Available online at www.sciencedirect.com

ScienceDirect

Games and Economic Behavior 65 (2009) 406–429

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Brown–von Neumann–Nash dynamics: The continuous strategy case

Josef Hofbauer [a], Jörg Oechssler [b], Frank Riedel [c,*]

[a] *Department of Mathematics, University of Vienna, Austria*
[b] *Department of Economics, University of Heidelberg, Germany*
[c] *Institute of Mathematical Economics, University of Bielefeld, Germany*

Received 7 September 2006

Available online 4 April 2008

Abstract

Brown and von Neumann introduced a dynamical system that converges to saddle points of zero sum games with finitely many strategies. Nash used the mapping underlying these dynamics to prove existence of equilibria in general games. The resulting Brown–von Neumann–Nash dynamics are a benchmark example for myopic adjustment dynamics that, in contrast to replicator dynamics, allow for innovation, but require less rationality than the best response dynamics. This paper studies the BNN dynamics for games with infinitely many strategies. We establish Nash stationarity for continuous payoff functions. For negative semidefinite games (that include zero sum games), we generalize the results of Brown and von Neumann. In addition, we show that evolutionarily robust Nash equilibria are asymptotically stable. A complete stability analysis for doubly symmetric games is also obtained.

JEL classification: C70; 72

Keywords: Learning in games; Evolutionary stability; BNN

1. Introduction

This paper studies the Brown–von Neumann–Nash (BNN) dynamics for games with a continuum strategy space. Studying BNN dynamics is interesting for at least three reasons. First,

* Corresponding author at: University of Bielefeld, Institute of Mathematical Economics, Postfach 100131, Bielefeld, Germany.

E-mail address: friedel@uni-bielefeld.de (F. Riedel).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2008.03.006

Brown and von Neumann (1950) used it as an algorithm to find saddle points in zero sum games. We extend their results to continuum zero sum games and certain non zero sum games (negative semidefinite or stable games). We also show that the corresponding flow is continuous (in the appropriate weak topology). This shows that one can use BNN to find Nash equilibria in continuum games. This may prove to be very useful in games that are not easily solvable. An interesting numerical application might be the continuous strategy version of the Blotto game (Borel, 1921; for a recent application see Laslier and Picard, 2002).

Second, Nash (1951) used the mapping underlying the BNN dynamics to prove existence of equilibria that carry his name. It is now well known that Nash always had a dynamical interpretation of this mapping in mind (see Weibull, 1996). In fact, one might conjecture that the BNN dynamics always converge to a Nash equilibrium. This is not the case, as Berger and Hofbauer (2000, 2006) have shown. It is thus a challenge to find sufficient conditions that ensure convergence for as many games as possible. We prove that the BNN dynamics converge for a large class of games, including zero sum games, versions of Cournot and Bertrand oligopolies, coordination games, and the War of Attrition.

Third, the BNN dynamics are the mean dynamics for a stochastic learning process for which a plausible microfoundation can be given (see Appendix B for an argument based on Sandholm, 2005, 2006). We leave it to the reader to decide whether this adjustment process describes human behavior or not. (This is, in the end, an empirical question beyond the scope of this paper). The BNN dynamics are, however, the benchmark example of innovative, yet not fully rational dynamics: they are myopic adjustment dynamics (Swinkels, 1993) that lie in between the replicator dynamics, which require arguably too little rationality to do justice to human behavior, and the best response dynamics, which may require too much of it.

Our results are as follows. In the next section we show that the BNN dynamics are well defined for very general strategy spaces. We allow strategies to come from some measurable space. This makes it possible to apply BNN in games of incomplete information like auctions where strategies are usually mappings from some real interval to another real interval. Furthermore, we show that under some mild assumptions on the strategy space and for bounded and Lipschitz continuous payoff function, the semi-flow induced by the BNN dynamics is weakly continuous. For several reasons it is useful to know whether the semi-flow B is weakly continuous. First, from an applied perspective, a starting point $P(0)$ can only be known as a rough approximation. Thus, it would be reassuring to know that dynamics that start at nearby initial points, do not diverge from each other too much. Second, a continuous model, as we use it here, is only employed for convenience. A continuous model should always be a good approximation for a finite model if the number of strategies gets large. Weak continuity of the flow is a sufficient condition for such an approximation to persist over time. Finally, from a mathematical perspective, for S a compact metric space, B is a continuous semi-flow on a compact metric space Δ, for which we can then employ a large body of dynamical systems theory, in particular make use of ω-limits to describe the asymptotic behavior.

