

The folk theorem with imperfect public information in continuous time

BENJAMIN BERNARD

Department of Mathematics and Statistical Sciences, University of Alberta

CHRISTOPH FREI

Department of Mathematics and Statistical Sciences, University of Alberta

We prove a folk theorem for multiplayer games in continuous time when players observe a public signal distorted by Brownian noise. The proof is based on a rigorous foundation for such continuous-time multiplayer games. We study in detail the relation between behavior and mixed strategies, and the role of public randomization to move continuously across games within the same model.

KEYWORDS. Folk theorem, repeated games, continuous time, imperfect observability.

JEL CLASSIFICATION. C73.

1. INTRODUCTION

Folk theorems constitute a set of key results in the theory of discrete-time repeated games. These theorems state that the set of equilibrium payoffs expands to the set of all feasible and individually rational payoffs \mathcal{V}^* as players are increasingly patient. A central requirement for the standard folk theorem to hold is that profitable deviations from strategy profiles can be detected. Clearly, this is satisfied if players perfectly observe each other's actions. When players observe only a public outcome with noise, players' deviations need to be sufficiently identifiable, that is, deviations of different players can be statistically distinguished. Under suitable identifiability assumptions, [Fudenberg et al. \(1994\)](#) establish the folk theorem in these games of imperfect public information.

A continuous-time analogue to these important games of moral hazard has been introduced by [Sannikov \(2007\)](#). He studies a class of continuous-time games with two players, where the public signal is distorted by Brownian noise and players' actions affect the drift rate of the signal. Rather than focusing on a folk theorem, [Sannikov \(2007\)](#) characterizes the set of pure strategy equilibrium payoffs via a differential equation of

Benjamin Bernard: bernard1@ualberta.ca

Christoph Frei: cfrei@ualberta.ca

We would like to thank the editor and three anonymous referees for valuable comments helping us to significantly improve the structure and exposition of the paper. It is a pleasure to thank Tilman Klumpp, Semyon Malamud, and Yuliy Sannikov for valuable comments and interesting discussions. Financial support by the Natural Sciences and Engineering Research Council of Canada under Grant RGPIN/402585-2011 is gratefully acknowledged.

Copyright © 2016 Benjamin Bernard and Christoph Frei. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](https://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.
DOI: 10.3982/TE1687

its boundary $\partial\mathcal{E}_p(r)$ for any discount rate $r > 0$. His paper offers seminal insight into continuous-time games, but his techniques rely heavily on using an *ordinary* differential equation, which is only possible in two-player games.

The purpose of this paper is to prove the folk theorem in an extension of these games to any number of players. In the course of this, we have to construct continuous-time perfect public equilibria (PPE) that achieve nearly efficient payoffs. Our construction of equilibrium profiles resembles the iterative procedures of discrete time, in the sense that the equilibrium profiles are constructed using a local optimality condition on random time intervals $[0, \tau_1), [\tau_1, \tau_2), \dots$ for a suitable increasing sequence of stopping times $(\tau_\ell)_{\ell \geq 0}$ with $\tau_\ell \rightarrow \infty$. Similarly to Fudenberg et al. (1994), the proof of the folk theorem can be summarized in the following two main steps:

Step 1. Under suitable identifiability conditions, any smooth closed set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$ is locally self-generating, that is, for any $w \in \mathcal{W}$, there exists a neighborhood U_w and a discount rate r_w such that any payoff $v \in U_w \cap \mathcal{W}$ is attained by the same enforceable strategy profile for any $r \in (0, r_w)$ with a continuation value that remains in \mathcal{W} for a time interval $[\tau_\ell, \tau_{\ell+1})$ of positive length.

Step 2. For a sufficiently small discount rate, any compact locally self-generating set is self-generating, i.e., payoffs can be attained with a continuation value that remains in \mathcal{W} forever.

In continuous-time games, the continuation value of a strategy profile is characterized by a stochastic differential equation (SDE) similar as in Sannikov (2007). Step 1 thus requires us to find a solution to an SDE that remains in \mathcal{W} up to a suitable stopping time. The uniformity condition means that these solutions exist on a fixed probability space for the entire neighborhood U_w , that is, locally, these are strong solutions to the SDE. This is important when we concatenate these local solutions to a global solution in Step 2: by compactness of \mathcal{W} , we have to deal with only finitely many probability spaces and it is thus easy to define an enlarged probability space that contains the concatenation. However, requiring that these local solutions are strong solutions in the definition of local self-generation creates additional difficulties in proving Step 1, as the conditions for the existence of strong solutions are much more stringent than for weak solutions. Nevertheless, we are able to show that local self-generation holds under suitable identifiability conditions. These conditions are essentially the same as in discrete time, plus an additional condition to ensure that strong solutions exist in a neighborhood of coordinate payoffs, i.e., payoffs that maximize or minimize a player's payoff on \mathcal{W} .

Requiring strong solutions to the SDE in Step 1 above entails that the local strategy profiles are constant. The constructed equilibrium profiles are therefore constant on each of the intervals $[\tau_\ell, \tau_{\ell+1})$, which is a very desirable feature from both an implementation and an interpretation standpoint. Indeed, the finite variation property of the equilibrium profiles precludes strategies of unbounded oscillation, that is, agents do not switch actions infinitely often in finite time. Note also that agents adapt their strategies only at stopping times $(\tau_\ell)_{\ell \geq 0}$. This leads to the interpretation of a continuously repeated game as a discretely repeated game where the length of the periods is not fixed but random. Indeed, on each of these intervals of random length, equilibrium profiles are constant and the corresponding SDE has a strong solution. This is consistent with

discrete games, where the sampled random variable at time t is necessarily fixed on the interval $[t, t + \Delta)$.

Our methods indicate that rather than approximating the continuous-time world by discrete-time models with fixed time intervals, one should consider approximations where the length of a period is random. This construction avoids the “chattering phenomenon” that arises in optimal control when passing from discrete to continuous time. In comparison, consider a discrete-time game where players observe cumulative outcomes of a diffusion process at fixed times $\Delta, 2\Delta, \dots$. Because it is not possible to bound the change in the public signal with fixed times as it is with stopping times, strategy profiles in the limit as $\Delta \rightarrow 0$ will typically exhibit unbounded oscillation or one has to consider weaker forms of convergence as in [Staudigl and Steg \(2014\)](#). Thus, one can only obtain well implementable solutions in the above sense either by directly working in a continuous-time setting or by defining a suitable sequence of discrete games where the time periods are of random length.

Not only is this paper the first to formally model continuously repeated games with any finite number of players, but it is also the first paper to introduce continuous-time strategies in mixed actions in games of imperfect public monitoring and to establish the continuous-time analogue of Kuhn’s theorem (realization equivalence of behavior strategies and mixed strategies). Additional results include a square-root law, establishing that changes in the discount rate r , the drift rate m , or the volatility σ of the public signal do not affect the game as long as the informativeness of the signal relative to the players’ discounting $\sigma^\top(\sigma\sigma^\top)^{-1}m/\sqrt{r}$ remains constant. This is similar in spirit to a square-root law that is obtained in [Faingold and Sannikov \(2011\)](#) and the news dependence of equilibria in [Daley and Green \(2012\)](#). In contrast to their results, however, we show that equilibrium strategies are transformed with a time change, and that the time-changed equilibria give rise to exactly the same path of the continuation values, at a different speed. Moreover, we show that public randomization can be used to move continuously over games within the same class of games when $\sigma^\top(\sigma\sigma^\top)^{-1}m/\sqrt{r}$ is increased. Together, these two results imply that $\mathcal{E}(r)$ is monotonic in the discount rate if players have access to a public randomization device.

While we study directly the continuous-time situation rather than a discrete-time approximation, we briefly mention some recent literature on the connection between discrete and continuous time in relation to the folk theorem. [Fudenberg and Levine \(2007\)](#) analyze a specific example between one long-lived and one short-lived player, where efficient limit equilibria can be obtained as the length of the time period Δ shrinks to 0 if the long-lived player’s actions affect the volatility of the Brownian signal. [Sannikov and Skrzypacz \(2010\)](#) consider games between two long-lived players, where the public signal has both a Brownian and a Poisson component, and players’ actions affect the drift of the Brownian motion and the intensity of the Poisson jumps. Building upon the methods of [Fudenberg et al. \(1994\)](#), they show a folk theorem uniformly for small Δ and highlight the different impacts of Brownian and Poisson signals. [Osório \(2012\)](#) studies a model where players’ actions incur pairwise identifiable jumps in the signal otherwise given by a Brownian motion. As Δ goes to 0, the signal becomes perfectly informative

and a folk theorem is obtained. This is similar to discrete-time repeated games with perfect monitoring, where $r \rightarrow 0$ and $\Delta \rightarrow 0$ have the same effect.

The remainder of the paper is organized as follows. We introduce the continuous-time model with strategies in mixed actions in Section 2. Section 3 states and discusses our main results and highlights some of the similarities and differences between discrete and continuous time. The proofs of all main results are contained in Section 4. In Appendix A, we provide a mathematical framework for continuous-time strategies in mixed actions and prove the continuous-time analogue of Kuhn's theorem. Similarly to discrete time, the role of mixing is to weaken the conditions under which the folk theorem holds. When players are restricted to pure strategies, the same techniques can be used to construct equilibrium profiles, but the conditions need to be strengthened as we explain in the supplementary file available on the journal website, <http://econtheory.org/supp/1687/supplement.pdf>. Appendix B shows how we use public randomization to prove monotonicity of the equilibrium payoff set in the discount rate. Finally, Appendices C and D contain proofs of auxiliary results.

2. THE MULTIPLAYER SETTING

2.1 The model

We consider a multiplayer game, where agents $i = 1, \dots, n$ continuously take actions from the finite sets \mathcal{A}^i at each moment of time $t \in [0, \infty)$. Players may be allowed to mix their actions, in which case they continuously choose an element from $\Delta(\mathcal{A}^i)$, the set of distributions over \mathcal{A}^i . We denote by $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^n$ and $\Delta(\mathcal{A}) = \Delta(\mathcal{A}^1) \times \dots \times \Delta(\mathcal{A}^n)$ the spaces of pure and mixed action profiles, respectively. Agents cannot see their opponents' actions and observe only the outcome of a public signal $Y = \sigma Z$ instead, where the constant volatility matrix $\sigma \in \mathbb{R}^{d \times k}$ is of rank d and Z is a k -dimensional Brownian motion on some probability space (Ω, \mathcal{F}, P) . The arrival of public information is captured by the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$. It can be strictly larger than the filtration generated by Y to allow for public randomization, but this is not needed in our proof of the folk theorems.¹

DEFINITION 1. A (public) behavior strategy of player i is an \mathbb{F} -progressively measurable stochastic process $A^i : \Omega \times [0, \infty) \rightarrow \Delta(\mathcal{A}^i)$.

Distributions in $\Delta(\mathcal{A}^i)$ may be degenerate, so that behavior strategies contain pure strategies as special cases. A function $m : \mathcal{A} \rightarrow \mathbb{R}^d$ describes the impact that a chosen pure action profile has on the drift rate of the public signal. The drift rate is extended to any mixed action profile $\alpha \in \Delta(\mathcal{A})$ by multilinearity, that is,

$$\mu(\alpha) = \sum_{a \in \mathcal{A}} m(a) \alpha^1(a^1) \dots \alpha^n(a^n). \quad (1)$$

¹In continuously repeated games it is currently unknown whether the PPE payoff set $\mathcal{E}(r)$ is monotonic in the discount rate r . Hence, even when the folk result applies, $\mathcal{E}(r)$ may not expand *monotonically* to \mathcal{V}^* without public randomization. Using public randomization, we are able to show monotonicity of $\mathcal{E}(r)$ with an appropriate time change; see Theorem 3.

The choice of a strategy profile affects the future distribution of the public signal through a change of the probability measure. This is the natural analogue to taking expectations under conditional distributions in discrete time; see, for example, [Fudenberg et al. \(1994, p. 1000\)](#). Strategy profile A induces a family $Q^A = (Q_t^A)_{t \geq 0}$ of probability measures, defined via its density process

$$\frac{dQ_t^A}{dP} := \exp\left(\int_0^t \mu(A_s)^\top (\sigma \sigma^\top)^{-1} dY_s - \frac{1}{2} \int_0^t \mu(A_s)^\top (\sigma \sigma^\top)^{-1} \mu(A_s) ds\right), \quad (2)$$

under which players observe the game when strategy profile A is played. Under this family of probability measures, the public signal can be decomposed into

$$Y = \int \mu(A_s) ds + \sigma Z^A,$$

where $Z^A := Z - \int \sigma^\top (\sigma \sigma^\top)^{-1} m(A_s) ds$ is a Brownian motion on $[0, t]$ under Q_t^A by Girsanov's theorem. That is, from the players' perspective, the signal Y consists of a drift term given by $\mu(A)$ and Brownian noise.

REMARK 1. Because independence is always subject to a certain probability measure, a change of measure might affect the outcome of mixed actions. We assert in [Lemma 14](#) that mixed actions remain unaffected by the change of measure in (2).

[Anderson \(1984\)](#) and [Simon and Stinchcombe \(1989\)](#) demonstrate that seemingly simple strategies need not necessarily lead to unique outcomes in continuous time. This is not a problem in our model because actions taken by agents do not immediately generate information. Indeed, since the normal distribution has unbounded support, any realization of the signal is possible after play of any strategy profile. Therefore, these are games of full-support public monitoring. Restricting attention to public strategies also means that the probability space can be identified with the path space of all publicly observable processes. For a realized path $\omega \in \Omega$, a pure strategy profile A thus naturally leads to the unique outcome $A(\omega)$.² This is analogous to discrete-time repeated games with full-support public monitoring; see [Mailath and Samuelson \(2006\)](#) for a thorough exposition of discrete-time games.

Players $i = 1, \dots, n$ receive an expected flow payoff according to $g^i : \mathcal{A} \rightarrow \mathbb{R}$ that depends on the action profile a^{-i} of player i 's opponents only through $m(a)$.^{3,4} Again, g^i is extended to mixed action profiles by multilinearity at all times.

²For nondegenerate behavior strategy profiles, the outcome of strategy profile A additionally includes the outcome of players' mixing. See [Appendix A](#) for the construction of a unified probability space on which these outcomes live.

³This is the most general payoff structure in continuous-time games of imperfect public information. Compare this to discrete-time repeated games, where players receive an expected instantaneous payoff $g^i(A_t) = \mathbb{E}_{Q_t^A}[f^i(A_t^i, \tilde{Y}) | \mathcal{F}_t]$, where $\tilde{Y} = Y_{t+1} - Y_t$ is the change in the public signal. In continuous time, the infinitesimal change in the public signal dY_t has drift $m(A_t) dt$ under Q^A .

⁴Because Brownian information can only be used to transfer value linearly, we need to impose an affine payoff structure $g^i(a) = b^i(a^i)m(a) - c^i(a^i)$ to show that Pareto-efficient payoffs are enforceable. This special payoff structure is used in [Theorem 2](#), but not in our other main results. For pure strategy profiles,

DEFINITION 2. Let $r > 0$ be a common discount rate. Player i 's *discounted expected future payoff* under a behavior strategy profile A at time t is given by

$$W_t^i(A; r) := \int_t^\infty r e^{-r(s-t)} \mathbb{E}_{Q_s^A}[g^i(A_s) | \mathcal{F}_t] ds. \tag{3}$$

We will omit the argument r when there is no chance of confusion. Because discounted expected future payoffs are normalized, $W_t(A)$ lies in the set of feasible payoffs $\mathcal{V} := \text{conv}\{g(a) | a \in \mathcal{A}\}$ at all times $t \geq 0$ with probability 1.

DEFINITION 3. A behavior strategy profile A is a *perfect public equilibrium* (PPE) for discount rate r if for every player $i = 1, \dots, n$ and every $t \geq 0$, we have

$$W_t^i(A; r) \geq W_t^i(\tilde{A}; r) \quad \text{a.s.}$$

for all public behavior strategy profiles \tilde{A} with $\tilde{A}^{-i} = A^{-i}$ a.e.⁵ We denote the set of payoffs achievable by perfect public equilibria by

$$\mathcal{E}(r) := \{x \in \mathcal{V} | \text{there exists a PPE } A \text{ with } W_0(A; r) = x \text{ a.s.}\}.$$

Similarly to discrete time, any player has a public best response to any public strategy profile of his opponents; see Lemma 15. Therefore, in a PPE, any player $i = 1, \dots, n$ can ensure that his payoff rate dominates his minmax payoff

$$\underline{v}^i = \min_{\alpha^{-i}} \max_{a^i \in \mathcal{A}^i} g^i(a^i, \alpha^{-i})$$

at all times by myopically maximizing against his opponents' strategy profile. Therefore, $\mathcal{E}(r) \subseteq \mathcal{V}^*$, where $\mathcal{V}^* := \{w \in \mathcal{V} | w^i \geq \underline{v}^i \forall i\}$ denotes the set of feasible and individually rational payoffs.

2.2 Enforceable strategy profiles and self-generation

DEFINITION 4. An action profile α is said to be *enforceable* if there exist sensitivities $\beta = (\beta^1, \dots, \beta^n)^\top \in \mathbb{R}^{n \times d}$ to the public signal, such that for every player i , the sum of expected flow payoff $g^i(\alpha)$ and promised continuation rate $\beta^i \mu(\alpha)$ is maximized in α^i . That is, for $i = 1, \dots, n$ and every $a^i \in \mathcal{A}^i$,

$$g^i(\alpha) + \beta^i \mu(\alpha) \geq g^i(a^i, \alpha^{-i}) + \beta^i \mu(a^i, \alpha^{-i}) \quad \text{a.s.} \tag{4}$$

this payoff structure is the same as in Sannikov (2007), and one can show that $W_t(A)$ is the \mathcal{F}_t -conditional expectation under some probability measure Q_∞^A of

$$\int_t^\infty r e^{-r(s-t)} (b^i(A_s^i) dY_s - c^i(A_s^i) ds).$$

That is, b^i is the sensitivity of player i 's payoff to the public signal and c^i is a cost-of-effort function.

⁵Because players maximize their discounted expected future payoff, deviations with time measure 0 or probability 0 are irrelevant. Therefore, two strategy profiles lead to the same continuation value if they are $P \otimes$ Lebesgue-a.e. the same.

A behavior strategy profile is *enforceable* if there exists a progressively measurable process $(\beta_t)_{t \geq 0}$ such that (4) is satisfied a.e.⁶

It follows from (3) that the continuation game after time t is equivalent to the entire game. In discrete time, this feature gives rise to an iterative procedure over the time periods; in continuous time it leads to an SDE characterizing the infinitesimal change in the continuation value. The following lemma is the analogue of Theorem 1 in Sannikov (2007) adapted to the multiplayer setting with behavior strategies.

