ is constant and equals t_2,

(A.3) for every $t \in S(\pi)$, there exists $\xi_t = (\xi_t^i)_{i \in N} \in [0,1]^I$ such that

$$
\frac{\Delta\pi_t(a)}{1-t} = \xi_t(a), \quad \forall a \in A^*. \tag{4}
$$

(A.4) For every $t \in T(\pi) \setminus \{1\}$ we have $\dot{\pi}_t(a) = 0$ for every $a \in A_{\geq 2}^*$.

The set of absorption paths is denoted by $\mathbb{A}$.

Remarks 4.4 *Let $\pi \in \mathbb{A}$ be an AP.*

1. *For every $t \in S(\pi) \cup T(\pi)$, the quantity $\pi_t(a)$ should be thought of as the unconditional probability that the play is absorbed by the action profile a, until the moment in which the total probability of absorption is t.*

10

2. Elements $t \in S(\pi)$ correspond to play in discrete time, and for such t, ξ_t is the mixed action profile the players play at t, and t is the total probability of absorption up to t. This explains (A.3).

3. Elements $t \in T(\pi) \setminus \{1\}$ correspond to play in continuous time. In intervals $(t_1, t_2) \in T(\pi)$, the time at which a player quits is a continuous random variable. Therefore, players cannot quit simultaneously with a positive probability. This explains (A.4).

4. If (t, t') is a connected component of $[0, 1] \setminus (S(\pi) \cup T(\pi))$, then $t \in S(\pi)$ and $t' = t + (1 - t)p(\xi_t)$, where ξ_t is defined in Eq. (4). This interval corresponds to the increase in probability due to play in discrete time.

5. Since, for all $a \in A^*$, $s \mapsto \pi_s(a)$ is nondecreasing, π is continuous at t if and only if $\widehat{\pi}$ is continuous at t, for every $t \in [0, 1]$. It follows from (A.2) that on each connected component of $[0, 1] \setminus (S(\pi) \cup T(\pi))$ the process π is constant.

6. Let $t \in S(\pi)$. Since π is càdlàg and nondecreasing, we get from (A.1) that $\widehat{\pi}_t > t$, and from (A.2) that $\widehat{\pi}_s = \widehat{\pi}_t$ for every $s \in [t, \widehat{\pi}_t)$. In particular $\widehat{\pi}_{\widehat{\pi}_t-} = \widehat{\pi}_t$.

7. For every $t \in [0, 1]$, both $\widehat{\pi}_{t-}$ and $\widehat{\pi}_t$ belong to $T(\pi) \cup S(\pi)$.

8. From (A.2) and Remark 4.4(7), we deduce that $[0, 1)$ is partitioned to a countable number of intervals $U = [t_1, t_2)$, with, either $U \subset T(\pi)$, or $t_1 \in S(\pi)$ and $t_2 = \widehat{\pi}_{t_1}$. On each of these intervals, π is continuous, with $\widehat{\pi}_t = t$ if $t \in U \subset T(\pi)$, and $\widehat{\pi}_t = t_2$ otherwise.

9. The function π is continuous at $t = 1$: indeed, since, for all $t \in [0, 1]$, $t \leq \widehat{\pi}_t \leq 1$, we have $\widehat{\pi}_1 = \lim_{t \nearrow 1} \widehat{\pi}_t = 1$.

10. For every $a \in A^*_{\geq 2}$, the function $t \mapsto \pi_t(a)$ is piecewise constant.

11. The reader may wonder why we defined $\dot{\pi}$ with liminf and not with limsup. This choice is crucial to ensure that the set of AP's is sequentially compact. Indeed, we show in Proposition 4.11 below that, if some function π is the limit of a sequence of AP's, then $\liminf_{s \nearrow t} \frac{\pi_s(a) - \pi_t(a)}{s - t} = 0$ for every $t \in T(\pi)$ and $a \in A^*_{\geq 2}$. This result does not hold when changing liminf into limsup in the definition of $\dot{\pi}$.

Example 4.5 *Consider the limit behavior in continuous time that is described in Section 4.1: first Player 1 quits in continuous time, until the total probability that she*

11

quits is $\frac{1}{2}$, then Player 2 quits in continuous time, until the total probability that she quits is $\frac{1}{2}$, then Player 3 quits in continuous time, until the total probability that she quits is $\frac{1}{2}$, and, if the play has not terminated, the players repeat this behavior. The AP that corresponds to this behavior is displayed in Figure 2. Player 1 quits first with total probability $\frac{1}{2}$, and therefore $\pi(Q^1, C^2, C^3)$ increases linearly from 0 at $t = 0$ to $\frac{1}{2}$ at $t = \frac{1}{2}$. Player 2 quits afterwards with total probability $\frac{1}{2}$, hence the probability that the play terminates until Player 2 is done with quitting is $\frac{3}{4}$. Therefore, $\pi(C^1, Q^2, C^3)$ increases linearly from 0 at $t = \frac{1}{2}$ to $\frac{1}{4}$ at $t = \frac{3}{4}$. Since Player 3 quits after Player 2 with total probability $\frac{1}{2}$, the probability that the play terminates until Player 3 is done with quitting is $\frac{7}{8}$. It follows that, $\pi(C^1, C^2, Q^3)$ increases linearly from 0 at $t = \frac{3}{4}$ to $\frac{1}{8}$ at $t = \frac{7}{8}$. Afterwards, Player 1 again quits in continuous time with total probability $\frac{1}{2}$, hence the probability that the play terminates when Player 1 is done with the next round of termination is $\frac{15}{16}$. As a result, $\pi(Q^1, C^2, C^3)$ increases linearly from $\frac{1}{2}$ at $t = \frac{7}{8}$ to $\frac{9}{16}$ at $t = \frac{15}{16}$, and so on.
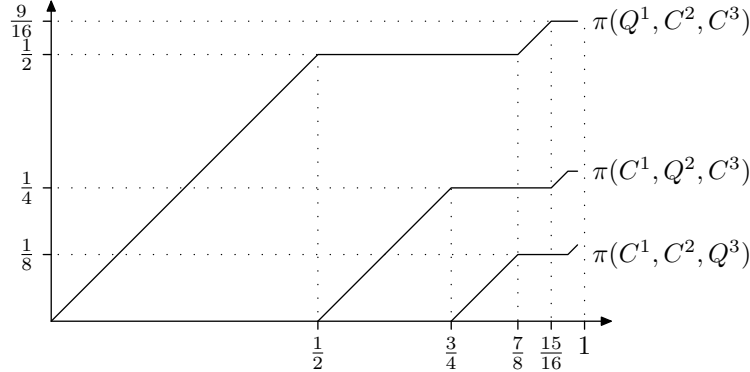

Figure 2: The AP in Example 4.5.

Example 4.6 *Figure 3 displays a more generic AP π for the case $|I| = 2$. This AP corresponds to the following behavior: first the players play the mixed action profile $\xi_0 = (\frac{1}{3}, \frac{1}{4})$ (that is, Player 1 (resp. Player 2) quits with probability $\frac{1}{3}$ (resp. $\frac{1}{4}$)), then they play the mixed action profile $\xi_{1/2} = (\frac{1}{2}, 0)$, and then they quit in continuous time until the game terminates, with Player 1 quitting at a double rate than Player 2.*

Indeed, $S(\pi) = \{0, \frac{1}{2}\}$ and $T(\pi) = [\frac{3}{4}, 1]$, hence the players play twice in discrete time (at $t = 0, \frac{1}{2}$) and then in continuous time (at $t \in [\frac{3}{4}, 1]$).

To find $\xi_0 = (\xi_0^1, \xi_0^2)$, we recall that it is the unique mixed action profile that satisfies $\frac{3}{12} = \Delta\pi_0(Q^1, C^2) = \xi_0^1(1 - \xi_0^2)$, $\frac{2}{12} = \Delta\pi_0(C^1, Q^2) = (1 - \xi_0^1)\xi_0^2$, and $\frac{1}{12} = \Delta\pi_0(Q^1, Q^2) = \xi_0^1\xi_0^2$, hence it is $\xi_0^1 = \frac{1}{3}$ and $\xi_0^2 = \frac{1}{4}$.

12

Similarly, $\xi_{1/2} = (\xi_{1/2}^1, \xi_{1/2}^2)$ is the unique mixed action profile that satisfies $\frac{1}{2} = \frac{3}{12}/\frac{1}{2} = \Delta\pi_{1/2}(Q^1, C^2) = \xi_{1/2}^1(1 - \xi_{1/2}^2)$, $0 = \Delta\pi_{1/2}(C^1, Q^2) = (1 - \xi_{1/2}^1)\xi_{1/2}^2$, and $0 = \Delta\pi_{1/2}(Q^1, Q^2) = \xi_{1/2}^1\xi_{1/2}^2$, hence it is $\xi_{1/2}^1 = \frac{1}{2}$ and $\xi_{1/2}^2 = 0$.

