Center for
Mathematical Economics
Working Papers

525

September 2014

On Repeated games with imperfect public monitoring: From discrete to continuous time

Mathias Staudigl and Jan-Henrik Steg

Center for Mathematical Economics (IMW)
Bielefeld University
Universitätsstraße 25
D-33615 Bielefeld · Germany

e-mail: imw@uni-bielefeld.de
http://www.imw.uni-bielefeld.de/wp/
ISSN: 0931-6558

On Repeated games with imperfect public monitoring: From discrete to continuous time

Mathias Staudigl* Jan-Henrik Steg†

September 16, 2014

Abstract

Motivated by recent path-breaking contributions in the theory of repeated games in continuous time, this paper presents a family of discrete-time games which provides a consistent discrete-time approximation of the continuous-time limit game. Using probabilistic arguments, we prove that continuous-time games can be defined as the limit of a sequence of discrete-time games. Our convergence analysis reveals various intricacies of continuous-time games. First, we demonstrate the importance of correlated strategies in continuous-time. Second, we attach a precise meaning to the statement that a sequence of discrete-time games can be used to approximate a continuous-time game.

Keywords: Continuous-time game theory, Stochastic optimal control, Weak convergence

1. Introduction

In this paper we study a class of repeated games with imperfect public monitoring which has been introduced in a continuous-time framework in the important contribution of Sannikov (2007). As has been convincingly shown in that paper and its followers, continuous-time games are in certain ways analytically more tractable than their discrete-time counterparts. Nevertheless, there are still open methodological questions how the continuous-time game model fits into the perceived game-theoretic literature. In particular, it is often heuristically argued that continuous-time games are limit cases of sequences

*Center for Mathematical Economics, Bielefeld University, Germany. e-mail: mathias.staudigl@uni-bielefeld.de; website: mwpweb.eu/MathiasStaudigl.

†Center for Mathematical Economics, Bielefeld University, Germany. e-mail: jsteg@uni-bielefeld.de

of discrete-time games. The purpose of this paper is to see in which sense this interpretation is possible. We construct a class of discrete-time games which provides a consistent approximation to the class of repeated games studied in Sannikov (2007), and show that the basic features of the game-theoretic model can indeed be interpreted as limits of sequences of discrete-time games. However, other features of the game model are less clear to interpret. The main obstacle one faces in giving a full-fledged interpretation of continuous-time games as limits of discrete-time games is the notion of a continuous-time strategy. Once passing to the limit from discrete to continuous time, strategy spaces must be enlarged in order to talk about convergence of strategies.[1] As a consequence, limit strategies must be interpreted as correlated strategies (or, abstractly speaking, measure-valued processes), and the limit of these correlated strategies are not easily interpretable in the continuous-time limit game. Nevertheless, we are still able to assign payoffs to the limit game and prove that the limit dynamics of the signal process are diffusions whose drift is controlled by the players' strategies. These facts clearly show in what sense continuous-time games can be interpreted as limits of discrete-time games, and in which sense not.

We feel that an honest study of the foundations of continuous-time game theory is important from a theoretical as well as from an applied perspective. From a theoretical point of view, we are the first who present a rigorous convergence analysis for the aforementioned class of continuous-time games. From an applied point of view, it is important to know that continuous-time games can be derived from discrete-time games, just because all real-world observations are discrete (though may be observed at high frequencies). A continuous-time framework of the strategic interaction can then be regarded as an idealized model, which can be analyzed with the help of powerful analytic tools.

As an independent methodological contribution, we show in this paper that there is a very specific way how repeated games in continuous time can be obtained from the more traditional discrete-time game framework. Our proof of convergence is purely probabilistic, using weak-convergence arguments. This is the weakest form of convergence of stochastic processes one can think of and is the natural one to study strategic interactions in the continuous-time limit. Weak-convergence arguments are nowadays standard in the stochastic control literature (Kushner, 1990), though have not received that much attention in strategic settings.

Many previous papers in the literature dealing with limits of discrete-time games choose a path-wise approach, which may lead to pathological, or not well-defined out-

[1] A classical example for this is the well-known chattering phenomenon. See Example 4.1 and also Example 4.2 in Section 4.2.

comes. This is illustrated very clearly in the fundamental contribution of Simon and Stinchcombe (1989), and has been recently discussed in some depth in Alos-Ferrer and Ritzberger (2008). Our weak-convergence approach does not suffer from these problems, just because we do not care about particular outcomes; all that the weak-convergence approach allows us to do is to make predictions on the probability distribution over the paths of play. Using this weaker notion of convergence, we are able to present a rigorous convergence theorem for a fairly large class of repeated games, and thereby we are able to give continuous-time games a clear meaning.

1.1 Related Literature

Besides the already mentioned references, there are several recent papers trying to shed some light on the connection between discrete-time and continuous-time game theory. Faingold (2008) is concerned with reputation models, whereas we consider repeated games with imperfect public monitoring. Fudenberg and Levine (2007, 2009) study repeated binary choice games between a single long-run player and a sequence of short-run players. They investigate the continuous-time limit of this family of games and work out some conditions under which "non-trivial" equilibria (meaning action profiles which are not stage-game Nash equilibria) can survive when passing to the continuous-time limit.

Contract theory and problems in corporate finance were among the first applications of stochastic analysis in game theory (see Holmström and Milgrom (1987) and Schättler and Sung (1993)). Contract theory in continuous time is now a very active field of research, starting with the paper by Sannikov (2008). For an excellent overview on contract theory in continuous time we refer to Cvitanić and Zhang (2013). There have been several attempts to provide limit theorems for principal-agent models as well, among which we would like to highlight Hellwig and Schmidt (2002) and Biais et al. (2007).

In the theory of zero-sum games there are several recent papers investigating the connection between discrete-time and continuous-time games. Cardaliaguet et al. (2013) and Gensbittel (2013) deal with repeated games with incomplete information on one side, and prove the convergence of the value of the family of discrete-time games to the value of the limit continuous-time game.

Closest to this paper is the recent work by Neyman (2013) and Sannikov and Skrzypacz (2010). Neyman (2013) considers stochastic games with finitely many states where the transition probability is a function of the time lag between two consecutive moves. He makes precise in what sense dynamic games converge to a continuous-time game, and some of his arguments are closely related to ours. The two main differences between his paper and ours is that (i) Neyman (2013) studies finite stochastic games, whereas we consider repeated games, and (ii) the dynamic is completely different. Neyman (2013) assumes that the limit dynamic is a Markov jump process in which the players control the transition rates. We instead consider the diffusion case. What we have in common with the work by Neyman (2013) is the concept of convergence of strategies. We will say more on this point later in the paper.[2]

Sannikov and Skrzypacz (2010) study a family of repeated games with imperfect public monitoring where the public signal process follows a discrete approximation of a jump diffusion. Using geometric arguments similar to Fudenberg et al. (1994), they report a uniform bound (in the limit of frequent actions) on the set of public perfect equilibrium payoffs. While we are not able to say much about the general connection between the sets of sequential equilibrium payoffs of the discrete-time and continuous-time games,[3] we rigorously prove the weak convergence of the game dynamics and the total expected payoffs of the players to corresponding objects of the limit continuous-time game with imperfect public monitoring as introduced in Sannikov (2007).

This paper is structured as follows: Section 2 introduces the class of repeated games with imperfect public monitoring in continuous time. Section 3 describes the approximating family of discrete-time games. Since the proofs of the main results of this paper are rather technical, we spend some time in motivating them. This motivating discussion starts in Section 4. The main results themselves are collected in Section 4.3. The proofs of the main results are organized in various intermediate lemmas which are stated in Section 6. Technical proofs and additional background information on the concepts we use are collected in extra appendices.

2. The continuous-time game

We start with a formal description of the repeated game in continuous time which we intend to approximate via a sequence of discrete-time games with frequent moves. This continuous-time game, denoted by Γ, is a straightforward extension of the model presented in Sannikov (2007) to the N-player case.[4]

[2] For an extension of our arguments to stochastic games see Staudigl and Steg (2014).

[3] Which is a completely open problem in general. See, however, Staudigl (2014b).

[4] We keep the presentation to a minimum; additional information, in particular in view towards characterization of sequential equilibrium payoffs, is given in the companion papers Staudigl (2014b) and Staudigl (2014a).

2.1 The game dynamics in continuous time

We are given a stochastic basis $(\mathbf{W}^{d+1}, \mathcal{W}^{d+1}, \mathsf{P})$, where $\mathbf{W}^{d+1} \triangleq \mathbf{C}(\mathbb{R}_+;\mathbb{R}^{d+1})$ is the space of continuous functions taking values in $\mathbb{R}^{d+1}$, endowed with the metric of uniform convergence on compact sets, and $\mathcal{W}^{d+1}$ is the σ-algebra generated by the cylinder sets.[5] A generic element of the space $\mathbf{W}^{d+1}$ is denoted by $\mathbf{w} = (w_0, \dots, w_d)$. We endow the measure space $(\mathbf{W}^{d+1}, \mathcal{W}^{d+1})$ with the Wiener measure P. Let $B = (B_0, \dots, B_d)^\top = (B_0, B_{(1)})^\top$ be the canonical projection mappings on $\mathbf{W}^{d+1}$, defined by

$$B_0(t, \mathbf{w}) = w_0(t), \quad B_{(1)}(t, \mathbf{w}) = (w_1(t), \dots, w_d(t)).$$

Hence, under P, the process $(B_0, B_{(1)})$ is $(d+1)$-dimensional standard Brownian motion.

N players $i \in I \triangleq \{1, \dots, N\}$ continuously monitor the evolution of a public signal (X_t, Y_t). This public signal consists of a (cumulative) signal process $X_t \in \mathbb{R}^d$, and a public correlation device $Y_t \in \mathbb{R}$. The signal X has continuous sample paths and is given as the unique strong solution to the stochastic differential equation $\mathrm{d}X_t = C\mathrm{d}B_{(1)}(t)$, $X_0 = CB_{(1)}(0)$, where C is a non-singular deterministic matrix. Similarly, the process Y has continuous sample paths and is given by $Y_t = B_0(t)$ for all t. We denote by $\{\mathcal{X}_t^\circ\}_{t\geq 0}$ and $\{\mathcal{Y}_t^\circ\}_{t\geq 0}$ the filtrations generated by the processes X and Y, respectively. Additionally, we denote by $\mathcal{X}_t \triangleq \mathcal{X}_{t+}^\circ$ and $\mathcal{Y}_t \triangleq \mathcal{Y}_{t+}^\circ$ the right-continuous augmentations of the respective σ-algebras. Let

$$\mathcal{W}_t^\circ \triangleq \sigma(\mathcal{X}_t^\circ, \mathcal{Y}_t^\circ),\ \mathcal{W}_\infty^\circ \triangleq \sigma\left(\bigcup_{0\leq t<\infty} \mathcal{W}_t^\circ\right) \text{ and } \mathcal{W}_t \triangleq \sigma(\mathcal{X}_t, \mathcal{Y}_t).$$

It is important to observe that $\mathcal{X}_t^\circ = \sigma(B_{(1)}(s); s \leq t)$ and $\mathcal{Y}_t^\circ = \sigma(B_0(s); s \leq t)$. In repeated games with imperfect public monitoring the players' strategies are formulated as functionals over the sample paths of the public signal process (X, Y), satisfying the usual "adaptedness" condition in dynamic games.

Definition 2.1. *An N-tuple of stochastic processes $\alpha = (\alpha^i)_{i\in I}$ is a* public strategy *profile if each $\alpha^i = \{\alpha_t^i; t \geq 0\}$ is a $\{\mathcal{W}_t\}_{t\geq 0}$-progressively measurable process taking values in the finite set of actions A^i available to player i. The set of public strategies for player i is defined as*

$$\mathcal{A}^i \triangleq \left\{\alpha^i : \mathbb{R}_+ \times \mathbf{W}^{d+1} \to A^i | (\forall T > 0) :\ \alpha^i|_{[0,T]} \textit{ is } \mathcal{B}([0,T]) \otimes \mathcal{W}_T\textit{-measurable}\right\}.$$

[5] This metric is given by

$$\rho(f,g) \triangleq \sum_{n\geq 1} 2^{-n} \min\{1, \|f-g\|_{[0,n]}\}, \text{ where } \|f-g\|_{[0,T]} \triangleq \sup_{0\leq t\leq T} |f(t) - g(t)|.$$

Remark 2.2. Public strategies, even if defined as processes on the probability space $\mathbf{W}^{d+1}$, can now be understood as functionals of the sample paths of the signal process (X, Y) because $B = (B_0, B_{(1)})^\top = (Y, C^{-1}X)^\top$ is the canonical projection mapping on $\mathbf{W}^{d+1}$. Hence, in the following we write $\alpha_t = \tilde{\alpha}_t(X, Y)$ if we want to indicate the functional character of the public strategy profile.

An important technical reason to choose the (canonical) sample path space as the underlying probability space is that the following change of measure using the Girsanov transformation is quite intricate with an infinite time horizon. On the one hand, one has to take care that there even exists a new probability measure with the desired properties on the given set-up (which works for the path space). On the other hand, this new measure will generally not even be absolutely continuous with respect to the initial measure, whence one cannot take the usual completion of the filtration by null sets. The two measures will only be equivalent when restricted to any $\mathcal{W}_t$ for fixed t.