In Section 3 we prove that, as in the finite case, the rest points of the dynamics coincide with the Nash equilibria (this property is called Nash stationarity) if the payoff function is continuous and the strategy space is a compact metric space. We also show by example that one cannot weaken the condition of continuity in general. Nash stationarity is a useful property. For example, if a trajectory of BNN converges to some point, this point must be a Nash equilibrium.

There turn out to be some important differences between the case of a finite number of pure strategies (as studied by Berger and Hofbauer, 2000 and Hofbauer, 2000) and the continuous strategy case studied here. Probably most important is the fact that strict equilibria are not necessarily (Lyapunov) stable in the continuous case.[1] We demonstrate this in Section 4 through an example with a quadratic payoff function. Interestingly, static stability concepts originally developed for the replicator dynamics (such as a "continuously stable strategy" (CSS), as in Eshel and Motro, 1981 and Evolutionary Robustness, as in Oechssler and Riedel, 2002) become relevant for stability with respect to the BNN dynamics. These concepts are introduced in Section 5.

In Sections 6 and 7 we deal with two classes of games, namely doubly symmetric games and negative semidefinite games, for which we have nearly complete results with respect to stability. For doubly symmetric games, we have the Fundamental Law of Natural Selection: average fitness is increasing over time. This leads to several stability results. We show that evolutionary robust Nash equilibria are asymptotically stable for BNN dynamics. If a game has only one Nash equilibrium, this one is asymptotically stable. An important class of examples are the quadratic games (including linear Cournot and Bertrand oligopolies). These games can be transformed into doubly symmetric games with the same excess payoff, and hence, same BNN trajectories. We show that for these games, asymptotic stability is equivalent to the fact that the slope of the best reply function is less than 1. This property has been called continuous stability (CSS) by Eshel and Motro (1981). For discrete BNN dynamics, asymptotic stability in doubly symmetric games is equivalent to ESS. As evolutionary robustness (ER) is the natural generalization of ESS for continuum games, one might conjecture that asymptotic stability for continuum BNN is equivalent to ER. Our quadratic example shows that this conjecture is not true.

Finally, we establish stability results for negative semidefinite games. This class includes zero sum games, contests, and the War of Attrition. Here, a function that measures the quadratic (positive) excess payoff is a Lyapunov function. Under suitably compactness and continuity conditions, we conclude that BNN dynamics converge to the (convex) set of Nash equilibria. A strict Nash equilibrium—which is then unique by negative semi-definiteness—is globally asymptotically stable. For example, the unique strict Nash equilibrium of contests is globally asymptotically stable for BNN. Finally, we treat at length the War of Attrition because of its importance. Here, the payoff function is not continuous, and our general results do not apply. However, we are able to prove global convergence to the unique Nash equilibrium by using arguments similar to Brown and von Neumann (1950). Some proofs and an attempt of a micro-foundation for the BNN dynamics are relegated to the Appendix.

2. The BNN dynamics

We consider symmetric two-player games with (pure) strategy set S. Let $\mathcal{A}$ be a σ-algebra on S and μ be a finite measure on $(S, \mathcal{A})$. Let $f : S \times S \to \mathbb{R}$ be a bounded measurable function, where $f(x, y)$ is the payoff for player 1 when he plays x and player 2 plays y. An interesting special case which will be treated in more detail is when S is a compact interval $S \subset \mathbb{R}$ with the Lebesgue measure.

A *population* is identified with the aggregate play of its members and is described by a probability measure P on the measurable space $(S, \mathcal{A})$. We denote by Δ the set of all populations (probability measures or mixed strategies) on S. Since Δ is not a vector space, we shall work with the linear span of Δ, that is the space $\mathcal{M}^e(S, \mathcal{A})$ of all finite and signed measures. Recall that ν is a finite signed measure on $(S, \mathcal{A})$ if there are two finite measures μ^1 and μ^2 such that for all sets $A \in \mathcal{A}$, $\nu(A) = \mu^1(A) - \mu^2(A)$.

[1] This is also the case for other dynamics like the replicator dynamics, see e.g. Oechssler and Riedel (2001, 2002).