In the interval $[3/4, 1]$ the slope of $\pi_t(Q^1, C^2)$ is twice the slope of $\pi_t(C^1, Q^2)$, reflecting the rates at which the players quit in the last phase of the game.
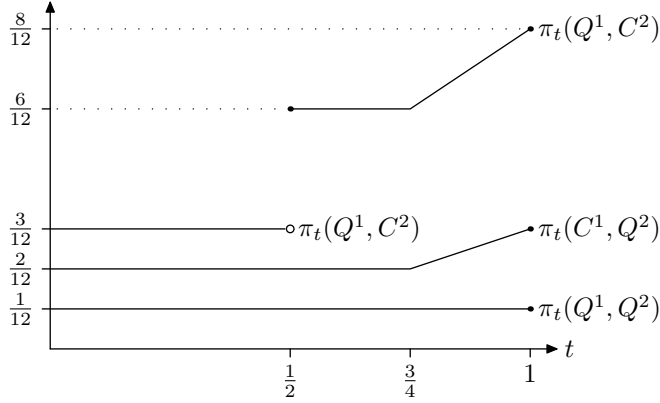

Figure 3: The AP in Example 4.6.

Remark 4.7 By Remark 4.4(3), in intervals that belong to $T(\pi)$, $\pi(a)$ increases only for $a \in A_1^*$. Moreover, in those intervals $\sum_{i \in I} \dot{\pi}_t(Q^i, C^{-i}) = 1$, and as seen in Example 4.6, $\dot{\pi}_t(Q^i, C^{-i})$ is equal to the ratio between the rate at which player i quits at t and the sum of rates at which all players quit at t.

The following result states that the set of all π^x, where x ranges over all absorbing strategy profiles, is sequentially dense in the set of AP's. Thus, the set of AP's is a compactification of the set of absorbing strategy profiles.

Proposition 4.8 For every AP π there is a sequence of absorbing strategy profiles $(x^k)_{k \in \mathbb{N}}$ such that $\pi^{x^k} \Rightarrow \pi$.

To prove Proposition 4.8 we need the following technical lemma, which states that a correlated action profile that (a) absorbs with low probability and (b) provided absorption occurs, absorbs mainly by single quittings, can be well approximated by an (independent) mixed action profile.

Lemma 4.9 Let $\varepsilon \in (0, \frac{1}{2}]$ be sufficiently small, and let $y \in \Delta(A)$ be a distribution that satisfies $p(y) := 1 - y(\vec{C}) \leq \varepsilon$ and $y(a) \leq \varepsilon y(Q^i, C^{-i})$ for each $i \in I$ and every

13

$a \in A_{\geq 2}^*$ such that $a^i = Q^i$. Then there exists a mixed action profile $\xi \in [0,1]^I$ such that $p(\xi) = p(y)$ and

$$|\xi(a) - y(a)| < 2^{|I|} \cdot \varepsilon p(y), \quad \forall a \in A^*, \tag{5}$$

where $\xi(a)$ is defined in Eq. (1).

Proof. If $y^i(Q^i, C^{-i}) = 0$ for some $i \in I$, then the assumptions imply that $y(a) = 0$ for every a such that $a^i = Q^i$. In such a case we set $\xi^i = 0$, and then Eq. (5) holds for every a such that $a^i = Q^i$. We therefore assume from now on that $y^i(Q^i, C^{-i}) > 0$ for each $i \in I$. Denote $\delta := \min_{i \in I} y^i(Q^i, C^{-i}) \in (0, \varepsilon]$.

To construct a mixed action profile ξ that satisfies the conditions we will construct a vector field φ over the set $\Xi := [0,1]^I$, prove that it has at least one zero, and prove that all its zeros satisfy the conditions in the lemma. It will be useful to require in addition that for every zero ξ of the vector field the ratio $\frac{\xi(Q^i, C^{-i})}{y(Q^i, C^{-i})}$ is the same for all $i \in I$.

Step 1: Definition of a vector field φ over Ξ.

For every $\xi \in \Xi$ and every $i \in I$ define

$$\varphi_0^i(\xi) := \frac{2}{\delta}\left(p(y) - p(\xi)\right) + \left(\frac{1}{|I|}\sum_{j \in I}\frac{\xi(Q^j, C^{-j})}{y(Q^j, C^{-j})} - \frac{\xi(Q^i, C^{-i})}{y(Q^i, C^{-i})}\right).$$

As we will see, at every zero ξ of the vector field that is yet to be defined, both summands in the definition of φ_0 will vanish, hence the properties we need will hold. The coefficient $\frac{2}{\delta}$ of the first summand ensures that the contribution of the first term is larger than that of the second term, so that $\varphi_0^i(\xi) < 0$ whenever $\xi^i = 1$. Define

$$\varphi^i(\xi) := \mathbf{1}_{\{\varphi_0^i(\xi) \geq 0\}} \cdot \varphi_0^i(\xi) + \mathbf{1}_{\{\varphi_0^i(\xi) < 0\}} \cdot \xi^i \cdot \varphi_0^i(\xi).$$

As we will see in Step 2, the multiplication by ξ^i on $\{\varphi_0^i(\xi) < 0\}$ ensures that $\varphi^i(\xi) \geq 0$ whenever $\xi^i = 0$.

Step 2: The vector field has a zero in Ξ.

By Brouwer's fixed point theorem, since Ξ is convex and compact, to prove that the vector field has a zero in Ξ it is sufficient to establish three properties: (a) φ is continuous, (b) $\varphi^i(\xi) \leq 0$ whenever $\xi^i = 1$, and (c) $\varphi^i(\xi) \geq 0$ whenever $\xi^i = 0$.

Property (a) follows from the definition of φ. We turn to prove Property (b). If $\xi^i = 1$ then $p(\xi) = 1$. Since for every $j \in I$ and every $\xi \in \Xi$ we have $\frac{\xi(Q^j, C^{-j})}{y(Q^j, C^{-j})} \in [0, \frac{1}{\delta}]$, and since $\varepsilon < \frac{1}{2}$, we have in this case

$$\varphi_0^i(\xi) \leq \frac{2}{\delta}(\varepsilon - 1) + \frac{1}{\delta} \leq 0.$$

14

Property (c) holds since on $\{\varphi_0^i(\xi) < 0\} \cap \{\xi^i = 0\}$ we have $\varphi^i(\xi) = 0$.

Step 3: For every zero ξ of φ we have (i) $p(y) = p(\xi)$ and (ii) $\frac{\xi(Q^i, C^{-i})}{y(Q^i, C^{-i})}$ is the same for all $i \in I$.

Suppose that $\varphi(\xi) = \vec{0}$. We will distinguish between three cases: $p(y) = p(\xi)$, $p(y) > p(\xi)$, or $p(y) < p(\xi)$.

Suppose first that $p(y) = p(\xi)$. It follows that $\xi^j > 0$ for at least one $j \in I$. If $\xi^i = 0$ for some $i \in I$, then $\varphi_0^i(\xi) > 0$ and hence $\varphi^i(\xi) > 0$, a contradiction. Hence $\xi^i > 0$ for every $i \in I$. But then the second summand in the definition of $\varphi_0^i(\xi)$ is 0 for all $i \in I$, which implies that $\frac{\xi(Q^i, C^{-i})}{y(Q^i, C^{-i})}$ is indepdnent of i.

Suppose now that $p(y) > p(\xi)$. Then the first summand in the definition of $\varphi_0^i(\xi)$ is positive for all $i \in I$. This implies that the second summand must be negative for all $i \in I$. But the sum of the second summand over all $i \in I$ is 0, a contradiction.

Last of all, suppose that $p(y) < p(\xi)$. In particular, $p(\xi) > 0$, and hence there is $i \in I$ such that $\xi^i > 0$. Let $i \in I$ be an index such that $\frac{\xi(Q^i, C^{-i})}{y(Q^i, C^{-i})}$ is maximal (and in particular positive). For this i, the second summand in the definition of $\varphi_0^i(\xi)$ is negative, but then $\varphi^i(\xi) < 0$, a contradiction.

We note that since $\frac{\xi(Q^i, C^{-i})}{y(Q^i, C^{-i})}$ is the same for all $i \in I$, so is the sign of $\xi(Q^i, C^{-i}) - y(Q^i, C^{-i})$.

Step 4: Every zero ξ of φ satisfies Eq. (5).

For each $a \in A_{\geq 2}^*$ we have

$$\xi(a) = \left(\prod_{\{i:\, a^i = Q^i\}} \xi^i \right) \cdot \left(\prod_{\{i:\, a^i = C^i\}} (1 - \xi^i) \right) \leq \prod_{\{i:\, a^i = Q^i\}} \xi^i \leq (p(\xi))^2 \leq \varepsilon p(y),$$

where the penultimate inequality holds since $\#\{i:\, a^i = Q^i\} \geq 2$. For each such a we also have $y(a) \leq \varepsilon y(Q^i, C^{-i}) \leq \varepsilon p(y)$, where i is any index such that $a^i = Q^i$. Hence, Eq. (5) holds for $a \in A_{\geq 2}^*$.