Public strategies affect the probability distribution over signal paths via a change of measure. Let $b : A \triangleq \prod_{i \in I} A^i \to \mathbb{R}^d$ be a given bounded vector-valued mapping, and denote by $f : A \to \mathbb{R}^{d+1}$ the map

$$f(a) \triangleq (0, C^{-1}b(a))^\top \triangleq (0, \mu(a))^\top \qquad \forall a \in A.$$

Hence, the progressively measurable process $(t, \mathbf{w}) \mapsto f(\alpha_t(\mathbf{w}))$ is bounded, so that the stochastic exponential

$$M_t^\alpha \triangleq \exp\left(\int_0^t f(\alpha_s) \cdot \mathrm{d}B(s) - \frac{1}{2}\int_0^t \|f(\alpha_s)\|^2 \,\mathrm{d}s\right)$$

is a true martingale on the stochastic basis $(\mathbf{W}^{d+1}, \mathcal{W}^{d+1}, \{\mathcal{W}_t\}_{t\geq 0}, \mathsf{P})$. By the Cameron-Martin-Girsanov theorem (see e.g. Rogers and Williams, 2000, Theorem IV.38.9), there exists a unique probability measure P^α on $(\mathbf{W}^{d+1}, \mathcal{W}_\infty^\circ)$ defined by[6]

$$\mathsf{P}^\alpha(\Gamma) \triangleq \mathsf{E}^{\mathsf{P}}[\mathbb{1}_\Gamma M_t^\alpha] \qquad \forall \Gamma \in \mathcal{W}_{t+}^\circ,\ t \geq 0.$$

Hence, for every t, the measure $\mathsf{P}^\alpha|_{\mathcal{W}_{t+}^\circ}$ is equivalent to $\mathsf{P}|_{\mathcal{W}_{t+}^\circ}$ (but P^α, a probability mea-

[6]In the remainder of this paper we use the following notation consistently. For any separable metric space X, we denote by $\Delta(X)$ the set of Borel probability measures on it. Given $m \in \Delta(X)$ we denote by E^m the integral operator with respect to m. Hence, for any m-integrable function $f : X \to \mathbb{R}$ we write

$$\mathsf{E}^m(f) \triangleq \int_X f \mathrm{d}m = \int_X f(x)\mathrm{d}m(x) = \int_X f(x)m(\mathrm{d}x).$$

sure on the measure space $(\mathbf{W}^{d+1}, \mathcal{W}^\circ_\infty)$, need not be equivalent to P), and the process

$$B_0^\alpha(t) = B_0(t), \quad B_{(1)}^\alpha(t) \triangleq B_{(1)}(t) - \int_0^t \mu(\alpha_s)\mathrm{d}s \tag{1}$$

is a standard $(d+1)$-dimensional Brownian motion with respect to $\{\mathcal{W}^\circ_{t+}\}_{t\geq 0}$ under P^α. Then the sample paths of the signal process (X, Y) satisfy

$$\begin{aligned} X_t &= \int_0^t b(\alpha_s)\mathrm{d}s + CB_{(1)}^\alpha(t), \\ Y_t &= B_0^\alpha(t) \qquad \forall t \geq 0. \end{aligned}$$

Hence, under P^α the public signal process X becomes a Brownian motion with drift $b(\alpha)$ and volatility C. In terms of stochastic analysis the tuple

$$(\mathbf{W}^{d+1}, \mathcal{W}^{d+1}, \{\mathcal{W}_t\}_{t\geq 0}, \mathsf{P}^\alpha), \ (X, Y, B^\alpha)$$

defines a *weak solution* to the stochastic differential equation

$$\begin{cases} \mathrm{d}X(t) = b(\alpha_t)\mathrm{d}t + C\mathrm{d}W_{(1)}(t), \ X(0) = \mathbf{0}, \\ \mathrm{d}Y(t) = \mathrm{d}W_0(t), \ Y(0) = 0, \end{cases} \tag{2}$$

where we recall that α can be interpreted as a (progressive) functional of the paths of (X, Y) and $W = (W_0, W_{(1)})^\top$ is a standard Brownian motion. The main point of interest of a weak solution is the induced distribution on the path space, since one does not fix the stochastic basis (and, in particular, the driving Brownian motion process) a priori.[7] This is the right solution concept when it is not possible to insist on the filtration being the one generated by the driving Brownian motion, which is the natural situation for games with imperfect monitoring by a public signal (and not the noise process).

Remark 2.3. When we discretize the game later on and pass to the limit, we need to account for even less measurability and consider also α as part of the "solution" because one will *not* be able to regard it as a given functional of the signal paths. Instead one can understand it as a control process to be chosen, in which case (2) becomes a *controlled* SDE. We demonstrate this concept first in Section 2.3. ◆

Remark 2.4. The reader will have observed that the construction of the game dynamics can be easily generalized to allow for a state-dependent volatility matrix $C(x)$. The con-

[7] In particular, once the existence of a weak solution to a given SDE is settled, there is always a "canonical" version of it. For an exceptionally clear presentation of the distinction between strong and weak solutions in stochastic control problems see Yong and Zhou (1999).

struction of weak solutions would be entirely analogous to the one described above if one assumes that the data are uniformly elliptic: Assume that there exists a constant $c > 0$ such that

$$\sum_{i,j=1}^{d} (CC^{\top})_{ij}(x)\gamma_i\gamma_j \geq c \|\gamma\|^2$$

for all $x, \gamma \in \mathbb{R}^d$. ◆

2.2 Payoff processes in continuous-time

In continuous time players receive a continuous flow of payoffs given by[8]

$$R_t^i(\alpha^i) \triangleq \int_0^t \phi^i(\alpha_s^i) \cdot \mathrm{d}X_s + \int_0^t \psi^i(\alpha_s^i)\mathrm{d}s. \tag{3}$$

The goal of the players is to maximize their normalized discounted infinite-horizon utility

$$U^i(\alpha) \triangleq \mathsf{E}^{\mathsf{P}^\alpha}\left[\int_0^\infty re^{-rt}\mathrm{d}R_t^i(\alpha^i)\right],$$

which, by Fubini's Theorem, is equal to

$$U^i(\alpha) = \mathsf{E}^{\mathsf{P}^\alpha}\left[\int_0^\infty re^{-rt}g^i(\alpha_t)\mathrm{d}t\right].$$

The function g^i is the expected utility rate function, defined by

$$g^i(a^i, a^{-i}) \triangleq \phi^i(a^i) \cdot b(a^i, a^{-i}) + \psi^i(a^i) \qquad \forall a^i \in A^i,\ a^{-i} \in A^{-i} = \prod_{j \neq i} A^j. \tag{4}$$

2.3 A projection result

Our class of admissible strategies in the game Γ contains implicitly a public correlation device (the independent Brownian motion B_0). In order to appreciate the notion of convergence of discrete-time games on which our arguments are based, it will be useful to understand the structure of the continuous-time game when the public correlation device is "averaged out". This will be made precise in this section.

[8] As mentioned in Sannikov and Skrzypacz (2010), this is the most general specification for the flow payoff function in games with imperfect public monitoring.

For a given public strategy profile $\alpha \in \mathcal{A} = \prod_{i \in I} \mathcal{A}^i$ we define a $\Delta(A)$-valued $\{\mathcal{X}_t\}_{t \geq 0}$-progressively measurable process $\lambda^\alpha = \{\lambda_t^\alpha\}_{t \geq 0}$ satisfying[9]

$$\lambda_t^\alpha(a) = \mathsf{P}^\alpha\left[\alpha_t = a | \mathcal{X}_t\right] \qquad \forall a \in A,\ t \geq 0. \tag{5}$$

Formally, one can interpret this process as an extensive-form correlated strategy (Forges, 1986; Lehrer, 1992). However, since there are no periods in a continuous-time game, we avoid this interpretation and simply call the process λ^α a *correlated strategy process.*

For fixed t, λ_t^α is easily constructed as the regular conditional probability distribution of the random variable α_t, given the information $\mathcal{X}_t$.[10] If we took only this definition, we would obtain an $\mathcal{X}_t$-adapted process, defined up to sets of measure 0 that depend on t. Since t is a continuous variable, there are uncountably many such null-sets, which eventually causes problems in the definition of the process $\lambda^\alpha = \{\lambda_t^\alpha; t \geq 0\}$. Nevertheless, we can rely on results from filtering theory (see Theorem 2.2.1 in Bain and Crisan (2000)), which guarantee that λ^α exists as a $\{\mathcal{X}_t\}_{t \geq 0}$-progressively measurable process.[11] We will use the correlated strategy process λ^α to construct weak solutions of the controlled SDE

$$\mathrm{d}X(t) = \sum_{a \in \mathcal{A}} b(a)\lambda_t(a)\mathrm{d}t + C\mathrm{d}W(t), \tag{6}$$

where $W(\cdot)$ is a standard d-dimensional Brownian motion relative to a filtration $\{\mathcal{G}_t\}_{t \geq 0}$, and $\{\lambda_t\}_{t \geq 0}$ is progressively measurable with respect to $\{\mathcal{G}_t\}_{t \geq 0}$ and $\Delta(A)$-valued. Setting $\Lambda(a, t) \triangleq \int_0^T \lambda_t(a)\mathrm{d}t$, we say that the pair (Λ, W) are the *driving forces* in this equation, whose solution is the signal process X. Λ is a measure-valued random variable on the space $A \times [0, \infty)$, with the properties that $\Lambda(A, T) = T$ for all $T \geq 0$, and the random variable $\Lambda(a, T)$ is $\mathcal{G}_T$-measurable for every T. Measures with this property are known as *relaxed controls* in optimal control theory, and we denote the space of all (deterministic) relaxed controls by $\mathcal{R}(A \times [0, \infty)) \equiv \mathcal{R}$.[12] All these processes are realized on some underlying probability space $(\Omega, \mathcal{G}, \mathsf{P})$, which is part of the solution. The next proposition explicitly constructs one weak solution of this controlled SDE.

[9]The sigma algebra $\mathcal{X}_t$ is the augmented version of the $\mathcal{X}_t^\circ$.

[10]Since A is finite and the underlying sample space is Polish, the regular conditional probability can be defined P-almost everywhere. On the null set where it may not be defined one can modify this process to be the Dirac measure on some arbitrarily chosen action profile $a \in A$. This does not affect the results.

[11]We just remark that if players ignore the public correlation device in forming their strategy, then the correlated strategy process is trivial and given by

$$\lambda_t^\alpha(a) = \delta_a(\alpha_t) \qquad \forall t \geq 0,\ a \in A$$

P^α-almost surely.

[12]More information on relaxed controls is given in Appendix A.1.

Proposition 2.5. *On the set-up* $(\mathbf{W}^{d+1}, \mathcal{W}^{d+1}, \mathcal{X} = \{\mathcal{X}_t\}_{t\geq 0}, \mathsf{P}^\alpha)$ *the signal process* X *is a weak solution to the stochastic differential equation*

$$\mathrm{d}X_t = \sum_{a\in A} b(a)\lambda_t^\alpha(a)\mathrm{d}t + C\mathrm{d}\bar{B}_t^\alpha,$$

where

$$\bar{B}_t^\alpha \triangleq B_{(1)}^\alpha(t) - \int_0^t (\mu_s^{\lambda^\alpha} - \mu(\alpha_s))\mathrm{d}s = B_{(1)}(t) - \int_0^t \mu_s^{\lambda^\alpha}\mathrm{d}s, \tag{7}$$

and $\mu_t^{\lambda^\alpha} \triangleq \sum_{a\in A}\mu(a)\lambda_t^\alpha(a)$. *The process* $\bar{B}^\alpha$ *is a d-dimensional Brownian motion on this set-up.*

Proof. By the construction (7), we only have to show that $\bar{B}^\alpha$ is a standard Brownian motion on the given set-up. $\bar{B}^\alpha$ is $\mathcal{X}$-adapted since X is so and λ^α progressively measurable. For any $t > s$, we compute that

$$\begin{aligned}\mathsf{E}^{\mathsf{P}^\alpha}[\bar{B}_t^\alpha - \bar{B}_s^\alpha|\mathcal{X}_s] &= \mathsf{E}^{\mathsf{P}^\alpha}\left[\int_s^t \left(\mu_r^{\lambda^\alpha} - \mu(\alpha_r)\right)\mathrm{d}r|\mathcal{X}_s\right] \\ &\quad + \mathsf{E}^{\mathsf{P}^\alpha}[B_{(1)}^\alpha(t) - B_{(1)}^\alpha(s)|\mathcal{X}_s] \\ &= 0.\end{aligned}$$

The last equality follows from the following considerations: First, for all $s \leq r \leq t$, the law of iterated expectations gives us

$$\mathsf{E}^{\mathsf{P}^\alpha}[\mu(\alpha_r)|\mathcal{X}_s] = \mathsf{E}^{\mathsf{P}^\alpha}\left\{\mathsf{E}^{\mathsf{P}^\alpha}[\mu(\alpha_r)|\mathcal{X}_r]|\mathcal{X}_s\right\} = \mathsf{E}^{\mathsf{P}^\alpha}\left[\mu_r^{\lambda^\alpha}|\mathcal{X}_s\right].$$

Second, since $B_{(1)}^\alpha$ is a P^α standard Wiener process, measurable with respect to the filtration $\{\mathcal{W}_t\}_{t\geq 0}$, which includes the filtration $\{\mathcal{X}_t\}_{t\geq 0}$, we obtain

$$\mathsf{E}^{\mathsf{P}^\alpha}\left[B_{(1)}^\alpha(t) - B_{(1)}^\alpha(s)|\mathcal{X}_s\right] = \mathsf{E}^{\mathsf{P}^\alpha}\left\{\mathsf{E}^{\mathsf{P}^\alpha}\left[B_{(1)}^\alpha(t) - B_{(1)}^\alpha(s)|\mathcal{W}_s\right]|\mathcal{X}_s\right\} = 0.$$

Hence, $\bar{B}^\alpha$ is an $\{\mathcal{X}_t\}_{t\geq 0}$-martingale under P^α. Its quadratic variation process under P^α is the same as that of $B_{(1)}^\alpha$:

$$[\bar{B}^\alpha, \bar{B}^\alpha]_t = t \quad \mathsf{P}^\alpha\text{-a.s.}$$

This shows that $\bar{B}^\alpha$ is a $\{\mathcal{X}_t\}$-Wiener process under the measure P^α. ◻

Given the process λ^α and its induced relaxed control measure $\Lambda^\alpha(a, T) \triangleq \int_0^T \lambda_t^\alpha(a)\mathrm{d}t$, Proposition 2.5 states that $(\mathbf{W}^{d+1}, \mathcal{W}^{d+1}, \mathcal{X}, \mathsf{P}^\alpha), (X, \Lambda^\alpha, \bar{B}^\alpha)$ is a weak solution of the controlled SDE (6).