LEMMA 1. *For an n -dimensional process W and a behavior strategy profile A , the following statements are equivalent:*

- (a) *The process W is the discounted expected future payoff under A .*
- (b) *The process W is a bounded semimartingale that satisfies for $i = 1, \dots, n$ that*

$$dW_t^i = r(W_t^i - g^i(A_t)) dt + r\beta_t^i(\sigma dZ_t - \mu(A_t) dt) + dM_t^i \quad (5)$$

for a martingale M^i (strongly) orthogonal to σZ with $M_0^i = 0$ and a progressively measurable process β^i with $\mathbb{E}_{Q_T^A}[\int_0^T |\beta_t^i|^2 dt] < \infty$ for all $T \geq 0$.

Moreover, a behavior strategy profile A is a PPE if and only if $\beta = (\beta^1, \dots, \beta^n)^\top$ related to $W(A)$ by (5) enforces A .

DEFINITION 5. A set $\mathcal{W} \subseteq \mathbb{R}^n$ is *self-generating* for discount rate $r > 0$ if for every $w \in \mathcal{W}$ there exists a solution (W, A, β, Z, M) to (5) such that β enforces A , $W_0 = w$ a.s., and $W_\tau \in \mathcal{W}$ a.s. for every stopping time τ .

LEMMA 2. *The set $\mathcal{E}(r)$ is the largest bounded self-generating set.*

This result is the equivalent of Theorem 1 in Abreu et al. (1990). It follows from Lemma 1 that any self-generating set is contained in $\mathcal{E}(r)$. The idea for the proof of the converse is that a PPE is subgame perfect, hence $W_t \in \mathcal{E}(r)$ a.s. since the continuation game is equivalent to the repeated game. For a formal proof one needs to deal with some measurability issues that we discuss in Section 4.1.

2.3 Enforceability and identifiability

The folk theorem will follow from Lemma 2 once we find suitable conditions, under which any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$ is self-generating for a sufficiently small discount rate. This means that we need to construct enforceable strategy profiles whose continuation values do not escape \mathcal{W} . In this section we will motivate some necessary conditions for that to be possible. By examining (5), we see that on $\partial\mathcal{W}$, both of the following statements have to be satisfied:

⁶It is enough to consider deviations to pure strategies, since any behavior strategy has a realization equivalent mixed strategy by Theorem 4, and a deviation to a mixed strategy can only be profitable if it has at least one profitable pure strategy in its support.




FIGURE 1. At every point $w \in \partial\mathcal{W}$, the tangential diffusion $r\beta_t(\sigma dZ_t - \mu(\alpha) dt)$ leads to an outward-pointing drift of order $r^2 \text{tr}(\beta\sigma\sigma^\top \beta^\top) dt$. For sufficiently small r , this term is dominated by the inward-pointing drift $r(W_t - g(A_t)) dt$.

- (i) The drift points inward, that is, $N_w^\top(g(\alpha) - w) > 0$, where N_w is the outer-pointing normal vector at $w \in \partial\mathcal{W}$.
- (ii) The volatility is tangential to \mathcal{W} .

Indeed, otherwise the continuation value would immediately escape \mathcal{W} ; see also Figure 1. The first condition translates to enforceability of action profiles with extremal payoffs. The second condition is achieved using enforceability on hyperplanes and the related concept of orthogonal enforceability.

DEFINITION 6.

- (i) Let $T \in \mathbb{R}^{n \times (n-1)}$ be a matrix whose column vectors T_1, \dots, T_{n-1} span the hyperplane $H \subseteq \mathbb{R}^n$. An action profile α is *enforceable on the hyperplane H* if there exists a matrix $B \in \mathbb{R}^{(n-1) \times d}$ such that α is enforced by $\beta = TB$.
- (ii) Let N be a vector in \mathbb{R}^n . A matrix $\beta \in \mathbb{R}^{n \times d}$ *enforces α orthogonal to N* if it enforces α and satisfies $N^\top \beta = 0$.

Observe that the two notions of enforceability are equivalent, i.e., α is enforceable on a hyperplane H if and only if it is enforceable orthogonal to the normal vector N of H .⁷ While enforceability on hyperplanes has the nice interpretation of transferring future value among players, it is often easier to work with the related concept of orthogonal enforceability because the normal vector to the smooth hypersurface $\partial\mathcal{W}$ is unique. We distinguish two types of hyperplanes.

DEFINITION 7. A hyperplane H is said to be *coordinate* if it is orthogonal to a coordinate axis. A hyperplane H is *regular* if it is not coordinate.

For an enforceable action profile α , the additional requirement to be enforceable on a coordinate hyperplane means that the corresponding player does not make any

⁷Indeed, if $\beta = TB$, then $N^\top \beta = 0$. Conversely, if $N^\top \beta = 0$, then all column vectors β_j lie in H , which means they can be written as linear combinations of the T_j . This is equivalent to $\beta = TB$.

transfers. Indeed, the systems (4) have a solution with $\beta^i = 0$ if and only if α is a best response for player i . For ease of reference we state it as a lemma.

LEMMA 3. *An enforceable action profile α is enforceable on a hyperplane orthogonal to the i th coordinate axis if and only if α satisfies the best response property for player i , that is, $g^i(\alpha) \geq g^i(a^i, \alpha^{-i})$ for all $a^i \in \mathcal{A}^i$.*

For an action profile α to be enforceable on regular hyperplanes, players' impacts on the distribution of the public signal need to be sufficiently identifiable. Let $M^i(\alpha)$ denote the $(d \times |\mathcal{A}^i|)$ -dimensional matrix, whose column vectors $\mu(a^i, \alpha^{-i}) - \mu(\alpha)$, $a^i \in \mathcal{A}^i$ are given by the impact on the drift rate of the public signal that player i 's deviation from α^i to a^i has. Observe that $\text{rank } M^i(\alpha) \leq |\mathcal{A}^i| - 1$ since

$$\sum_{a^i \in \mathcal{A}^i} \alpha^i(a^i) (\mu(a^i, \alpha^{-i}) - \mu(\alpha)) = 0$$

by multilinearity of μ . Denote by $\Lambda^i(\alpha) := \text{span } M^i(\alpha)$ the space of all those possible impacts that player i 's deviations may have on the distribution of the public signal.

DEFINITION 8.

- (i) An action profile α has *individual full rank for player i* if $M^i(\alpha)$ has rank $|\mathcal{A}^i| - 1$. If this is true for every player $i = 1, \dots, n$, then α has *individual full rank*.
- (ii) An action profile α is said to have *pairwise full rank for players i and j* if the matrix $M^{ij}(\alpha) = [M^i(\alpha), M^j(\alpha)]$ has rank $|\mathcal{A}^i| + |\mathcal{A}^j| - 2$. An action profile has *pairwise full rank* if this is true for all pairs of players $j \neq i$.
- (iii) A mixed action profile α is *pairwise identifiable* if for any two players i and j ,

$$\Lambda^i(\alpha) \cap \Lambda^j(\alpha) = \{0\}.$$

Having individual full rank implies that the system of inequalities (4) can be solved with equality, thus any action profile with individual full rank is enforceable; see [Lemma 16](#) for details. Pairwise identifiability means that deviations of any two players lead to linearly independent impacts on the drift rate of the public signal. Therefore, deviations of any two players can be statistically distinguished. Finally, note that having pairwise full rank is equivalent to having individual full rank and pairwise identifiability.

The next result is the analogue of Lemma 5.5 in [Fudenberg et al. \(1994\)](#). The proof shows that under the assumption of pairwise identifiability, any two players' incentives are isolated by an orthogonal decomposition. The proof also sheds some light on the source of the additional assumption of the continuous-time folk theorem.

LEMMA 4. *Suppose that an enforceable action profile α is pairwise identifiable. Then it is enforceable on all regular hyperplanes.*

PROOF. We show that α is enforceable orthogonal to the normal vector N of the hyperplane. Because the hyperplane is regular, N has at least two nonzero entries and we will assume that these are the first two. Let $\beta \in \mathbb{R}^{n \times d}$ enforce α . Pairwise identifiability implies that $\mathbb{R}^d = (\Lambda^i(\alpha) \cap \Lambda^j(\alpha))^\perp = \Lambda^i(\alpha)^\perp + \Lambda^j(\alpha)^\perp$ for all $i \neq j$, hence we can decompose β^i into $\beta^i = \beta^{\perp,i} + \tilde{\beta}^i$, where $\beta^{\perp,i} \perp \Lambda^i(\alpha)$ for all i and

$$\tilde{\beta}^1 \perp \Lambda^2(\alpha) \quad \text{and} \quad \tilde{\beta}^i \perp \Lambda^1(\alpha), \quad i = 2, \dots, n.$$

Let $G^i(\alpha)$ denote the row vector of losses $g^i(\alpha) - g^i(a^i, \alpha^{-i})$ in player i 's expected flow payoff when he switches from α^i to a^i , so that α is enforceable if and only if the inequality $G^i(\alpha) \geq \beta^i M^i(\alpha)$ holds componentwise for every player $i = 1, \dots, n$. We construct a matrix $B = (B^1, \dots, B^n)^\top$ enforcing α on the hyperplane H by setting

$$B^1 = \tilde{\beta}^1 - \sum_{i=2}^n \frac{N^i}{N^1} \tilde{\beta}^i, \quad B^2 = \tilde{\beta}^2 - \frac{N^1}{N^2} \tilde{\beta}^1 \quad \text{and} \quad B^i = \tilde{\beta}^i, \quad i = 3, \dots, n. \quad (6)$$

Indeed, since $\beta^{\perp,1}$ and $\tilde{\beta}^2, \dots, \tilde{\beta}^n$ are orthogonal to $\Lambda^1(\alpha)$, it follows that

$$B^1 M^1(\alpha) = \tilde{\beta}^1 M^1(\alpha) - \sum_{i=2}^n \frac{N^i}{N^1} \tilde{\beta}^i M^1(\alpha) = (\beta^1 - \beta^{\perp,1}) M^1(\alpha) \leq G^1(\alpha).$$

The inequalities for players $i = 2, \dots, n$ are verified in the same manner. Finally, $N^\top B = 0$ by construction. \square

Lemmas 3 and 4 tell us how to construct incentives on any hyperplane as a function of the chosen action profile A_t and the normal vector N_{W_t} .⁸ It follows from Itô's formula applied to (5) that the tangential volatility leads to an outward-pointing drift of order $r^2 \text{tr}(\beta \sigma \sigma^\top \beta^\top)$; see also Figure 1. Therefore, for W to remain in \mathcal{W} it is necessary that β is bounded, so that the outward-pointing drift is dominated by the inward-pointing drift $r(W - g(A))$ for r small enough. As we can see from (6), this may be tricky where the tangent hyperplane changes from being regular to coordinate. The following lemma states various conditions such that β is locally bounded. Since the construction in (6) is bounded where it is Lipschitz continuous, and we need Lipschitz continuity later on, we assert both properties. Observe that this issue does not exist in discrete time, since transitions are discrete and hence attained payoffs are always either coordinate or bounded away from being coordinate.

LEMMA 5. *Let $N \in \mathbb{R}^n \setminus \{0\}$ and let α be an enforceable action profile. Suppose that one of the following conditions holds true:*

- (i) *The profile α is pairwise identifiable and N is not parallel to any coordinate axis.*
- (ii) *The profile α is pairwise identifiable and enforceable orthogonal to N .*

⁸To be precise, the normal vector at the point $\pi(W_t)$ for a suitable projection π of \mathbb{R}^n onto $\partial\mathcal{W}$.

- (iii) The profile α is enforceable orthogonal to e_i and α^i is a unique best response to α^{-i} , that is, $\alpha^i = a^i$ for some $a^i \in \mathcal{A}^i$ and $g^i(a^i, \alpha^{-i}) > g^i(\tilde{a}^i, \alpha^{-i})$ for every $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$.
- (iv) The profile α is a static Nash equilibrium.

Then there exist a neighborhood U_N of N and a bounded, Lipschitz continuous map $\beta_\alpha : U_N \rightarrow \mathbb{R}^{n \times d}$ such that $\beta_\alpha(x)$ enforces α orthogonal to x .

3. MAIN RESULTS AND DISCUSSION

3.1 The folk results

Let $\mathcal{A}^{(i)} \subseteq \mathcal{A}$ denote the pure action profiles that maximize player i 's payoff over \mathcal{A} . The continuous-time folk theorem differs from its discrete-time counterpart in the additional assumption that for every player i , one element of $\mathcal{A}^{(i)}$ and the minmax profile $\underline{\alpha}_i$ against player i are either pairwise identifiable or satisfy the unique best response property for player i .⁹ While the argument in the previous section leaves open the question whether there exists a different construction of incentives that is bounded without these assumptions, we show in Section 3.3 that some form of identifiability condition is necessary when the unique best response property fails.

THEOREM 1 (Minmax folk theorem). *Suppose that the following conditions hold:*

- (i) Every pure action profile has individual full rank.
- (ii) For every pair of players i and j , there exists an action profile α_{ij} with pairwise full rank for these players.
- (iii) For every player i , there exists an action profile $a_i^* \in \mathcal{A}^{(i)}$ that is either pairwise identifiable or satisfies the unique best response property for player i .
- (iv) For every player i , best responses to the minmax profile $\underline{\alpha}_i^{-i}$ are unique.

Then for any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.¹⁰

Observe that the second condition is satisfied if there exists at least one pairwise identifiable pure action profile because then it has pairwise full rank by condition (i). A sufficient condition for the folk theorem to hold is that all pure action profiles have pairwise full rank. Then conditions (i), (ii), and (iii) are clearly fulfilled and the fourth condition can be circumvented by footnote 9. The condition that all pure action profiles have pairwise full rank is also sufficient to establish a folk theorem in pure strategies; see the Supplement for details.

⁹To be precise, we need that an approximation of the minmax profile either is pairwise identifiable or has the unique best response property. While the unique best response property of the minmax profile carries over to approximations by linearity of the expectation, pairwise identifiability of the approximation requires pairwise full rank of all pure action profiles; see Lemma 18 for details.

¹⁰A set \mathcal{W} is called *smooth* if it is a closed and convex subset of \mathbb{R}^n with nonempty interior and a twice differentiable boundary.

If one is not interested in generating all payoffs in the interior of \mathcal{V}^* , but only payoffs that dominate the payoff of a stage-game Nash equilibrium α_e , we prove a weaker Nash-threat folk theorem similar to Theorem 6.1 in Fudenberg et al. (1994).

THEOREM 2 (Nash-threat folk theorem). *Let \mathcal{V}^0 be the convex hull of $g(\alpha_e)$ and the Pareto-efficient payoff vectors Pareto-dominating $g(\alpha_e)$. Suppose that either g is affine in m and every Pareto-efficient pure action profile is pairwise identifiable, or that following statements hold:*

- (i) *For every pair of players i and j , there exists at least one profile α_{ij} with pairwise full rank for that pair of players.*
- (ii) *For every player i , there exists an enforceable action profile $a_i^* \in \mathcal{A}^{(i)}$ that is either pairwise identifiable or satisfies the best response property for player i .*

Then for any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^0$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.

Because pairwise identifiability of action profiles is essential for the folk result to hold, it is worth mentioning a special class of games where this assumption is always satisfied. A game is said to be of a *product structure* if the impacts of players' deviations on the drift are orthogonal, that is, $\Lambda^i(a) \perp \Lambda^j(a)$ for all $i \neq j$ and all pure action profiles $a \in \mathcal{A}$. Clearly, this implies pairwise identifiability of all pure action profiles. Therefore, a Nash-threat folk theorem holds for any game, in which Pareto-efficient action profiles are enforceable.

COROLLARY 1. *Consider a game with a product structure such that g is an affine function of m . For any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}_0$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

Since pairwise identifiability and individual full rank are equivalent to having pairwise full rank, we obtain the minmax folk theorem for games with a product structure in the following form.

COROLLARY 2. *Suppose that in a game with a product structure, every pure action profile has individual full rank. Then for any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

The special case of these corollaries, where $n = 2$ players are restricted to pure strategies, has already been established in Sannikov (2007) as a consequence of his precise characterization of $\partial \mathcal{E}_p(r)$.

3.2 Monotonicity of $\mathcal{E}(r)$

In this class of continuously repeated games it is currently unknown whether the PPE payoff set is monotonic in the discount rate r without public randomization. Therefore,

$\mathcal{E}(r)$ may not expand monotonically to \mathcal{V}^* even when the folk theorems apply. We elaborate in this section on how players can use public randomization to achieve payoffs also under any smaller discount rate.

Previously, the drift function m and the volatility σ were fixed. In this section and [Appendix B](#) we will view the PPE payoff set $\mathcal{E}(r, m, \sigma)$ as an object depending on all three game primitives. The following square-root law tells us that scaling the discount rate by λ has the same effect on equilibrium payoffs as dividing the signal-to-noise ratio $\sigma^\top(\sigma\sigma^\top)^{-1}m$ by $\sqrt{\lambda}$. This result is similar in spirit to [Corollary 1](#) in [Faingold and Sannikov \(2011\)](#), where a square-root law is obtained for continuous-time games between one long-lived player and a continuum of short-lived players.

LEMMA 6. *Let $\tilde{m} : \mathcal{A} \rightarrow \mathbb{R}^d$ and $\tilde{\sigma} \in \mathbb{R}^{d \times k}$ be such that $\tilde{\sigma}\tilde{\sigma}^\top$ is invertible and $\tilde{\sigma}^\top(\tilde{\sigma}\tilde{\sigma}^\top)^{-1}\tilde{m} = \sqrt{\lambda}\sigma^\top(\sigma\sigma^\top)^{-1}m$ for some $\lambda > 0$. Then a strategy profile A is a PPE for the game primitives (r, m, σ) if and only if $(A_{\lambda t})_{t \geq 0}$ is a PPE with respect to the game primitives $(\lambda r, \tilde{m}, \tilde{\sigma})$. Moreover, for every $t \geq 0$,*

$$\tilde{W}_t((A_{\lambda t})_{t \geq 0}; \lambda r, \tilde{m}, \tilde{\sigma}) = W_{\lambda t}(A; r, m, \sigma) \quad \text{a.s.}, \quad (7)$$

where \tilde{W} is the discounted expected future payoff with respect to the time-changed filtration $(\mathcal{F}_{\lambda t})_{t \geq 0}$. In particular, the equilibrium payoff set depends on (r, m, σ) only through the ratio $\sigma^\top(\sigma\sigma^\top)^{-1}m/\sqrt{r}$.