The average payoff of a measure P against a measure Q is defined as

$$E(P,Q)=\int_S\int_S f(x,y)Q(\mathrm{d}y)P(\mathrm{d}x). \tag{1}$$

Let

$$\sigma(x,P):=E(\delta_x,P)-E(P,P)$$

denote the difference between the payoff of strategy $x\in S$ (identified with the Dirac measure δ_x on x) and the average population payoff. The *excess payoff* of pure strategy x when matched against population P is defined as

$$\sigma_+(x,P):=\max\bigl(\sigma(x,P),0\bigr).$$

We now define the *Brown–von Neumann–Nash dynamics* on the measure space $(S,\mathcal{A},\mu)$ as the differential equation on Δ

$$\dot P(A)=\int_A \sigma_+(x,P)\mu(\mathrm{d}x)-P(A)\int_S \sigma_+(x,P)\mu(\mathrm{d}x), \tag{2}$$

for all $A\in\mathcal{A}$. Let $\Sigma(P):=\int_S \sigma_+(s,P)\mu(\mathrm{d}s)$ denote the *total excess*. If $\Sigma(P)>0$, then the *relative excess* for a subset $A\in\mathcal{A}$ is denoted by $R^P(A):=\frac{1}{\Sigma(P)}\int_A\sigma_+(x,P)\mu(\mathrm{d}x)$ and defines a probability measure on $(S,\mathcal{A})$, absolutely continuous with respect to μ, with density function $r^P(x)=\frac{1}{\Sigma(P)}\sigma_+(x,P)$. Then (2) can be rewritten as

$$\dot P(A)=\Sigma(P)\bigl(R^P(A)-P(A)\bigr). \tag{3}$$

Hence, under the BNN dynamics, a population P moves toward its relative excess measure R^P, and the speed of motion is proportional to the total excess. For later reference note that by construction of R^P we have that

$$E\bigl(R^P,P\bigr)\geqslant E(P,P),\quad \forall P, \tag{4}$$

and the inequality is strict whenever total excess does not vanish. This property implies that the dynamics is weakly compatible in the sense of Friedman (1991) as it moves towards relatively fitter strategies.

We first show existence and uniqueness of solutions of the differential equation (2) by interpreting it as a differential equation in a suitable Banach space.

Theorem 1. *For each $P=P(0)\in\Delta$ there is a unique solution $P(t)\in\Delta$ of the ordinary differential equation* (2) *for $t\in[0,\infty[$.*

See Appendix A for a proof and precise meaning of this statement.

Given that a unique solution to the BNN dynamics exists, we can define the semi-flow

$$B:\Delta\times[0,\infty[\,\to\Delta,$$

where $B(P,t)=P(t)$ denotes the population at time t when the BNN dynamics start in $P=P(0)$.

In most applications, S is a metric space and then the weak topology on Δ is a natural choice. For several reasons it is useful to know whether the semi-flow B is weakly continuous. First, from an applied perspective, a starting point $P(0)$ can only be known as a rough approximation. Thus, it would be reassuring to know that dynamics that start at nearby initial points, do not diverge from each other too much. Second, a continuous model as we use it here is only employed for convenience. A continuous model should always be a good approximation for a finite model if the number of strategies gets large. Weak continuity of the flow is a sufficient condition for such an approximation to persist over time (cf. Oechssler and Riedel, 2002). Finally, from a mathematical perspective, for S a compact metric space, B is a continuous semi-flow on a compact metric space Δ, for which we can then employ a large body of dynamical systems theory, in particular make use of ω-limits to describe the asymptotic behavior.

Theorem 2. *Let S be a separable metric space and f be bounded and Lipschitz continuous. Then the semi-flow B is continuous with respect to the weak topology of measures.*

Proof. See Appendix A.

3. Nash stationarity

A nice property of the BNN dynamics is that for continuous f (and thus, in particular, for the finite strategy case) the rest points of the dynamics coincide with the Nash equilibria.[2] The total excess $\Sigma(P)$ vanishes if $\sigma(x, P) \leqslant 0$ or $E(\delta_x, P) \leqslant E(P, P)$ for μ-almost all $x \in S$, in particular, if P is a Nash equilibrium. For continuous payoff functions f the reverse holds also.

Proposition 1. *Let S be a compact metric space, μ a finite Borel measure on S with full support. Suppose f is continuous. Then P is a rest point of the BNN dynamics if and only if (P, P) is a Nash equilibrium.*

Proof. If P is a best reply to itself, then $\sigma_+(x, P) = 0$ for all x, and stationarity follows.