To show that Eq. (5) holds for $a \in A_1^*$, note that

$$\left| \sum_{a \in A_1^*} \xi(a) - \sum_{a \in A_1^*} y(a) \right| = \left| \sum_{a \in A^*} \xi(a) - \sum_{a \in A^*} y(a) - \sum_{a \in A_{\geq 2}^*} \xi(a) + \sum_{a \in A_{\geq 2}^*} y(a) \right|$$

$$= \left| \sum_{a \in A_{\geq 2}^*} \xi(a) - \sum_{a \in A_{\geq 2}^*} y(a) \right| < 2^{|I|} \varepsilon p(y).$$

Since the sign of $\xi(Q^i, C^{-i}) - y(Q^i, C^{-i})$ is independent of i, this implies that Eq. (5) holds for every $a \in A_1^*$ as well. ∎

15

Proof of Proposition 4.8. The idea of the proof is to discretize $[0,1]$; that is, for every $k \in \mathbb{N}$, we define a countable set $S^k = (s_n^k)_{n \in \mathbb{N}} \subset [0,1]$ and a strategy profile x^k in such a way that x_n^k approximates the behavior under π between the n'th and $(n+1)$'st point of S^k. The set S^k contains (a) the points t in $S(\pi)$ with high jumps: $\Delta \widehat{\pi}_t \geq \frac{1-t}{k}$, and (b) covers $[0,1]$ minus the corresponding intervals $[t, \widehat{\pi}_t)$ with well chosen points s_n^k such that $s_{n+1}^k \leq \frac{1}{k}(1 - s_n^k)$, i.e., the conditional probability of absorption in $[s_n^k, s_{n+1}^k)$ is less than $\frac{1}{k}$. If s_n^k satisfies the condition in (a), we take $x_n^k = \xi_{s_n^k}$; otherwise, we use Lemma 4.9 to approximate the behavior of π in the interval (s_n^k, s_{n+1}^k).

We turn to the formal construction. Fix an AP $\pi \in \mathbb{A}$ and $k \geq 2$. Let

$$S_0^k := \{t \in S(\pi) : \widehat{\pi}_t - t \geq \tfrac{1-t}{k}\} = \{t \in S(\pi) : p(\xi_t) \geq \tfrac{1}{k}\}.$$

Define the set $S^k = (s_n^k)_{n \in \mathbb{N}} \subset [0,1]$ as follows:

- $s_1^k := 0$.

- For $n \in \mathbb{N}$, define inductively $s_{n+1}^k := \sup\left(\left((S(\pi) \cup T(\pi)) \cap [0, s_n^k + \tfrac{1-s_n^k}{k}] \right) \cup \{\widehat{\pi}_{s_n^k}\} \right)$.
  In words, if $s_n^k \in S_0^k$ then $s_{n+1}^k = \widehat{\pi}_{s_n^k}$, and if $s_n^k \notin S_0^k$, then s_{n+1}^k is the maximal point in $S(\pi) \cup T(\pi)$ smaller than $s_n^k + \tfrac{1-s_n^k}{k}$.

For every $n \in \mathbb{N}$ define a correlated action profile $y_n^k \in \Delta(A)$ by

$$y_n^k(a) := \begin{cases} \dfrac{\pi_{s_{n+1}^k -}(Q^i, C^{-i}) - \pi_{s_n^k -}(Q^i, C^{-i})}{1 - s_n^k}, & a \in A^*, \\[2ex] 1 - \sum_{a' \neq \vec{C}} y(a') = \dfrac{1 - s_{n+1}^k}{1 - s_n^k}, & a = \vec{C}. \end{cases}$$

We argue that $y(a) \leq \frac{1}{k} y(Q^i, C^{-i})$ for every $i \in I$ and every $a \in A_{\geq 2}^*$ such that $a^i = Q^i$. Indeed, this inequality holds since for every $t \in [s_n^k, s_{n+1}^k)$, if $t \in S(\pi)$ then $\xi_t^i \leq \frac{1}{k}$ for each $i \in I$, while if $t \in T(\pi)$ than $\dot{\pi}_t(a) = 0$ for each such action profile a. We can then apply Lemma 4.9 to y_n^k, and obtain a mixed action profile $\widehat{\xi}_n^k$ that satisfies (i) $p(\widehat{\xi}_n^k) = p(y_n^k) = \frac{s_{n+1}^k - s_k^n}{1 - s_n^k}$ and (ii) Eq. (5).

Define a strategy profile x^k as follows:

(D.1) If $s_n^k \in S_0^k$, set $x_n^k := \xi_{s_n^k}$.

(D.2) If $s_n^k \notin S_0^k$, set $x_n^k := \widehat{\xi}_n^k$.

The convergence $\pi^{x^k} \Rightarrow \pi$ will follow as soon as we show that

$$\|\pi^{x^k}_{s^k_n-} - \pi_{s^k_n-}\|_\infty \le s^k_n \cdot 2^{|I|}/k, \quad \forall k \in \mathbb{N}, \forall n \in \mathbb{N}. \tag{6}$$

For every fixed $k \in \mathbb{N}$ we prove this inequality by induction over n. For $n = 1$ Eq. (6) trivially holds, because $s^k_1 = 0$, hence both sides of Eq. (6) vanish.

We shall suppose that the relation is true for some $n \in \mathbb{N}$ and prove that it holds for $n+1$. (D.1) and (i) ensure that $\widehat{\pi}^{x^k}_{s^k_{n+1}-} - \widehat{\pi}^{x^k}_{s^k_n-} = \widehat{\pi}_{s^k_{n+1}-} - \widehat{\pi}_{s^k_n-}$: for every $n \in \mathbb{N}$, the probability of absorption at stage n under the AP π^{x^k}, is the same as under the original AP π in $[s^k_n, s^k_{n+1})$. This implies that $\widehat{\pi}^{x^k}_{s^k_n-} = \widehat{\pi}_{s^k_n-}$ for every $n \in \mathbb{N}$.

If $s^k_n \in S^k_0$, then (D.1) implies that $s^k_{n+1} = \widehat{\pi}_{s^n_k-}$ and $\pi^{x^k}_{s^k_{n+1}-}(a) - \pi^{x^k}_{s^k_n-}(a) = \pi_{s^k_{n+1}-}(a) - \pi_{s^k_n-}(a)$ for every $a \in A^*$, and therefore Eq. (6) holds for $n+1$.

Suppose now that $s^k_n \notin S^k_0$. By (i) and (ii),

$$
\begin{aligned}
|\pi^{x^k}_{s^k_{n+1}-}(a) - \pi_{s^k_{n+1}-}(a)| & \le s^k_n \cdot 2^{|I|}/k + (1 - s^k_n) \cdot \frac{s^k_{n+1} - s^k_n}{1 - s^k_n} \cdot 2^{|I|}/k \\
& = s^k_{n+1} \cdot 2^{|I|}/k,
\end{aligned}
$$

as desired. ∎

Remark 4.10 *The behavior "Player 1 quits with probability 1, and all other players continue throughout the game" may be translated in many ways to AP's. Here are some examples:*

- *Player 1 quits with probability 1 in the first stage of the game. In this case, we have $T(\pi) = \{1\}$ and $S(\pi) = \{0\}$ (Figure 4(a)).*

- *Player 1 quits with probability $\frac{1}{2}$ in each stage. In this case, we have $T(\pi) = \{1\}$ and $S = \{0, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}, \cdots\}$ (Figure 4(b)).*

- *Player 1 "quits continuously". Here $S(\pi) = \emptyset$, $T(\pi) = [0,1]$, and $\pi_t(Q^1, C^{-1}) = t$, for every $t \in [0,1]$ (Figure 4(c)).*

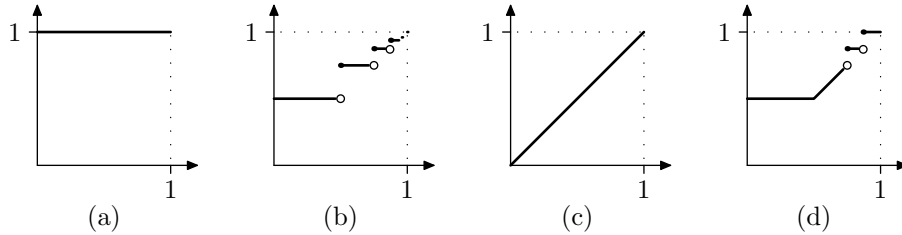- *And we may have combinations of the above (Figure 4(d)).*

Figure 4: Four possibilities for the function $\pi_t(Q^1, C^{-1})$ in Remark 4.10.

17

Proposition 4.11 *The set $\mathbb{A}$ of AP's is sequentially compact: for every sequence $(\pi^k)_{k\in\mathbb{N}}$ of AP's there exists $\pi \in \mathbb{A}$ and a subsequence, still denoted $(\pi^k)_{k\in\mathbb{N}}$, which converges weakly to π. Moreover, this subsequence can be chosen in such a way that for every $t \in S(\pi)$, there are two sequences $(t_k)_{k\in\mathbb{N}} \subset [0,1]$ and $(\xi^k)_{k\in\mathbb{N}} \subset [0,1]^I$ such that for every $k \in \mathbb{N}$ we have $t_k \in S(\pi^k)$ and Eq. (4) holds for π^k and ξ^k at t_k, and such that $\lim_{k\to\infty} t_k = t$, $\lim_{k\to\infty} \pi^k_{t_k} = \pi_t$, and $\lim_{k\to\infty} \xi^k = \xi_t$. Furthermore the limit ξ_t satisfies Eq. (4) for π.*

As described in Section 4.1, even when all AP's $(\pi^k)_{k\in\mathbb{N}}$ are defined by strategy profiles, the limit AP need not be defined by a strategy profile. Furthermore, in this case, the sequence of AP's that was constructed in the proof of Proposition 4.8 and converges to π does not need to coincide with the original sequence $(\pi^k)_{k\in\mathbb{N}}$.