For $\lambda = 1$, the result says that the continuation value depends on m and σ only through the informativeness of the signal, the signal-to-noise-ratio $\sigma^\top(\sigma\sigma^\top)^{-1}m$. This is very intuitive since the induced probability measure in (2) depends on m and σ only through that quantity. Therefore, such a transformation leads to the same distribution over possible signals. For $\lambda < 1$, [Lemma 6](#) says that players becoming more patient has the same effect as increasing the informativeness of the signal. As time becomes less valuable to the players, a longer interval of observations of the public signal becomes available at the same cost, hence players can better estimate m . Observe that we consider players as being more patient when their discount rate is lower simply because they value future payoffs more. In this class of games, being more patient can be taken very literally, by executing the same strategy profile at a slower speed. Similarly for $\lambda > 1$, players being less patient has the same effect on the game as a decrease in the informativeness of the signal. Observe that for any value of λ , the same action profiles are enforceable since \tilde{m} is an invertible linear transformation of m . Therefore, (4) is solvable for m if and only if it is solvable for \tilde{m} .

The quantity $\gamma = \sigma^\top(\sigma\sigma^\top)^{-1}m/\sqrt{r}$ is a measure of the informativeness of the public signal adjusted for the patience of players. It follows from (7) at time 0 that the equilibrium payoff set depends on the game primitives (r, m, σ) only through γ . This is similar to [Daley and Green \(2012\)](#), where the quality of news, a quantity corresponding to $\gamma^\top\gamma$, plays a central role in the equilibrium analysis. When the signal becomes more informative relative to players' discounting, that is, $\hat{\gamma} \geq \gamma$ in every component, the following result establishes that a PPE A for γ can be transformed to a PPE \hat{A} for $\hat{\gamma}$ by adding noise through public randomization.

THEOREM 3. *Let A be a PPE with respect to (r, m, σ) . Let $\hat{\mathbb{F}} = (\hat{\mathcal{F}}_t)_{t \geq 0}$ denote the filtration generated by the public filtration \mathbb{F} and the addition of a public randomization device. Denote by $\hat{W}_t(\cdot)$ the discounted expected future payoff conditional on $\hat{\mathcal{F}}_t$. Then by using the public randomization suitably, A can be transformed to a PPE \hat{A} for different game parameters in the following cases:¹¹*

- (a) *For any symmetric $\Lambda \in \mathbb{R}^{k \times k}$ with eigenvalues in $[-1, 1]$ such that $\ker(\sigma\Lambda) = \ker(\sigma)$, there exists a PPE \hat{A} , differing from A only through public randomization, with $\hat{W}_t(\hat{A}; r, m, \sigma\Lambda) = W_t(A; r, m, \sigma)$ a.s.*
- (b) *For any symmetric $\Lambda \in \mathbb{R}^{d \times d}$ with eigenvalues in $[-1, 1] \setminus \{0\}$ such that $\Lambda\sigma\sigma^\top = \sigma\sigma^\top\Lambda$, there exists a PPE \hat{A} , differing from A only through public randomization, with $\hat{W}_t(\hat{A}; r, \Lambda^{-1}m, \sigma) = \hat{W}_t(\hat{A}; r, m, \Lambda\sigma) = W_t(A; r, m, \sigma)$ a.s.*
- (c) *For any $\lambda \in (0, 1)$, there exists a PPE \hat{A} , differing from $(A_{\lambda t})_{t \geq 0}$ only through public randomization, with $\hat{W}_t(\hat{A}; \lambda r, m, \sigma) = W_{\lambda t}(A; r, m, \sigma)$ a.s.*

This is a remarkable result because public randomization allows one to transform PPE continuously over games with a higher quality of the signal, achieving exactly the same path of continuation values. An equivalent result does not hold in discrete time.¹² Returning to the question of how to achieve equilibrium payoffs also under a smaller discount rate, we see in (c) that this is done by performing a time change and then adding public randomization to the time-changed strategy profile. Monotonicity of $\mathcal{E}(r)$ in r is thus a direct corollary to (c) in [Theorem 3](#).

COROLLARY 3. *If players have access to a public randomization device, then it follows that $\mathcal{E}(r) \subseteq \mathcal{E}(r')$ for any $0 < r' < r$.*

3.3 Motivation for the proof of the folk theorem

In this section we highlight some of the differences and similarities between the proofs of the folk theorem in discrete and continuous time. In this motivation, some technical details are omitted; see [Section 4](#) for the full proof. In both discrete and continuous

¹¹By adding public randomization to a strategy profile A , we mean that disregarding the information of the public randomization device in \hat{A} is identical to playing A . Formally, ${}^\circ\hat{A} = A$, where ${}^\circ(\cdot)$ is the optional projection onto the public filtration \mathbb{F} without public randomization. See [Appendix A](#) for a short introduction to the optional projection.

¹²Even though an equivalent result to our [Theorem 3](#) does not hold in discrete time, one can still observe the time change in (c). Section 6 of [Abreu et al. \(1990\)](#) or, equivalently, equation (4.2) in [Fudenberg et al. \(1994\)](#), shows how the continuation value u of the public signal Y has to be adjusted to decompose the same payoff v when the discount factor changes from δ_1 to $\delta_2 \in (\delta_1, 1)$. Namely,

$$u(Y; \delta_2) = (1 - \lambda)v + \lambda u(Y; \delta_1) = (1 - \lambda)(1 - \delta_1)g(a) + (\lambda + (1 - \lambda)\delta_1)u(Y; \delta_1),$$

where $\lambda = \delta_1(1 - \delta_2)/(\delta_2(1 - \delta_1)) \in (0, 1)$. Observe that $u(Y; \delta_2)$ is a convex combination of $u(Y; \delta_1)$ and the current payoff $g(a)$. This can be interpreted that time moves more slowly because the current-period payoff lags into the continuation payoff. At the same time, the current-period action profile remains unchanged, which is consistent with a slower play of the same strategy profile.




FIGURE 2. The left panel shows the decomposition of the payoff $w = \delta g(\alpha) + (1 - \delta)\mathbb{E}[u(Y)]$ on the boundary $\partial\mathcal{W}$ in discrete time. The continuation payoff $u(Y)$ lies on a hyperplane parallel to the tangent hyperplane \mathcal{S}_w for all realizations y of Y . The right panel shows that payoffs close to w are decomposable with respect to the same hyperplane.

time, the folk theorem is established by showing that any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$ is self-generating for a sufficiently small discount rate r and therefore has to be a subset of the PPE payoff set by Lemma 2. To show that \mathcal{W} is self-generating, we construct equilibrium strategies with continuation values in \mathcal{W} in the following steps:

Step 1. For any payoff $w \in \mathcal{W}$, there exists an enforceable strategy profile A^w with initial payoff w and continuation value in \mathcal{W} for a short but positive amount of time τ_w . The discount rate r_w may depend on the payoff w .

Step 2. There exist a neighborhood U_w of w and $\tilde{r} > 0$ such that for any $r \in (0, \tilde{r})$, the time τ and the strategy profile A can be chosen uniformly across U_w .

Step 3. By compactness we can concatenate these local solutions to a global solution.

In both discrete and continuous time, payoffs in the interior of \mathcal{W} are attainable by a static Nash equilibrium for a sufficiently low discount rate. We will thus focus on payoffs $w \in \partial\mathcal{W}$ in this motivation.

Step 1. In discrete time, Fudenberg et al. (1994) decompose any payoff $w \in \partial\mathcal{W}$ into a current-period payoff $g(\alpha)$ outside of \mathcal{W} and a continuation promise in the interior of \mathcal{W} . In their decomposition, the continuation promise is parallel to the tangent hyperplane \mathcal{S}_w at w , that is, α is enforceable on \mathcal{S}_w ; see also Figure 2. A payoff set \mathcal{W} is called *decomposable on tangent hyperplanes* if such a decomposition is possible for any $w \in \partial\mathcal{W}$, which is sufficient for \mathcal{W} to be self-generating in discrete time.

As we mentioned in Section 2.3, in continuous time it is not enough that action profile α is enforceable on \mathcal{S}_w only. Instead, α has to be enforceable on all nearby hyperplanes so that the movement of the continuation value can be continuously adjusted to follow $\partial\mathcal{W}$. For self-generation in continuous time, a payoff set \mathcal{W} has to be *uniformly decomposable on tangent hyperplanes*, that is, for any $w \in \partial\mathcal{W}$, there exists an enforceable action profile α with $g(\alpha)$ strictly separated from \mathcal{W} by \mathcal{S}_w so that (α, N_w) satisfies one of the four conditions of Lemma 5, where N_w is the unique outer-pointing normal vector to $\partial\mathcal{W}$ at w . Then $w \mapsto \beta_\alpha^w$ is locally Lipschitz continuous on $\partial\mathcal{W}$, i.e., if an action profile is enforceable on a given hyperplane, then it can be enforced on nearby hyperplanes without changing the volatility significantly. Indeed, since \mathcal{W} is smooth, $w \mapsto N_w$

is Lipschitz continuous, hence the concatenation with the locally Lipschitz continuous map from Lemma 5 is Lipschitz continuous on a suitable neighborhood U_w of w . It is for this difference that the continuous-time folk theorem has the additional conditions (iii) and (iv).¹³ The stopping time τ_w can be chosen as the time when the conditions of Lemma 5 are no longer satisfied, i.e., $\tau_w = \inf\{t \geq 0 | W_t^w \notin B_\varepsilon(w)\}$ for a suitable $\varepsilon > 0$.

Step 2. Because continuation payoffs are strictly separated from the hyperplane S_w in discrete time, it is possible to decompose payoffs v close enough to w on a translate to S_w by moving the continuations by a small (and constant) amount; see also Figure 2. Because the variance of the continuation payoff is decreasing in δ , any payoff decomposable for $\tilde{\delta}$ is also decomposable for $\delta \in (\tilde{\delta}, 1)$ by convexity of \mathcal{W} .

In continuous time, payoffs $v \in U_w$ close to w can be decomposed by the same action profile α and Lipschitz continuous function $v \mapsto \beta_\alpha^v$ for sufficiently small \tilde{r} by uniform decomposability. The main difficulty in this step is to ensure that the stopping times τ_v are uniformly bounded from below by a strictly positive stopping time, that is, $\inf_{v \in U_w} \tau_v > 0$ a.s. For this step it is important that the SDE (5) locally admits a strong solution so that the probability space Ω and the Brownian motion Z can be fixed on the entire neighborhood U_w . We show that the SDE is “sufficiently nice” so that solutions W^v to (5) with initial value $v \in U_w$ have the continuous flow property, that is, $v \mapsto W^v$ is continuous for almost every $\omega \in \Omega$. This implies that $\inf_{v \in U_w} \tau_v > 0$ a.s. if U_w is sufficiently small; see the proof of Lemma 9 for details.

Step 3. The neighborhoods in Step 2 form an open cover of \mathcal{W} , hence by compactness, there exists a finite subcover $\mathcal{U} = \{U_1, \dots, U_N\}$. It follows that $r_{\mathcal{U}} := \min_{U \in \mathcal{U}} \tilde{r}_U$ is strictly positive. In discrete time, the action profiles found in the one-period decomposition are concatenated at every step to form an enforceable strategy profile with continuations that remain in \mathcal{W} forever.

In continuous time, we also take the minimum over the stopping times, i.e., we set $\tau_{\mathcal{U}} := \min_{U \in \mathcal{U}} \tau_U$, which is strictly positive almost surely. For this step it is crucial again that (5) locally admits strong solutions, so that we have to deal with only finitely many

¹³In some games, slightly weaker conditions for enforceability on coordinate hyperplanes may be sufficient. However, if α does not have the unique best response property for player i , then some sort of identifiability conditions has to be satisfied for all deviations $a^i \in A^i$ with $g^i(a^i, \alpha^{-i}) = g^i(\alpha)$ to ensure enforceability on hyperplanes infinitesimally close to being coordinate. Indeed, let $j \neq i$ be a player for whom α^j is not a best response to α^{-j} . Any neighborhood of e_i contains vectors $N(\varepsilon) := \varepsilon e_j + \sqrt{1 - \varepsilon^2} e_i$ for ε arbitrarily close to zero. From $N(\varepsilon)^\top \beta(\varepsilon) = 0$ it follows that

$$\beta^i(\varepsilon) = \frac{-\varepsilon}{\sqrt{1 - \varepsilon^2}} \beta^j(\varepsilon).$$

For $a^j \in A^j$ with $g^j(a^j, \alpha^{-j}) - g^j(\alpha) = \delta > 0$, enforceability implies $\beta^j(\varepsilon)(\mu(\alpha) - \mu(a^j, \alpha^{-j})) \geq \delta$ and hence $\beta^j(\varepsilon)$ is not orthogonal to $\mu(\alpha) - \mu(a^j, \alpha^{-j})$ for any ε close enough to zero. The enforceability condition for player i , imposes

$$0 \leq \varepsilon \beta^j(\varepsilon)(\mu(a^i, \alpha^{-i}) - \mu(\alpha)).$$

It follows that either the right hand side is identically zero, that is, $\beta^j(\varepsilon)$ is orthogonal to $\mu(a^i, \alpha^{-i}) - \mu(\alpha)$, or $\beta^j(\varepsilon)(\mu(a^i, \alpha^{-i}) - \mu(\alpha))$ takes opposite signs for positive and negative ε , respectively. Both imply that $\mu(a^i, \alpha^{-i}) - \mu(\alpha)$ is linearly independent of $\mu(a^j, \alpha^{-j}) - \mu(\alpha)$. That is, some sort of identifiability condition has to be satisfied.

probability spaces and filtrations. These can be enlarged at the beginning of the concatenation to contain all the necessary information. Therefore, $\tau_{\mathcal{U}}$ is indeed a stopping time, i.e., is measurable with respect to the enlarged filtration. Observe that $\tau_{\mathcal{U}}$ depends on the specific subcover \mathcal{U} that is chosen. By construction, the stopping time $\tau_{\mathcal{U}}$ is a functional of the public signal such that the continuation value W^v does not escape \mathcal{W} on $[0, \tau_{\mathcal{U}})$ regardless of the starting value v . Since $W_{\tau_{\mathcal{U}}}^v \in \mathcal{W}$, we use the same procedure to find a solution to (5) on $[\tau_{\mathcal{U}}, \tau_{\mathcal{U},2})$ starting at $W_{\tau_{\mathcal{U}}}^v$, where $\tau_{\mathcal{U},2} - \tau_{\mathcal{U}}$ is independent of $\tau_{\mathcal{U}}$ and identically distributed as $\tau_{\mathcal{U}}$ by independence and stationarity of the increments of Brownian motion.

An iteration of this procedure leads to a sequence of stopping times $(\tau_{\mathcal{U},\ell})_{\ell \geq 0}$ and solutions $(W^\ell, A^\ell, \beta^\ell, Z^\ell)_{\ell \geq 0}$ to (5) on $[\tau_{\mathcal{U},\ell}, \tau_{\mathcal{U},\ell+1})$ such that $\tau_{\mathcal{U},\ell+1} - \tau_{\mathcal{U},\ell}$ are independent and identically distributed (i.i.d.) as $\tau_{\mathcal{U}}$.¹⁴ This implies $\tau_{\mathcal{U},\ell} \rightarrow \infty$ a.s. Indeed, any sequence $(\tau_{\mathcal{U},\ell})_{\ell \geq 0}$ of random variables with strictly positive and i.i.d. increments $\tau_{\mathcal{U},\ell+1} - \tau_{\mathcal{U},\ell}$ diverges to ∞ a.s. by the strong law of large numbers; see Lemma 11 for details. Therefore, a countable concatenation of the solutions $(W^\ell, A^\ell, \beta^\ell, Z^\ell)$ to (5) yields a solution on $[0, \infty)$. This shows that \mathcal{W} is self-generating and hence $\mathcal{W} \subseteq \mathcal{E}(r)$. Observe that the global solution is only a weak solution to (5) because the underlying Brownian motion was also concatenated at $\tau_{\mathcal{U},\ell}$, depending on what element of the cover $W_{\tau_{\mathcal{U},\ell}}$ fell into; hence the Brownian motion cannot be fixed a priori. We elaborate on some technical difficulties that arise with weak solutions in Section 4.1.

3.4 Finite-variation property of equilibrium profiles

Because the constructed equilibrium profiles are concatenations of locally constant strategy profiles at i.i.d. copies of a positive stopping time $\tau_{\mathcal{U}}$, the resulting equilibrium profiles exhibit finitely many changes on every finite time interval. This is a very desirable feature for implementation because it seems unrealistic that agents can adapt their strategy profiles arbitrarily often. In this section, we present an example of such a strategy profile and compare it to the techniques used in Sannikov (2007). Consider the two-player partnership example of Section 2 in Sannikov (2007), reproduced in Figure 3 for the sake of exposition. To illustrate that the finite-variation property does not depend on the players' ability of mixing, we restrict the example to pure strategies.

Figure 3 shows a possible cover for a smooth payoff set \mathcal{W} in the interior of \mathcal{V}^* , such that on each element of the cover, the SDE (5) admits a strong solution. To ensure that the stopping times can be chosen strictly positive uniformly on each element of the cover, payoffs in a band of width ε around the element of the cover need to be decomposable with respect to the same pure action profile. The strategy profile is changed only when the continuation value leaves this band, ensuring that the strategy profile remains constant for a small but positive amount of time.¹⁵

¹⁴We omit mentioning M^ℓ as a part of the solution because $M^\ell \equiv 0$ in a strong solution to (5).

¹⁵The stopping times $\tau_{\mathcal{U},\ell}$ constructed in the proof are functionals of the signal such that the continuation value would not escape the band of width ε if it were to start at any point of *any* neighborhood. For all practical purposes (including this example), one may think of these times as the times $\tilde{\tau}_\ell$ at which the continuation value leaves the band of width ε around $U_{W_{\tilde{\tau}_\ell-1}}$. Because $\tilde{\tau}_\ell \geq \tau_{\mathcal{U},\ell}$ a.s. for any ℓ , the concatenation still extends to ∞ .