Let P^* be a stationary point of (6), that is

$$\int_A \sigma_+(x, P^*)\mu(\mathrm{d}x) = P^*(A)\Sigma(P^*) \tag{5}$$

for all Borel sets A. We distinguish two cases, $\Sigma(P^*) = 0$ and $\Sigma(P^*) > 0$.

Case 1: $\Sigma(P^*) = 0$. In this case, for μ-almost every x, we have

$$\sigma_+(x, P^*) = 0.$$

$\sigma_+(x, P^*)$ inherits continuity from f. As μ has full support, it follows that $\sigma_+(x, P^*) = 0$ holds true for all $x \in S$. This is equivalent to

$$E(\delta_x, P^*) \leqslant E(P^*, P^*),$$

and it follows that P^* is a best reply to itself.

Case 2: $\Sigma(P^*) > 0$. Since P^* is a stationary point of (6), we get from (5) that P^* has a density p^* with respect to Lebesgue measure and

$$p^*(x) = \frac{\sigma_+(x, P^*)}{\Sigma(P^*)}$$

[2] Sandholm (2005) calls this property "Nash stationarity."

for P^*-almost every x. For every x with $p^*(x) > 0$, we have thus

$$\sigma_+(x, P^*) > 0,$$

or

$$E(\delta_x, P^*) > E(P^*, P^*).$$

By integrating, we get

$$E(P^*, P^*) = \int_{\{x: p^*(x)>0\}} E(\delta_x, P^*) p^*(x) \mu(\mathrm{d}x) > E(P^*, P^*),$$

a contradiction. Hence, we cannot have $\Sigma(P^*) > 0$ for a stationary point P^*. This concludes the proof. □

Example 1. In general, one cannot weaken the continuity assumption made in the previous proposition. Take $S = [0, 1]$, $\mu = \mathrm{d}x$, and let the payoff function $f(x, y) = 0$ for all $x < 1$ and all $y \in [0, 1]$. Let $f(1, y) = 1$ for all $y \in [0, 1]$. In this game, the strategy $x^* = 1$ is strictly dominant. The unique Nash equilibrium is thus $(1, 1)$. However, every distribution P on $[0, 1]$ with a strictly positive Lebesgue density $p(x) > 0$ for all $x \in [0, 1]$ has $\Sigma(P) = 0$. Thus, all such P are rest points of BNN, but not Nash equilibria.

4. An example: quadratic games

In the previous section we saw that all symmetric Nash equilibria are rest points of the BNN dynamics. However, some of those Nash equilibria may turn out to be unstable. One is used to think of strict Nash equilibria as particularly stable with respect to all kinds of dynamics. And indeed, in the case of finite strategy sets S it is straightforward to show that strict Nash equilibria are asymptotically stable with respect to the BNN dynamics (see e.g. Berger and Hofbauer, 2000). The following simple example shows that this is not the case anymore for general S.

For this example we shall assume that S is a compact interval in $\mathbb{R}$ endowed with the Lebesgue measure. Thus, (2) can be written as

$$\dot{P}(A) = \int_A \sigma_+(x, P)\,\mathrm{d}x - P(A) \int_S \sigma_+(x, P)\,\mathrm{d}x. \tag{6}$$

Example 2. Let $S \subset \mathbb{R}$ be an interval around 0 and $f(x, y) = -x^2 + axy$ be a linear-quadratic game with $a > 0$. For all parameters a, $(0, 0)$ is a strict Nash equilibrium. However, for $a > 2$, this strict Nash equilibrium is unstable with respect to the BNN dynamics.[3] For $a < 2$, BNN dynamics globally converge to the strict Nash equilibrium as both mean and variance converge to 0 along any solution of BNN.

The statements in the example are formally proven by Proposition 4 below. Here we shall present an informal argument to illustrate the issue.