In terms of the correlated strategy process $\{\lambda_t^\alpha, \mathcal{X}_t; t \geq 0\}$, we can compute the expected discounted payoff of player i as follows: Using Fubini's theorem and the law of iterated expectations, one sees that

$$\begin{aligned}
U^i(\alpha) &= \mathsf{E}^{\mathsf{P}^\alpha}\left[\int_0^\infty r g^i(\alpha_t) e^{-rt} \mathrm{d}t\right] \\
&= \mathsf{E}^{\mathsf{P}^\alpha}\left[\sum_{a\in A} g^i(a)\left(r\int_0^\infty \delta_a(\alpha_t) e^{-rt}\mathrm{d}t\right)\right] \\
&= \sum_{a\in A} g^i(a)\int_0^\infty \mathsf{E}^{\mathsf{P}^\alpha}[\delta_a(\alpha_t)] r e^{-rt}\mathrm{d}t \\
&= \sum_{a\in A} g^i(a)\int_0^\infty \mathsf{E}^{\mathsf{P}^\alpha}\{\mathsf{E}^{\mathsf{P}^\alpha}[\delta_a(\alpha_t)|\mathcal{X}_t]\} r e^{-rt}\mathrm{d}t \\
&= \mathsf{E}^{\mathsf{P}^\alpha}\left\{\sum_{a\in A} g^i(a)\int_0^\infty \lambda_t^\alpha(a) r e^{-rt}\mathrm{d}t\right\}.
\end{aligned}$$

This expression will be crucial for arguing that the sequence of repeated game payoffs converges to the payoffs in the continuous time model.

3. The discrete-time game

The family of discrete-time games, denoted by Γ^h, is modeled in the standard way.[13] Since we are only concerned with repeated games with public signals and public strategies, we can describe the information sets of the players completely in terms of sequences of realizations of the signal process. Let $\{\xi_n\}_{n\in\mathbb{N}}$ be an i.i.d. sequence of random variables with law $\rho \in \Delta(\mathbb{R}^d)$, satisfying the following assumptions:

Assumption 3.1. *The probability law ρ satisfies the following conditions:*

$$\int_{\mathbb{R}^d} x \mathrm{d}\rho(x) = 0, \tag{8}$$

$$\int_{\mathbb{R}^d} x x^\top \mathrm{d}\rho(x) = \mathrm{Id}, \tag{9}$$

$$\mathrm{supp}(\rho) = \mathbb{R}^d. \tag{10}$$

The first two assumptions are normalizations, while the full support assumption of the measure ρ is needed to support the hypothesis that players only have imperfect observa-

[13] See Mailath and Samuelson (2006) for an excellent survey of the discrete-time literature.

tions of the actions of their opponents and cannot infer them directly from the realization of the new signal. Formally, we will construct the repeated game dynamics such that the support of the controlled signal distributions is independent of the controls (see eq. (12) and the ensuing discussion). We now come to the explicit construction of the discrete-time game.

3.1 The game dynamics in discrete time

Players monitor the evolution of the cumulative signal process

$$X_n^h = \sqrt{h}C\sum_{k=0}^{n-1}\xi_{k+1},\ X_0^h = \mathbf{0}, \tag{11}$$

so that the realization of the signal ξ_{n+1} determines the jump $\Delta X_n^h \triangleq X_{n+1}^h - X_n^h$ of the cumulative signal process. Let $E = (\mathbb{R}^d)^{\mathbb{N}}$ denote the sequence space, and $\mathcal{B}(E)$ its Borel σ-algebra. Further, we let $\mathbb{P}$ denote the probability law $\rho^{\otimes\mathbb{N}}$ and $\mathcal{F}_n^h$ the filtration generated by the process X^h, i.e.

$$\mathcal{F}_n^h \triangleq \sigma(X_0^h,\dots,X_n^h),\ \mathcal{F}_0^h = \{\varnothing, E\}.$$

Hence, on the set-up $(E,\mathcal{B}(E),\{\mathcal{F}_n^h\}_n,\mathbb{P})$ the process X^h is a Markov chain with stationary transition probabilities and moments

$$\mathsf{E}^{\mathbb{P}}[\Delta X_n^h|\mathcal{F}_n^h] = \mathbf{0},\ \mathsf{Var}^{\mathbb{P}}[\Delta X_n^h|\mathcal{F}_n^h] = hCC^\top.$$

The public histories in the discrete-time game Γ^h are partial sequences of signal realizations, $(x_0,\dots,x_n) \in (\mathbb{R}^d)^{n+1},\ n \geq 0$. A pure public strategy is a measurable map from public histories to pure actions. Formally, a discrete-time pure public strategy for player i is a collection of measurable functions $\underline{a}_n^{i,h} : (\mathbb{R}^d)^{n+1} \to A^i$. Given the previous signal realizations $x_{[n]} = (x_0,\dots,x_n)$, player i responds with the action $\underline{a}_n^{i,h}(x_0,\dots,x_n) \in A^i$. A collection of measurable mappings $\{\underline{a}_n^{i,h}\}_{n\in\mathbb{N}_0}$ is called a *pure public strategy*. As a convenient notational device we define $\{\mathcal{F}_n^h\}_n$-adapted processes $\{\alpha_n^{i,h}\}_n$ by

$$\alpha_n^{i,h} \triangleq \underline{a}_n^{i,h}(X_0^h,\dots,X_n^h) \quad \forall n \in \mathbb{N}_0,$$

and associate strategy profiles with $\alpha_n^h \triangleq (\alpha_n^{1,h},\dots,\alpha_n^{N,h})$. The space of pure public strategies in the discrete-time game Γ^h is denoted by $\mathcal{A}^{i,h}$. The set of pure public strategy profiles is the Cartesian product $\mathcal{A}^h \triangleq \prod_{1\leq i\leq N} \mathcal{A}^{i,h}$.[14]

For every action profile $a \in A$, define the probability law ν_a^h on $(\mathbb{R}^d, \mathcal{R}^d)$ by

$$\nu_a^h(\Gamma) \triangleq \rho((\theta_a^h)^{-1}(\Gamma)) \quad \forall \Gamma \in \mathcal{R}^d, \tag{12}$$

where

$$\theta_a^h(u) \triangleq \theta^h(a,u) \triangleq \sqrt{h}C^{-1}b(a) + u \quad \forall (a,u) \in A \times \mathbb{R}^d.$$

Given Assumption 3.1, the measure ν_a^h is seen to satisfy the moment conditions

$$\int_{\mathbb{R}^d} z \mathrm{d}\nu_a^h(z) = \mathsf{E}^{\mathbb{P}}[\theta_a^h(\xi_1)] = \sqrt{h}C^{-1}b(a) \triangleq \sqrt{h}\mu(a),$$
$$\int_{\mathbb{R}^d} (z - \sqrt{h}C^{-1}b(a))(z - \sqrt{h}C^{-1}b(a))^\top \mathrm{d}\nu_a^h(z) = \mathsf{Var}^{\mathbb{P}}[\theta_a^h(\xi_1)] = \mathrm{Id}\,.$$

By construction, the support of each measure ν_a^h is the support of ρ and hence independent of the action profile as in Abreu et al. (1990).[15] To understand this construction better, observe that if we take the probability measure $\mathbb{P}_a^h \triangleq (\nu_a^h)^{\otimes \mathbb{N}}$ on the sequence space $(E, \mathcal{B}(E))$ and run the iid process $\{\xi_n\}$ on the set-up $(E, \mathcal{B}(E), \mathbb{P}_a^h)$ (recall that these mappings are the canonical projections), then we see that the law of the individual projections is changed from ρ to ν_a^h. Hence, when running the process X^h on this set-up, we see that its increments ΔX_n^h are still iid, but with mean $hb(a)$ and covariance $hCC^\top$. Now we generalize this setting to general public strategies, not only those which constantly play a single action profile. Hence, fix a strategy profile in feedback form $\{\underline{a}_n^h\}_{n\geq 0}$ and define a sequence of stochastic kernels $\{\nu_n^h\}_{n\in\mathbb{N}_0}$ by

$$\nu_n^h(\Gamma|x_{[n]}) \triangleq \nu^h_{\underline{a}_n^h(x_{[n]})}(\Gamma) \qquad \forall n \geq 0,\, \Gamma \in \mathcal{B}(\mathbb{R}^d),\, x_{[n]} = (x_0, \dots, x_n) \in (\mathbb{R}^d)^{n+1}. \tag{13}$$

Clearly, each measure $\nu^h_{\underline{a}_n^h(x_{[n]})}(\cdot)$ is a probability on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ given $x_{[n]} \in (\mathbb{R}^d)^{n+1}$. Us-

[14] Public strategies are the usual class of strategies in games with public monitoring. Clearly enough, it is a restrictive class of strategies. However, it has some nice properties. First, it is closed under best-replies, meaning that if the (N-1) other players use public strategies, there is a best-reply for the remaining player in public strategies as well. Second, public strategies allow us to use recursive techniques to analyze equilibrium payoff processes. Third, the probability distribution over the game tree induced by pure private strategies (i.e. strategies in which player i may condition on his own past decisions, unobservable to the opponents), can be copied by pure public strategies. Hence, they give rise to the same expected payoff. See Mailath and Samuelson (2006) for further discussion.

[15] If the support of the probability measures ν_a^h shifted with the action profile, the players could in principle infer the action profile from the distribution function. The constant support hypothesis makes such a statistical inference impossible.

ing our shorthand notation α^h for the feedback strategy profile $\underline{a}^h \in \mathcal{A}^h$, we let $\mathbb{P}^{\alpha^h}$ denote the induced probability measure on $(E, \mathcal{B}(E))$, which is uniquely defined by the Ionescu-Tulcea Theorem (see e.g. Ethier and Kurtz, 1986). Under this measure, the cumulative signal process X^h is a non-stationary Markov chain whose increments satisfy the moment conditions ($\mathbb{P}$-a.s.)

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}[\Delta X_n^h | \mathcal{F}_n^h] = hb(\alpha_n^h), \tag{14}$$

$$\mathsf{Var}^{\mathbb{P}^{\alpha^h}}[\Delta X_n^h | \mathcal{F}_n^h] = hCC^\top. \tag{15}$$

3.2 Viewing Γ^h as a continuous-time game

We think of the discrete-time game Γ^h as a strategic interaction in continuous time $t \in \mathbb{R}_+$, but where changes in the environment happen only at multiples of a common mesh size $h > 0$. Hence, the time points where players observe a new signal and may adapt their actions are $t_n^h = nh$ for $n \in \mathbb{N}_0$. Accordingly, the dynamic variables of the game are continuous-time processes with piecewise constant sample paths. Specifically, we denote by $\bar{X}^h$ the piecewise constant interpolation of the discrete-time signal process, i.e.

$$\bar{X}_t^h = X_n^h \quad \forall t \in [t_n^h, t_{n+1}^h),\, n \geq 0.$$

The interpolated filtration is defined as $\mathcal{F}_t^h \triangleq \mathcal{F}_{\lfloor t/h \rfloor}^h$ for all $t \geq 0$. The process $\alpha_t^h = \alpha_n^h$ for $t \in [t_n^h, t_{n+1}^h)$ denotes the interpolated strategy process. Next, we identify the driving noise process of the piecewise-constant dynamic. Let

$$\bar{B}_t^h \triangleq C^{-1} \sum_{n=0}^{\lfloor t/h \rfloor - 1} (\Delta X_n^h - hb(\alpha_n^h)) = C^{-1} \sum_{n=0}^{\lfloor t/h \rfloor - 1} \varepsilon_n^h, \tag{16}$$

where

$$\varepsilon_n^h \triangleq \Delta X_n^h - hb(\alpha_n^h) \tag{17}$$

is a zero-mean martingale difference under the measure $\mathbb{P}^{\alpha^h}$. Hence, $\{\bar{B}_t^h\}_{t \geq 0}$ is a càdlàg martingale with respect to the filtration $\{\mathcal{F}_t^h\}_{t \geq 0}$ with zero mean and covariance $h\,\mathrm{Id}$, as one can easily check using eq. (14). In terms of this martingale noise process we obtain the sample path representation of the signal process $\bar{X}^h$ as

$$\bar{X}^h_{t^h_{n+1}} = \bar{X}^h_{t^h_n} + hb(\alpha^h_{t^h_n}) + C(\bar{B}^h_{t^h_{n+1}} - \bar{B}^h_{t^h_n}).$$

This is the stochastic difference equation solved by the signal process X^h when the players play the repeated game with the public strategy profile α^h. The process $\bar{B}^h$ is the driving martingale term of this stochastic difference equation, and thus will be our natural candidate to approximate a standard d-dimensional Wiener process. It is important to emphasize that we have not changed the definition of the stochastic process X^h itself. This process is still generated by the equation (11), but its distribution is changed since the public strategies of the players influence the law of the per-period signal random variables ξ_n.