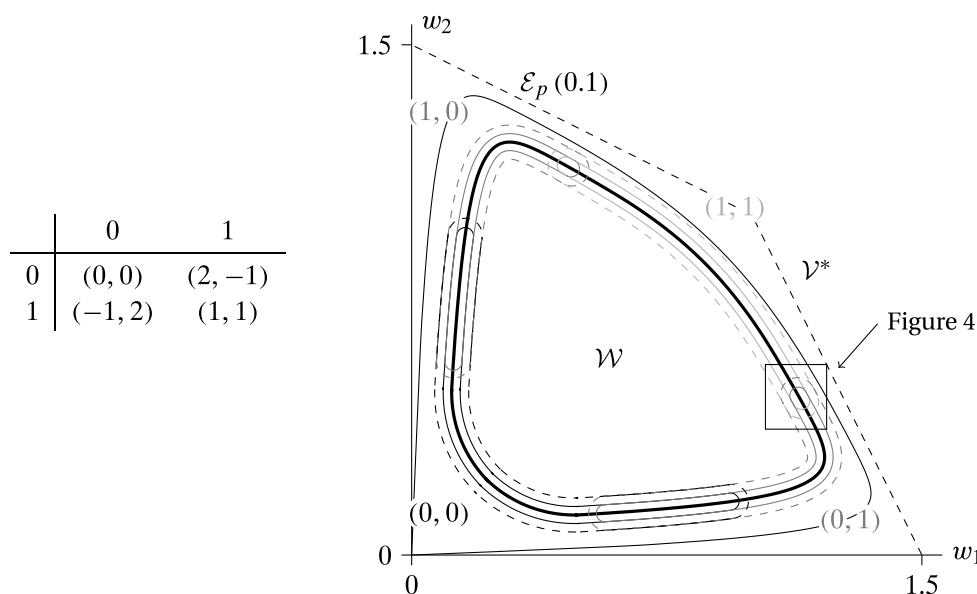


FIGURE 3. The matrix of static payoffs (w_1, w_2) is shown to the left, and the right panel shows a cover of $\partial\mathcal{W}$ (bold black line) into four overlapping sets (solid lines in gray scale), such that payoffs in a band of width ε around the sets (dashed lines in gray scale) can be decomposed with respect to the same pure action profile for discount rate $r = 0.1$. The cover of \mathcal{W} is completed by playing the static Nash equilibrium in the interior of \mathcal{W} . Also depicted is $\partial\mathcal{E}_p(0.1)$ (thin black line) constructed with the techniques in Sannikov (2007).

In comparison, the construction of equilibrium profiles in Sannikov (2007) works even on the boundary $\partial\mathcal{E}_p(r)$, where the constructed strategy profiles are constant up to a finite number of “switching points.” However, due to unbounded variation of Brownian motion, the players will switch between action profiles an infinite number of times during a finite time interval when the continuation value crosses a switching point; see also Figure 4. While our approach of constructing equilibrium profiles is more general in the sense that it is applicable to any finite number of players, this example shows that it can have advantages even in two-player games.

If players are not restricted to pure strategies, the realizations of their strategies are drawn continuously. Therefore, players switch actions infinitely often on finite time intervals even for constant (but mixed) strategy profiles. However, mixing is done individually for each player, and because of the multilinearity in (1), the public signal is not affected by the different realizations of a player’s mixed action as long as his strategy profile remains constant. The strategies of a player’s opponents are therefore not affected by the realizations of his mixed strategy; hence a change of actions within a constant strategy profile is a less complicated operation than a change of strategy profile. Moreover, because m and g are extended to mixed action profiles by multilinearity, continuous-time mixing may be interpreted as a division of effort among the pure actions in its support. This is a common formulation in continuous-time games of strategic experimentation; see, for example, Bolton and Harris (1999) and Keller and Rady (2010).




FIGURE 4. The left panel shows the simulation of the continuation value of a PPE in a zoom-in of Figure 3. Lines in light gray, dark gray, and black mean that action profiles (1, 1), (0, 1), and (0, 0), respectively, are played; see Figure O.1 in the Supplement for a colored version. When the continuation value leaves the band around the cover of $\partial\mathcal{W}$, the static Nash equilibrium is played until the boundary of \mathcal{W} is reached. The upper right panel shows the corresponding strategy profile. The lower right panel shows a strategy profile constructed with the techniques in Sannikov (2007) with unbounded variation when the continuation value crosses the switching point S (left panel).

4. PROOFS OF THE MAIN RESULTS

4.1 Weak definition of $\mathcal{E}(r)$ and self-generation

In Section 2.1 we introduced our main object of study, the set $\mathcal{E}(r)$ of payoffs that are achievable by public perfect equilibria. For the proof of our results it is necessary to elaborate on what it means for a payoff $x \in \mathcal{V}$ to be achieved by some PPE A . As we outlined in Section 3.3, we will construct weak solutions to the SDE (5) achieving x . This means that the public signal, the filtrations, and the whole probability space may depend on x . We arrive at the specification

$$\mathcal{E}(r) := \left\{ x \in \mathcal{V} \left| \begin{array}{l} \text{there exists } (\Omega, \mathcal{F}, \mathbb{F}, P) \text{ containing an } (\mathbb{F}, P) \text{ Brownian} \\ \text{motion } Z \text{ and a PPE } A \text{ with } W_0(A) = x \text{ } P\text{-a.s.} \end{array} \right. \right\}.$$

We will also call $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$ a stochastic framework for A .

REMARK 2. Technically, this weak definition is necessary to ensure existence of the solutions. But even aside from the technical advantages, the weak solution concept is appropriate here. From an interpretation standpoint, the difference between a strong solution and a weak solution to an SDE lies in the causality of the noise. If the noise is defined exogenously and not affected by players' actions, this corresponds to a strong solution. In games of imperfect information, however, the noise is induced by players' strategies and thus cannot be fixed at the beginning. This is in line with discrete-time games, where we only care about the distribution of the public signal and not on what probability space the distribution is realized.

As illustrated in Section 3.3, we will piece together local solutions to obtain a global solution. At any stopping time τ , the value $W_\tau(A)$ is a random variable. Because of the weak formulation, the probability space depends on the point $x \in \mathcal{E}(r)$, and hence it is not clear what measurability conditions a random variable in $\mathcal{E}(r)$ should satisfy. This is clarified by the following lemma, whose rather technical proof is contained in Appendix C.

LEMMA 7. *Let X be an \mathcal{F}_0 -measurable random variable in a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$. Then the following statements are equivalent:*

- (a) *We have $X \in \mathcal{E}(r)$ a.s.*
- (b) *There exists a PPE A with $W_0(A) = X$ a.s.*

This means we can only achieve random variables by a PPE on a fixed probability space. At first glance, this might seem like a major restriction because we are dealing with weak solutions. However, we will need this result only to concatenate locally strong solutions to a global solution, at which point we only have finitely many probability spaces by compactness. These probability spaces can be enlarged at the beginning of the concatenation, after which it remains fixed. From Lemmas 1 and 7 and we obtain the following stochastic characterization of $\mathcal{E}(r)$.

LEMMA 8. *The following statements are equivalent for an \mathcal{F}_0 -measurable random variable X in a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$:*

- (a) *We have $X \in \mathcal{E}(r)$ a.s.*
- (b) *There exists a strategy profile A , a square-integrable progressively measurable process β , a martingale M orthogonal to σZ , and a bounded semimartingale W such that β enforces A , $W_0 = X$ a.s., and A, β, W, Z , and M satisfy (5).*

We conclude this section with the proof of self-generation. The argument sheds some first insight into the necessity of weakly defining $\mathcal{E}(r)$.

PROOF OF LEMMA 2. By Lemma 1, any bounded self-generating set \mathcal{W} is contained in $\mathcal{E}(r)$. Since $\mathcal{E}(r)$ is bounded, it remains to show that $\mathcal{E}(r)$ is self-generating. Take $x \in \mathcal{E}(r)$ so that Lemma 1 yields the existence of a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$, a behavior strategy profile A enforced by β , a martingale M orthogonal to σZ , and a bounded semimartingale W satisfying (5) with $W_0 = x$ a.s. We now fix a stopping time τ and show that $W_\tau \in \mathcal{E}(r)$ a.s. To do so, we set $\tilde{X} = W_\tau$, $\tilde{\mathcal{F}}_t = \mathcal{F}_{t+\tau}$, $\tilde{Z}_t = Z_{\tau+t} - Z_\tau$, $\tilde{M}_t = M_{\tau+t} - M_\tau$, $\tilde{W}_t = W_{\tau+t}$, $\tilde{\beta}_t = \beta_{\tau+t}$, and $\tilde{A}_t = A_{\tau+t}$. Because the tilde processes and filtrations satisfy condition (b) in Lemma 8, we obtain that $W_\tau = \tilde{X} \in \mathcal{E}(r)$ a.s. \square

4.2 Construction of continuous-time equilibria

In Sections 2.3 and 3.3, we motivated the condition of uniform decomposability on tangent hyperplanes for a payoff set \mathcal{W} to be self-generating. In this section we will show

that this condition is also sufficient. In a first step, we show that any uniformly decomposable set \mathcal{W} is locally self-generating.

DEFINITION 9. A set $\mathcal{W} \subseteq \mathbb{R}^n$ is called *locally self-generating* if for every point $w \in \mathcal{W}$, there exist an open neighborhood U_w of w , a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$, an enforceable strategy profile A , a martingale M orthogonal to σZ , and $\tilde{r} > 0$ such that for every discount rate $r \in (0, \tilde{r})$ there exists a stopping time $\tau > 0$ such that for all $v \in U_w \cap \mathcal{W}$, there exist W^v with $W_0^v = v$ a.s. and β^v such that $v \mapsto W^v$ and $v \mapsto \beta^v$ are Borel measurable and on $0 \leq t \leq \tau$, the processes W^v, β^v, A, M, Z are related by (5), β_t^v enforces A_t a.s., and $W_t^v \in \mathcal{W}$ a.s.

LEMMA 9. *Suppose that a smooth set $\mathcal{W} \subseteq \mathcal{V}^*$ is uniformly decomposable on tangent hyperplanes. Then \mathcal{W} is locally self-generating.*

PROOF. Suppose first that w is in the interior of \mathcal{W} . Let $U_w = B_\varepsilon(w)$, where $\varepsilon > 0$ is chosen such that the open ball $B_{2\varepsilon}(w)$ is contained in \mathcal{W} . For a static Nash equilibrium α_e , the constant strategy profile $A \equiv \alpha_e$ is enforced by $\beta \equiv 0$. For any $r > 0$ and any $v \in U_w$, let W^v be a strong solution to

$$dW_t^v := r(W_t^v - g(\alpha_e)) dt$$

with initial condition $W_0^v = v$ a.s. The explicit solution $W_t^v = v + (e^{rt} - 1)(v - g(\alpha_e))$ is measurable in v . Let $t_0(r) := \log(1 + \varepsilon/(\|w - g(\alpha_e)\| + \varepsilon))/r$. Then for any $v \in B_\varepsilon(w)$ it follows that

$$\|W_{t_0}^v - v\| \leq (e^{rt_0} - 1)(\varepsilon + \|w - g(\alpha_e)\|) \leq \varepsilon$$

on $[0, t_0]$. Therefore, in the interior of \mathcal{W} , we can support any discount rate $r > 0$ by choosing $U_w = B_\varepsilon(w)$ and the deterministic time $\tau \equiv t_0(r) > 0$.

For any $w \in \partial\mathcal{W}$, denote by N_w the outward unit normal to $\partial\mathcal{W}$ in w and denote by S_w the tangent hyperplane to $\partial\mathcal{W}$ in w . By smoothness of \mathcal{W} , both of these are unique and continuous in $w \in \partial\mathcal{W}$. Fix now a payoff $w \in \partial\mathcal{W}$. It will be convenient to work in a coordinate system with origin in w and a basis consisting of an orthonormal basis of S_w and N_w , where we choose the n th coordinate in the direction of N_w .

Since $\partial\mathcal{W}$ is a C^2 submanifold, we can locally parametrize it by a twice differentiable function φ . Let $\hat{x} = (x_1, \dots, x_{n-1})$ denote the projection onto the first $n - 1$ components so that the boundary $\partial\mathcal{W}$ is locally given by $(\hat{x}, \varphi(\hat{x}))$. By assumption, there exists an enforceable action profile α such that $g(\alpha)$ is strictly separated from \mathcal{W} by S_w . Let β_α be the locally Lipschitz continuous function from Lemma 5, which assigns to any vector $x \in \mathbb{R}^n$ a matrix β enforcing α orthogonal to x . Choose an $\varepsilon > 0$ such that the following statements hold:

- (i) We have $N_v^\top N_w > 0$ for all $v \in B_{2\varepsilon}(w) \cap \partial\mathcal{W}$.
- (ii) For all $v \in B_{2\varepsilon}(w)$, $\|\nabla\varphi(\hat{v})\| \leq p_1$ and $|\Delta_{ij}\varphi(\hat{v})| \leq p_2$ for $i, j = 1, \dots, n$ and constants $p_1, p_2 > 0$, where $\Delta_{ij}\varphi$ denotes the second partial derivative of φ with respect to \hat{v}_i and \hat{v}_j .

- (iii) The map $\beta_\alpha(-\nabla\varphi(\hat{\cdot}), 1)$ is Lipschitz continuous on $B_{2\varepsilon}(w)$ and $|(\beta_\alpha\sigma\sigma^\top\beta_\alpha^\top)_{ij}| \leq B$ for $i, j = 1, \dots, n$ and a constant $B < \infty$.
- (iv) We have $c(\varepsilon) := N_w^\top(g(\alpha) - w) - 2\varepsilon(1 + p_1) - p_1\|g(\alpha) - w\| > 0$.

The first condition makes sure that a local parametrization φ exists with bounded gradient as in condition (ii). Since $\partial\mathcal{W}$ is assumed to be C^2 , the first two derivatives of φ are continuous, hence locally bounded. In particular, by letting ε small enough we get the first two conditions to hold. For the third condition, observe that φ is continuous with bounded derivative by condition (ii), hence is Lipschitz continuous. Since the projection $\hat{\cdot}$ is Lipschitz continuous with Lipschitz constant 1 and the composition of Lipschitz continuous functions is Lipschitz again, the third condition holds in a small neighborhood of $N_w = (-\nabla\varphi(\hat{w}), 1)$ by Lemma 5. Finally, $N_w^\top(g(\alpha) - w)$ in condition (iv) is positive by strict separation of $g(\alpha)$ from \mathcal{W} . Because $\nabla\varphi$ is continuous and $\nabla\varphi(\hat{w}) = 0$, p_1 can be made arbitrarily small by choosing a small ε . This implies that $c(\varepsilon) > 0$ for sufficiently small ε .

Fix a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$ and an ε satisfying all of the above conditions. Denote $U_w := B_\varepsilon(w)$, fix a discount rate $r \leq 2c(\varepsilon)/((n - 1)^2 p_2 B) =: \tilde{r}$, and let $A \equiv \alpha$. For all $v \in B_\varepsilon(w)$, let W^v denote the strong solution to

$$dW_t^v = r(W_t^v - g(\alpha)) dt + r\beta_\alpha(-\nabla\varphi(\hat{W}_t^v), 1)(\sigma dZ_t - \mu(\alpha) dt)$$

on $\llbracket 0, \tau_v \rrbracket$ with $W_0 = v$ a.s., where $\tau_v := \inf\{t > 0 | W_t^v \notin B_{2\varepsilon}(w)\}$. Using that $\beta_\alpha \circ (-\nabla\varphi(\hat{\cdot}), 1)$ is uniformly bounded and Lipschitz continuous on $B_{2\varepsilon}(w)$ by condition (iii), a strong solution to this SDE exists by Theorem 5.2.1 of Øksendal (1998).¹⁶ Note that the process $\beta_t := \beta_\alpha(-\nabla\varphi(\hat{W}_t^v), 1)$ is progressively measurable on $\llbracket 0, \tau_v \rrbracket$ as a concatenation of a progressively measurable process with a Borel measurable function. Moreover, it enforces A and it is bounded on $\llbracket 0, \tau_v \rrbracket$ by Lemma 5, hence is locally square integrable.

Let $D_t^v := W_t^{v,n} - \varphi(\hat{W}_t^v)$ measure the distance from W_t to $\partial\mathcal{W}$ in the direction of N_w as shown in Figure 5. By Itô's formula,

$$\begin{aligned} dD_t^v &= (-\nabla\varphi(\hat{W}_t^v), 1)^\top dW_t^v - \frac{1}{2} \sum_{i,j=1}^{n-1} \frac{\partial^2\varphi(\hat{W}_t^v)}{\partial x_i \partial x_j} d\langle W^{v,i}, W^{v,j} \rangle_t \\ &= r \left((-\nabla\varphi(\hat{W}_t^v), 1)^\top (W_t^v - g(\alpha)) - \frac{r}{2} \sum_{i,j=1}^{n-1} \frac{\partial^2\varphi(\hat{W}_t^v)}{\partial x_i \partial x_j} (\beta_t \sigma \sigma^\top \beta_t^\top)_{ij} \right) dt \\ &\quad + r(-\nabla\varphi(\hat{W}_t^v), 1)^\top \beta_\alpha(-\nabla\varphi(\hat{W}_t^v), 1)(\sigma dZ_t - \mu(\alpha) dt) \\ &\leq r \left(\frac{r}{2} (n - 1)^2 p_2 B - c(\varepsilon) \right) dt, \end{aligned}$$

¹⁶Uniform boundedness implies the linear growth condition needed for the existence result. The function $f(x) = r(x - g(\alpha))$ is linear, hence Lipschitz continuous and of linear growth.




FIGURE 5. The distance of W_t^v to $\partial\mathcal{W}$ is measured by D_t^v in the direction of N_w for all v in a neighborhood of w . The continuation values W_t^v remain in \mathcal{W} if and only if $D_t^v \leq 0$. The sensitivity β_α of W^v to Z is chosen orthogonal to $(-\nabla\varphi(\hat{W}_t^v), 1)$, the normal vector to $\partial\mathcal{W}$ in the projection $(\hat{W}_t^v, \varphi(\hat{W}_t^v))$ of W_t^v onto $\partial\mathcal{W}$.

where we used that $x^\top \beta_\alpha(x) = 0$ for all $x \in B_{2\varepsilon}(w)$ and that conditions (ii) and (iv) imply

$$\begin{aligned} & ((-\nabla\varphi(\hat{W}_t^v), 0) + N_w)^\top (W_t^v - w + w - g(\alpha)) \\ & \leq \|\nabla\varphi(\hat{W}_t^v)\| (2\varepsilon + \|g(\alpha) - w\|) + 2\varepsilon - N_w^\top (g(\alpha) - w) \leq -c(\varepsilon). \end{aligned}$$

This implies that for any $r \in (0, \tilde{r})$, D^v is absolutely continuous with $dD_t^v/dt \leq 0$ on $\llbracket 0, \tau_v \rrbracket$, where τ_v depends on r . Since $D_0^v \leq 0$ for all $v \in U_w \cap \mathcal{W}$, it follows that $D_t^v \leq 0$ on $\llbracket 0, \tau_v \rrbracket$. Next, we show that the stopping times τ_v are uniformly bounded from below by a stopping time $\tau > 0$. The idea is that this SDE is sufficiently nice such that the flow $v \mapsto W^v$ is continuous and thus W^v can be approximated by $W^{\bar{v}}$ for \bar{v} close to v . This leads to a cover of $B_\varepsilon(w)$ with a finite subcover, over which the minimum of stopping times is still positive. Denote $V_t^v := e^{-rt}(W_t^v - v)$ and derive from the product rule that it is the solution to the SDE

$$dV_t^v = re^{-rt}(v - g(\alpha)) dt + re^{-rt} \beta_\alpha(-\nabla\varphi(v + \widehat{e^{rt}V_t^v}), 1)(\sigma dZ_t - \mu(\alpha) dt).$$

Fix a time horizon $T > 0$ to make e^{rt} bounded and Lipschitz continuous. To apply Theorem V.37 of Protter (2005), we write V^v in its integrated form

$$V_t^v = r(1 - e^{-rt})(v - g(\alpha)) + \int_0^t F(V^v)_s (\sigma dZ_s - \mu(\alpha) ds), \quad t \leq T,$$

where $F(V^v)_s = re^{-rs} \beta_\alpha(-\nabla\varphi(v + \widehat{e^{rs}V_s^v}), 1)$. Both the finite variation part and F are Lipschitz,¹⁷ hence Theorem V.37 of Protter (2005) applies and we deduce that the flow

¹⁷The result requires that F is functional Lipschitz, which is satisfied for any operator induced by a Lipschitz function; see Protter (2005, p. 251).