Note first that the game with payoff function $f(x, y) = -x^2 + axy$ is strategically equivalent to the doubly symmetric game with payoff function $f(x, y) = -x^2 + axy - y^2 = f(y, x)$. The

[3] For a formal definition of stability see Definition 1 below.

behavior of BNN is the same under both payoff functions since $\sigma_+(x, P)$ is the same for both payoff functions. Let $P_i := \int_S x^i P(\mathrm{d}x)$ denote the ith moment of P. Then

$$E(\delta_x, P) = \int_S \bigl(-x^2 + axy - y^2\bigr) P(\mathrm{d}y) = -x^2 + axP_1 - P_2 \tag{7}$$

and

$$\begin{aligned} E(P, P) &= \int_S \int_S \bigl(-x^2 + axy - y^2\bigr) P(\mathrm{d}x) P(\mathrm{d}y) \\ &= aP_1^2 - 2P_2. \end{aligned} \tag{8}$$

Therefore

$$\sigma_+(x, P) = \bigl[x(aP_1 - x) - aP_1^2 + P_2\bigr]_+ \tag{9}$$

and the density $r^P(x, P) = \sigma_+(x, P)/\Sigma(P)$ of the relative excess measure is the positive part of a quadratic function[4] in x. Thus, the density of R^P is a symmetric function about the maximizer $x = \frac{a}{2} P_1$. It follows that the mean of the excess measure is also $\frac{a}{2} P_1$. From Eq. (3), it follows that the mean value of the population changes under BNN according to

$$\dot{P}_1 = \Sigma(P)\bigl(R_1^P - P_1\bigr),$$

where R_1^P denotes the mean of R^P. Applying this to our example, we obtain

$$\dot{P}_1 = \Sigma(P)\Bigl(\frac{a}{2} - 1\Bigr) P_1. \tag{10}$$

Thus, for $a < 2$, $P_1(t) \to 0$, whereas for $a = 2$, $P_1(t) \equiv P_1(0)$, and for $a > 2$, $P_1(t)$ moves away from 0. Hence, for $a > 2$, δ_0 is unstable with respect to the BNN dynamics.

To prove asymptotic stability for $a < 2$, note that (8) implies

$$P_2 = \frac{1}{2} a P_1^2 - \frac{1}{2} E(P, P).$$

Thus,

$$\dot{P}_2 = aP_1 \dot{P}_1 - \frac{1}{2} \frac{\mathrm{d}}{\mathrm{d}t} E(P, P).$$

Since the game is doubly symmetric, we have furthermore that

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} E\bigl(P(t), P(t)\bigr) &= E(\dot{P}, P) + E(P, \dot{P}) = 2E(\dot{P}, P) \\ &= 2\Sigma(P) E\bigl(R^P - P, P\bigr) \geqslant 0, \end{aligned} \tag{11}$$

where the last inequality follows from (4). By (10) we obtain

$$\begin{aligned} \dot{P}_2 &= aP_1 \Sigma(P)\Bigl(\frac{a}{2} - 1\Bigr) P_1 - 2\Sigma(P) E\bigl(R^p - P, P\bigr) \\ &= \Sigma(P)\Bigl[aP_1^2\Bigl(\frac{a}{2} - 1\Bigr) - 2E\bigl(R^p - P, P\bigr)\Bigr] \leqslant 0. \end{aligned}$$

[4] We ignore here boundary effects, assuming essentially $S = \mathbb{R}$. For a compact interval S the result follows from the analysis in Section 6, see Proposition 4.

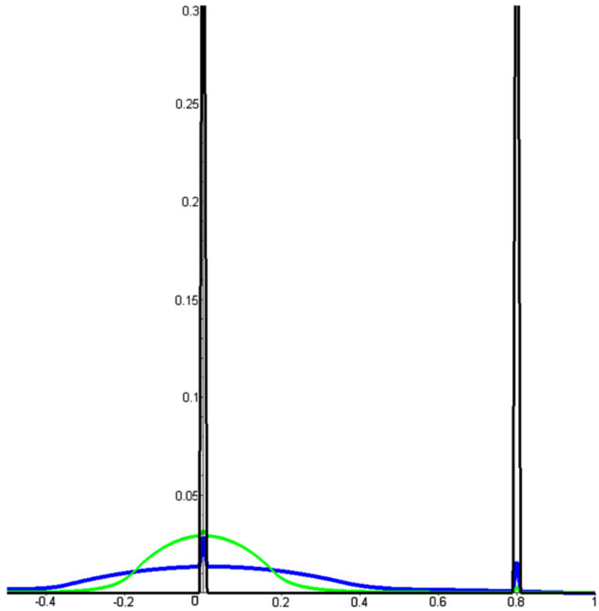

Fig. 1. Simulated BNN dynamics for the payoff function $f(x, y) = -x^2 - 2xy$. The initial population has mass 0.5 on the two points 0 and 0.8. After a few steps of the discretized BNN dynamics, the point masses have decreased drastically and the distribution is dispersed between −0.4 and 0.4. The gray curve shows that after 100 steps the distribution starts concentrating around the long run equilibrium 0. In fact, we have convergence with respect to the weak topology, but no convergence with respect to the strong topology.