Proof. Let $(\pi^k)_{k\in\mathbb{N}}$ be a sequence of AP's. Since $\mathbf{F}$ is sequentially compact, there exists a subsequence, still denoted $(\pi^k)_{k\in\mathbb{N}}$, and $\pi \in \mathbf{F}$, such that $\pi^k \Rightarrow \pi$. We have to show that $\pi \in \mathbb{A}$.

Since $\pi^k_t \to \pi_t$ for a.e. $t \in [0,1]$, it follows that $\widehat{\pi}^k_t \to \widehat{\pi}_t$ for a.e. $t \in [0,1]$. Together with the weak monotonicity of $t \mapsto \widehat{\pi}$ this implies that $\widehat{\pi}_t \geq t$ for all $t \in [0,1]$.

To show that (A.2) holds for π, let U be a connected component of $[0,1] \setminus (T(\pi) \cup S(\pi))$. Fix $t \in U$. Since π is continuous at t, we have $\pi_t = \lim_{k\to\infty} \pi^k_t$. Since $\widehat{\pi}_t > t$, for every $\varepsilon \in (0, \widehat{\pi}_t - t)$, there exists $k_0 \in \mathbb{N}$ such that for every $k \geq k_0$ we have $\widehat{\pi}^k_t > \widehat{\pi}_t - \varepsilon > t$. Since π^k belong to $\mathbb{A}$, it is constant on $[t, \widehat{\pi}_t - \varepsilon)$. It follows that π is also constant on $[t, \widehat{\pi}_t - \varepsilon)$. Since this is true for every $\varepsilon > 0$ sufficiently small, π is constant on $[t, \widehat{\pi}_t)$, and is equal to π_t.

We turn to prove that (A.3) holds for π. Fix $t \in S(\pi)$. There exists a subsequence of $(\pi^k)_{k\in\mathbb{N}}$, still denoted $(\pi^k)_{k\in\mathbb{N}}$, and a sequence $(s_k)_{k\in\mathbb{N}} \subset [0,1]$ such that $\lim_{k\to\infty} s_k = t$ and $\lim_{k\to\infty} \pi^k_{s_k} = \pi_t$. For each k, set $t_k := \min\{s \leq s_k, \pi^k_s = \pi^k_{s_k}\}$, where the infimum is attained because of the right continuity of π^k. Since $t \in S(\pi)$ we have $\widehat{\pi}_t > t$, hence $\widehat{\pi}^k_{t_k} > t_k$ for every k sufficiently large. By the definition of t_k and (A.2), it follows that $t_k \in S(\pi^k)$.

We argue that $\lim_{k\to\infty} t_k = t$. Let $\widetilde{t}$ be an accumulation point of $(t_k)_{k\in\mathbb{N}}$. Since $t_k \leq s_k \to t$, we have $\widetilde{t} \leq t$. If $\widetilde{t} < t$, consider $s \in [\widetilde{t}, t)$ such that $\pi^k_s \to \pi_s$. Then, for every $\varepsilon > 0$ and every k large enough, we have

$$\widehat{\pi}_t - \varepsilon \leq \widehat{\pi}^k_{s_k} = \widehat{\pi}^k_{t_k} \leq \widehat{\pi}_s + \varepsilon \leq \widehat{\pi}_{t-} + \varepsilon,$$

which is impossible for $\varepsilon < (\widehat{\pi}_t - \widehat{\pi}_{t-})/2$.

Since $\lim_{k\to\infty} t_k = t$, every accumulation point of $(\pi^k_{t_k-})_{k\in\mathbb{N}}$ belongs to the set $\{\pi_{t-}, \pi_t\}$, and, since $\lim_{k\to\infty} \widehat{\pi}^k_{t_k-} = \lim_{k\to\infty} s_k = t < \widehat{\pi}_t$ (originally, it is (s_k) which

18

converges to t), it follows that $\lim_{k\to\infty} \pi^k_{t_k-} = \pi_{t-}$, which implies that $\lim_{k\to\infty} \Delta \pi^k_{t_k} = \Delta \pi_t$.

For each $k \in \mathbb{N}$, since $t_k \in S(\pi^k)$, there exists $\xi^k \in [0,1]^I$ such that

$$\Delta \pi^k_{t_k}(a) = (1 - t_k) \left(\prod_{\{i:\, a^i = Q^i\}} \xi^{k,i} \right) \left(\prod_{\{i:\, a^i = C^i\}} (1 - \xi^{k,i}) \right), \quad a \in A^*. \tag{7}$$

We can find a subsequence of $(t_k)_{k\in\mathbb{N}}$ and $\xi \in [0,1]^I$, such that $\lim_{k\to\infty} \xi^{k,i} = \xi^i$ for all $i \in I$. Taking the limit as $k \to \infty$ in Eq. (7) we get

$$\Delta \pi_t(a) = (1 - t) \left(\prod_{\{i:\, a^i = Q^i\}} \xi^i \right) \left(\prod_{\{i:\, a^i = C^i\}} (1 - \xi^i) \right), \quad a \in A^*.$$

This proves that (A.3) holds, as well as the existence of the sequences $(t_k)_{k\in\mathbb{N}}$ and $(\xi^k)_{k\in\mathbb{N}}$ for every $t \in S(\pi)$ as described in the statement of the proposition.

We finally prove that (A.4) holds as well. Fix $t \in T(\pi) \setminus \{1\}$, so that $\widehat{\pi}_t = t$. We have to show that $\dot{\pi}_t(a) = 0$ for every $a \in A^*_{\geq 2}$. Since $t \in T(\pi)$, there is a nonincreasing sequence $(t_k)_{k\in\mathbb{N}}$ that converges to t such that $\widehat{\pi}_{t_k-} = t_k$ for every k. For the same reason, for every $\varepsilon > 0$ there is $k_0 \in \mathbb{N}$ and $\delta > 0$ such that for every $k \geq k_0$ and every $t' \in [t_k, t_k + \delta) \cap S(\pi^k)$ we have $p(\xi^k_{t'}) < \varepsilon$. Indeed, otherwise there is $\varepsilon > 0$ such that for every $k_0 \in \mathbb{N}$ and every $\delta > 0$ there is $k \geq k_0$ and $t' \in [t_k, t_k + \delta) \cap S(\pi^k)$ for which $p(\xi^k_{t'}) \geq \varepsilon$. But then, letting k_0 go to infinity and δ go to 0, we deduce that $t \in S(\pi)$ and $p(\xi_t) \geq \varepsilon$, a contradiction.

For every mixed action profile ξ that satisfies $p(\xi) < \varepsilon$, we have $\xi^i < \varepsilon$ for every i, and therefore

$$\xi(a) = \left(\prod_{\{i:\, a^i = Q^i\}} \xi^i \right) \cdot \left(\prod_{\{i:\, a^i = C^i\}} (1 - \xi^i) \right) \leq \frac{\varepsilon}{1 - \varepsilon} p(\xi), \quad \forall a \in A^*_{\geq 2}.$$

We deduce that for every $\varepsilon > 0$ there is $k_0 \in \mathbb{N}$ and $\delta > 0$ such that for every $k \geq k_0$ and every $t' \in (t_k, t_k + \delta) \cap S(\pi^k)$, we have $\xi_{t'}(a) \leq \frac{\varepsilon}{1-\varepsilon} p(\xi_{t'})$ for every $a \in A^*_{\geq 2}$. This implies that for every $t' \in (t_k, t_k + \delta) \cap (T(\pi^k) \cup S(\pi^k))$,

$$\pi^k_{t'-}(a) - \pi^k_{t_k-}(a) \leq (t' - t_k) \frac{\varepsilon}{1 - \varepsilon}, \quad \forall a \in A^*_{\geq 2}, \forall t' \in (t_k, t_k + \delta) \cap (T(\pi^k) \cup S(\pi^k)).$$

Since this inequality holds for every $\varepsilon > 0$, we deduce that $\dot{\pi}_t(a) = 0$ for every $a \in A^*_{\geq 2}$. $\blacksquare$

19

4.3   The Payoff Path

Let π be an AP. For every $0 \le t < 1$ and every $a \in A^*$, the difference $\pi_1(a) - \pi_t(a)$ is the probability of absorption by the action profile a in the interval $(t, 1]$. Since the total probability of absorption in $[t, 1]$ is $1 - \widehat{\pi}_t$, the expected payoff after absorption probability t is given by

$$\gamma_t(\pi) := \begin{cases} \frac{\sum_{a \in A^*} (\pi_1(a) - \pi_t(a)) r(a)}{1 - \widehat{\pi}_t}, & \text{if } \widehat{\pi}_t < 1, \\ \vec{0}, & \text{if } \widehat{\pi}_t = 1. \end{cases} \qquad (8)$$

We call the function $\gamma(\pi) : [0, 1] \to \mathbb{R}^I$ the *payoff path*.

Remarks 4.12   *1. Payoff paths take their values in $[-M, M]^I$, where $M = \|r(a)\|_\infty$.*

2. The quantity $\gamma_{0-}(\pi) = \sum_{a \in \mathbb{A}^*} \pi_1(a) r(a)$ is the expected payoff under π in the game. The definition of $\gamma_t(\pi)$ when $\widehat{\pi}_t = 1$ is irrelevant, because in this case the game is already over at t.

3. For every absorbing strategy profile x, we have

$$\gamma_{t_n-}(\pi^x) = \gamma_n(x), \quad \forall n \in \mathbb{N},$$

where π^x is the AP defined by x and $t_n = \mathbf{P}_x(\theta < n)$. This equality reflects the equivalence between each strategy profile x and the AP π^x.