$v \mapsto V^v(\omega)$ is continuous for almost all ω .¹⁸ Let $\sigma_v := \inf\{t > 0 \mid e^{rt} \|V_t^v\| > \varepsilon/2\}$, which is strictly positive by continuity of V^v . For any $v, \bar{v} \in \overline{B_\varepsilon(w)}$, define the auxiliary process $\Lambda_t^{v, \bar{v}} := e^{rt} V_t^v 1_{\llbracket 0, \sigma_{\bar{v}} \rrbracket} + e^{r\sigma_{\bar{v}}} V_{\sigma_{\bar{v}}}^v 1_{\llbracket \sigma_{\bar{v}}, \infty \rrbracket}$. For fixed $\bar{v} \in \overline{B_\varepsilon(w)}$, the map $v \mapsto \Lambda^{v, \bar{v}}$ is continuous for almost all ω . Because $\|\Lambda^{v, \bar{v}}\| \leq \varepsilon/2$, there exists an $\varepsilon_{\bar{v}} > 0$ such that $\|\Lambda^{v, \bar{v}}\| < \varepsilon$ for all $v \in B_{\varepsilon_{\bar{v}}}(\bar{v})$. This implies $\|W^v - v\| < \varepsilon$ on $\llbracket 0, \sigma_{\bar{v}} \rrbracket$ for all $v \in B_{\varepsilon_{\bar{v}}}(\bar{v}) \cap B_\varepsilon(w)$ and thus $\tau_v > \sigma_{\bar{v}}$ a.s. Since $\overline{B_\varepsilon(w)}$ is compact, there exists a finite subcover $B_{\varepsilon_{\bar{v}_1}}(\bar{v}_1), \dots, B_{\varepsilon_{\bar{v}_m}}(\bar{v}_m)$ of $\overline{B_\varepsilon(w)}$ and thus $\tau := \inf_{\ell=1, \dots, m} \sigma_{\bar{v}_\ell}$ is strictly positive. Since $v \mapsto W^v$ is continuous, it is Borel measurable and because β is a continuous functional of W , so is $v \mapsto \beta^v$. \square

Having locally constructed strong solutions to (5), we need to piece them together to construct a global weak solution. The following lemma tells us that this is possible if the corresponding payoff set is compact.

LEMMA 10. *Let $\mathcal{W} \subseteq \mathbb{R}^n$ be a compact locally self-generating set. Then there exists a discount rate \tilde{r} such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for any $r \in (0, \tilde{r})$.*

PROOF. The family of open neighborhoods $(U_w)_{w \in \mathcal{W}}$ forms an open cover of \mathcal{W} ; hence by compactness there exists a finite subcover $(U^k)_{k=1, \dots, N}$. By making these sets disjoint, we obtain a finite, Borel measurable cover of \mathcal{W} . On each of these (now disjoint) sets U^k we have a stochastic framework $(\Omega^k, \mathcal{F}^k, \mathbb{F}^k, P^k, Z^k)$. Let $(\Omega, \mathcal{F}, \mathbb{F}, P)$ be the product space, that is, $\Omega := \Omega^1 \times \dots \times \Omega^N$, $\mathcal{F} := \mathcal{F}^1 \otimes \dots \otimes \mathcal{F}^N$, and similarly for \mathcal{F}_t , $t \geq 0$, and define $P(B) := P(B^1) \dots P(B^N)$ for $B = B^1 \times \dots \times B^N \in \mathcal{F}$. Choose any discount rate r smaller than $\tilde{r} = \min(\tilde{r}^1, \dots, \tilde{r}^N) > 0$. Then, dependent on r , for every U_k there exist a strategy profile A^k , a martingale M^k orthogonal to σZ^k , and a stopping time $\tau^k > 0$ such that for every $v \in U^k \cap \mathcal{W}$, there exist β^v, W^v satisfying the appropriate conditions. Let now $\tau(\omega^1, \dots, \omega^N) := \min(\tau^1(\omega^1), \dots, \tau^N(\omega^N), 1)$, which is positive P -a.s.

Fix $v \in \mathcal{W}$ and let κ be the index such that $v \in U^\kappa$. On $\llbracket 0, \tau \rrbracket$ we set $Z = Z^\kappa$, $A = A^\kappa$, $M = M^\kappa$, $W = W^v$, and $\beta = \beta^v$, and we know that they have the desired properties. In particular, $W_\tau \in \mathcal{W}$; hence we can concatenate the solution with the processes related to the neighborhood U^k that W_τ falls into. More precisely, on $\llbracket \tau, \tilde{\tau} \rrbracket$, where $\tilde{\tau} - \tau$ is identically distributed as τ , we set

$$Z_t(\omega) := Z_{\tau(\omega)}^\kappa(\omega^\kappa) + \sum_{k=1}^N Z_{t-\tau(\omega)}^k(\omega^k) 1_{\{W_{\tau(\omega)} \in U^k\}}$$

and similarly for M . It follows from the strong Markov property of each Z^k that Z is an (\mathbb{F}, P) -Brownian motion. Since τ is bounded, the optional stopping theorem implies

¹⁸Here, $V^v(\omega)$ is to be understood as an element of the space \mathcal{D}^n of càdlàg functions from $[0, \infty)$ to \mathbb{R}^n with the topology of uniform convergence on compacts. A compatible metric is given by

$$d(f, g) := \sum_{k=1}^{\infty} \frac{1}{2^k} \left(1 \wedge \sup_{0 \leq s \leq k} \|f(s) - g(s)\| \right);$$

see Protter (2005, p. 220).

that M is a (P, \mathbb{F}) -martingale. Because P is defined as the product measure and each M^k is orthogonal to σZ^k , M is orthogonal to σZ . Further, define

$$A_t(\omega) := \sum_{k=1}^N A_{t-\tau(\omega)}^k(\omega^k) 1_{\{W_{\tau(\omega)} \in U^k\}},$$

which is $\Delta(\mathcal{A})$ -valued and \mathbb{F} -progressively measurable, hence a valid strategy profile. Finally, let $\beta_t(\omega) := \beta_{t-\tau(\omega)}^{W_{\tau(\omega)}}(\omega)$ and $W_t(\omega) = W_{t-\tau(\omega)}^{W_{\tau(\omega)}}(\omega)$, which are \mathbb{F} -progressively measurable since $v \mapsto \beta^v$ and $v \mapsto W^v$ are measurable by assumption. By construction, β enforces A on $\llbracket 0, \tilde{\tau} \rrbracket$ and all the processes are related by (5).

An iteration of this procedure thus leads to a sequence of stopping times $(\tau_\ell)_{\ell \geq 0}$ and solutions $(W^\ell, A^\ell, \beta^\ell, Z^\ell, M^\ell)_{\ell \geq 0}$ to (5) on $\llbracket \tau_\ell, \tau_{\ell+1} \rrbracket$ such that $\tau_{\ell+1} - \tau_\ell$ are independent and identically distributed as τ . By the subsequent Lemma 11, τ_ℓ diverges to ∞ a.s.; hence a countable concatenation of the solutions $(W^\ell, A^\ell, \beta^\ell, Z^\ell, M^\ell)$ yields a solution to (5) on $[0, \infty)$ attaining v . Since v was arbitrary, \mathcal{W} is self-generating for discount rate r and hence $\mathcal{W} \subseteq \mathcal{E}(r)$ by Lemma 2. The statement follows since $r \in (0, \tilde{r})$ was arbitrary. \square

LEMMA 11. *Let $(\tau_\ell)_{\ell \geq 0}$ be a sequence of random variables with $\tau_0 = 0$ such that $\tau_{\ell+1} - \tau_\ell$ are strictly positive and i.i.d. Then $\tau_\ell \rightarrow \infty$ a.s.*

PROOF. Let $\tilde{\tau}_\ell := \tau_\ell - \tau_{\ell-1}$ for $\ell \geq 1$. Then $\tilde{\tau}_\ell$ are i.i.d. and $\tau_\ell = \sum_{k=1}^{\ell} \tilde{\tau}_k$. Therefore, the strong law of large numbers implies that

$$\frac{1}{\ell} \tau_\ell = \frac{1}{\ell} \sum_{k=1}^{\ell} \tilde{\tau}_k \rightarrow \mathbb{E}[\tau_1] \quad \text{a.s.}$$

That is, for all $\varepsilon > 0$ there exists ℓ_0 such that for any $\ell \geq \ell_0$, $|\tau_\ell/\ell - \mathbb{E}[\tau_1]| < \varepsilon$ a.s. Letting $\varepsilon = \mathbb{E}[\tau_1]/2$, we obtain $\tau_\ell > \ell \mathbb{E}[\tau_1]/2$ a.s. for all $\ell \geq \ell_0$. This lower bound, and hence also τ_ℓ , diverges to ∞ a.s. since $\mathbb{E}[\tau_1] > 0$ because $\tau_1 = \tau_1 - \tau_0 > 0$ a.s. \square

4.3 Proofs of the folk theorems

The proofs of Theorems 1 and 2 are completed by showing that the given conditions on the game primitives imply that any closed, smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$ is uniformly decomposable on tangent hyperplanes. Auxiliary results in the spirit of Lemmas 6.1–6.3 of Fudenberg et al. (1994) are contained in Appendix D. Figures 6 and 7 show how these results come together in the proofs of the folk theorems.

PROOF OF THEOREM 2. The statement follows from Lemmas 9 and 10 once we show that \mathcal{W} is uniformly decomposable on tangent hyperplanes. Suppose first that all Pareto-efficient pure action profiles a_1, \dots, a_N are pairwise identifiable and that the expected flow payoff is affine in m . Then the profiles a_1, \dots, a_N are enforceable by Lemma 19. Since \mathcal{W} is contained in the interior of $\text{conv}(g(\alpha_e, a_1, \dots, a_N))$, at any point $w \in \partial \mathcal{W}$ there exists an enforceable action profile $\alpha \in \{\alpha_e, a_1, \dots, a_N\}$ such that $g(\alpha)$ is separated from




FIGURE 6. Solid arrows indicate the use in the proof; dashed arrows represent sufficient conditions for a result to apply in the proof of Theorem 1. To establish a Nash-threat folk theorem, Lemma 18 is not needed; hence we can either assume conditions (ii) and (iv) and proceed as above or assume that Pareto-efficient pure action profiles are pairwise identifiable and use the shortcut as in Figure 7.

\mathcal{W} by \mathcal{S}_w . Suppose first that \mathcal{S}_w is regular. Then α and N_w satisfy conditions (i) or (iv) of Lemma 5.

If \mathcal{S}_w is coordinate to the i th axis, then w either maximizes or minimizes player i 's payoff on \mathcal{W} by convexity. If w maximizes player i 's payoff, then a_i^p from Corollary 4 maximizes g^i over \mathcal{A} and it is pairwise identifiable by assumption; hence condition (ii) of Lemma 5 is fulfilled. If w minimizes player i 's payoff, then $g(\alpha_e)$ is separated from \mathcal{W} by \mathcal{S}_w by assumption and $(\alpha_e, -e_i)$ satisfy condition (iv) of Lemma 5.

Suppose now that conditions (i) and (ii) are satisfied instead. In this case, we can actually decompose smooth payoff sets \mathcal{W} in the interior of the slightly larger payoff set $\mathcal{V}^\dagger := \{v \in \mathcal{V}^* \mid v^i \geq g^i(\alpha_e)\}$. Denote by $\tilde{a}_1, \dots, \tilde{a}_K$ pure action profiles with extremal payoffs such that \mathcal{V}^\dagger is contained in $\text{conv}(g(\alpha_e, \tilde{a}_1, \dots, \tilde{a}_K))$. For each \tilde{a}_ℓ there exists a mixed action profile α_ℓ with pairwise full rank by Lemma 17, such that $g(\alpha_\ell)$ is arbitrarily close to $g(\tilde{a}_\ell)$ for all ℓ . Moreover, by Lemma 16, $\alpha_1, \dots, \alpha_K$ are all enforceable and pairwise identifiable. Because \mathcal{W} is contained in the interior of \mathcal{V}^\dagger , we can choose $\alpha_1, \dots, \alpha_K$ such that \mathcal{W} is contained in the interior of $\text{conv}(g(\alpha_e, \alpha_1, \dots, \alpha_K))$. Therefore, for every $w \in \mathcal{W}$, there exists an enforceable action profile such that its expected payoff is strictly separated from \mathcal{W} by \mathcal{S}_w .




FIGURE 7. Condition (v) enables a shortcut to establish a Nash-threat folk theorem. Lemmas 9, 10 and their preliminary results are still needed, but are not reproduced in this figure to avoid repetition.

If S_w is regular or if w minimizes the payoff of a player i on \mathcal{W} , the statement works in the same way as before. If w maximizes the payoff of player i over \mathcal{W} , we use a_i^* instead of a_i^p . It follows from Lemma 3 that a_i^* is enforceable orthogonal to e_i and (a_i^*, e_i) satisfies condition (ii) or (iii) of Lemma 5 by assumption. \square

PROOF OF THEOREM 1. Because of condition (ii), any extremal payoff can be approximated by the payoff of a mixed action profile with pairwise full rank. Hence all points w where S_w is regular can be dealt with as in the proof of Theorem 2 under conditions (i) and (ii). In the case where w maximizes player i 's payoff on \mathcal{W} , condition (i) and Lemmas 3 and 16 show that a_i^* is enforceable orthogonal to e_i . Because of condition (iii), (a_i^*, e_i) satisfies either condition (ii) or (iii) of Lemma 5. If w minimizes player i 's payoff on \mathcal{W} , condition (iv) and Lemma 18 ensure that there exists an enforceable action profile α_i with best response property for player i such that $g(\alpha_i)$ is strictly separated from \mathcal{W} by S_w and $(\alpha_i, -e_i)$ satisfies condition (iii) of Lemma 5. \square

5. CONCLUSION

In this paper, we study continuous-time multiplayer games with imperfect information, where the signal is distorted by Brownian noise. Establishing a rigorous theory of continuous-time repeated games for any number of players, we provide a mathematically sound basis of how to model strategies in mixed actions. An important concept introduced in this paper is the notion of *uniform decomposability on tangent hyperplanes* of a payoff set \mathcal{W} . For such a payoff set, there exists a stopping time τ , such that equilibrium profiles with continuation values in \mathcal{W} need to be adapted only at independent copies of τ . Using this technique, we are able to construct surprisingly simple continuous-time equilibria attaining nearly efficient outcomes, thereby establishing the folk theorem in continuous time. The simple nature of these equilibria is also very desirable for implementation of these strategies.

The techniques of this paper suggest that if a payoff set \mathcal{W} is smooth and uniformly decomposable on tangent hyperplanes, then there exists a discrete-time game with periods of random length, such that the equilibrium payoff set of the discrete-time game

coincides with \mathcal{W} and a canonical embedding of the discrete-time equilibria into continuous time leads to continuous-time equilibria. By considering a sequence of smooth and uniformly decomposable inner approximations of $\mathcal{E}(r)$, it is possible to obtain a sequence of such games that approximate the continuous-time game such that the equilibrium payoff sets of the discrete-time games converge to $\mathcal{E}(r)$. This leads to interesting questions for future research on the connection between equilibria in discrete- and continuous-time models.

APPENDIX A: MIXING IN CONTINUOUS TIME

In a continuous-time setting, realizations of a mixed action have to be drawn continuously. Suppose that player i plays a fixed mixed action over an interval $[s, t]$ and samples from his mixing distribution only at discrete intervals. An opponent who samples more frequently may realize this after a couple of his own samples, and henceforth play a best response to the already sampled action of player i . To avoid such a scenario, sampling has to be done continuously, where the realizations are drawn from a continuum of independent events. This means, though, that a behavior strategy could involve a continuum of probability spaces on which the distributions $\Delta(\mathcal{A}^i)$ are evaluated. In this appendix, we show the existence of a unified probability space containing all the public information as well as the outcomes of behavior strategies. In this framework it is also possible to formulate private strategies, and we show that players always have a public best reply to any public strategy profile of their opponents. Moreover, we provide a definition of public mixed strategies (as opposed to behavior strategies) and prove a continuous-time analogue to Kuhn's theorem.

Fix a mixed action profile $\alpha \in \Delta(\mathcal{A})$. We call an \mathcal{A} -valued random variable γ on some probability space an *instantiation* of α if γ has distribution α . An outcome of α is then identified with $\gamma(\omega)$ for some $\omega \in \Omega$. We call $(\gamma_t)_{t \geq 0}$ an *instantiation* of a strategy profile A if at each point in time t , the \mathcal{F}_t -conditional distribution of γ_t equals A_t . Then $\gamma(\omega)$ is the outcome of the strategy profile A . In continuous time, an instantiation of a strategy profile A is constructed by setting

$$\gamma_t = \left(\sum_{a^1 \in \mathcal{A}^1} a^1 1_{\Xi_t^1(a^1)}, \dots, \sum_{a^n \in \mathcal{A}^n} a^n 1_{\Xi_t^n(a^n)} \right), \quad t > 0, \tag{8}$$

where Ξ_t^i is a partition of Ω independent of the public filtration \mathbb{F} , of all partitions Ξ_s^i for $s < t$, and of all partitions Ξ_t^j for $j \neq i$ such that $A_t^i(a^i) = P(\Xi_t^i(a^i) | \mathcal{F}_t)$. The following lemma justifies that we can indeed do this.