3.3 Payoff processes in the discrete-time game

Given a public strategy profile α^h, the cumulative payoff to player i at the beginning of stage $n \in \mathbb{N}_0$ of the game Γ^h is defined as

$$R^i_n(\alpha^{i,h}) \triangleq \sum_{k=0}^{n-1} \phi^i(\alpha^{i,h}_k) \cdot \Delta X^h_k + h \sum_{k=0}^{n-1} \psi^i(\alpha^{i,h}_k). \tag{18}$$

The functions ϕ^i, ψ^i are the same as in the continuous-time payoff process (3). Players discount future payoffs at the fixed positive rate $r > 0$, so that the weight of a period n payoff rate is $\int_{t^h_n}^{t^h_{n+1}} re^{-rt}\mathrm{d}t = \delta^n(1-\delta)$, where $\delta = \delta^h \triangleq e^{-rh}$. The scalar $\delta \in (0,1)$ can therefore be interpreted as the discount factor between two consecutive periods in the discrete-time game. Hence, the random discounted (normalized) payoff of player i in the discrete-time game is

$$\sum_{n=0}^{\infty} (1-\delta)\delta^n h^{-1} \Delta R^i_n(\alpha^{i,h}).$$

Recall the definition of the function g^i

$$g^i(a^i, a^{-i}) \triangleq \phi^i(a^i) \cdot b(a^i, a^{-i}) + \psi^i(a^i),$$

which corresponds to the expected stage game payoff rate of player i when the increment distribution of the public signal process X^h is $\nu^h_a(\cdot)$. By Fubini's theorem, the normalized discounted payoff of player i over the infinite time horizon can be expressed as

$$U^i(\alpha^{i,h},\alpha^{-i,h}) \triangleq \mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\sum_{n=0}^{\infty}(1-\delta)\delta^n g^i(\alpha^{i,h}_{t^h_n},\alpha^{-i,h}_{t^h_n})\right]. \tag{19}$$

4. Convergence of the repeated game dynamics

4.1 Preliminary Discussion

The proof of convergence of the family of repeated games Γ^h proceeds in two steps: First, we investigate the compactness properties of interpolated data (constructed in Section 3.2) describing the discrete-time interaction. This analysis will show that the signal process, as well as its driving martingale noise process, are tight sequences of random variables taking values in the space of right-continuous functions having left limits, denoted by $\mathbf{D}^d$.[16] We endow $\mathbf{D}^d$ with the Skorokhod topology under which it becomes a complete separable metric space (see e.g. Jacod and Shiryaev, 2002, for more details). Its induced Borel σ-algebra (i.e. the Borel sigma-algebra generated by the open sets of the Skorokhod topology) is denoted by $\mathcal{D}^d$. Relative compactness allows us to extract subsequences which converge in distribution. Moreover, a standard probabilistic construction, known as the Skorokhod representation, allows us to realize the stochastic processes on one common probability space (independent of the mesh size h), carrying random variables which are equal in law to the process $(\bar{X}^h,\bar{B}^h)$, but which converge not only in law but even almost surely.[17]

A delicate point in our convergence analysis is to understand the limit properties of the family of interpolated discrete-time public strategies. Since the space of pure public strategies is not weakly compact, we have to embed this space into a larger one, having the necessary compactness properties. This compactification of repeated game strategies is a well known tool in optimal control theory (Warga, 1972; Kushner and Dupuis, 2001; El Karoui et al., 1987). We stick to the control-theoretic terminology and call the result-

[16]Let $(\Omega,\mathcal{F},\mathbb{P})$ be some probability space on which we have defined a sequence of random variables $\{Y_n\}_{n\in\mathbb{N}}$ taking values in some separable metric space (E,d). This sequence is tight if the induced family of probability measures $\mathbb{P}\circ Y_n^{-1}=P_n$ is tight. A family of probability measures $\{P_n\}_n$ is tight if for each $\varepsilon>0$ there exists a compact set $K\subset E$ such that $\inf_{n\in\mathbb{N}}P_n(K)\geq 1-\varepsilon$. A family of probability measures $\{P_n\}_n$ on a separable metric space (E,d) is tight if and only if it is relatively compact under the topology of weak convergence of measures on $\Delta(E)$. More details can be found in Ethier and Kurtz (1986) or Billingsley (1999).

[17]The use of the Skorokhod representation in our convergence analysis is completely justified, since only the distributional properties of the repeated game dynamics are what can be analyzed with weak-convergence arguments. The stochastic basis on which the limit processes are realized is not of any importance.

ing objects *relaxed control measures*. Game-theoretically, a relaxed control measure is the appropriate continuous-time concept of a repeated game correlated strategy.[18]

Once the convergence of the interpolated data has been settled, we show that the resulting limit processes solve a SDE identical to the projected SDE (6). Recall from our discussion in Section 2.3 that this SDE represents the evolution of the cumulative signal process X once we average over the public correlation device, and the laws of the public strategies are transported to relaxed control measures. As it will turn out, we will be able to compute expected payoffs of the players knowing only this projected system. It will become apparent that exactly this projected system is what can be approximated by sequences of discrete time games. Some implications of this result on the analysis of equilibria in continuous-time games are discussed in the concluding section of this paper.

4.2 Relaxed Control Measures

Given a public strategy profile in the discrete-time game α^h, define

$$\lambda_t^h(a) \triangleq \delta_a(\alpha_t^h) \qquad \forall a \in A,\, t \geq 0, \tag{20}$$

and

$$\Lambda^h(a, T) \triangleq \int_0^T \lambda_t^h(a)\mathrm{d}t \qquad \forall a \in A,\, T > 0. \tag{21}$$

Observe that $\Lambda^h(A, T) = T$ for every $T \geq 0$, so that the marginal distribution on the time dimension of any induced random measure Λ^h is Lebesgue measure. Hence, for every h the family of random variables Λ^h defines a random element of the space $\mathcal{R}(A \times [0, \infty)) \equiv \mathcal{R}$, which we already encountered in Section 2.3.[19] Using the language introduced in that section, each random variable Λ^h is a relaxed control. The triple $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ completely describes the game dynamics of the discrete-time game Γ^h when the public strategy α^h is used. It is a random element of the space

$$\tilde{\mathcal{Z}} = \mathbf{D}^d \times \mathcal{R} \times \mathbf{D}^d.$$

[18] Formally, our notion of convergence is what Neyman (2013) called w^* convergence of Markov strategies (which is just weak* convergence in suitably defined function spaces). Our strategies are more complex objects than just Markov, which is the reason why we have to use the concept of weak convergence of relaxed controls. We share with Neyman (2013) the interpretation of the limit objects as correlated strategies.

[19] Since players are assumed to use pure public strategies in the game Γ^h, the measure Λ^h is just the occupation measure of the strategy process.

The following examples illustrate clearly why we have to use relaxed control measures in our convergence theorem.

Example 4.1. Correlation in the limit. There are two players with common action set $A^i = \{0,1\}$. For every $h > 0$ assume that the players use the following pure public strategy

$$\alpha_n^{i,h} = \begin{cases} 0 & \text{if } n \text{ is odd,} \\ 1 & \text{if } n \text{ is even.} \end{cases}$$

The resulting profile of public strategies α^h is clearly admissible in the game Γ^h. However, defining the piecewise-constant interpolation $\alpha_t^{i,h}$ as above and looking at sample paths when the mesh size h goes to 0, shows that this function does not have a classical limit. This example is well known in optimal control theory, where it is known as the "chattering phenomenon". It would be unnatural to rule out such chattering controls. Hence, we have to use a weaker concept for a limit. The classical remedy to the chattering problem is to enlarge the space of admissible strategies. Hence, instead of the sequence of public strategies α^h, we consider their "measure" representation (20) and the induced relaxed control measures Λ^h (21). It follows from classical results (summarized in Appendix A) that every sequence of relaxed control measures has a weakly convergent subsequence. It can be checked that the weak limit of the relaxed control measure induced by this pair of strategies has a "density" given by

$$\lambda_t(a) = \frac{1}{2}\delta_{(0,0)}(a) + \frac{1}{2}\delta_{(1,1)}(a),$$

and it is clear that this measure in $\Delta(A)$ cannot be represented via independent randomization of the players. Only a correlated move can generate this distribution. Hence, there is correlation in the limit of frequent moves.[20] ◆

Example 4.2. Measurability in the limit. Another aspect why we obtain more general objects in the limit than standard strategies of the continuous-time game is that information becomes continuous and one cannot decide certain events of the approximating sequence in the limit anymore. As an example fix some time $t \in \mathbb{R}_+$ and two profiles a and a' with different drifts $b(a) \neq b(a')$. Suppose that for any $h > 0$ the profile a is played for

[20] This "problem" has already been observed in Fudenberg and Levine (1986) where they discuss the problem of approximating general open-loop strategies via a sequence of piecewise constant strategies. Since their aim was to assign a natural limit to a sequence of repeated games where players use conventional strategies in the limit (i.e., there is no public correlation device), they are forced to conclude that not all sequences of discrete-time games have a "well-defined limit". Including a public correlation device in the limit game and using the notion of weak convergence allows us to assign a well-defined limit to the sequence of repeated games.

all $n < \lfloor t/h \rfloor$ and also for $n \geq \lfloor t/h \rfloor$ if $(\Delta X^h_{\lfloor t/h \rfloor -1} - hb(a)) \cdot \mathbf{1} \leq 0$. Else the profile a' is played. The probability distribution of the increment process ξ_n under this strategy is therefore ν^h_a for $n < \lfloor t/h \rfloor$ and $\nu^h_{a'}$ for $n \geq \lfloor t/h \rfloor$ if $(\Delta X^h_{\lfloor t/h \rfloor -1} - hb(a)) \cdot \mathbf{1} \leq 0$ holds true. Using the definition of these probability laws given by eq. (12), it is clear that

$$\nu^h_a \left(\{x \in \mathbb{R}^d | (\sqrt{h}Cx - hb(a)) \cdot \mathbf{1} > 0\} \right) = \rho(\{x \in \mathbb{R}^d | Cx \cdot \mathbf{1} > 0\}),$$

and thus is independent of h. Below we will see that the public signal $\bar{X}^h$ converges to a continuous (diffusion) process in the limit, that only changes its drift at the fixed time t with the given probability. Hence, it is not decidable by observing the limit signal process whether the switch has actually occurred. ◆

4.3 Main Theorems

Following our previous discussion, we can identify the discrete-time repeated game with the triple $(\bar{X}^h, \Lambda^h, \bar{B}^h)$, constituting the information process, the relaxed control measure, and the driving martingale process, respectively. In terms of this triple, we now state the main results of this paper. The first theorem settles the question in which sense the family of discrete time games Γ^h provides a foundation for the continuous-time game Γ.

Theorem 4.3 (Weak convergence of game dynamics). *Let $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ be a triple of càdlàg processes obtained from the sequence of discrete-time games Γ^h by piecewise constant interpolation. Then we have the following:*

(i) Every such family of processes is tight in the topologies of their respective sample path spaces;

(ii) For a given subsequence $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ (we do not relabel) there exist a subsubsequence and a probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathsf{P}})$ such that convergence w.p.1 to a limit $(\bar{X}, \Lambda, \bar{B})$ holds. This limit triple has the following properties:

(ii.a) Λ is a relaxed control measure with almost sure decomposition

$$\Lambda(\{a\} \times \mathrm{d}t) = \lambda_t(a) \otimes \mathrm{d}t;$$

(ii.b) Let $\bar{\mathcal{G}}_t \triangleq \sigma(\bar{X}_s, \Lambda(s), \bar{B}_s, s \leq t)$, then $\bar{B}$ is a.s. continuous and is a standard Wiener process;

(ii.c) The process $\bar{X}$ is a.s. continuous and the tuple $(\bar{\Omega}, \bar{\mathcal{G}}, \{\bar{\mathcal{G}}_t\}_{t \geq 0}, \bar{\mathsf{P}}, \bar{X}, \Lambda = \lambda_t \otimes \mathrm{d}t, \bar{B})$

is a weak solution of the stochastic differential equation

$$\text{(22)} \qquad \mathrm{d}X(t) = \sum_{a\in A} b(a)\lambda_t(a)\mathrm{d}t + C\mathrm{d}W(t).$$

This theorem is proven by a sequence of intermediate propositions in Section 6. Tightness of the sequence of relaxed control measures is shown in Section 6.1. The tightness for the processes $\bar{X}^h$ and $\bar{B}^h$ is demonstrated in Proposition 6.3. Part (ii) of the theorem is the well-known Skorokhod representation. (ii.a) is shown in Proposition 6.2. (ii.b) and (ii.c) are shown in Section 6.2, following Proposition 6.3.