4. When $T(\pi) = [0, 1]$, the expression for the payoff path simplifies to

$$\gamma_t(\pi) = \frac{\sum_{i \in I} \left(\pi_1(Q^i, C^{-i}) - \pi_t(Q^i, C^{-i}) \right) r(Q^i, C^{-i})}{1 - t}, \quad \forall t \in [0, 1). \qquad (9)$$

We then have for every $0 \le s < t < 1$,

$$(1 - t)\gamma_t = (1 - s)\gamma_s + \sum_{i \in I} \left(\pi_s(Q^i, C^{-i}) - \pi_t(Q^i, C^{-i}) \right) r(Q^i, C^{-i}).$$

Hence, the function $t \mapsto \gamma_t$ is a solution of the differential equation

$$(1 - t)\dot{\gamma}_t = \gamma_t - \sum_{i \in I} \dot{\pi}_t(Q^i, C^{-i}) r(Q^i, C^{-i}), \ t \in [0, 1). \qquad (10)$$

5. Let $(\pi^k)_{k \in \mathbb{N}}$ be a sequence of AP's that converges to a limit π. Then,

$$\gamma_t(\pi) = \lim_{k \to \infty} \gamma_t(\pi^k),$$

for every $t \in [0, 1)$ where π is continuous.

We now adapt the definition of sequential ε-perfectness to AP's.

Definition 4.13 *Let $\varepsilon \geq 0$. Player i is* sequentially ε-perfect *at the AP π if the following conditions hold:*

(SP.1) *For every $t \in S(\pi)$ such that $\widehat{\pi}_t < 1$, player i is ε-perfect at the mixed action profile ξ_t in the strategic-form game $G_\Gamma(\gamma_t(\pi))$, where ξ_t satisfies Eq. (4).*

(SP.2) *For every $t \in T(\pi) \setminus \{1\}$,*

> (a) $\gamma_t^i(\pi) \geq r^i(Q^i, C^{-i}) - \varepsilon$, *and*
> (b) *if $\dot{\pi}_t(Q^i, C^{-i}) > 0$, then $\gamma_t^i(\pi) \leq r^i(Q^i, C^{-i}) + \varepsilon$.*

An AP π is sequentially ε-perfect, *denoted ε-AP, if all players are sequentially ε-perfect at π.*

In words, an AP is sequentially 0-perfect (resp. sequentially ε-perfect) if (i) whenever the players play in discrete time ($t \in S(\pi)$), the mixed action that they play is a Nash equilibrium (resp. ε-perfect) in the one-shot game induced by the continuation payoff, and (ii) whenever the players play in continuous time ($t \in T(\pi)$), every player who quits with a positive rate is indifferent (resp. indifferent up to ε) between continuing and quitting, and no player who quits with rate 0 can profit (resp. can profit more than ε) by quitting.

It follows from Definition 4.1 that player i is sequentially ε-perfect at an absorbing strategy profile x if and only if she is sequentially ε-perfect at the AP π^x.

We shall see now that standard continuity arguments imply that a limit of sequentially ε-perfect AP's as ε goes to 0 is a 0-AP.

Proposition 4.14 *Let $(\pi^k)_{k \in \mathbb{N}}$ be a sequence of AP's that converges to a limit π, let $(\varepsilon^k)_{k \in \mathbb{N}}$ be a sequence of non-negative reals that converges to 0, and let $i \in I$. If for every $k \in \mathbb{N}$ player i is sequentially ε^k-perfect at the AP π^k, then player i is sequentially 0-perfect at the AP π.*

Proof. Fix $t \in S(\pi)$. We prove that in this case (SP.1) holds with $\varepsilon = 0$. Since $\pi^k \Rightarrow \pi$, following Proposition 4.11 we can find a sequence $(t_k)_{k \in \mathbb{N}}$, with $t_k \in S(\pi^k)$ for all $k \in \mathbb{N}$, that converges to t, $\lim_{k \to \infty} \pi^k_{t_k} = \pi_t$ and $\xi_t = \lim_{k \to \infty} \xi^k$, where ξ^k satisfies Eq. (4) at t_k for π^k, for all $k \in \mathbb{N}$. Since $\pi^k_{t_k}$ converges to π_t, we have also $\gamma_{t_k}(\pi^k) \to \gamma_t(\pi)$.

By definition, if player i is sequentially ε^k-perfect at π^k, then she is ε^k-perfect at the mixed action profile ξ^k in the strategic-form game $G_\Gamma(\gamma_{t_k}(\pi^k))$. As discussed

in Section 3.1, it follows that player i is 0-perfect at ξ_t in the strategic-form game $G_\Gamma(\gamma_t(\pi))$, i.e., (SP.1) holds with $\varepsilon = 0$.

Now let $t \in T(\pi) \setminus \{1\}$. We will prove that (SP.2.a) holds with $\varepsilon = 0$. Let $(t_k)_{k \in \mathbb{N}}$ be a nonincreasing sequence of times converging to t, such that $\pi^k_{t_{k-}} \to \pi_t$. Since $t \in T(\pi)$, we can choose a subsequence of $(t_k)_{k \in \mathbb{N}}$, still denoted $(t_k)_{k \in \mathbb{N}}$, such that $\widehat{\pi}_{t_{k-}} = t_k$ for every k. This implies that $\gamma_t(\pi) = \lim_{k \to \infty} \gamma_{t_k}(\pi^k)$.
Following Remark 4.4(7), this implies that, for each $k \in \mathbb{N}$ there are only two possibilities: either $t_k \in T(\pi^k)$ or $t_k \in S(\pi^k)$.

Suppose first that $t_k \in T(\pi^k)$ for every $k \in \mathbb{N}$ large enough. Then (SP.2.a), applied to π^k, yields
$$\gamma^i_{t_k}(\pi^k) \geq r^i(Q^i, C^{-i}) - \varepsilon^k,$$
and, letting k go to $+\infty$, we obtain that (SP.2.a) with $\varepsilon = 0$ holds for π at t.

Next let us suppose the existence of a subsequence of $(\pi^k)_{k \in \mathbb{N}}$ such that $t_k \in S(\pi^k)$ for every $k \in \mathbb{N}$. By assumption we have

$$r^i(Q^i, \xi^{k,-i}) \leq \gamma^i_{t_k}(\pi^k) + \varepsilon^k. \tag{11}$$

Since $t \in T(\pi) \setminus \{1\}$, necessarily $\lim_{k \to \infty} p(\xi^k_{t_k}) = 0$.

The result follows by letting $k \to \infty$ in Eq. (11).

The proof that (SP.2.b) holds with $\varepsilon = 0$ is similar, hence (SP.2) holds for every $t \in T(\pi)$ such that $\pi^k_t \to \pi_t$. For t such that π^k_t does not converge to π_t, (SP.2) holds by the right-continuity of π. ■

The following result relates the concepts of ε-equilibria in discrete-time games and 0-AP's.

Theorem 4.15 *Let Γ be a quitting game that for every $\varepsilon > 0$ sufficiently small possesses neither an ε-equilibrium under which the game terminates with probability 1 in the first stage, nor an ε-equilibrium where all players always continue. Then Γ admits an ε-equilibrium for every $\varepsilon > 0$, if and only if there is a 0-AP.*

Theorem 4.15 highlights the significance of AP's in simplifying the study of ε-equilibria in quitting games. The set of strategy profiles is compact in the product topology, yet in this topology the payoff function is not continuous and the set of ε-equilibria is not compact. Moreover, since players may not be indifferent among their actions along the equilibrium, it may be difficult to identify and to characterize ε-equilibria. On the other hand, in the weak topology the sets of AP's and 0-AP's are compact, the payoff is continuous, and along a 0-AP for every $t \in [0, 1)$ players are indifferent between actions they play with positive probability or rate. Therefore,

the study of 0-AP's seems to be simpler than that of ε-equilibria, yet, in view of Theorem 4.15, it may suffice to answer various questions on ε-equilibria. We will see such a case in the next section.

Proof. Theorem 3.4, Proposition 4.11, and Proposition 4.14 imply that if the game admits an ε-equilibrium for every $\varepsilon > 0$, then there is a 0-AP. Regarding the converse implication, let π be a 0-AP. By Proposition 4.8, there exists a sequence $(x^k)_{k \in \mathbb{N}}$ of strategy profiles such that $\pi^{x^k} \Rightarrow \pi$. For every $k \in \mathbb{N}$ let $(s_n^k)_{n \in \mathbb{N}}$ be the sequence of real numbers defined in the proof of Proposition 4.8 for x^k. We then have $\lim_{k \to \infty} \sup_{n \in \mathbb{N}} \| \gamma_{s_n^k}(\pi^{x^k}) - \gamma_{s_n^k}(\pi) \|_{\infty} = 0$, which implies that x^k is an ε^k-equilibrium for every $k \in \mathbb{N}$, for some sequence $(\varepsilon^k)_{k \in \mathbb{N}}$ that converges to 0. ∎

Theorem 4.15 is related to Gobbino and Simon (2020), who separated the dynamics of the sequence $(\gamma_n(x))_{n \in \mathbb{N}}$, where x is an absorbing sequentially ε-perfect strategy profile, into "large" motion (the discrete part of the AP) and "small" motion (the continuous part of the AP).