LEMMA 12. *There exist independent filtrations $\mathbb{M}^i = (\mathcal{M}_t^i)_{t \geq 0}$, $i = 1, \dots, n$, that are independent of the public filtration, such that at all $t > 0$, \mathcal{M}_t^i contains finite partitions of Ω of arbitrary size that are independent of \mathcal{M}_s^i for all $s < t$.*

PROOF. Fix a player i . We start by constructing a process $(U_t^i)_{t \geq 0}$ such that each U_t^i is standard uniformly distributed and independent of U_s^i for $s < t$. Indeed, its finite-dimensional distributions satisfy

$$P(U_{t_1}^i \leq c_1, \dots, U_{t_m}^i \leq c_m) = \prod_{j=1}^m P(U_{t_j}^i \leq c_j) = \prod_{j=1}^m c_j$$

for all $t_j \in [0, \infty)$ such that $t_j \neq t_\ell$ for $j \neq \ell$, all $c_j \in [0, 1]$, and all $m \in \mathbb{N}$. Since this family of finite-dimensional distributions is consistent, Kolmogorov's existence theorem (see, for example, Theorem 36.2 of Billingsley 1986) tells us that such a process indeed exists. Independent partitions of the appropriate size can now be found as the pre-image of a partition of $[0, 1]$ under U^i . Therefore, the filtration generated by U^i will serve as \mathbb{M}^i . Clearly we can do this construction finitely many times in an independent way and independently of \mathbb{F} . \square

The filtration \mathbb{M}^i is the personal source of randomness that player i has available for mixing. Because these filtrations are independent, neither do players learn anything about the signal from their personal source of randomness nor can they predict the outcome of their opponent's mixing. For $i = 1, \dots, n$, let \mathbb{F}^i denote the augmented filtration generated by \mathbb{F} and \mathbb{M}^i .

LEMMA 13. *A stochastic process γ is the instantiation of a behavior strategy profile if and only if $\gamma^i : \Omega \times [0, \infty) \rightarrow \mathcal{A}^i$ is \mathbb{F}^i -progressively measurable for every player i .*

PROOF. For a behavior strategy profile A , define an instantiation γ through (8), where the partitions Ξ_t^i exist by Lemma 12. Since \mathbb{M}^i are defined as the filtrations generated by Ξ^i , γ^i has the necessary measurability properties. For the converse, let γ be a stochastic process such that γ^i is \mathbb{F}^i -progressively measurable with values in \mathcal{A}^i . For any $a \in \mathcal{A}$, define $A(a) := \mathcal{O}(1_{\{\gamma=a\}})$ as the \mathbb{F} -optional projection of $1_{\{\gamma=a\}}$. It is the unique \mathbb{F} -optional process X such that

$$\mathbb{E}[1_{\{\gamma_\tau=a\}} 1_{\{\tau < \infty\}} | \mathcal{F}_\tau] = X_\tau 1_{\{\tau < \infty\}} \quad \text{a.s.}$$

for every \mathbb{F} stopping time τ ; see Section VI.2 of Dellacherie and Meyer (1982) for further details on the optional projection.¹⁹ Observe that $A(a)$ is \mathbb{F} -progressively measurable and that $\sum_{a \in \mathcal{A}} A(a) = 1$ a.e. Because $\mathbb{M}^1, \dots, \mathbb{M}^n$ are independent of each other,

$$A_t(a) = \mathbb{E}[1_{\{\gamma_t^1=a^1\}} \cdots 1_{\{\gamma_t^n=a^n\}} | \mathcal{F}_t] = P(\gamma_t^1 = a^1 | \mathcal{F}_t) \cdots P(\gamma_t^n = a^n | \mathcal{F}_t) \quad \text{a.s.}$$

¹⁹One may wonder why we use the optional projection and not its alternative, the predictable projection defined in (VI.43.2) of Dellacherie and Meyer (1982). First, in discrete time the optional projection reduces to the standard conditional expectations $\mathcal{O}(m(A))_\ell = \mathbb{E}[m(A_\ell) | \mathcal{F}_\ell]$ for $\ell \in \mathbb{N}$ by Remark VI.44.b of Dellacherie and Meyer (1982), which means taking the step-by-step average over the mixing, whereas the predictable projection corresponds to $\mathbb{E}[m(A_\ell) | \mathcal{F}_{\ell-1}]$, which would mean to average over not just the mixing but also other new information. Second, without mixing (\mathbb{M}^i trivial), the predictable projection does not need to correspond to the original process, as opposed to the optional projection. In particular, if actions are correlated in an unpredictable way by using information orthogonal to Z , the predictable projection would filter this out.

This means that the players' distributions are conditionally independent, given the public information. Therefore, A is indeed a behavior strategy profile. \square

Observe that $\mu(A) = \mathcal{O}(m(\gamma))$ for any instantiation γ of A . Because $\mu(A)$ is \mathbb{F} -progressively measurable, so is the density process dQ^A/dP defined in (2). Therefore, we immediately obtain the following consistency result.

LEMMA 14. *For any behavior strategy profile A , the following statements hold:*

- (i) *The family Q^A of probability measures agrees with P on $\mathbb{M}^1, \dots, \mathbb{M}^n$.*
- (ii) *The filtrations $\mathbb{F}, \mathbb{M}^1, \dots, \mathbb{M}^n$ are independent under Q^A .*
- (iii) *The \mathbb{F} -optional projections under Q^A and P coincide.*

The first property says that a change of measure does not affect the weight a player assigns to any pure action. By the second property, a mixed action profile remains a mixed action profile under Q^A . Finally, the last statement implies that the infinitesimal average is not affected by a change of measure.

LEMMA 15. *Suppose that player i 's opponents play a public strategy profile A^{-i} . Then player i has a best response in public strategies.*

PROOF. Suppose that player i has additional information available in a filtration \mathbb{G}^i and let γ be an instantiation of A . Similar to the proof of Lemma 13, γ^i is progressively measurable with respect to the augmented filtration of \mathbb{F}, \mathbb{G}^i , and \mathbb{M}^i , where $\mathbb{M}^1, \dots, \mathbb{M}^n$ are independent of \mathbb{G}^i . Then a public best response is given by

$$\tilde{A}^i(a^i) := a^i \mathcal{O}_{\mathbb{F}^i}(1_{\{\gamma^i=a^i\}}), \quad a^i \in \mathcal{A}^i,$$

where $\mathcal{O}_{\mathbb{F}^i}(\cdot)$ denotes the optional projection onto \mathbb{F}^i . Indeed, for any \mathbb{F} stopping time τ we obtain, by \mathcal{F}_τ -conditional independence of \mathcal{M}_τ^i and \mathcal{M}_τ^j for $j \neq i$,

$$\begin{aligned} \mu(\tilde{A}^i, A^{-i})_\tau 1_{\{\tau < \infty\}} &= \mathbb{E}[m(\tilde{\gamma}_\tau^i, \gamma_\tau^{-i}) 1_{\{\tau < \infty\}} | \mathcal{F}_\tau] \\ &= \sum_{a \in \mathcal{A}} m(a) \mathbb{E}[\mathcal{O}_{\mathbb{F}^i}(1_{\{\gamma^i=a^i\}})_\tau 1_{\{\tau < \infty\}} | \mathcal{F}_\tau] \prod_{j \neq i} \mathcal{O}(1_{\{A^j=a^j\}})_\tau \\ &= \mu(A^i, A^{-i})_\tau 1_{\{\tau < \infty\}}, \end{aligned}$$

where $\tilde{\gamma}^i$ is an instantiation of \tilde{A}^i . Therefore, $Q^{(\tilde{A}^i, A^{-i})} = Q^A$ and it follows from (3) that $W(\tilde{A}^i, A^{-i}) = W(A)$ a.e. by projectivity of the conditional expectation. \square

A mixed strategy (as opposed to a behavior strategy) is a mixture over pure strategies. Formally, it is a probability measure on the space of all pure strategies. In the absence of public randomization, the public filtration is generated by Z ; hence $\Omega = C[0, \infty)^d$ without loss of generality, where $C[0, \infty)^d$ denotes the space of all continuous functions from $[0, \infty)$ to \mathbb{R}^d . Therefore, a pure strategy for player i can be seen as a progressively measurable mapping $\gamma^i : C[0, \infty)^d \times [0, \infty) \rightarrow \mathcal{A}^i$.

DEFINITION 10. A *mixed strategy* for player i is a probability measure κ^i on the set

$$\mathcal{P}^i = \{\gamma^i : C[0, \infty)^d \times [0, \infty) \rightarrow \mathcal{A}^i \mid \gamma^i \text{ is progressively measurable}\}^{20}$$

A *mixed strategy profile* κ is given by the product measure $\kappa = \kappa^1 \otimes \dots \otimes \kappa^n$ on the product σ -algebra on $\mathcal{P} = \mathcal{P}^1 \times \dots \times \mathcal{P}^n$, where κ^i is a mixed strategy for player $i = 1, \dots, n$. Player i 's discounted expected future payoff of a mixed strategy profile is given by

$$W_t^i(\kappa) = \int_{\mathcal{P}} W_t^i(\gamma) d\kappa(\gamma) = \int_{\mathcal{P}^1} \dots \int_{\mathcal{P}^n} W_t^i(\gamma^1, \dots, \gamma^n) d\kappa^1(\gamma^1) \otimes \dots \otimes d\kappa^n(\gamma^n).$$

Since mixing over strategies is a more complicated procedure than mixing over actions, it is usually easier to work with behavior strategies than with mixed strategies. In the remainder of this appendix we show in an analogue to Kuhn's theorem (see [Kuhn 1953](#)) that the two notions are equivalent. A mixed strategy profile κ and a behavior strategy profile A are *realization equivalent* if they lead to the same distribution over outcomes, that is, for any $a \in \mathcal{A}$,

$$A(a) = \kappa(\{\gamma \in \mathcal{P} \mid \gamma = a\}) \quad \text{a.e.}$$

THEOREM 4 (Analogue of Kuhn's theorem). *Every mixed strategy profile is realization equivalent to some behavior strategy profile. Conversely, every behavior strategy profile has a realization equivalent mixed strategy profile.*

PROOF. Let κ be a mixed strategy profile. Fix a player i and define for any $a^i \in \mathcal{A}^i$ and any $(\omega, t) \in C[0, \infty)^d \times [0, \infty)$,

$$A_t^i(a^i; \omega) := \kappa^i(\{\gamma^i \in \mathcal{P}^i \mid \gamma_t^i(\omega) = a^i\}).$$

It can be deduced that $A^i(a^i)$ is almost everywhere well defined and a progressively measurable process for all $a^i \in \mathcal{A}^i$. Indeed, the sets

$$S(a^i) = \{(\gamma^i, \omega, t) \in \mathcal{P}^i \times C[0, \infty)^d \times [0, \infty) \mid \gamma_t^i(\omega) = a^i\}$$

are elements of the product σ -algebra of $\sigma\mathcal{P}^i$ (see [footnote 20](#)) and the progressive σ -algebra on $C[0, \infty)^d \times [0, \infty)$. Since $\{\gamma^i \in \mathcal{P}^i \mid \gamma_t^i(\omega) = a^i\}$ are the (ω, t) sections of $S(a^i)$, it follows from measurable induction that the mapping

$$(\omega, t) \rightarrow \kappa^i(\{\gamma^i \in \mathcal{P}^i \mid \gamma_t^i(\omega) = a^i\})$$

is progressively measurable, which means that $A^i(a^i)$ is progressively measurable. Moreover, the processes $A^i(a^i)$ are nonnegative and their sum over $a^i \in \mathcal{A}^i$ is 1. Since $\kappa = \kappa^1 \otimes \dots \otimes \kappa^n$, it follows that $A(a) := A^1(a^1) \dots A^n(a^n)$ indeed defines a realization equivalent behavior strategy profile.

²⁰Formally, κ^i is defined not on \mathcal{P}^i itself, but on the σ -algebra $\sigma\mathcal{P}^i$ on \mathcal{P}^i generated by the coordinate maps $\pi_t : \mathcal{P}^i \rightarrow (C[0, t]^d \rightarrow \mathcal{A}^i)$ given by $\pi_t(\gamma^i) := \gamma_t^i$; see also [Billingsley \(1986, p. 509\)](#).

Let now A be a behavior strategy profile and let U^1, \dots, U^n be independent processes with standard uniform marginals as in the proof of Lemma 12 on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$. For any player i , enumerate $A^i = \{a^i_1, \dots, a^i_{K_i}\}$ and define

$$\gamma^i_t(\tilde{\omega}, \omega) = \sum_{j=1}^{K_i} a^i_j \mathbf{1}_{\{\sum_{\ell=1}^{j-1} A^i_\ell(a^i_\ell; \omega) \leq U^i_t(\tilde{\omega}) < \sum_{\ell=1}^j A^i_\ell(a^i_\ell; \omega)\}},$$

which we consider as a mapping in $\tilde{\omega}$ from $\tilde{\Omega}$ to the set \mathcal{P}^i of progressively measurable processes $C[0, \infty)^d \times [0, \infty) \rightarrow A^i$. Using the σ -algebra from footnote 20, this mapping becomes measurable; hence we can define a probability measure κ^i on \mathcal{P}^i as the pre-image of γ^i under \tilde{P} , that is, $\kappa^i = \tilde{P} \circ (\gamma^i)^{-1}$. Therefore, $\kappa = \kappa^1 \otimes \dots \otimes \kappa^n$ indeed defines a realization equivalent mixed strategy profile. \square

APPENDIX B: TIME-CHANGED PPEs AND MONOTONICITY OF $\mathcal{E}(r)$

PROOF OF LEMMA 6. For a strategy profile A in a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$, we define the time-changed processes

$$\tilde{A}_t := A_{\lambda t}, \quad \tilde{Z}_t := \frac{1}{\sqrt{\lambda}} Z_{\lambda t}, \quad \tilde{Y}_t := \tilde{\sigma} \tilde{Z}_t.$$

Observe that \tilde{A} is progressively measurable with respect to the time-changed filtration $\tilde{\mathbb{F}} = (\tilde{\mathcal{F}}_t)_{t \geq 0}$, where $\tilde{\mathcal{F}}_t := \mathcal{F}_{\lambda t}$, and \tilde{Z} is an $\tilde{\mathbb{F}}$ -Brownian motion by the scaling property of Brownian motion. The family $(\tilde{Q}_t^{\tilde{A}})_{t \geq 0}$ of probability measures induced by \tilde{A} with respect to \tilde{m} , $\tilde{\sigma}$, and \tilde{Y} is defined as

$$\frac{d\tilde{Q}_t^{\tilde{A}}}{dP} := \exp\left(\int_0^t \tilde{\mu}(\tilde{A}_s)^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} d\tilde{Y}_s - \frac{1}{2} \int_0^t \tilde{\mu}(\tilde{A}_s)^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{\mu}(\tilde{A}_s) ds\right).$$

Abbreviate $\hat{m} := \tilde{m}/\sqrt{\lambda}$ and observe that $\tilde{\sigma}^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \hat{m} = \sigma^\top (\sigma \sigma^\top)^{-1} m$ by assumption. With the substitution $d\tilde{s} = \lambda ds$ we arrive at

$$\begin{aligned} \frac{d\tilde{Q}_t^{\tilde{A}}}{dP} &= \exp\left(\int_0^t \hat{\mu}(A_{\lambda s})^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{\sigma} dZ_{\lambda s} - \frac{1}{2} \int_0^t \hat{\mu}(A_{\lambda s})^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \hat{\mu}(A_{\lambda s}) \lambda ds\right) \\ &= \exp\left(\int_0^{\lambda t} \mu(A_{\tilde{s}})^\top (\sigma \sigma^\top)^{-1} \sigma dZ_{\tilde{s}} - \frac{1}{2} \int_0^{\lambda t} \mu(A_{\tilde{s}})^\top (\sigma \sigma^\top)^{-1} \mu(A_{\tilde{s}}) d\tilde{s}\right) = \frac{dQ_{\lambda t}^A}{dP} \end{aligned}$$

and hence $\tilde{Q}_t^{\tilde{A}}$ coincides with $Q_{\lambda t}^A$ on $\tilde{\mathcal{F}}_t = \mathcal{F}_{\lambda t}$. Observe that the expected flow payoff $g^i(a) = f^i(a^i, m(a)) = f^i(a^i, \sigma \tilde{\sigma}^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{m}(a)/\sqrt{\lambda})$ still depends on a^{-i} only through $\tilde{m}(a)$. Substituting $d\tilde{s} = \lambda ds$ again, we obtain for every $t \geq 0$,

$$\begin{aligned} \tilde{W}_t^i(\tilde{A}; \lambda r, \tilde{m}, \tilde{\sigma}) &:= \int_t^\infty \lambda r e^{-\lambda r(s-t)} \mathbb{E}_{\tilde{Q}_s^{\tilde{A}}} [g^i(\tilde{A}_s) | \tilde{\mathcal{F}}_t] ds \\ &= \int_{\lambda t}^\infty r e^{-r(\tilde{s}-\lambda t)} \mathbb{E}_{Q_{\tilde{s}}^A} [g^i(A_{\tilde{s}}) | \mathcal{F}_{\lambda t}] d\tilde{s} = W_{\lambda t}^i(A; r, m, \sigma) \quad \text{a.s.} \end{aligned} \tag{9}$$

Because all unilateral deviations of $(A_{\lambda t})_{t \geq 0}$ correspond to unilateral deviations of A , it follows from (9) that $(A_{\lambda t})_{t \geq 0}$ is a PPE with respect to $(\lambda r, \tilde{m}, \tilde{\sigma})$ if and only if A is a PPE with respect to (r, m, σ) . \square

PROOF OF THEOREM 3. To show (a), we need to show that a PPE A with respect to (r, m, σ) can be transformed to a PPE with respect to $(r, m, \sigma\Lambda)$. Let $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$ be the stochastic framework of A , and let β, Z , and M be the processes from Lemma 1 that satisfy (5) for A and $W = W(A)$ with respect to μ and σ . Let Z^\perp be a k -dimensional Brownian motion orthogonal to both Z and M , and denote by $\hat{\mathbb{F}}$ the augmented filtration generated by \mathbb{F} and Z^\perp . Because Λ is symmetric, we can write $\Lambda = Q^\top D Q$ for an orthogonal matrix Q and a diagonal matrix

$$D = \begin{pmatrix} \mathbb{I}_\ell & 0 & 0 \\ 0 & -\mathbb{I}_m & 0 \\ 0 & 0 & \tilde{D} \end{pmatrix}$$

such that \tilde{D} has entries in $(-1, 1)$. Define $\tilde{\Lambda} := Q^\top \sqrt{\mathbb{I}_k - D^2} Q$ so that $\Lambda^2 + \tilde{\Lambda}^2 = \mathbb{I}_k$. Then $\hat{Z} := \Lambda Z + \tilde{\Lambda} Z^\perp$ is a Brownian motion with respect to $\hat{\mathbb{F}}$. Set

$$\hat{\Lambda} := Q^\top \begin{pmatrix} 0 & 0 \\ 0 & (\sqrt{\mathbb{I}_{k-\ell-m} - \tilde{D}^2})^{-1} \end{pmatrix} Q$$

and define $\hat{Z}^\perp := \hat{\Lambda}(Z - \Lambda \hat{Z})$. It follows from

$$d\langle \hat{Z}^\perp, \hat{Z}^\perp \rangle_t = \begin{pmatrix} 0 & 0 \\ 0 & \mathbb{I}_{k-\ell-m} \end{pmatrix} dt$$

and $d\langle \hat{Z}^\perp, \hat{Z} \rangle_t = 0$ that \hat{Z}^\perp is a martingale orthogonal to \hat{Z} . This gives us the decomposition $Z = \Lambda \hat{Z} + \tilde{\Lambda} \hat{Z}^\perp$ and hence

$$\begin{aligned} dW_t &= r(W_t - g(A_t)) dt + r\beta_t(\sigma dZ_t - \mu(A_t) dt) + dM_t \\ &= r(W_t - g(A_t)) dt + r\beta_t(\sigma\Lambda d\hat{Z}_t - \mu(A_t) dt) + r\beta_t\sigma\tilde{\Lambda} d\hat{Z}_t^\perp + dM_t. \end{aligned}$$