Some discussion on its implications is in order. The game dynamics which we obtain as a weak limit from the interpolated data $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ are defined on a probability space which is larger than the one we have used in the explicit construction of the continuous-time process in Section 2. Neither does the information embodied in this set-up correspond to the information the players have in the continuous-time game, nor do the random variables "look" like the ones we have used there. They only describe some continuous-time game in *distribution*. Hence, the proper way of interpreting Theorem 4.3 is not one of "strong convergence" in which the driving data of the repeated game literally converge to a limit repeated game, but rather that the limit objects we obtain are able to describe a continuous-time repeated game in *distribution*.[21] In particular, the process Λ which we obtain as the limit of the discrete-time relaxed control Λ^h does not correspond to a relaxed control measure as the ones we have constructed in Section 2.3 (recall that the objects studied there are $\mathcal{X}_t$-adapted, while the limit is $\bar{\mathcal{G}}_t$-adapted, a filtration which is strictly larger). This said, one may have serious doubts how relevant our theorem actually is. We can only provide a pragmatic reply to this concern: The expected payoff the players obtain in the game only depends on the distributions induced on the path space by the public strategy. As such, Theorem 4.3 says that the family of discrete-time game dynamics converges in distribution and thus allows us to assign a limit value to each family of discrete-time games. Viewed from this angle, our convergence analysis really shows what is possible, in terms of the expected payoffs, when passing to the limit. In fact, this leads us to the second main result of this paper.

Theorem 4.4 (Weak convergence of expected payoffs)**.** *Let* $\{(\bar{X}^h, \Lambda^h, \bar{B}^h); h \in (0,1)\}$ *be a convergent family of random elements of the space* $\tilde{\mathcal{Z}}$ *obtained from the sequence of discrete-time*

[21] This implies that our approach is not able to say something about the relation between the dynamics of the game in discrete and continuous time. This is also the reason why it does not make sense in our framework to perform an equilibrium analysis, since concepts like sequential rationality cannot be verified when passing from discrete to continuous time. A different approach has to be chosen to tackle this question. See Staudigl (2014b).

games Γ^h. In terms of the Skorokhod representation $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathsf{P}})$, $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ (we do not relabel) we have

$$\lim_{h\to 0} U(\alpha^h) = \lim_{h\to 0} \mathsf{E}^{\bar{\mathsf{P}}}\left[\int_0^\infty re^{-rt} g(\alpha_t^h)\mathrm{d}t\right] = \mathsf{E}^{\bar{\mathsf{P}}}\left[r\int_0^\infty e^{-rt}\sum_{a\in\mathcal{A}} g^i(a)\lambda_t(a)\mathrm{d}t\right].$$

5. Conclusion

We have developed a consistent discrete-time approximation scheme for a class of repeated games in continuous time in which the players receive information through a diffusion signal driven by Brownian motion. Convergence of the family of discrete-time games happens rather generally, when only insisting on weak convergence of the interpolated game dynamics. As we have argued in the main text, weak convergence is sufficient when one is interested in obtaining a consistent approximation of the total life-long utility of a player. There are two issues which deserve more attention for future research: First, it is an open problem to get a general interpretation theorem for the limit set of relaxed controls in terms of continuous-time correlated strategies. Second, with our weak convergence arguments one is not able to say much about limits of public perfect equilibria of the game in discrete time. This is a complex topic which has, to our best knowledge, not been studied at all. A more refined analysis of the continuation value process is needed, which is initiated in Staudigl (2014b). In that note a general martingale decomposition for the continuation value process in discrete time is performed, which is the closest possible representation of the corresponding continuous-time processes. While a complete convergence proof is out of reach at the moment, the reader will see that the martingale approach reveals a clean and transparent connection between discrete and continuous-time game theory. In fact, we strongly belief that within this martingale framework one is able to prove the convergence of the continuation payoff process to the corresponding continuous-time process, and we regard this as an important question for future research.

We also remark that the weak convergence approach of this paper can be extended to stochastic games with public signals. The arguments become, however, much more involved and can be found in Staudigl and Steg (2014).

As a final remark we would like to point out that our analysis shows that continuous-time games can be "interpreted" as the limit objects of sequences of discrete-time games with frequent actions only if there is a public correlation device available to the players. Public correlation is frequently used in economics either to facilitate the analysis, or to capture exogenous uncertainty ("sunspots") (see e.g. Stokey and Lucas, 1989; Duffie et al., 1994). Public correlation is also used in Staudigl (2014a) where a general control-theoretic framework is developed to compute the set of public perfect equilibrium payoffs in a continuous-time game. In that note a novel geometric characterization of the set of perfect public equilibrium payoffs in terms of PDEs and the principal curvatures of a self-generating set is derived.

6. Proofs

6.1 Compactness of relaxed controls

The main (technical) reason why we are interested in relaxed controls is the following fundamental property.

Proposition 6.1. *The sequence of relaxed control measures $\{\Lambda^h\}_{h>0}$ is a tight family of random measures in $\mathcal{R}$. Hence, every subsequence has a weakly convergent subsubsequence.*

Proof. See Appendix A.1. □

Given weak convergence we can make use of the Skorokhod representation theorem (see e.g. Ethier and Kurtz, 1986). This theorem states that we can find a probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathsf{P}})$ which carries a sequence of measure-valued random variables $\bar{\Lambda}^h$, each of which is equal in law to Λ^h, but which converges with probability 1 (i.e. $\bar{\mathsf{P}}$-almost everywhere) to the limit random measure $\bar{\Lambda}$ having the same law as Λ. Since all our statements are just concerned with weak convergence of processes, we usually do not make a notational distinction between the w.p.1 converging random variables guaranteed by the Skorokhod representation theorem and the original ones. For the limit Λ we have the following disintegration result.

Proposition 6.2. *Let $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathsf{P}})$ be a probability space and $\bar{\Lambda}$ a random element of $\mathcal{R}(A \times [0, \infty))$. Then there exists $\lambda : [0, \infty) \times \bar{\Omega} \times \mathcal{B}(A) \to [0, 1]$ such that:*

(i) $\lambda(t, \omega, \cdot) \equiv \lambda_t(\cdot|\omega) \in \Delta(A)$ *for all* $t \geq 0$, $\omega \in \bar{\Omega}$;

(ii) For each $a \in A$, $\lambda_\cdot(a|\cdot)$ *is* $\mathcal{B}([0, \infty)) \otimes \bar{\mathcal{F}}$*-measurable;*

(iii) For all sets $C \in \mathcal{B}(A \times [0, \infty)) = \mathcal{B}(A) \otimes \mathcal{B}([0, \infty))$, *the decomposition*

$$\Lambda(C) = \int_{[0,\infty)} \int_A \mathbb{1}_C(a, t) \lambda_t(\mathrm{d}a) \mathrm{d}t$$

holds $\bar{\mathsf{P}}$-almost everywhere.

Proof. See Appendix A.1. □

This result ensures that any random measure Λ, defined as the almost sure limit under the Skorokhod representation of the sequence of relaxed control measures $\{\Lambda^h\}_{h>0}$, admits the decomposition

$$\Lambda(\{a\} \times B) = \int_B \lambda_t(a)\mathrm{d}t \quad \forall B \in \mathcal{B}([0,\infty))$$

by *(iii)*, where λ is a measurable process by *(ii)* with values in $\Delta(A)$ by *(i)*. The subsequent sections will relate this limiting process to a correlated strategy process of the continuous-time repeated game Γ.[22]

6.2 Convergence of the signal process and the driving martingale noise

The first result we need for the weak convergence analysis is a compactness property of the processes $\{(\bar{X}^h, \bar{B}^h), h \in (0,1)\}$. To establish this, we fix a sequence of public strategies $\{\alpha^h; h \in (0,1)\}$, where α^h is the public strategy profile used in the discrete-time game Γ^h.

Proposition 6.3. *The family of interpolated processes $\{\bar{X}^h\}_{h>0}$ and $\{\bar{B}^h\}_{h>0}$ is tight.*

Proof. We use Aldou's tightness criterion to prove the claim (see e.g. Jacod and Shiryaev, 2002, Theorem 4.5, pp. 356). Let $K_b \triangleq \max_{a\in A} |b(a)|$. Then, using $(a+b)^2 \leq 2a^2 + 2b^2$, it follows that

$$\begin{aligned}
\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|\bar{X}^h_t\right\|^2\right] &= \mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|h\sum_{k=0}^{\lfloor t/h\rfloor -1} b(\alpha^h_k) + C\bar{B}^h_t\right\|^2\right] \\
&\leq 2\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|h\sum_{k=0}^{\lfloor t/h\rfloor -1} b(\alpha^h_k)\right\|^2\right] + 2\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|C\bar{B}^h_t\right\|^2\right] \\
&\leq 2K_b^2 t^2 + 2t\,\mathrm{tr}(CC^\top) \triangleq K_t.
\end{aligned}$$

In the step to the last equality we have used an explicit computation of the norm $\left\|C\bar{B}^h_t\right\|^2$,

[22] As the reader can see from the statement of the Proposition, the disintegration result holds for much more general action spaces than the finite ones we are considering in this paper. Indeed, our convergence theorem holds also if each individual action set A^i is a compact separable metric space. More details are given in Appendix A.1.

which will be given below. Using Chebyshev's inequality, we see that for every $\eta > 0$

$$\mathbb{P}^{\alpha^h}\left[\left\|\bar{X}_t^h\right\| \geq \eta\right] \leq \frac{1}{\eta^2}\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|\bar{X}_t^h\right\|^2\right] \leq \frac{K_t}{\eta^2}.$$

Hence, for every $\delta > 0$ we can choose $\eta > \sqrt{\frac{K_t}{\delta}}$ so that

$$\sup_{h>0}\mathbb{P}^{\alpha^h}\left[\left\|\bar{X}_t^h\right\| \geq \eta\right] \leq \delta.$$

This gives the first condition in Aldou's tightness criterion. For the second condition we have to introduce some more notation. Let $\mathcal{T}_T^h$ be the set of $\mathcal{F}_t^h$ stopping times τ which are less or equal to T with probability 1. As in Billingsley (1999) (Theorem 16.9, pp. 177), we have to show that for every η, ε, h there exist δ_0 and h_0 such that if $\delta \leq \delta_0$ and $h \leq h_0$ and $\tau \in \mathcal{T}_T^h$, then

$$\mathbb{P}^{\alpha^h}\left[\left\|\bar{X}_{\tau+\delta}^h - \bar{X}_\tau^h\right\| \geq \eta\right] \leq \varepsilon.$$

This follows if we can show that

$$\lim_{\delta\to 0^+}\limsup_{h\to 0}\sup_{\tau\in\mathcal{T}_T^h}\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[1 \wedge \left\|\bar{X}_{\tau+\delta}^h - \bar{X}_\tau^h\right\|\right] = 0.$$

Using Jensen's inequality and the strong Markov property of the process $\bar{X}^h$, the latter follows from the estimate

$$\begin{aligned}\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[1 \wedge \left\|\bar{X}_{\tau+\delta}^h - \bar{X}_\tau^h\right\|\right] &\leq \sqrt{\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|\bar{X}_{\tau+\delta}^h - \bar{X}_\tau^h\right\|^2\right]} \\ &\leq \sqrt{K_\delta} \to 0 \text{ as } \delta \to 0^+.\end{aligned}$$

The tightness of the process $\bar{B}^h$ is established in the same way. Since some of the bounds used in the argument will be useful later on to give a characterization of the weak limits of this process, we nevertheless give a sketch of the proof. Recall that

$$\bar{B}_t^h = C^{-1}\sum_{n=0}^{\lfloor t/h\rfloor - 1}\varepsilon_n^h,$$

where $\varepsilon_n^h = \Delta X_n^h - hb(\alpha_n^h)$ a martingale difference sequence under the measure $\mathbb{P}^{\alpha^h}$ with

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\varepsilon_n^h(\varepsilon_n^h)^\top|\mathcal{F}_n^h\right] = hCC^\top. \tag{23}$$