5   Continuous Equilibria

An AP π is *continuous* if it does not contain discrete-time aspects; that is, if $T(\pi) = [0, 1]$. When π is continuous, $\sum_{a \in A_{\geq 2}^*} \pi_1(a) = 0$, yet the converse need not hold. Such equilibria only depend on $(r(Q^i, C^{-i}))_{i \in I}$ and not on the whole payoff function. In this sense, continuous equilibria are (partially) detailed-free, and, perhaps, more robust to misspecification of payoffs. To simplify terminology, we use the term *continuous equilibria* for sequentially 0-perfect continuous AP's.

In this section we provide a sufficient condition for the existence of a continuous equilibrium. This sufficient condition uses the concept of linear complementarity problems, which encompassed linear programming and quadratic programming, see, e.g., Cottle and Dantzig (1968), Balinski (1978), and Murty (1988). To link our sufficient condition to linear complementarity problems, we find it convenient to normalize the payoffs and assume w.l.o.g. that $r^i(Q^i, C^{-i}) = 0$ for each $i \in I$.

Definition 5.1 *Let R be an $(n \times n)$-matrix, and let $q \in \mathbb{R}^n$. For each i, $1 \leq i \leq n$, denote by R^i the i'th column of R. The* linear complementarity problem $\mathrm{LCP}(R, q)$

is the following problem:

$$\text{Find} \quad w \in \mathbb{R}_+^n, \ and \ z = (z_0, z_1, \ldots, z_n) \in \Delta(\{0, 1, \ldots, n\}),$$

$$\text{such that} \quad w = z_0 q + \sum_{i=1}^n z_i R^i, \tag{12}$$

$$z_i = 0 \ or \ w_i = 0, \quad \forall i \in \{1, 2, \ldots, n\}.$$

A matrix R is a Q-*matrix* if for every $q \in \mathbb{R}$ the problem $\mathrm{LCP}(R, q)$ has at least one solution.

Let Γ be a quitting game, and denote by $R(\Gamma)$ the $(|I| \times |I|)$ matrix $(r^i(Q^j, C^{-j}))_{i,j \in I}$. Solan and Solan (2020) proved that if $R(\Gamma)$ is not a Q-matrix, then Γ has a stationary 0-equilibrium. Here we study another family of matrices.

Definition 5.2 *We say that a matrix R is a $\overline{Q}$-matrix if R as well as all its principal minors are Q-matrices.*

Remark 5.3 *A (1×1)-matrix $R = (R_{11})$ is a $\overline{Q}$-matrix if and only if $R_{11} \geq 0$. A (2×2)-matrix $R = (R_{ij})$ is a $\overline{Q}$-matrix if and only if there is a non-negative row whose diagonal entry is 0. A (3×3)-matrix R is a $\overline{Q}$-matrix if and only if, up to a conjugation with a permutation matrix, one of the following conditions holds:*

- *The sign structure of R is $\begin{pmatrix} 0 & ? & ? \\ \geq & 0 & ? \\ \geq & \geq & ? \end{pmatrix}$, where ? means that the sign of the entry is irrelevant.*

- *The sign structure of R is $\begin{pmatrix} 0 & \leq & \geq \\ \geq & 0 & \leq \\ \leq & \geq & 0 \end{pmatrix}$ and the determinant of R is non-negative.*

The following result identifies a new class of quitting games were ε-equilibria exist.

Theorem 5.4 *If $R(\Gamma)$ is a $\overline{Q}$-matrix, then Γ admits a continuous equilibrium.*

Remark 5.5 *1. Theorem 5.4 is not tight: there may be continuous equilibria when its condition is not satisfied. Indeed, it may be that the restriction of $R(\Gamma)$ to a subset J of players satisfies the condition of Theorem 5.4, and therefore there is a continuous equilibrium π for the subgame that involves those players (when*

players not in J are restricted to always continue), and it may further happen that the players not in J obtain non-negative payoffs along this AP. In such a case, all players are sequentially 0-perfect at π. Yet the rows of $R(\Gamma)$ that correspond to players not in J may be arbitrary, hence $R(\Gamma)$ need not be a $\overline{Q}$-matrix.

We do not know whether the existence of a continuous equilibrium along which all players quit with positive probability implies that $R(\Gamma)$ is a $\overline{Q}$-matrix.

2. *We are not aware of a characterization of Q-matrices or of $\overline{Q}$-matrices, yet we can point at a family of matrices that are $\overline{Q}$.*

Recall that a matrix is P_0 if all its principal minors are non-negative. One family of matrices that are included in the set $\overline{Q}$ is the set of all P_0-matrices whose diagonal entries are 0. To see that this inclusion holds, we need to recall the set of P-matrices, which are matrices R for which all principal minors are positive. It is well known that for P-matrices R, the linear complementarity problem $\mathrm{LCP}(R, q)$ has exactly one solution for every $q \in \mathbb{R}^n$, see, e.g., Murty (1988, Chapter 3). The set of P_0-matrices is the closure of the set of P-matrices (since if $\det(R) \geq 0$ then $\det(R + \varepsilon I) > 0$ for every $\varepsilon > 0$, and therefore if R is a P_0-matrix, then $R + \varepsilon I$ is a P-matrix), which implies that every P_0-matrix is a Q-matrix.

We note that there are $\overline{Q}$-matrices that are not in P_0, for example, the following (3×3)-matrix:

$$\begin{pmatrix} 0 & 1 & 1 \\ -1 & 0 & 1 \\ -1 & 1 & 0 \end{pmatrix}.$$

3. *The* standard linear complementarity problem *is the problem (12), where z is not required to be in $\Delta(\{0, 1, \ldots, n\})$, but rather to satisfy $z_0 = 1$ and $z_i \geq 0$ for every $i \in \{1, 2, \ldots, n\}$. A matrix R is a Q-matrix according to Definition 5.1 if (a) for every $q \in \mathbb{R}^n$ the standard linear complementarity problem for R and q has a solution, or (b) there is a convex combination Rz of the columns R such that $(Rz)_i = 0$ whenever $z_i > 0$. A matrix R that satisfies (a) is a Q-matrix w.r.t. the standard linear complementarity problem, and a matrix R all of whose principal minors are Q-matrices w.r.t. the standard linear complementarity problem is a completely-Q matrix w.r.t. the standard linear complementarity problem. Such matrices have been studied by, e.g., Cottle (1980), who proved that the family of completely-Q matrices coincide with the family of strictly semi-monotone matrices.*

Proof of Theorem 5.4.

Step 1: Convex combinations in the non-negative orthant.

We will show here that for every nonempty subset $J \subseteq I$ of players there is a probability distribution $z \in \Delta(J)$ that satisfies

$$\sum_{i \in J} z_i r^j(Q^i, C^{-i}) \geq 0, \quad \forall j \in J, \tag{13}$$

$$\sum_{i \in J} z_i r^j(Q^i, C^{-i}) = 0 \text{ for at least one } j \in J. \tag{14}$$

The assumption that $R = R(\Gamma)$ is a $\overline{Q}$-matrix is used only in this step of the proof.

Fix $i_0 \in J$ and let $\widehat{q} \in \mathbb{R}^J$ be the vector that is defined by

$$\widehat{q}_{i_0} := -1, \quad \widehat{q}_i := 0 \quad \forall i \in J \setminus \{i_0\}.$$

The matrix $\widehat{R} := (r^i(Q^j, C^{-j}))_{i,j \in J}$ is a principal minor of R. Therefore, the linear complementarity problem $\text{LCP}(\widehat{R}, \widehat{q})$ has a solution $(\widehat{w}, \widehat{z})$. Since $\widehat{q}_{i_0} < 0$, it cannot be that $\widehat{z}_0 = 1$. If i_0 is the only player $i \in J$ such that $\widehat{z}_i > 0$, then, since $r^{i_0}(Q^{i_0}, C^{-i_0}) = 0$ and $\widehat{q}_{i_0} < 0$, we have $\widehat{z}_{i_0} = 1$. Otherwise, there is $i_1 \in J \setminus \{i_0\}$ such that $\widehat{z}_{i_1} > 0$, and consequently $\widehat{w}_{i_1} = 0$.

Define $z_i := \frac{\widehat{z}_i}{1 - \widehat{z}_0}$ for each $i \in J$. Since $\widehat{w}_i \geq 0$ and $\widehat{q}_i \leq 0$ for every $i \in J$, and since $\widehat{w}$ is a convex combination of $\widehat{q}$ and $\sum_{i \in J} z_i r(Q^i, C^{-i})$, it follows that Eq. (13) holds. If $z_{i_0} = 1$, then Eq. (14) holds with $j = i_0$. Otherwise, since $\widehat{w}_{i_1} = \widehat{q}_{i_1} = 0$, we have $\sum_{i \in J} z_i r^{i_1}(Q^i, C^{-i}) = 0$, and Eq. (14) holds with $j = i_1$.

Step 2: Viability theory.

For every $z \in \Delta(I)$ denote $z \cdot R := \sum_{i \in I} z_i R^i$, and let Y be the boundary of $\mathbb{R}^I_+$. For every $q \in Y$, set

$$F(q) := \{z \in \Delta(I) : z_i > 0 \Rightarrow q_i = 0, \quad (z \cdot R)_i \geq 0 \text{ whenever } q_i = 0\}.$$

Note that $F(q)$ depends only on the set $\{i \in I : q_i = 0\}$. We will show that there exist measurable functions $z : [t_0, 1] \to \Delta(I)$ and $q : [t_0, 1] \to \mathbb{R}^I$ such that for every $t \in [t_0, 1]$ we have (a) $q(t) \in Y$ and (b) $z(t) \in F(q(t))$.