Since $\ker(\sigma\Lambda) = \ker(\sigma)$ by assumption, it follows that there exists a matrix $\Delta \in \mathbb{R}^{d \times d}$ such that $\Delta\sigma = \sigma\Lambda$.²¹ Therefore, $\sigma\Lambda\hat{Z} = \Delta^2\sigma Z + \sigma\Lambda\tilde{\Lambda}Z^\perp$ is orthogonal to M and hence also to $\hat{M} := M + \int r\beta_s\sigma\tilde{\Lambda} d\hat{Z}_s^\perp$. It follows that W also fulfills (5) for $\sigma\Lambda$ with processes β, \hat{Z} , and \hat{M} . Since β enforces A , A is also a PPE in the continuous-time game with volatility $\sigma\Lambda$ by Lemma 1. Note, however, that A is a PPE as an $\hat{\mathbb{F}}$ -progressively measurable process, that is, when players use the new orthogonal information suitably. We derive from Itô's formula that

$$d(e^{-rt}W_t) = -re^{-rt}g(A_t) dt + re^{-rt}\hat{\beta}_t(\sigma\Lambda d\hat{Z}_t - \mu(A_t) dt) + e^{-rt}d\hat{M}_t. \quad (10)$$

²¹Indeed, σ is an isomorphism from $\ker(\sigma)^\perp$ to \mathbb{R}^d with inverse $\sigma^\top(\sigma\sigma^\top)^{-1}$ and hence the vectors $b_i := \sigma^\top(\sigma\sigma^\top)^{-1}e_i$ for $i = 1, \dots, d$ form a basis of $\ker(\sigma) = \ker(\sigma\Lambda)$, where e_i denotes the i th standard basis vector in \mathbb{R}^d . Define $\delta_i := \sigma\Lambda b_i$ for $i = 1, \dots, d$ and set $\Delta = (\delta_1, \dots, \delta_d)$. By construction, we have $\Delta\sigma b_i = \delta_i = \sigma\Lambda b_i$. Since (b_1, \dots, b_d) can be completed to a basis of \mathbb{R}^n with any basis of $\ker(\sigma\Lambda)$ and $\ker(\sigma) = \ker(\sigma\Lambda)$, it follows that $\Delta\sigma = \sigma\Lambda$.

Observe that $\ker(\sigma\Lambda) = \ker(\sigma)$ implies $\text{rank}(\sigma\Lambda) = d$; hence we can define a family $(\tilde{Q}_t^A)_{t \geq 0}$ of probability measures as in (2) with respect to m and $\sigma\Lambda$. Integrating (10) from t to T and taking \tilde{Q}_T^A -conditional expectations on $\hat{\mathcal{F}}_t$ yields

$$W_t = \int_t^T e^{-r(s-t)} \mathbb{E}_{\tilde{Q}_s^A}[g(A_s)|\hat{\mathcal{F}}_t] ds + e^{-r(T-t)} \mathbb{E}_{\tilde{Q}_T^A}[W_T|\hat{\mathcal{F}}_t].$$

Taking the limit as $T \rightarrow \infty$ yields $W_t(A; r, m, \sigma) = \hat{W}_t(A; r, m, \sigma\Lambda)$ a.s. since W is bounded. This concludes the proof of (a).

For statement (b), let first $\tilde{\sigma} = \Lambda\sigma$. Define the $k \times k$ matrix $\Lambda' := \sigma^\top(\sigma\sigma^\top)^{-1}\Lambda\sigma$ so that $\sigma\Lambda' = \Lambda\sigma$. Note that $\Lambda'^\top = \sigma^\top\Lambda(\sigma\sigma^\top)^{-1}\sigma = \Lambda'$, i.e., Λ' is symmetric. Every vector in the kernel of σ is an eigenvector of Λ' with eigenvalue 0. Let λ be an eigenvalue of Λ' to an eigenvector v that is not in the kernel of σ . Then $\Lambda\sigma v = \sigma\Lambda'v = \lambda\sigma v$, that is, λ is an eigenvalue of Λ for eigenvector σv . Since this applies to all eigenvalues of Λ' outside the kernel of σ , the eigenvalues of Λ' lie in $[-1, 1]$ and $\ker(\sigma\Lambda') = \ker(\sigma)$. Moreover, $\sigma\Lambda' = \Lambda\sigma$ has rank d because Λ is invertible and σ has rank d . The statement now follows by applying (a) to Λ' . Observe that the change $\tilde{m} = \Lambda^{-1}m$ is completely equivalent since

$$\sigma^\top\Lambda^\top(\Lambda\sigma\sigma^\top\Lambda^\top)^{-1}m = \sigma^\top\Lambda\Lambda^{-1}(\sigma\sigma^\top)^{-1}\Lambda^{-1}m = \sigma^\top(\sigma\sigma^\top)^{-1}\Lambda^{-1}m$$

and hence the induced probability measures coincide at all times. For statement (c), Lemma 6 implies $\tilde{W}_t((A_{\lambda t})_{t \geq 0}, \lambda r, m, \sigma/\sqrt{\lambda}) = W_{\lambda t}(A, r, m, \sigma)$ a.s. for $\lambda \in (0, 1)$. The statement follows from (a) for the matrix $\Lambda = \text{diag}_k(\sqrt{\lambda})$ applied to $\sigma/\sqrt{\lambda}$. \square

APPENDIX C: PROOFS OF LEMMAS 1, 5, AND 7

The statement of Lemma 1 is rather intuitive, since (3) can be rewritten as

$$W_t^i(A) = re^{rt} \left(\lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^A} \left[\int_0^u e^{-rs} g^i(A_s) ds \middle| \mathcal{F}_t \right] - \int_0^t e^{-rs} g^i(A_s) ds \right)$$

and hence $dW_t^i = rW_t^i dt - rg^i(A_t) dt + d$ “martingale” by the product rule. However, the limiting probability measure Q_∞^A is not equivalent to P on \mathcal{F}_∞ ; hence we cannot immediately apply a martingale representation result.²²

PROOF OF LEMMA 1. To show (a) \implies (b), observe first that $W^i := W^i(A)$ is bounded, as it remains in \mathcal{V} at all times. Fix $T > 0$ and derive from (3) that

$$\begin{aligned} w_T^i &:= W_T^i - r \int_0^T (W_t^i - g^i(A_t)) dt \\ &= W_T^i + r \int_0^T g^i(A_t) dt - r \int_0^\infty \int_0^{s \wedge T} re^{-r(s-t)} \mathbb{E}_{Q_s^A}[g^i(A_s)|\mathcal{F}_t] dt ds \end{aligned} \tag{11}$$

²²The probability measure Q_∞^A that coincides with Q_t^A on \mathcal{F}_t for every $t \geq 0$. It is not obtained as the limit in (2), but its existence is asserted by Proposition I.7.4 of Karatzas and Shreve (1998).

is a bounded \mathcal{F}_T -measurable random variable. By Corollary 1 to Theorem IV.37 of Protter (2005), any square-integrable Q_T^A -martingale can be decomposed uniquely into a stochastic integral with respect to $\sigma Z - \int \mu(A_s) ds$ and a square-integrable martingale M orthogonal to Z . Applying this to $\mathbb{E}_{Q_T^A}[w_T^i | \mathcal{F}_t]$, we obtain an \mathcal{F}_0 -measurable c_T^i , a progressively measurable process $(\beta_{t,T}^i)_{0 \leq t \leq T}$ with $\mathbb{E}_{Q_T^A}[\int_0^T |\beta_{t,T}^i|^2 dt] < \infty$, and a Q_T^A -martingale $(M_{t,T}^i)_{0 \leq t \leq T}$ orthogonal to Z with $M_{0,T}^i = 0$ such that

$$w_T^i = c_T^i + \int_0^T r \beta_{t,T}^i (\sigma dZ_t - \mu(A_t) dt) + M_{T,T}^i.$$

To prove that (b) holds, we need to show that c_T^i , $\beta_{t,T}^i$, and $M_{t,T}^i$ do not depend on T . Let $\tilde{T} \leq T$ and take in (11) conditional expectations on $\mathcal{F}_{\tilde{T}}$ under $Q_{\tilde{T}}^A$ to deduce that

$$\begin{aligned} \mathbb{E}_{Q_{\tilde{T}}^A}[w_T^i | \mathcal{F}_{\tilde{T}}] - w_{\tilde{T}}^i &= \mathbb{E}_{Q_{\tilde{T}}^A}[w_T^i | \mathcal{F}_{\tilde{T}}] - W_{\tilde{T}}^i + r \int_{\tilde{T}}^T \mathbb{E}_{Q_t^A}[g^i(A_t) | \mathcal{F}_{\tilde{T}}] dt \\ &\quad - r \int_{\tilde{T}}^{\infty} \int_{\tilde{T}}^{s \wedge T} r e^{-r(s-t)} \mathbb{E}_{Q_s^A}[g^i(A_s) | \mathcal{F}_{\tilde{T}}] dt ds \\ &= \mathbb{E}_{Q_{\tilde{T}}^A}[w_T^i | \mathcal{F}_{\tilde{T}}] - W_{\tilde{T}}^i - \int_{\tilde{T}}^{\infty} r e^{-r(s-\tilde{T})} \mathbb{E}_{Q_s^A}[g^i(A_s) | \mathcal{F}_{\tilde{T}}] ds \\ &\quad + \int_{\tilde{T}}^{\infty} r e^{-r(s-\tilde{T})} \mathbb{E}_{Q_s^A}[g^i(A_s) | \mathcal{F}_{\tilde{T}}] ds \\ &= 0, \end{aligned}$$

using that $\mathbb{E}_{Q_{\tilde{T}}^A}[X | \mathcal{F}_{\tilde{T}}] = \mathbb{E}_{Q_s^A}[X | \mathcal{F}_{\tilde{T}}]$ for \mathcal{F}_s -measurable X and $s \leq T$. Taking $\tilde{T} = 0$, this shows that $c_T^i = W_0^i$ does not depend on T . It also implies

$$w_{\tilde{T}}^i = \mathbb{E}_{Q_{\tilde{T}}^A}[w_T^i | \mathcal{F}_{\tilde{T}}] = W_0^i + \int_0^{\tilde{T}} r \beta_{t,T}^i (\sigma dZ_t - \mu(A_t) dt) + M_{\tilde{T},T}^i,$$

which yields $\beta_{\cdot,T}^i = \beta_{\cdot,\tilde{T}}^i$ a.e. and $M_{\tilde{T},T}^i = M_{\tilde{T},\tilde{T}}^i$ a.s. by the uniqueness of the orthogonal decomposition. Taking \mathcal{F}_t -conditional expectations, we deduce $M_{t,\tilde{T}}^i = M_{t,T}^i$ a.s. for every t ; hence $W(A)$ satisfies (b).

To prove (b) \implies (a), we derive from Itô's formula that

$$d(e^{-rt} W_t^i) = -r e^{-rt} g^i(A_t) dt + r e^{-rt} \beta_{t,T}^i (\sigma dZ_t - \mu(A_t) dt) + e^{-rt} dM_t^i. \quad (12)$$

Integrating (12) from t to T and taking Q_T^A -conditional expectations on \mathcal{F}_t thus yields

$$W_t^i = \int_t^T r e^{-r(s-t)} \mathbb{E}_{Q_s^A}[g^i(A_s) | \mathcal{F}_t] ds + e^{-r(T-t)} \mathbb{E}_{Q_T^A}[W_T^i | \mathcal{F}_t].$$

Since W is bounded, the second summand converges to zero a.s. as T tends to ∞ ; hence W_t^i is indeed the discounted expected future value of A .

For the last statement, fix a player i and a time t , and let \tilde{A} be a strategy profile with $\tilde{A}^{-i} = A^{-i}$. For β related to $W(A)$ by (5), we obtain from (12) for $u \geq t$ that

$$W_t^i(A) = e^{-r(u-t)}W_u^i(A) - \int_t^u re^{-r(s-t)}(\beta_s^i(\sigma dZ_s - \mu(A_s) ds) - g^i(A_s) ds - dM_s^i).$$

As we let $u \rightarrow \infty$, the term $e^{-r(u-t)}W_u^i(A)$ vanishes since $W^i(A)$ is bounded. Since M is a martingale up to time u also under $Q_u^{\tilde{A}}$, we obtain

$$\begin{aligned} W_t^i(\tilde{A}) &= \lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^{\tilde{A}}} \left[\int_t^u re^{-r(s-t)} g^i(\tilde{A}_s) ds \middle| \mathcal{F}_t \right] \\ &= W_t^i(A) + \lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^{\tilde{A}}} \left[\int_t^u re^{-r(s-t)} ((g^i(\tilde{A}_s) - g^i(A_s)) ds \right. \\ &\qquad \qquad \qquad \left. + \beta_s^i(\sigma dZ_s - \mu(A_s) ds)) \middle| \mathcal{F}_t \right] \quad \text{a.s.} \end{aligned}$$

Because w_T^i defined in (11) is bounded, it follows from the construction of β that the process $\int_t^u re^{-r(s-t)} \beta_s^i(\sigma dZ_s - \mu(A_s) ds)$ is up to any time $u \in (t, \infty)$ a bounded mean oscillation (BMO) martingale under the probability measure $Q_u^{\tilde{A}}$. This implies by Theorem 3.6 of Kazamaki (1994) that $\int_t^u re^{-r(s-t)} \beta_s^i(\sigma dZ_s - \mu(\tilde{A}_s) ds)$ is a BMO martingale under $Q_u^{\tilde{A}}$. Together with Fubini's theorem this implies

$$\begin{aligned} W_t^i(\tilde{A}) - W_t^i(A) &= \int_t^\infty re^{-r(s-t)} \mathbb{E}_{Q_s^{\tilde{A}}} [g^i(\tilde{A}_s) - g^i(A_s) + \beta_s^i(\mu(\tilde{A}_s) - \mu(A_s)) | \mathcal{F}_t] ds \quad \text{a.s.} \end{aligned} \tag{13}$$

If β enforces A , the above conditional expectation is nonpositive; hence A is a PPE. To show the converse, assume toward a contradiction that there exists a player i and a set $\Xi \subseteq \Omega \times [0, \infty)$ with $P \otimes \text{Lebesgue}(\Xi) > 0$, such that for some other strategy \hat{A}^i ,

$$g^i(\hat{A}^i, A^{-i}) - g^i(A) + \beta^i(\mu(\hat{A}^i, A^{-i}) - \mu(A)) > 0 \quad \text{on } \Xi.$$

Set $\tilde{A}^i := \hat{A}^i 1_\Xi + A^i 1_{\Xi^c}$. Because β is progressively measurable, we can and do choose Ξ and \hat{A} to be progressively measurable as well. In particular, \tilde{A}^i is a behavior strategy for player i . For such an \tilde{A} , the expectation in (13) is strictly positive for $t = 0$, a contradiction. \square

PROOF OF LEMMA 5. Statement (i). Let $\beta_\alpha^i(N) = B^i$ as in (6), which is locally Lipschitz continuous in N . The statement holds by choosing U_N such that the first two coordinates are bounded away from zero.

Statement (ii). The case where N is not parallel to a coordinate axis is shown in statement (i); hence suppose now that $N = e_1$. Let $\tilde{\beta}^1, \dots, \tilde{\beta}^n$ be defined as in the proof of Lemma 4 and set

$$\beta_\alpha^1(x) := - \sum_{i=2}^n \frac{x^i}{x^1} \tilde{\beta}^i, \quad \beta_\alpha^i(x) = \tilde{\beta}^i, \quad i = 2, \dots, n. \tag{14}$$

Along the lines of the proof of [Lemma 4](#), it follows that $\beta_\alpha(x)$ enforces α orthogonal to x if $x_1 \neq 0$. The statement follows by choosing U_N bounded away from $\{x^1 = 0\}$.

Statement (iii). Suppose $N = e_1$ and that $\alpha^1 = a^1 \in \mathcal{A}^1$ is a unique best response to α^{-1} . Let $\beta_\alpha^i(x)$ as in (14), except that $\tilde{\beta}^i$ are replaced by β^i . Then, clearly, (4) is fulfilled for players $i = 2, \dots, n$. Because of the unique best response property, there exists an $\varepsilon > 0$ such that $g^1(\alpha^1) \geq g^1(\tilde{a}^1, \alpha^{-1}) + \varepsilon$ for every $\tilde{a}^1 \in \mathcal{A}^1 \setminus \{a^1\}$. Let us set $B = \max_{i=2, \dots, n} \max_{\tilde{a}^1 \in \mathcal{A}^1} |\beta^i(\mu(\tilde{a}^1, \alpha^{-1}) - \mu(\alpha))|$, which is finite as β is fixed. If $B = 0$, then α is a Nash equilibrium and the result holds by statement (iv). Suppose, therefore, that $B > 0$. Then for all x in

$$U_{e_1} := \left\{ x \in \mathbb{R}^n \mid \|x - e_1\| \leq \frac{\varepsilon}{B(n-1) + \varepsilon} \right\},$$

x^1 is bounded away from 0, and hence $|x^i|/x^1 \leq \varepsilon/(B(n-1))$. It follows that

$$|\beta_\alpha^1(x)(\mu(\tilde{a}^1, \alpha^{-1}) - \mu(\alpha))| = \sum_{i=2}^n \frac{|x^i|}{x^1} |\beta^i(\mu(a^1, \alpha^{-1}) - \mu(\alpha))| \leq \varepsilon$$

for every $\tilde{a}^1 \in \mathcal{A}^1$. Together with the unique best response property for player 1, this shows that $\beta_\alpha(x)$ enforces α for $x \in U_{e_1}$. For all $x \in U_{e_1}$, $\beta_\alpha(x)^\top x = 0$ by construction and β_α is Lipschitz continuous and bounded since x_1 is bounded away from 0.