By definition of the martingale noise process $\bar{B}^h$ we have

$$\begin{aligned}
\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|\bar{B}_t^h\right\|^2\right] &= \mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[C^{-1}\sum_{n=0}^{\lfloor t/h\rfloor-1}\varepsilon_n^h \cdot C^{-1}\sum_{n=0}^{\lfloor t/h\rfloor-1}\varepsilon_n^h\right] \\
&= \mathsf{E}^{\mathbb{P}^{\alpha^h}}\left\{\operatorname{tr}\left[C^{-1}\left(\sum_{n=0}^{\lfloor t/h\rfloor-1}\varepsilon_n^h\right)\left(\sum_{n=0}^{\lfloor t/h\rfloor-1}\varepsilon_n^h\right)^\top (C^{-1})^\top\right]\right\} \\
&= \operatorname{tr}\left[C^{-1}\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left\{\left(\sum_{n=0}^{\lfloor t/h\rfloor-1}\varepsilon_n^h\right)\left(\sum_{n=0}^{\lfloor t/h\rfloor-1}\varepsilon_n^h\right)^\top\right\}(C^{-1})^\top\right] \\
&= \operatorname{tr}\left[C^{-1}\sum_{n,k=0}^{\lfloor t/h\rfloor-1}\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left\{\varepsilon_n^h(\varepsilon_k^h)^\top\right\}(C^{-1})^\top\right].
\end{aligned}$$

A straightforward computation, using (23) and the law of iterated expectations, shows that for all n

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}[\varepsilon_n^h(\varepsilon_n^h)^\top] = hCC^\top,$$

and for all $k \neq n$

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}[\varepsilon_n^h(\varepsilon_k^h)^\top] = O \text{ the zero matrix in } \mathbb{R}^{d\times d}.$$

Hence,

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|\bar{B}_t^h\right\|^2\right] = h\lfloor t/h\rfloor \operatorname{tr}(\mathrm{Id}) \leq td. \tag{24}$$

This gives the first bound needed in Aldou's tightness criterion, and the second bound is obtained as in the case of the process $\bar{X}^h$. □

Together with Proposition 6.1 this result shows that the triple $\{(\bar{X}^h, \Lambda^h, \bar{B}^h); h \in (0,1)\}$ is tight as a family of random elements in the path space $\tilde{\mathcal{Z}} = \mathbf{D}^d \times \mathcal{R} \times \mathbf{D}^d$. Under our choice of norms, the space $\tilde{\mathcal{Z}}$ is separable and endowed with a metric which makes it complete. Under these topological conditions, tightness is equivalent to relative compactness, which means that for every subsequence of $\{(\bar{X}^h, \Lambda^h, \bar{B}^h); h \in (0,1)\}$ there is a subsubsequence, still denoted by $\{(\bar{X}^h, \Lambda^h, \bar{B}^h); h \in (0,1)\}$, converging in distribution to a limit $(\bar{X}, \Lambda, \bar{B})$. Take the converging subsequence $\{(\bar{X}^h, \Lambda^h, \bar{B}^h); h \in (0,1)\}$ and its weak limit $(\bar{X}, \Lambda, \bar{B})$ as given for the remainder of this section. To complete the proof of Theorem 4.3 it remains to show that

(a) $\bar{B}$ is a standard Wiener process with respect to $\{\bar{\mathcal{G}}_t\}$, where $\bar{\mathcal{G}}_t \triangleq \sigma(\bar{X}_s, \Lambda(s), \bar{B}_s, s \leq t)$;

(b) $\bar{X}$ is a.s. continuous;

(c) the tuple $(\bar{\Omega}, \bar{\mathcal{G}}, \{\bar{\mathcal{G}}_t\}_{t\geq 0}, \bar{\mathsf{P}}, \bar{X}, \Lambda, \bar{B})$ is a weak solution of the stochastic differential equation (22).

For every $h > 0$ the process $\bar{B}^h$ is the driving noise of the public signal process $\bar{X}^h$. To show that the laws of the processes $\{\bar{B}^h\}_{h>0}$ converge weakly to Wiener law on any $[0, T]$, we can appeal to Theorem 8 in (Gihman and Skorohod, 1979, p. 197). By (23) we have for every $n \geq 0$

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[(\bar{B}^h_{t^h_{n+1}} - \bar{B}^h_{t^h_n})(\bar{B}^h_{t^h_{n+1}} - \bar{B}^h_{t^h_n})^\top \Big| \mathcal{F}^h_{t^h_n}\right] = \mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[(C^{-1}\varepsilon^h_n)(C^{-1}\varepsilon^h_n)^\top \Big| \mathcal{F}^h_n\right] = h\,\mathrm{Id}\,.$$

Therefore, we only have to show that the following Lindeberg-type condition holds under Assumption 3.1: For any $\gamma > 0$,

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\sum_{k=0}^{\lfloor T/h \rfloor -1} \mathbf{1}_{\left\{\|C^{-1}\varepsilon^h_k\|^2 \geq \gamma\right\}} \left\|C^{-1}\varepsilon^h_k\right\|^2\right] \to 0 \qquad \text{as } h \to 0. \tag{25}$$

Consider the k-th summand in (25), or actually its conditional expectation wrt. $\mathcal{F}^h_k$ as we are inside $\mathsf{E}^{\mathbb{P}^{\alpha^h}}[\cdot]$. Recalling $\varepsilon^h_k = \Delta X^h_k - hb(\alpha^h_k) = \sqrt{h}C\xi_{k+1} - hb(\alpha^h_k)$ and that the law $\nu^h_{\alpha^h_k}$ of ξ_{k+1} under $\mathbb{P}^{\alpha^h}$ conditioned on $\mathcal{F}^h_k$ is that of $\xi_{k+1} + \sqrt{h}C^{-1}b(\alpha^h_k)$ under $\mathbb{P}$ (defined as $\rho^{\otimes\mathbb{N}}$) we have

$$\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\mathbf{1}_{\left\{\|C^{-1}\varepsilon^h_k\|^2 \geq \gamma\right\}} \left\|C^{-1}\varepsilon^h_k\right\|^2 \Big| \mathcal{F}^h_k\right] = \int_{\mathbb{R}^d} \mathbf{1}_{\left\{\|\sqrt{h}x\|^2 \geq \gamma\right\}} \left\|\sqrt{h}x\right\|^2 \mathrm{d}\rho(x)$$

$\mathbb{P}^{\alpha^h}$-a.s. Hence, the left-hand side of (25) is

$$\lfloor T/h \rfloor \int_{\mathbb{R}^d} \mathbf{1}_{\left\{\|\sqrt{h}x\|^2 \geq \gamma\right\}} \left\|\sqrt{h}x\right\|^2 \mathrm{d}\rho(x) \leq T \int_{\mathbb{R}^d} \mathbf{1}_{\left\{\|x\|^2 \geq \frac{\gamma}{h}\right\}} \|x\|^2 \,\mathrm{d}\rho(x).$$

The last term vanishes as $h \searrow 0$ for fixed $\gamma > 0$, since the law ρ has finite second moment; cf. (9).

Now $\bar{B}$ has almost surely continuous paths since its distribution, the Wiener law, puts probability 1 on $\mathbf{W}^d = \mathbf{C}(\mathbb{R}_+; \mathbb{R}^d) \subset \mathbf{D}^d$. Next we show that with probability 1

$$\bar{X}(t) = \int_0^t \sum_{a \in A} b(a)\lambda_s(a)\mathrm{d}s + C\bar{B}(t) \qquad \forall t \geq 0, \tag{26}$$

where $\lambda = \{\lambda_t; t \geq 0\}$ is taken from the disintegration of Λ according to Proposition 6.2. This will prove that $(\bar{X}, \Lambda, \bar{B})$ satisfy (22) and that also $\bar{X}$ has almost surely continuous paths. To establish (26) we make use of the Skorokhod representation theorem. Hence, we continue working on a new probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathsf{P}})$ on which we have defined a family $(\tilde{X}^h, \tilde{\Lambda}^h, \tilde{B}^h)$ which, for every $h \in (0,1)$, is a random element of the set $\tilde{\mathcal{Z}}$, and whose $\bar{\mathsf{P}}$-law equals the $\mathbb{P}^{\alpha^h}$-law of the original triple $(\bar{X}^h, \Lambda^h, \bar{B}^h)$. As the specific probability space is of no importance for our arguments we stick to the original notation and simply assume that $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ converge almost surely to $(\bar{X}, \Lambda, \bar{B})$.

$\bar{X}^h$ and $\bar{B}^h$ converge in the Skorokhod topology while the measures Λ^h converge weakly to the measure Λ as a random element of the space $\mathcal{R}$. Then, as b is bounded and $t \mapsto \Lambda(a, [0,t))$ uniformly continuous, the RHS of

$$\bar{X}^h_t = \int_0^{\lfloor t/h \rfloor} \sum_{a \in A} b(a)\lambda^h_s(a)\mathrm{d}s + C\bar{B}^h_t \qquad \forall t \geq 0$$

converges also in the Skorokhod topology[23] to the RHS of (26) (as functions of $t \in \mathbb{R}_+$). Hence we have (26) with probability 1.

Finally we show that the weak limit $\bar{B}$ of the processes $\bar{B}^h$ is a martingale wrt. $\{\bar{\mathcal{G}}_t\}$, where $\bar{\mathcal{G}}_t \triangleq \sigma(\bar{X}_s, \Lambda(s), \bar{B}_s, s \leq t)$. We establish this by proving that for fixed times $t, s > 0$

[23]The integral converges even uniformly on any compact interval $[0, T]$. Indeed, since the drift is uniformly bounded, we can estimate

$$\left\| \int_t^u \sum_{a \in \mathcal{A}} b(a)(\lambda^h_s(a) - \lambda_s(a))\mathrm{d}s \right\| \leq 2dK_b(u-t) \qquad \forall u > t \geq 0.$$

If we now fix $\varepsilon > 0$ and set $\delta = \varepsilon/4dK_b$, then $\left\{ \int_0^t \sum_{a \in A} b(a)(\lambda^h_s(a) - \lambda_s(a))\mathrm{d}s; t = k\delta \in [0,T], k \in \mathbb{N}_0 \right\}$ converges to 0 uniformly with probability 1 as $h \to 0$, i.e., for all h sufficiently small the set is bounded by $\varepsilon/2$. By the previous estimate and the choice of δ we can extend the set to all $t \in [0,T]$ and it will be bounded by ϵ.

in the limit

$$\mathsf{E}^{\bar{\mathsf{P}}}\left[\Psi_t\left(\bar{B}_{t+s}-\bar{B}_t\right)\right]=0,$$

where Ψ_t is an arbitrary member of a class of functions generating the σ-field $\bar{\mathcal{G}}_t$.

For any integers p, q, let $\varphi_j(\cdot) \in \mathbf{C}_b(A \times \mathbb{R}_+; \mathbb{R}),\ j \leq p$, be some bounded continuous function from $A \times \mathbb{R}_+$ to $\mathbb{R}$, and define the duality pairing between the sets $\mathcal{R}$ and $\mathbf{C}(A \times \mathbb{R}_+; \mathbb{R})$ by[24]

$$(\varphi, m)_t \triangleq \int_0^t \sum_{a \in A} \varphi(a,t) m(a, \mathrm{d}t) \quad \forall t \geq 0.$$

Further, let $t_i, i \leq q$ be time points in $[0,t]$, and let $\Psi(\cdot)$ be a bounded continuous function of the inputs

$$\left((\varphi_j, m)_{t_i}, \bar{X}^h_{t_i}, \bar{B}^h_{t_i}, j \leq p, i \leq q\right).$$

To simplify notation set

$$\Psi^h_t \triangleq \Psi\left((\varphi_j, \Lambda^h)_{t_i}, \bar{X}^h_{t_i}, \bar{B}^h_{t_i}, j \leq p, i \leq q\right).$$

Then

$$\mathsf{E}^{\bar{\mathsf{P}}}\left[\Psi^h_t\left(\bar{B}^h_{t+s}-\bar{B}^h_t\right)\right]=0,$$

because $\bar{B}^h$ is an $\{\mathcal{F}^h_t\}$-martingale and Ψ_t is $\mathcal{F}^h_t$-measurable.

Thanks to our use of the Skorohod representation theorem, we know that with probability 1 $\bar{X}^h$ and $\bar{B}^h$ converge in the Skorokhod topology and the limit paths $\bar{X}_\cdot$ and $\bar{B}_\cdot$ are continuous. Thus also $\bar{X}^h_u$ and $\bar{B}^h_u$ converge to $\bar{X}_u$ and $\bar{B}_u$, resp., with probability 1 for any $u \in \mathbb{R}_+$. Further, with probability 1 the measures Λ^h converge weakly to the measure Λ on $\mathcal{R}(A \times \mathbb{R}_+)$. Hence, $(\varphi_j, \Lambda^h)_{t_i}$ converge to $(\varphi_j, \Lambda)_{t_i}$ for all $j \leq p$, $i \leq q$ a.s., implying the convergence $\Psi^h_t \to \Psi_t \triangleq \Psi\left((\varphi_j, \Lambda)_{t_i}, \bar{X}_{t_i}, \bar{B}_{t_i}, j \leq p, i \leq q\right)$ a.s.