The set-valued function F is upper semi-continuous with convex values, and by Step 1 it has nonempty values. For every $q \in Y$ denote by $T_Y(q)$ the tangent cone at q:

$$T_Y(q) := \{d \in \mathbb{R}^I : q + \delta d \in Y \text{ for all } \delta > 0 \text{ small}\}.$$

A careful analysis of the tangent cone shows that $\frac{\delta}{t} z \cdot R + (1 - \frac{\delta}{t})q \in Y$ for every z satisfying Eqs. (13)–(14) and $\delta > 0$ small enough, where $J = \{i \in I : q_i = 0\}$.

26

Fix $(q_0, t_0) \in Y \times (0, 1)$. For every measurable function $z : [t_0, 1] \to \Delta(I)$, consider the following controlled dynamic:

$$\begin{cases} \dot{q}(t) = \frac{1}{t}(z(t) \cdot R - q(t)), & \forall t \in [t_0, 1], \\ q(t_0) = q_0. \end{cases} \tag{15}$$

The set Y is closed, and the set-valued function F is upper-semicontinuous with nonempty, closed, and convex values. By the classical Viability Theorem (Aubin, 1991, Theorem 3.3.4) it follows that there exist measurable functions $z : [t_0, 1] \to \Delta(I)$ and $q : [t_0, 1] \to \mathbb{R}^I$ such that (a) and (b) above hold for every $t \in [t_0, 1]$.

Step 3: Constructing a continuous equilibrium.

Fix an arbitrary $q_0 \in Y$. For every $n \in \mathbb{N}$ let (z^n, q^n) be a solution of Eq. (15) with $q_0^n = q_0$ and $t_0 = \frac{1}{n}$, such that $q^n(t) \in Y$ and $z^n(t) \in F(q^n(t))$ for every $t \in [\frac{1}{n}, 1]$. Define $\pi^n \in \mathbb{A}$ by

$$\dot{\pi}_t^n(Q^i, C^{-i}) = z_i^n(1 - t), \quad \forall t \in [0, 1 - \tfrac{1}{n}), \quad \forall i \in I, \tag{16}$$

and an arbitrary continuous evolution on $[1 - \frac{1}{n}, 1]$. By definition, π^n is a continuous AP. Eq. (15) implies that, for all $0 \le t \le 1 - \frac{1}{n}$,

$$(1 - t)q^n(1 - t) - \frac{1}{n}q_0 = \int_{\frac{1}{n}}^{1-t} z^n(s)\mathrm{d}s \cdot R = \int_t^{1-\frac{1}{n}} z^n(1 - s)\mathrm{d}s \cdot R.$$

In addition, for every $t \in [0, 1 - \frac{1}{n}]$,

$$\begin{aligned} \gamma_t(\pi^n) &= \frac{1}{1 - t} \int_t^1 z^n(1 - s)\mathrm{d}s \cdot R \\ &= \frac{1}{1 - t} \int_t^{1-1/n} z^n(1 - s)\mathrm{d}s \cdot R + \frac{1}{1 - t} \int_{1-1/n}^1 z^n(1 - s)\mathrm{d}s \cdot R \\ &= q^n(1 - t) - \frac{q_0}{(1 - t)n} + \frac{1}{1 - t} \int_{1-1/n}^1 z^n(1 - s)\mathrm{d}s \cdot R. \end{aligned}$$

It follows that

$$\|\gamma_t(\pi^n) - q^n(1 - t)\|_\infty \le \frac{2\|R\|_\infty}{(1 - t)n}, \quad \forall n \in \mathbb{N}, \forall t \in [0, 1 - \tfrac{1}{n}].$$

Let π be an accumulation point of (π^n), and assume w.l.o.g. that $\pi^n \Rightarrow \pi$. We will prove that π is a continuous equilibrium. Since π^n is continuous, so is π. Consequently,

27

for every $t \in [0,1)$ the limit $\lim_{n\to\infty} q^n(1-t)$ exists and is equal to $\gamma_t(\pi)$. Since $q^n(1-t) \in Y$ for every $t \in [0, \frac{1}{n}]$, we deduce that $\gamma_t(\pi) \in Y$ for every $t \in [0,1)$, and therefore (SP.2.a) with $\varepsilon = 0$ holds for each $i \in I$.

We turn to prove that (SP.2.b) holds as well. Fix $i \in I$ and let $t \in [0,1)$ be such that $\dot{\pi}_t(Q^i, C^{-i}) > 0$. Then there exists a sequence $(t_n)_{n\in\mathbb{N}}$ such that $\lim_{n\to\infty} t_n = t$ and $\dot{\pi}^n_{t_n}(Q^i, C^{-i}) > 0$ for every n sufficiently large. This implies that for every n sufficiently large we have $z_i^n(1-t_n) > 0$, and therefore $q_i^n(1-t_n) = 0$. By taking the limit as n goes to infinity we deduce that $\gamma_t^i(\pi) = 0$, and (SP.2.b) indeed holds.

Since Condition (SP.2) of Definition 4.13 holds for π, and since i is arbitrary, π is sequentially 0-perfect. ∎

When π is a continuous equilibrium, we can assign to each $t \in [0,1)$ the set of players who quit with positive rate at t. In the next two examples, $[0,1)$ is divided into countably many intervals, and a single player quits with positive rate in each interval. We therefore describe π by a totally ordered index set $\mathcal{K}$, where each $k \in \mathcal{K}$ corresponds to an interval, such that $k < k'$ if and only if the interval that corresponds to k precedes the interval that corresponds to k', and a list of pairs $(i_k, p_k)_{k\in\mathcal{K}}$, where for each $k \in \mathcal{K}$, i_k is the player who quits with positive rate along the interval that corresponds to k, and $p_k \in (0,1]$ is the probability by which player i_k quits, given that the game did not terminate before. For instance, when $\mathcal{K}$ is well-ordered, as happens in Examples 5.6 and 5.7 below, under π player i_0 quits in the interval $[0, p_0)$, player i_1 quits in the interval $[p_0, p_0 + (1 - p_0)p_1)$, and so on. In this case, it is w.l.o.g. to assume that these intervals are maximal in the sense that no player quits in two consecutive intervals, that is, $i_k \neq i_{k+1}$ for each $k \in \mathcal{K}$. Example 5.8 below illustrates a case where $\mathcal{K}$ is not well-ordered. Finally, we note that since the play eventually absorbs, $\sum_{k\in\mathcal{K}} p_k = \infty$.

Example 5.6 *As in the example in Section 4.1, suppose that there are three players, and*

$$R(\Gamma) = \begin{pmatrix} 0 & -1 & 2 \\ 2 & 0 & -1 \\ -1 & 2 & 0 \end{pmatrix}.$$

The matrix R is a $\overline{Q}$-matrix, hence a continuous equilibrium exists. One such equilibrium is the one were the sequence $(i_k, p_k)_k$ is:

$$(1, \frac{1}{2}), (2, \frac{1}{2}), (3, \frac{1}{2}), (1, \frac{1}{2}), (2, \frac{1}{2}), (3, \frac{1}{2}), (1, \frac{1}{2}), (2, \frac{1}{2}), (3, \frac{1}{2}), \dots . \tag{17}$$

In fact, Flesch, Thuijsman, and Vrieze (1997) showed that all continuous equilibria in this example can be obtained from the one in Eq. (17) by starting the period at any $t \in [0, \frac{7}{8}]$ (instead of at $t = 0$).

The following example shows that continuous equilibria, even when periodic, may exhibit a wild behavior.

Example 5.7 *Suppose that there are five players and*

$$
R(\Gamma) = \begin{pmatrix}
0 & -\frac{1}{2} & 2 & -1 & 2 \\
2 & 0 & -\frac{1}{2} & -2 & \frac{7}{2} \\
-\frac{1}{2} & 2 & 0 & -3 & \frac{47}{8} \\
1 & 1 & 1 & 0 & \frac{5}{2} \\
-1 & -1 & -1 & \frac{10}{7} & 0
\end{pmatrix}.
$$

It is a bit tedious but not difficult to show that the corresponding matrix $R(\Gamma)$ is a $\overline{Q}$-matrix, and therefore a continuous equilibrium exists.[3]

In this example there are many periodic continuous equilibria $(i_k, p_k)_{k \in \mathbb{N}}$. In fact, for every $l \in \mathbb{N}$ there is such an equilibrium with period $3l + 2$, where the sequence $(i_k)_{k \in \mathbb{N}}$ is an infinite repetition of $(\underbrace{1, 2, 3, 1, 2, 3, \ldots, 1, 2, 3}_{l \text{ cycles of length } 3}, 4, 5)$. There is also a con-tinuous equilibrium that has this structure for $l = \infty$; that is, the index set is $\mathcal{K} = \mathbb{N}^2$ (with the lexicographic order) and $(i_k, p_k)_{k \in \mathcal{K}}$ is an infinite repetition of

$$
\underbrace{(1, \frac{1}{4}), (2, \frac{1}{6}), (3, \frac{1}{20}), (1, \frac{1}{76}), (2, \frac{1}{300}), (3, \frac{1}{598}), \ldots,}_{\text{countably many cycles of length } 3} (4, \frac{1}{2}), (5, \frac{1}{2}),
$$

where the total probability of absorption in each repetition is strictly less than 1. In particular, the sequence $(i_k, p_k)_{k \in \mathcal{K}}$ is well-ordered but not order-equivalent to the set $\mathbb{N}$. We do not know whether there exist games where there is a continuous equilibrium but none that is periodic with a finite period.