Statement (iv). This is clear since $\beta_\alpha(x) = 0$ for all $x \in \mathbb{R}^n$. \square

PROOF OF LEMMA 7. To show that (a) \implies (b), let $X \in \mathcal{E}(r, m)$ a.s. Although we may have different probability spaces in (a) for each realization $X = x$, we can use the fact that the models all share the same path space to construct a regular conditional probability on that space. The path space of a behavior strategy A and its stochastic framework is given by $\mathcal{D} := \Delta(\mathcal{A})^{[0, \infty)} \times C[0, \infty)^d$, and $C[0, \infty)^d$ is the space of continuous functions $[0, \infty) \rightarrow \mathbb{R}^d$. One can show that $\Omega = \mathcal{V} \times \mathcal{D}$ is complete and separable,²³ hence by Theorem V.3.19 in [Karatzas and Shreve \(1998\)](#) there exists a regular conditional probability $P_x(F) : \mathcal{V} \times \mathcal{F} \rightarrow [0, 1]$, which means that it has the following properties:

- (i) For each $x \in \mathcal{V}$, P_x is a probability measure on (Ω, \mathcal{F}) .
- (ii) For each $F \in \mathcal{F}$, the mapping $x \mapsto P_x(F)$ is $\mathcal{B}(\mathcal{V})$ -measurable.
- (iii) For each $F \in \mathcal{F}$, $P_x(F) = P(F|X = x)$ for ν -a.e. $x \in \mathcal{V}$, where ν is the distribution of X .

We know that for each $x \in \mathcal{E}(r, m)$, there exists a PPE A^x achieving x . Let A now be the process defined pointwise by A^x on $\{X = x\}$. It follows from the properties of a regular conditional probability that A is a PPE achieving X . Indeed, for any player i and any

²³ \mathcal{V} is complete and separable as a closed subset of \mathbb{R}^n , $C[0, \infty)^d$ is both complete and separable with respect to the uniform metric by Theorem 43.6 of [Munkres \(2000\)](#) and $\Delta(\mathcal{A})^{[0, \infty)}$ is compact by Tychonov's theorem and hence complete and separable.

behavior strategy profile \tilde{A} with $\tilde{A}^{-i} = A^{-i}$,

$$\begin{aligned} P(W_0^i(A) \geq W_0^i(\tilde{A})) &= \int_{\mathcal{E}(r,m)} P(W_0^i(A) \geq W_0^i(\tilde{A})|X = x) \, d\nu(x) \\ &= \int_{\mathcal{E}(r,m)} P_x(W_0^i(A^x) \geq W_0^i(\tilde{A})) \, d\nu(x) = 1 \end{aligned}$$

and in the same way $P(W_0(A) = X) = \int_{\mathcal{E}(r,m)} P_x(W_0(A^x) = x) \, d\nu(x) = 1$.

To show the implication (b) \implies (a), suppose that $X \notin \mathcal{E}(r, m)$ on an \mathcal{F}_0 -measurable set Ξ with $\nu(\Xi) > 0$. Since there are only finitely many players, this implies the existence of an \mathcal{F}_0 -measurable set $\tilde{\Xi}$ with $\nu(\tilde{\Xi}) > 0$ such that some player i can improve his strategy to $\tilde{A}^{x,i}$ for $x \in \tilde{\Xi}$. Letting $\hat{A}^i := A^i 1_{\tilde{\Xi}^c}(x) + \tilde{A}^{x,i} 1_{\tilde{\Xi}}(x)$, it follows that

$$P(W_0^i(A) \geq W_0^i(\hat{A}^i, A^{-i})) = \int_{\mathcal{V}} P_x(W_0^i(A) \geq W_0^i(\tilde{A}^{x,i}, A^{-i})) \, d\nu(x) < 1,$$

contradicting the assumption that A is a PPE. □

APPENDIX D: AUXILIARY RESULTS

In this appendix we provide some auxiliary results related to enforceability and pairwise identifiability. These results are needed in the proofs of the folk theorems as indicated in Figures 6 and 7. They differ from the results in Fudenberg et al. (1994) in that we model the change of the signal's distribution through a change of its drift.

LEMMA 16. *An action profile α is enforceable if for every player i one of below conditions holds:*

- (i) *The profile α has individual full rank for player i .*
- (ii) *The action α^i is a best response of player i to α^{-i} .*

The enforceability condition (4) imposes n systems of linear inequalities, one for each player i . Because $\text{rank } M^i(\alpha) \leq |\mathcal{A}^i| - 1$, we cannot simply solve the system i by applying the left inverse of $M^i(\alpha)$, but we need to additionally exploit that the linear dependence among the columns of $G^i(\alpha)$ and $M^i(\alpha)$ is the same.

PROOF OF LEMMA 16. Fix a player i and suppose first that 1. is satisfied for action profile α . Let $a^i \in \mathcal{A}^i$ be an action with $\alpha^i(a^i) > 0$ and enumerate $\mathcal{A}^i = \{a_1^i, \dots, a_{K_i}^i\}$ such that $a^i = a_{K_i}^i$ is the last element. For the sake of brevity, denote by $M_j^i(\alpha)$ the column of $M^i(\alpha)$ corresponding to action a_j^i . Because of the linear dependence among the columns of $M^i(\alpha)$, we obtain

$$M_{K_i}^i(\alpha) = - \sum_{j=1}^{K_i-1} \frac{\alpha(a_j^i)}{\alpha(a_{K_i}^i)} M_j^i(\alpha).$$




FIGURE 8. The right panel shows \mathcal{V}^Δ and \mathcal{V}^* for the stage game with payoffs given in the table to the left. In the proof of the folk theorems, we only need that we can approximate extremal points of \mathcal{V}^* and the minmax payoffs, which all lie in \mathcal{V}^Δ .

Condition 1 implies that there is no other linear dependence among the columns of $M^i(\alpha)$ and thus the $d \times (|A^i| - 1)$ -dimensional submatrix $\tilde{M}^i(\alpha)$ consisting of the first $|A^i| - 1$ columns has full column rank. In particular, $\tilde{M}^i(\alpha)$ has a left inverse $\tilde{M}_L^i(\alpha)$ and

$$\beta^i = G^i(\alpha) \begin{bmatrix} \tilde{M}_L^i(\alpha) \\ 0 \end{bmatrix}$$

solves the system for player i with equality. Indeed,

$$\begin{aligned} G^i(\alpha) \begin{bmatrix} \tilde{M}_L^i(\alpha) \\ 0 \end{bmatrix} M^i(\alpha) &= G^i(\alpha) \begin{bmatrix} \mathbb{I}_{K_i-1} & \sum_{j=1}^{K_i-1} \frac{\alpha(a_j^i)}{\alpha(a_{K_i}^i)} \tilde{M}_L^i(\alpha) \tilde{M}_j^i(\alpha) \\ 0 & 0 \end{bmatrix} \\ &= \left(G_1^i(\alpha), \dots, G_{K_i-1}^i(\alpha), - \sum_{j=1}^{K_i-1} \frac{\alpha(a_j^i)}{\alpha(a_{K_i}^i)} G_j^i(\alpha) \right), \end{aligned}$$

where we used that $\tilde{M}_L^i(\alpha) \tilde{M}_j^i(\alpha) = e_j$. The claim under condition 1 follows since

$$G_{K_i}^i(\alpha) = - \sum_{j=1}^{K_i-1} \frac{\alpha(a_j^i)}{\alpha(a_{K_i}^i)} G_j^i(\alpha).$$

Under condition 2, $\beta^i = 0$ solves the inequalities for player i . □

The following lemmas are the analogues of Lemmas 6.1–6.3 in Fudenberg et al. (1994) in our setting. Let \mathcal{V}^Δ denote the set of payoffs achievable in mixed actions. While it may be strictly smaller than \mathcal{V} for some games (see Figure 8), the extremal payoffs always correspond to pure action profiles and hence are contained in \mathcal{V}^Δ .

LEMMA 17. *Suppose that for every pair of players i, j , there exists a mixed action profile α^{ij} having ij -pairwise full rank. Then the set of payoffs C of action profiles with pairwise full rank for all pairs of players is dense in \mathcal{V}^Δ .*

PROOF. Let $E \subseteq \Delta(\mathcal{A})$ denote the set of mixed action profiles with pairwise full rank. By Lemma 6.2 of Fudenberg et al. (1994), E is dense in $\Delta(\mathcal{A})$. Since the map $g : \Delta(\mathcal{A}) \rightarrow \mathcal{V}^\Delta$ is continuous and surjective, $C \supseteq g(E)$ is dense in \mathcal{V}^Δ . □

LEMMA 18. *Suppose that every pure action profile has individual full rank. Then for any $\varepsilon > 0$ and any player i , there exists an enforceable action profile α with best response property for player i and $|g^i(\alpha) - \underline{v}^i| < \varepsilon$. Moreover, if every pure action is pairwise identifiable or the best response of player i to the minmax profile $\underline{\alpha}_{-i}^i$ is unique, then the pair $(\alpha, -e_i)$ satisfies condition (ii) or (iii) of Lemma 5, respectively.*

Note that the first part of the statement is identical to Lemma 6.3 of Fudenberg et al. (1994). However, we also need that the resulting action profile leads to a locally uniform decomposition as in Lemma 5.

PROOF OF LEMMA 18. Fix a player i and let $\underline{\alpha}_i^{-i}$ denote a minmax profile against player i . By assumption, (a^i, a^{-i}) has individual full rank for every $a^i \in \mathcal{A}^i$ and every $a^{-i} \in \mathcal{A}^{-i}$. Therefore, similar as in the proof of Lemma 6.2 in Fudenberg et al. (1994), one can find a sequence of profiles $(\underline{\alpha}_{(k)}^{-i})_{k \geq 1}$ converging to $\underline{\alpha}_i^{-i}$, such that $(a^i, \underline{\alpha}_{(k)}^{-i})$ has individual full rank for every $a^i \in \mathcal{A}^i$ and all k .

Let a_k^i be a best response for player i to $\underline{\alpha}_{(k)}^{-i}$. The profiles $(a_k^i, \underline{\alpha}_{(k)}^{-i})$ are enforceable orthogonal to $-e_i$ by Lemmas 16 and 3. Let $\underline{a}^i \in \mathcal{A}^i$ be an accumulation point of $(a_k^i)_{k \geq 1}$ and choose a subsequence $(k_\nu)_{\nu \geq 1}$ such that $a_{k_\nu}^i = \underline{a}^i$ for all $\nu \in \mathbb{N}$. Observe that \underline{a}^i is also a best response to $\underline{\alpha}_i^{-i}$ because

$$g^i(\underline{a}^i, \underline{\alpha}_i^{-i}) = \lim_{\nu \rightarrow \infty} g^i(\underline{a}^i, \underline{\alpha}_{(k_\nu)}^{-i}) \geq \lim_{\nu \rightarrow \infty} g^i(\tilde{a}^i, \underline{\alpha}_{(k_\nu)}^{-i}) = g^i(\tilde{a}^i, \underline{\alpha}_i^{-i}),$$

hence $\underline{a}^i \in \arg \max g^i(\cdot, \underline{\alpha}_i^{-i})$ and $g^i(\underline{a}^i, \underline{\alpha}_i^{-i}) = \underline{v}^i$. Therefore, for any $\varepsilon > 0$ we can find ν large enough such that $|g^i(\underline{a}^i, \underline{\alpha}_{(k_\nu)}^{-i}) - \underline{v}^i| < \varepsilon$.

Under condition 1, $(\underline{\alpha}_{(k)}^{-i})_{k \geq 1}$ can be chosen in a way that $(a^i, \underline{\alpha}_{(k)}^{-i})$ has pairwise full rank for every $a^i \in \mathcal{A}^i$ and all k . Therefore, $((\underline{a}^i, \underline{\alpha}_{(k_\nu)}^{-i}), -e_i)$ satisfies the second condition of Lemma 5. Under condition 2, it follows from multilinearity that there exists a ν large enough such that \underline{a}^i is also a unique best response to $\underline{\alpha}_{(k_\nu)}^{-i}$. Therefore, condition (iii) of Lemma 5 is fulfilled for the pair $((\underline{a}^i, \underline{\alpha}_{(k_\nu)}^{-i}), -e_i)$. \square

DEFINITION 11. An action profile α *Pareto-dominates* a profile $\tilde{\alpha}$ if $g^i(\alpha) \geq g^i(\tilde{\alpha})$ for every player i and $g^j(\alpha) > g^j(\tilde{\alpha})$ for at least one player j .

An action profile is *Pareto-efficient* if it is not Pareto-dominated by any other action profile.

LEMMA 19. *Suppose that $g^i(a) = b^i(a^i)m(a) - c^i(a^i)$. Then any Pareto-efficient pure action profile is enforceable.*

PROOF. Fix a Pareto-efficient pure action profile $a \in \mathcal{A}$. Because its payoff is on the “upper right” boundary of \mathcal{V} , there exists a direction $N \in \mathbb{R}^n$ with $N^i > 0$ for every i such that $g(a) = \arg \max_{v \in \mathcal{V}} N^\top v$. Then β with row vectors $\beta^i := \sum_{j \neq i} b^j(a^j)N^j / N^i$ enforces a .

Indeed, for every $\tilde{a}^i \in \mathcal{A}^i$, we have

$$\begin{aligned}
 g^i(\tilde{a}^i, a^{-i}) + \beta^i m(\tilde{a}^i, a^{-i}) &= g^i(\tilde{a}^i, a^{-i}) + \frac{1}{N^i} \sum_{j \neq i} (N^j g^j(\tilde{a}^i, a^{-i}) + N^j c^j(a^j)) \\
 &= \frac{1}{N^i} \sum_{j=1}^n N^j g^j(\tilde{a}^i, a^{-i}) + \frac{1}{N^i} \sum_{j \neq i} N^j c^j(a^j) \\
 &\leq \frac{1}{N^i} \sum_{j=1}^n N^j g^j(a) + \frac{1}{N^i} \sum_{j \neq i} N^j c^j(a^j) \\
 &= g^i(a) + \beta^i m(a). \quad \square
 \end{aligned}$$

Finally, we establish that for every player i there exists at least one Pareto-efficient action profile globally maximizing player i 's payoff. Hence, from Lemmas 3 and 19, it is enforceable on the corresponding coordinate hyperplane.

COROLLARY 4. *Suppose that g is affine in m . Then, for every player i , there exists an enforceable Pareto-efficient pure action profile a_i^P that maximizes g^i over \mathcal{A} . In particular, it is enforceable on the hyperplane coordinate to the i th axis.*

PROOF. Let $\mathcal{A}^{(i)} \subseteq \mathcal{A}$ denote the set of pure action profiles that maximize player i 's payoff over \mathcal{A} . Action profiles in $\mathcal{A}^{(i)}$ can be Pareto dominated only by other action profiles in $\mathcal{A}^{(i)}$ because Pareto dominance of some $\tilde{a} \in \mathcal{A}$ over $a \in \mathcal{A}^{(i)}$ entails $g^i(\tilde{a}) \geq g^i(a)$. Since the relation of Pareto dominance is transitive and irreflexive, there cannot be any circular relations on $\mathcal{A}^{(i)}$. Because there are only finitely many elements in $\mathcal{A}^{(i)}$, at least one element is not dominated, hence Pareto-efficient. Such an action profile a_i^P is enforceable by Lemma 19 and because it is a static best response for player i , it is enforceable on the hyperplane orthogonal to the i th axis due to Lemma 3. \square

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063. [417, 424]
- Anderson, Robert M. (1984), "Quick-response equilibrium." Working Papers in Economic Theory and Econometrics IP-323, Center for Research in Management, University of California, Berkeley. [415]
- Billingsley, Patrick (1986), *Probability and Measure*, second edition. Wiley, New York. [439, 441]
- Bolton, Patrick and Christopher Harris (1999), "Strategic experimentation." *Econometrica*, 67, 349–374. [428]
- Daley, Brendan and Brett Green (2012), "Waiting for news in the market for lemons." *Econometrica*, 80, 1433–1504. [413, 423]

Dellacherie, Claude and Paul-André Meyer (1982), *Probabilities and Potential, B: Theory of Martingales*. North-Holland Pub. Co., Amsterdam. [439]

Faingold, Eduardo and Yuliy Sannikov (2011), “Reputation in continuous-time games.” *Econometrica*, 79, 773–876. [413, 423]

Fudenberg, Drew, David K. Levine, and Eric Maskin (1994), “The folk theorem with imperfect public information.” *Econometrica*, 62, 997–1039. [411, 412, 413, 415, 419, 422, 424, 425, 435, 448, 449, 450]

Fudenberg, Drew and David K. Levine (2007), “Continuous time limits of repeated games with imperfect public monitoring.” *Review of Economic Dynamics*, 10, 173–192. [413]

Karatzas, Ioannis and Steven Shreve (1998), *Methods of Mathematical Finance*, volume 39 of Applications of Mathematics. Springer, New York. [444, 447]

Kazamaki, Norihiko (1994), *Continuous Exponential Martingales and BMO*, volume 1579 of Lecture Notes in Mathematics. Springer, Berlin. [446]

Keller, Godfrey and Sven Rady (2010), “Strategic experimentation with Poisson bandits.” *Theoretical Economics*, 5, 275–311. [428]

Kuhn, Harold W. (1953), “Extensive games and the problem of information.” In *Contributions to the Theory of Games* (Harold W. Kuhn and Albert W. Tucker, eds.), 193–216, Princeton University Press, Princeton. [441]

Mailath, George J. and Larry Samuelson (2006), *Repeated Games and Reputations*. Oxford University Press, New York, NY. [415]

Munkres, James R. (2000), *Topology*, second edition. Pearson, New York. [447]

Øksendal, Bernt (1998), *Stochastic Differential Equations*, fifth edition. Springer, New York. [432]

Osório, António M. (2012), “A folk theorem for games when frequent monitoring decreases noise.” *The B.E. Journal of Theoretical Economics*, 12, Art. 11, 27. [413]

Protter, Philip E. (2005), *Stochastic Integration and Differential Equations*, second edition, volume 21 of Stochastic Modelling and Applied Probability. Springer, New York. [433, 434, 445]

Sannikov, Yuliy (2007), “Games with imperfectly observable actions in continuous time.” *Econometrica*, 75, 1285–1329. [411, 412, 416, 417, 422, 427, 428, 429]

Sannikov, Yuliy and Andrzej Skrzypacz (2010), “The role of information in repeated games with frequent actions.” *Econometrica*, 78, 847–882. [413]

Simon, Leo K. and Maxwell B. Stinchcombe (1989), “Extensive form games in continuous time: Pure strategies.” *Econometrica*, 57, 1171–1214. [415]

Staudigl, Mathias and Jan-Henrik Steg (2014), "On repeated games with imperfect public monitoring: From discrete to continuous time." Preprint available at <http://ssrn.com/abstract=2497608>. [413]

Co-editor George J. Mailath handled this manuscript.

Submitted 2013-10-24. Final version accepted 2015-3-11. Available online 2015-3-12.