Analogously to (24) one can show that $\mathsf{E}^{\mathbb{P}^{\alpha^h}}\left[\left\|\bar{B}^h_{t+s}-\bar{B}^h_t\right\|^2\right] \leq sd$. Consequently, the family $\left\{\bar{B}^h_{t+s}-\bar{B}^h_t; 0<h<1\right\}$ with their respective laws $\mathbb{P}^{\alpha^h}$ is bounded in L^2 and thus

[24] The term "duality" is a misuse of terminology which is justified on the following grounds: One can identify the space of deterministic relaxed controls $\mathcal{R}$ as a subset of $L^\infty([0,\infty); \Delta(A))$. Then taking as test functions mappings $\phi \in L^1([0,\infty); \mathbf{C}(A; \mathbb{R}))$, the pairing $(\varphi, m)_t$ would be indeed a dual mapping between these two function spaces. We take below test functions which are additionally continuous in the time variable, which is more demanding than just being Lebesgue integrable.

uniformly integrable, and so is $\left\{\Psi_t^h\left(\bar{B}_{t+s}^h - \bar{B}_t^h\right); 0 < h \leq 1\right\}$ as $\Psi(\cdot)$ is bounded.

It follows that $\Psi_t^h\left(\bar{B}_{t+s}^h - \bar{B}_t^h\right)$ converges even in L^1. Hence,

$$\mathsf{E}^{\bar{\mathsf{P}}}\left[\Psi_t\left(\bar{B}_{t+s} - \bar{B}_t\right)\right] = \lim_{h\to 0} \mathsf{E}^{\bar{\mathsf{P}}}\left[\Psi_t^h\left(\bar{B}_{t+s}^h - \bar{B}_t^h\right)\right] = 0.$$

Since the number of functions φ_j and the chosen time points t_i were arbitrary, we have

$$\mathsf{E}^{\bar{\mathsf{P}}}\left[\bar{B}_{t+s} - \bar{B}_t | \bar{\mathcal{G}}_t\right] = 0.$$

The proof of Theorem 4.3 is now complete.

6.3 Convergence of payoffs

The convergence of payoffs is a simple corollary of the convergence of the signal process and the continuity of the problem data. Denote by $\{\Lambda^h\}_{h>0}$ a sequence of relaxed control measures with decomposition $\Lambda^h(a, \mathrm{d}t) = \lambda_t^h(a)\mathrm{d}t$ which converges weakly to a relaxed control measure Λ. Switching to the Skorokhod representation, there is a probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathsf{P}})$ and random variables $(\bar{X}^h, \Lambda^h, \bar{B}^h)$ (we do not relabel) taking values in the set $\tilde{\mathcal{Z}}$ and having the same law as the original triple, but which converge with probability 1 to the limit $(\bar{X}, \Lambda, \bar{B})$.

Then a straightforward manipulation of the expected utility of players gives

$$\begin{aligned} U(\alpha^h) &= \mathsf{E}^{\bar{\mathsf{P}}}\left[r\int_0^\infty e^{-rt} g(\alpha_t^h)\mathrm{d}t\right] \\ &= \mathsf{E}^{\bar{\mathsf{P}}}\left[r\int_0^\infty e^{-rt}\left(\sum_{a\in A} g(a)\lambda_t^h(a)\right)\mathrm{d}t\right] \\ &= \sum_{a\in A} g(a)\mathsf{E}^{\bar{\mathsf{P}}}\left[r\int_0^\infty \lambda_t^h(a)e^{-rt}\mathrm{d}t\right] \\ &= \mathsf{E}^{\bar{\mathsf{P}}}[\Lambda^h(\hat{g})], \end{aligned}$$

where $\hat{g}(a,t) \triangleq g(a)re^{-rt}$, and $\Lambda^h(\hat{g}) \triangleq \sum_{a\in A}\int_0^\infty g(a)\lambda_t^h(a)e^{-rt}\mathrm{d}t$. Since the map $\hat{g}$ is bounded and continuous, the convergence of the expected payoff vector $U(\alpha^h)$ is now a straightforward consequence of the weak convergence of the relaxed control measures Λ^h and the characterization of weak convergence for a sequence of random measures as described in Appendix A. As only distributions matter for the expectation, the convergence of expected payoffs is indeed independent of the spaces on which the Λ^h and Λ are defined, respectively. This completes the proof of Theorem 4.4.

Acknowledgments

This paper was partially written during a visit to the Hausdorff Research Institute for Mathematics at the University of Bonn in the framework of the Trimester Program Stochastic Dynamics in Economics and Finance. We are grateful to Frank Riedel and Max Stinchcombe for their very useful comments, and seminar participants of the workshop "Stochastic Dynamic Games with Applications in Economics and Finance". In particular we would like to thank Fabien Gensbittel for many stimulating discussions on this topic. Mathias Staudigl acknowledges financial support from the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017.

Appendix

A. Weak convergence of relaxed control measures

Recall that a sequence of probability measure $\{\mu_n\}$ on a Polish space E converges weakly to a limit measure $\mu \in \Delta(E)$, denoted by $\mu_n \overset{w}{\to} \mu$, if for all bounded continuous functions $f : E \to \mathbb{R}$

$$\lim_{n\to\infty} \mu_n(f) = \lim_{n\to\infty} \int_E f \mathrm{d}\mu_n = \int_E f \mathrm{d}\mu = \mu(f).$$

A.1 Relaxed control measures

In order to prove convergence of repeated game strategies, we have to embed them in a larger space which we will describe in this appendix. The results derived in this appendix hold at a very general level, which shows that many of our convergence arguments can be extended to larger families of games. In particular, the space of stage-game action profiles A can be an arbitrary compact metric space, so that its Borel σ-algebra $\mathcal{B}(A)$ is countably generated.

Let $\mathcal{R}(A \times [0,T])$ be the set of Borel measures Λ on $A \times [0,T]$ with the property that $\Lambda(A \times [0,T]) = T = \mathsf{Leb}([0,T])$ for all $T < \infty$. If $T = \infty$, we write $\mathcal{R}(A \times [0,\infty))$. With a slight, but convenient, abuse of notation we write a relaxed control measure $\Lambda \in \mathcal{R}(A \times [0,\infty))$ as a bimeasure of the form

$$\Lambda(C \times D) = \Lambda(C, D) \quad \forall C \in \mathcal{B}(A), D \in \mathcal{B}([0,\infty)).$$

This notation is justified by the fact (Morando's Theorem (Morando, 1969), see Appendix 8 in Ethier and Kurtz (1986)) that all relaxed control measures can be disintegrated in the form

$$\Lambda(C \times D) = \int_{A\times[0,\infty)} \mathbb{1}_{C\times D}(a,t)\lambda_t(\mathrm{d}a)\mathrm{d}t,$$

where the time derivative $\lambda : \mathbb{R}_+ \times A \to [0,1]$ is a map such that

(i) $\lambda_t \in \Delta(A) \quad \forall t \geq 0$, and

(ii) $t \mapsto \lambda_t(A)$ is $\mathcal{B}(\mathbb{R}_+)$-measurable for all $A \in \mathcal{B}(A)$.

The existence of a time derivative is essentially due to the fact that $t \mapsto \Lambda(B \times [0,t])$ is absolutely continuous with respect to Lebesgue measure.

Let $\Delta(\mathcal{A} \times [0,T])$ denote the space of Borel probability measures on $A \times [0,T]$, with Borel σ-algebra $\mathcal{B}(A \times [0,T])$. We denote by d^T the Prohorov distance on this set, i.e., for any two measures $\gamma, \rho \in \Delta(A \times [0,T])$

$$d^T(\gamma,\rho) \triangleq \inf\{\varepsilon > 0 | \gamma(C) \leq \rho(C^\varepsilon) + \varepsilon \text{ for all closed } C \in \mathcal{B}(A \times [0,T])\},$$

where $C^\varepsilon \triangleq \{x \in \mathbb{R} | \operatorname{dist}(x,C) \leq \varepsilon\}$. This distance defines a metric on $\Delta(A \times [0,T])$ which makes it a complete separable metric space (Ethier and Kurtz, 1986, Theorem 3.1.7). For $\Lambda \in \mathcal{R}(A \times [0,T])$ define

$$\bar{\Lambda}_{(T)}(C) \triangleq \frac{1}{T}\Lambda(C) \quad \forall C \in \mathcal{B}(A \times [0,T]). \tag{27}$$

By construction, $\bar{\Lambda}_{(T)}$ is a probability measure on $A \times [0,T]$. For $T = \infty$, we define

$$\bar{d}(\gamma,\rho) \triangleq \sum_{k=1}^{\infty} 2^{-k} d^n(\bar{\gamma}_{(k)}, \bar{\rho}_{(k)}) \quad \forall \gamma,\rho \in \mathcal{R}(A \times [0,\infty)).$$

It can be shown that the mapping $\bar{d}$ defines a metric for the space of relaxed control measures $\mathcal{R}(A \times [0,\infty))$, and that a sequence $\{\Lambda^n\}_{n=1}^\infty$ in $\mathcal{R}(A \times [0,\infty))$ converges to an element $\Lambda \in \mathcal{R}(A \times [0,\infty))$, if and only if $d^k(\bar{\Lambda}^n_{(k)}, \bar{\Lambda}_{(k)}) \to 0$ as $n \to \infty$, if and only if $\bar{\Lambda}^n_{(k)} \xrightarrow{w} \bar{\Lambda}_{(k)}$ for every $k \in \mathbb{N}$.

Lemma A.1 (Prokhorov)**.** *The space of relaxed control measures $\mathcal{R}(A \times [0,\infty))$ is weakly sequentially compact under the metric $\bar{d}$.*

Proof. Let $\{\Lambda^n\}_{n\geq 1}$ be a sequence in $\mathcal{R}(A \times [0,\infty))$. We will construct a weakly converging subsequence by a diagonalization argument.

For every $j \geq 1$, $\{\bar{\Lambda}^n_{(j)}\}$ is a sequence of Borel probability measures on the compact metric space $A \times [0,j]$. Hence, $\{\bar{\Lambda}^n_{(j)}\}_{n\geq 1}$ is tight and the direct half of Prokhorov's theorem (Billingsley, 1999, Theorem 1.5.1) guarantees that we can extract a weakly converging subsequence $\{\bar{\Lambda}^{nj}\}_{n\geq 1}$ such that $\bar{\Lambda}^{nj}_{(j)} \xrightarrow{w} \bar{\Lambda}_{(j)} \in \Delta(A \times [0,j])$ and therefore $\Lambda^{nj} \xrightarrow{w} \Lambda_{(j)} := j\bar{\Lambda}_{(j)}$ on $A \times [0,j])$. We use this insight to construct a weakly converging subsequence on $A \times [0,\infty)$ via a diagonal procedure. Indeed, let $\{\Lambda^{n1}\}_{n\geq 1}$ be a subsequence such that $\Lambda^{n1} \xrightarrow{w} \Lambda_{(1)}$ on $A \times [0,1]$ as $n \to \infty$. Continue with the sequence $\{\Lambda^{n1}\}_{n\geq 1} \subset \mathcal{R}(A \times [0,\infty))$, but focus on its restriction to the set $A \times [0,2]$ denoted by $\{\Lambda^{n1}\}_{n\geq 1}$. By Prokhorov's Theorem there exists a weakly converging subsequence $\{\Lambda^{n2}\}_{n\geq 1} \subset \mathcal{R}(A \times [0,\infty))$ such that $\Lambda^{n2} \to \Lambda_{(2)}$. Since $\{\Lambda^{n2}\}_{n\geq 1}$ is a subsequence of $\{\Lambda^{n1}\}_{n\geq 1}$, it follows that $\Lambda^{n2} \xrightarrow{w} \Lambda_{(1)}$ as $n \to \infty$. Proceeding in this way, the k-th subsequence

$\{\Lambda^{nk}\}_{n\geq 1}$ has the property that $\Lambda^{nk} \xrightarrow{w} \Lambda_{(j)}$ for all $j = 1, 2, \ldots, k$. Now we can construct a diagonal sequence $\{\Lambda^{nn}\}_{n\geq 1}$ with the property that $\{\Lambda^{nn}\}_{n\geq j}$ is a subsequence of the weakly converging subsequence $\{\Lambda^{nj}\}_{n\geq 1}$ for all $j \leq n$. Hence, we have now obtained a (sub)sequence $\{\Lambda^{n_k}\}$ such that $\Lambda^{n_k} \xrightarrow{w} \Lambda_{(j)}$ on $A \times [0, j]$ for all $j \in \mathbb{N}$ and therefore $\Lambda^{n_k} \to \Lambda := \lim_{j\to\infty} \Lambda_{(j)}$ under $\bar{d}$ (Λ is understood as the limit of the monotone sequence $(\Lambda_{(j)})$ where $\Lambda_{(j)}$ is extended to $A \times [0, \infty)$ by setting $\Lambda_{(j)}(B) = \Lambda_{(j)}(B \cap \mathcal{A} \times [0, j])$ for all $B \in \mathcal{B}(A \times [0, \infty))$, making Λ a Borel measure on $A \times [0, \infty)$).