Ashkenazi-Golan, Krasikov, Rainer, and Solan (2021) provide an algorithm for calculating the union of the range of all payoff paths that correspond to continuous equilibria with a well-ordered index set $\mathcal{K}$, where a single player quits in each interval, as well as a characterization of the set of such continuous equilibrium payoffs as a limsup of a certain sequence of sets. The idea of the algorithm is as follows. If along an

---

[3]The reader may note that in this example, the (5×5)-matrix $R(\Gamma)$ is not a P_0-matrix.

AP π players quit only in continuous time, then the AP is uniquely defined by the list of pairs $(i_k, p_k)_{k \in \mathcal{K}}$. When $\mathcal{K}$ is well-ordered, we call player i_{k+1} the *successor* of player i_k, for each $k \in \mathcal{K}$. For a fixed $i \in I$, let $\mathcal{E}_i^*$ be the set of such continuous equilibrium payoffs that can be attained when the play starts with player i. The collection of sets $(\mathcal{E}_i^*)_{i \in I}$ satisfies the following recursion: $w \in \mathcal{E}_i^*$ if and only if $w \in \mathbb{R}_+^I$ and there exists $p \in (0,1]$, $j \in I \setminus \{i\}$, and $v \in \mathcal{E}_j^*$ such that

$$w = pR^i(\Gamma) + (1-p)v, w_i = 0, \tag{18}$$

where $w \in \mathbb{R}_+^I$ captures (SP.2.a) and $w_i = 0$ corresponds to (SP.2.b). In Eq. (18), p corresponds to the length of the interval spent by player i quitting and, if $v \in \mathcal{E}_j^*$, then j is the successor of i. When the payoffs are generic, if j is a successor of i then

$$R_{ji} < 0 < R_{ij}. \tag{19}$$

Define a directed graph as follows: the set of vertices is I and there is a directed edge between i and j if they satisfy Eq. (19). Under some weak assumptions we show that the set of continuous equilibrium payoffs $\bigcup_{i \in I} \mathcal{E}_i^*$, is a fixed point of an operator that follows the geometry of this directed graph. For instance, in Example 5.6, as shown by Flesch, Thuijsman, and Vrieze (1997),

$$\mathcal{E}_i^* = \{w \in \mathbb{R}_+^3 : w_i = 0, w_{i+1} \neq 0, w_{i+1} + w_{i+2} = 1\},$$

where the addition is modulo 3, so that the set of continuous equilibrium payoffs with a well-ordered index set coincides with the boundary of the two-dimensional simplex.

We end this section by presenting a 0-AP for which the corresponding index set is not well-ordered.

Example 5.8 *Suppose that there are five players, and*

$$R(\Gamma) = \begin{pmatrix} 0 & 1 & -3 & -1 & 2 \\ -3 & 0 & 1 & 2 & 0 \\ 1 & -3 & 0 & -1 & 2 \\ 3 & 3 & 3 & 0 & -1 \\ -3 & -3 & -3 & 2 & 0 \end{pmatrix}.$$

The game admits a periodic continuous equilibrium $(i_k, p_k)_{k \in \mathcal{K}}$ that is defined over the index set $\mathcal{K} = \mathbb{N} \times (-\mathbb{N})$ (with the lexicographic order), and is given by an infinite

30

repetition of

$$\underbrace{\ldots, (3, \frac{2}{3^7 - 1}), (2, \frac{2}{3^6 - 1}), (1, \frac{2}{3^5 - 1}), (3, \frac{2}{3^4 - 1}), (2, \frac{2}{3^3 - 1}), (1, \frac{2}{3^2 - 1})}_{\textit{countably many cycles of length 3}}, (4, \frac{1}{2}), (5, \frac{1}{2}).$$

Under this equilibrium, denoted π, the total probability of absorption in each period is $\frac{5}{6}$, and $\gamma_0(\pi) = \gamma_{5/6}(\pi) = (0,0,0,1,0)$, $\gamma_{1/3}(\pi) = (0,1,0,0,1)$, and $\gamma_{1/2}(\pi) = (1,0,1,0,0)$.

6 Discussion

The behavior of players in dynamic games in general, and quitting games in particular, may be complex. It might be that in some stage, the players mix their actions, knowing that the set of players who will terminate the game will be random. It might also happen that some player wants to quit, but she wants to guarantee that no other player knows when she quits, to avoid the outcome where she quits with someone else. While in discrete time a player cannot guarantee that no other player will be able to quit with her, in continuous time this can be done. Equilibrium behavior in quitting games may exhibit both types of behavior: periods of discrete-time behavior, when players quit with positive probability, and periods of continuous-time behavior, when players quit at a given rate.

The concepts of discrete-time strategies and continuous-time strategies can capture only one of the two possible behaviors described above. The concept of AP allows to describe both behaviors. Though it is not known whether all quitting games have ε-equilibria, we showed that *if* an ε-equilibrium exists for every $\varepsilon > 0$, *then* there exists a 0-AP. This result shows that the reason for having games that possess ε-equilibria for every $\varepsilon > 0$ but no 0-equilibria, is that the nature of discrete time does not allow players to completely hide the stage in which they quit, thereby allowing other players to quit simultaneously with them (albeit with small probability) and make a low profit.

The space of AP's $\mathbb{A}$ is sequentially compact, and the function that assigns to every AP its payoff path is continuous. It is not difficult to show that $\mathbb{A}$ is contractible. We do not know whether these properties can be used to prove the existence of an ε-equilibrium in some family of quitting games.

31

References

[1] Ashkenazi-Golan G., Krasikov I. Rainer R., and Solan E. (2020) Algorithms for Continuous Equilibria in Quitting Games. In preparation.

[2] Aubin J.-P. (1991) *Viability Theory*, Birkhäuser.

[3] Balinski M.L. and Cottle R.W. (Eds.) (1978) *Complementarity and Fixed Point Problems*. Amsterdam: North-Holland.

[4] Cottle R.W. (1980) Compleyetly-Q Matrices, *Mathematical Programming*, **19**, 347–351.

[5] Cottle R.W. and Dantzig G.B. (1968) Complementary Pivot Theory of Mathematical Programming, *Linear Algebra and its Applications*, **1**, 103–125.

[6] Flesch J., Schoenmakers G., and Vrieze K. (2008) Stochastic Games on a Product State Space, *Mathematics of Operations Research*, **33**, 403–420.

[7] Flesch J., Schoenmakers G., and Vrieze K. (2009) Stochastic Games on a Product State Space: the Periodic Case, *International Journal of Game Theory*, **38**, 263–289.

[8] Flesch J., Thuijsman F. and Vrieze K. (1997) Cyclic Markov Equilibria in Stochastic Games, *International Journal of Game Theory*, **26**, 303–314

[9] Gobbino M. and Simon R.S. (2020) Discrete Viability Theory. In preparation.

[10] Mertens J.F. and Neyman A. (1981) Stochastic Games, *International Journal of Game Theory*, **10**, 53–66.

[11] Murty K.G. (1988) *Linear Complementarity, Linear and Non-linear Programming*, Heldermann Verlag.

[12] Simon R.S. (2007) The Structure of Non-Zero-Sum Stochastic Games, *Advances in Applied Mathematics*, **38**, 1–26.

[13] Simon R.S. (2012) A Topological Approach to Quitting Games, *Mathematics of Operations Research*, **37**, 180–195.

[14] Solan E. (1999) Three-Player Absorbing Games, *Mathematics of Operations Research*, **24**, 669–698.

[15] Solan E. (2003) The Dynamics of the Nash Correspondence and n-Player Stochastic Games, *International Game Theory Review*, **3**, 291–300.

[16] Solan E. and Solan O.N. (2020) Quitting Games and Linear Complementarity Problems, *Mathematics of Operations Research*, **45**, 434–454.

[17] Solan E., Solan O.N., and Solan R. (2020) Jointly Controlled Lotteries with Biased Coins, *Games and Economic Behavior*, **119**, 383–391.

[18] Solan E. and Vieille N. (2001) Quitting Games, *Mathematics of Operations Research*, **26**, 265–285.

[19] Solan E. and Vieille N. (2002) Quitting Games - an Example, *International Journal of Game Theory*, **31**, 365–381.

[20] Sorin S. and Vigeral G. (2020) Limit Optimal Trajectories in Zero-Sum Stochastic Games, *Dynamic Games and Applications*, **10**, 555–572.

[21] Vieille N. (1992). Weak Approachability, *Mathematics of Operations Research*, **17**, 781–791.

[22] Vieille N. (2000a) Equilibrium in 2-Person Stochastic Games I: A Reduction, *Israel Journal of Mathematics*, **119**, 55–91.

[23] Vieille N. (2000b) Equilibrium in 2-Person Stochastic Games II: The Case of Recursive Games, *Israel Journal of Mathematics*, **119**, 93–126.