The second term in the bracket has the required sign by (11). The first term is negative for $0 < a < 2$. The inequality is strict unless $P = \delta_0$, which proves that P_2 is a Lyapunov function for the BNN dynamics. □

The fact that the parameter a is decisive for stability, suggests that second derivatives of f may play an important role. The following sections show that this intuition is correct.

The above example also shows that δ_0 is unstable in the strong topology arising from the variational norm.[5] Even if the initial measure has some positive mass on 0, this will disperse into a smooth distribution of better replies near 0 and the mass at 0 will decrease to 0, see Fig. 1 for a numerical example. This is in contrast to the replicator dynamics. In the above example 0 is an "uninvadable" strategy (in the sense of Bomze, 1991), which implies that it is stable in the strong topology with respect to the replicator dynamics for every $a \in \mathbb{R}$, see Oechssler and Riedel (2001, Theorem 3).

5. Stability and the measure of closeness

Most, if not all, relevant strategy spaces carry an appropriate metric. For subsets of $\mathbb{R}^n$, there is the Euclidean distance. When considering Bayesian games, strategies are given by certain

[5] In the variational norm the distance between two probability measures P and Q is given by $\|P - Q\| = 2\sup_{A\in\mathcal{A}} |P(A) - Q(A)|$ (see e.g. Shiryaev, 1995).

classes of functions that also come with metrics. For this reason, we assume from now on that S carries a metric d.

The choice of topology is an important issue when defining dynamic stability as one has to specify what it means for a populations Q to be "close" to a given population P. See Oechssler and Riedel (2002) for an extensive discussion on this. For the reasons stated there, we find it most appropriate to use the topology of weak convergence to measure closeness of populations in evolution. Note also that the BNN dynamics tends to disperse the population as it introduces all strategies that are fitter than average even if we start with a homogeneous population. If μ is, e.g., Lebesgue measure, the distance between the trajectory and a possible stationary homogeneous state is always maximal in the strong topology. As a consequence, BNN dynamics cannot converge in the strong topology, see Fig. 1 for an example.

The weak topology is defined as follows. A sequence (P_n) converges weakly to P if $\int_S f\,\mathrm{d}P_n \to \int_S f\,\mathrm{d}P$ for every bounded, continuous real function f. The Prohorov metric can be used to measure the distance between populations. It is defined as (cf. Billingsley, 1968, p. 238)

$$\rho(P,Q) := \inf\bigl\{\varepsilon > 0\colon\ Q(A) \leqslant P\bigl(A^{\varepsilon}\bigr) + \varepsilon \text{ and } P(A) \leqslant Q\bigl(A^{\varepsilon}\bigr) + \varepsilon,\ \forall A \in \mathcal{A}\bigr\},$$

where $A^{\varepsilon} := \{x\colon\ \exists y \in A, d(y,x) < \varepsilon\}$. Thus, P_n converges weakly to P if and only if $\rho(P_n, P) \to 0$.

The weak topology captures the following notion of closeness. If $Q = (1-\varepsilon)\delta_u + \varepsilon\delta_x$ with $0 \leqslant \varepsilon \leqslant 1$, then $\rho(\delta_u, Q) = \min\{\varepsilon, d(u,x)\}$. Thus, population Q is close to the Dirac measure δ_u only if a small subpopulation deviates to a (possibly far away) pure strategy x or if a (possibly large) part of the population deviates to a nearby strategy x. In particular, the distance between two homogeneous populations agrees with the natural metric on the set of pure strategies, i.e. $\rho(\delta_u, \delta_x) = d(u,x)$, when u and x are close to each other.

The next definition specifies the *dynamic* stability concepts we will use in the following.