$\Lambda \in \mathcal{R}(A \times [0, \infty))$, because for any $T \in [0, \infty)$, $\Lambda^{n_k} \xrightarrow{w} \Lambda_{(\lceil T \rceil)}$ on $A \times [0, \lceil T \rceil]$, and if we approximate $\mathbb{1}_{t\leq T}$ by continuous functions f_m on $[0, \lceil T \rceil]$, then $\int_{A\times[0,\lceil T \rceil]} f_m(t) \mathrm{d}\Lambda^{n_k}(a, t)$ converges to T uniformly in k as $m \to \infty$, implying $\Lambda(A \times [0, T]) = T$. □

Corollary A.2. *The space of relaxed controls $\mathcal{R}(A \times [0, \infty))$ is compact.*

Proof. The space of relaxed controls is a metric space. By Lemma A.1 every sequence in this space has a converging subsequence. Thus, the space $\mathcal{R}(A \times [0, \infty))$ is sequentially compact and therefore compact. □

Now we extend the weak compactness properties of relaxed control measures to the stochastic setting. Consider a sequence of random elements $\{\Lambda^n\}_{n\geq 1}$ on $\mathcal{R}(A \times [0, \infty))$. We say that the sequence converges in distribution, denoted as $\Lambda^n \xrightarrow{d} \Lambda$, if for all bounded continuous functions $f : A \times [0, \infty) \to \mathbb{R}_+$ with compact support we have[25]

$$\Lambda^n(f) \xrightarrow{d} \Lambda(f).$$

Proposition A.3. *Let $\{\Lambda^n\}_{n=0}^{\infty}$ be a sequence of random relaxed control measures. There exists a subsequence which converges in distribution to a random element Λ taking values in the set $\mathcal{R}(A \times [0, \infty))$.*

Proof. The law of the family of random variables $\{\Lambda^n\}$ is a probability measure on $\mathcal{R}(A \times [0, \infty))$. By Corollary A.2 the space $\mathcal{R}(A \times [0, \infty))$ is a compact metric space, and therefore $\Delta(\mathcal{R}(A \times [0, \infty))$ is compact under the Prokhorov metric. Hence, the sequence $\{\Lambda^n\}_{n\geq 1}$ is tight, and therefore relatively compact. □

Consider a sequence of random measures $\{\Lambda^n\}_{n\geq 1}$ in $\mathcal{R}(A \times [0, \infty))$. After relabeling (if necessary) let $\{\Lambda^n\}_{n\geq 1}$ denote the weakly convergent subsequence with limit Λ. We know that any such limit measure has as marginal distribution on the time axis the Lebesgue measure. For the proper interpretation of the almost sure limit of a sequence of relaxed controls, we would like to ensure that there exists an *almost sure decomposition* of

[25]Cf. Theorem 16.16 in Kallenberg (2002) for equivalent notions of convergence of random measures.

the limit measure in the form $\Lambda(C\times D)=\int_{C\times D}\lambda_t(\mathrm{d}a)\mathrm{d}t$, for some measurable process $\lambda:[0,\infty)\times\Omega\to\Delta(A)$. The first step to obtain such a process is the following general fact, which we have learned from Stockbridge (1990). In order to be self-contained, we provide a proof of this result. Note that in the following Lemma the set A can be an arbitrary compact separable metric space.

Lemma A.4. *Let $(\Omega,\mathcal{F},\mathsf{P})$ be a probability space and Λ a random element of $\mathcal{R}(A\times[0,\infty))$. Then there exists $\lambda:[0,\infty)\times\Omega\times\mathcal{B}(A)\to[0,1]$ such that:*

(i) $\lambda(t,\omega,\cdot)\equiv\lambda_t(\cdot|\omega)\in\Delta(A)$ *for all* $t\geq 0,\omega\in\Omega$;

(ii) For each $B\in\mathcal{B}(A)$, $\lambda_\cdot(B|\cdot)$ *is* $\mathcal{B}([0,\infty))\otimes\mathcal{F}$*-measurable;*

(iii) For all sets $C\in\mathcal{B}(A\times[0,\infty))=\mathcal{B}(A)\otimes\mathcal{B}([0,\infty))$, *the decomposition*

$$\Lambda(C)=\int_{A\times[0,\infty)}\mathbb{1}_C(a,t)\lambda_t(\mathrm{d}a)\mathrm{d}t$$

holds w.p.1.

Proof. On the algebra

$$\mathcal{H}\triangleq\{H=\Gamma_1\times\Gamma_2\times\Gamma_3|\Gamma_1\in\mathcal{B}(A),\Gamma_2\in\mathcal{B}([0,\infty)),\Gamma_3\in\mathcal{F}\}$$

define the set-function

$$\rho(\Gamma_1\times\Gamma_2\times\Gamma_3)\triangleq\mathsf{E}\left[\int_{A\times[0,\infty)}\mathbb{1}_{\Gamma_1\times\Gamma_2}(u,t)\mathbb{1}_{\Gamma_3}e^{-t}\mathrm{d}\Lambda(a,t)\right].$$

Extend ρ to a measure on $\mathcal{A}\times[0,\infty)\times\Omega$, and observe that $\rho(A\times[0,\infty)\times\Omega)=1$. Set $\mathcal{Y}\triangleq[0,\infty)\times\Omega$ and $\mathcal{Y}\triangleq\mathcal{B}([0,\infty))\otimes\mathcal{F}$. Then the marginal measure $\mu(\cdot)\triangleq\rho(A\times\cdot)$ on $(\mathcal{Y},\mathcal{Y})$ is absolutely continuous with respect to the measure $\mathsf{Leb}\times\mathsf{P}$, and has Radon-Nikodym derivative $\mathrm{d}\mu=e^{-t}(\mathrm{d}\mathsf{Leb}\times\mathrm{d}\mathsf{P})$. By Morando's Theorem (see the beginning of this section) there exists a mapping $\lambda:\mathcal{Y}\times\mathcal{B}(A)\to[0,\infty)$ such that (i) $\lambda(t,\omega,\cdot)$ is a measure on $\mathcal{A}$, (ii) for all $B\in\mathcal{B}(A)$ the mapping $\lambda(\cdot,\cdot,B)$ is $\mathcal{B}([0,\infty))\otimes\mathcal{F}$ measurable, and (iii) for all $\Gamma_1\in\mathcal{B}(A),\Gamma_2\in\mathcal{B}([0,\infty)),\Gamma_3\in\mathcal{F}$ we have

$$\rho(\Gamma_1\times\Gamma_2\times\Gamma_3)=\mathsf{E}\left[\int_{A\times[0,\infty)}\mathbb{1}_{\Gamma_1\times\Gamma_2}(a,t)\mathbb{1}_{\Gamma_3}\lambda_t(\mathrm{d}a)e^{-t}\mathrm{d}t\right].$$

Hence,

$$\mathsf{E}\left[\int_{A\times[0,\infty)}\mathbb{1}_{\Gamma_1\times\Gamma_2}(u,t)\mathbb{1}_{\Gamma_3}e^{-t}\mathrm{d}\Lambda(a,t)\right]=\mathsf{E}\left[\int_{A\times[0,\infty)}\mathbb{1}_{\Gamma_1\times\Gamma_2}(a,t)\mathbb{1}_{\Gamma_3}\lambda_t(\mathrm{d}a)e^{-t}\mathrm{d}t\right]$$

for all sets Γ_i in the respective Borel σ-field. Therefore, by almost sure uniqueness of conditional expectations, it follows that

$$\Lambda(\Gamma_1 \times \Gamma_2) = \int_{A\times[0,\infty)} \mathbb{1}_{\Gamma_1\times\Gamma_2}(a,t)\lambda_t(\mathrm{d}a)\mathrm{d}t$$

holds w.p.1. This almost sure decomposition can be nicely expressed via the formal notation $\mathrm{d}\Lambda(a,t) = \lambda_t(\mathrm{d}a) \otimes \mathrm{d}t$. Since the time marginal of the random measure Λ is the Lebesgue measure, it follows that $\lambda_t(A|\omega) = 1$ is true P-a.e. $\square$

Applying these general facts to our repeated game contexts, we can draw the following conclusions: The previous lemma ensures that any random measure Λ on $\mathcal{R}(A \times [0,\infty))$ admits the almost sure decomposition

$$\Lambda(\{a\} \times B) = \int_B \lambda_t(a)\mathrm{d}t \quad \forall B \in \mathcal{B}([0,\infty)),\ a \in A,$$

where λ is a measurable process with values in $\Delta(A)$. This proves Proposition 6.2.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica Vol. 58*(5), 1041–1063.

Alos-Ferrer, C. and K. Ritzberger (2008). Trees and extensive forms. *Journal of Economic Theory 143*, 216–250.

Bain, A. and D. Crisan (2000). *Fundamentals of Stochastic Filtering*. Springer - Stochastic Modelling and Applied Probability.

Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet (2007). Dynamic security design: Convergence to continuous time and asset pricing implications. *The Review of Economic Studies 74*(2), pp. 345–390.

Billingsley, P. (1999). *Convergence of Probability Measures* (2nd ed.). Wiley Series in Probability and Statistics.

Cardaliaguet, P., C. Rainer, D. Rosenberg, and N. Vieille (2013). Markov games with frequent actions and incomplete information. *arXiv:1307.3365v1 [math.OC]*.

Cvitanić, J. and J. Zhang (2013). *Contract Theory in Continuous-Time Models*. Springer Finance.

Duffie, D., J. Geanakoplos, A. Mas-Colell, and A. McLennan (1994). Stationary markov equilibria. *Econometrica 62*(4), 745–781.

El Karoui, N., D. Nguyen, and M. Jeanblanc-Picqué (1987). Compactification methods in the control of degenerate diffusions: existence of an optimal control. *Stochastics 20*(3), 169–219.

Ethier, S. N. and T. G. Kurtz (1986). *Markov Processes: Characterization and Convergence*. New York: Wiley.

Faingold, E. (2008). Building a reputation under frequent decisions. Unpublished manuscript, Yale University.

Forges, F. (1986). An approach to communication equilibria. *Econometrica 54*(6), pp. 1375–1385.

Fudenberg, D. and D. K. Levine (1986). Limit games and limit equilibria. *Journal of Economic Theory 38*, 261–279.

Fudenberg, D. and D. K. Levine (2007). Continuous time limits of repeated games with imperfect public monitoring. *Review of Economic Dynamics 10*, 173–192.

Fudenberg, D. and D. K. Levine (2009). Repeated games with frequent signals. *Quarterly Journal of Economics 124*(1), 233–265.

Fudenberg, D., D. K. Levine, and E. Maskin (1994). The folk theorem with imperfect public information. *Econometrica Vol. 62, No. 5*, 997–1039.

Gensbittel, F. (2013). Continuous-time limit of dynamic games with incomplete information and a more informed player. *hal-00910970, version 1*.

Gihman, I. I. and A. V. Skorohod (1979). *The Theory of Stochastic Processes III*. Berlin Heidelberg New York: Springer-Verlag.

Hellwig, M. F. and K. M. Schmidt (2002). Discrete-time approximation of the Holmström-Milgrom model of intertemporal incentive provision. *Econometrica Vol. 70, No. 6*, 2225–2264.

Holmström, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica Vol. 55*(2), 303–328.

Jacod, J. and A. N. Shiryaev (2002). *Limit Theorems for Stochastic processes*. Springer Verlarg - Grundlehren der Mathematischen Wissenschaften, Volume 288.

Kallenberg, O. (2002). *Foundations of modern probability* (2nd ed.). New York [u.a.]: Springer.

Kushner, H. J. (1990). *Weak convergence methods and singularly perturbed stochastic control and filterin problems*. Birkhäuser: System& Control: Foundations & Applications.

Kushner, H. J. and P. Dupuis (2001). *Numerical Methods for Stochastic control problems in continuous time* (2nd ed.). New York: Springer.

Lehrer, E. (1992). Correlated equilibria in two-player repeated games with nonobservable actions. *Mathematics of Operations Research 17*(1), 175–199.

Mailath, G. J. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford University Press.

Morando, P. (1969). *Measures aleatoires*, pp. 190–229. Seminaire de Probabilites III, Lecture Notes in Mathematics. Spinger-Verlag.

Neyman, A. (2013). Stochastic games with short-stage duration. *Dynamic Games and Applications 3*, 236–278.

Rogers, L. and D. Williams (2000). *Diffusions, Markov processes and Martingales - Volume 2 Itô Calculus* (2nd ed.). Cambridge Mathematical Library.

Sannikov, Y. (2007). Games with imperfectly observable actions in continuous time. *Econometrica 75*(5), 1285–1329.

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies 75*, 957–984.

Sannikov, Y. and A. Skrzypacz (2010). The role of information in repeated games with frequent actions. *Econometrica Vo. 78*(3), 847–882.

Schättler, H. and J. Sung (1993). The first-order approach to the continuous-time principal agent problem with exponential utility. *Journal of Economic Theory 61*(2), 331 – 371.

Simon, L. K. and M. B. Stinchcombe (1989). Extensive form games in continuous time: Pure strategies. *Econometrica 57*(5), 1171–1214.

Staudigl, M. (2014a). A geometric characterization of perfect public equilibria in continuous time. In preparation.

Staudigl, M. (2014b). On repeated games with imperfect public monitoring: Characterization of continuation payoff processes. Bielefeld University.

Staudigl, M. and J.-H. Steg (2014). A limit theorem for stochastic games. Technical report, Center for Mathematical Economics (IMW), University Bielefeld. In preparation.

Stockbridge, R. H. (1990). Time-average control of martingale problems: Existence of a stationary solution. *The Annals of Probability 18*(1), pp. 190–205.

Stokey, N. L. and R. E. Lucas (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Warga, J. (1972). *Optimal Control of differential and functional equations*. Academic Press.

Yong, J. and X. Y. Zhou (1999). *Stochastic Controls- Hamiltonian Systems and HJB equations*. Springer.