Definition 1. Let $Q^\star$ be a rest point of the BNN dynamics. Then

- $Q^\star$ is called (*Lyapunov*) *stable* if for all $\varepsilon > 0$ there exists an $\eta > 0$ such that $\rho(Q(0), Q^\star) < \eta \Rightarrow \rho(Q(t), Q^\star) < \varepsilon$ for all $t > 0$.
- $Q^\star$ is called *asymptotically stable* if additionally there exists $\varepsilon > 0$ such that $\rho(Q(0), Q^\star) < \varepsilon \Rightarrow \rho(Q(t), Q^\star) \to 0$.

Dynamic stability can be related to a number of static stability concepts which have the advantage that they can easily be checked given the payoff function. Since strictness of Nash equilibrium is not sufficient for dynamic stability, stronger concepts are required. As it turns out, concepts originally developed for the continuous version of the replicator dynamics in evolutionary biology like CSS (Eshel and Motro, 1981) and Evolutionary Robustness (Oechssler and Riedel, 2002) become relevant for the BNN dynamics as well.

The classical definition of an evolutionary stable strategy (ESS) (Maynard Smith, 1974) requires that for all mutant populations R there exists an invasion barrier ε such that the original population P does better against the mixed population $(1-\eta)P + \eta R$ than R does for all $\eta \leqslant \varepsilon$. In this definition some invasion barrier exists for each R.

Eshel and Motro (1981) introduced the following definition for $S \subset \mathbb{R}$.

Definition 2 *(CSS).* A strategy u is a *continuously stable strategy (CSS)* if (1) it is an ESS and (2) there exists an $\varepsilon > 0$ such that for all $v \in S$ with $|v - u| < \varepsilon$ there exists an $\eta > 0$ such that for all $x \in S$ with $|v - x| < \eta$

$$f(v, x) > f(x, x) \quad \text{if and only if} \quad |v - u| < |x - u|. \tag{12}$$

As shown by Eshel (1983), if f is twice differentiable, a necessary condition for an interior ESS u to be a CSS is that

$$f_{xx}(u, u) + f_{xy}(u, u) \leqslant 0. \tag{13}$$

Condition (13) is sufficient if the weak inequality is replaced by a strict one.

The following condition was introduced by Oechssler and Riedel (2002) and is stronger than CSS.

Definition 3. A population $P^* \in \Delta(S)$ is *evolutionarily robust* if there exists $\varepsilon > 0$ such that for all $Q \neq P^*$ with $\rho(Q, P^*) < \varepsilon$ we have

$$E(P^*, Q) > E(Q, Q). \tag{14}$$

When (14) holds for all $Q \neq P$, P is called *globally evolutionarily robust*. If f is twice differentiable, a necessary condition for an interior, homogeneous δ_u to be evolutionary robust is that (cf. Oechssler and Riedel, 2002)

$$f_{xx}(u, u) + 2 f_{xy}(u, u) \leqslant 0. \tag{15}$$

6. Doubly symmetric games

Games in which all players have the same payoff function f and which have a symmetric payoff function, $f(x, y) = f(y, x)$ for all $x, y \in S$, are called doubly symmetric. Doubly symmetric games or games that can be transformed into the symmetric form (as the one in Example 1) have the property that the mean payoff $E(P, P)$ is increasing along every solution of BNN.

Lemma 1. *Let $(S, \mathcal{A}, \mu)$ be a measure space with μ a finite measure on S. Consider a doubly symmetric game. Then the mean payoff $E(P, P)$ is monotonically increasing along every solution of BNN, and strictly increasing along every nonstationary solution. A (strict) local maximizer of mean payoff is (asymptotically) stable under BNN. If (S, d) is a compact metric space, μ a finite Borel measure on S with full support, and f is Lipschitz continuous, then the set of limit points of any trajectory is a non-empty connected compact set of Nash equilibria.*

Proof. The fact that mean payoff $E(P, P)$ is monotonically increasing along every solution of BNN follows directly from (11), as we have

$$\dot{\Lambda}\big(P(t)\big) = -\frac{\mathrm{d}}{\mathrm{d}t} E\big(P(t), P(t)\big) = -2\Sigma\big(P(t)\big) E\big(R^{P(t)} - P(t), P(t)\big) \leqslant 0,$$

and the inequality is strict for nonstationary trajectories. If P^* is a (strict) local maximizer of mean payoff, $\Lambda(Q) := E(P^*, P^*) - E(Q, Q)$ is a Lyapunov function. From a suitable generalization of Lyapunov's theorem (see e.g. Oechssler and Riedel, 2002, Appendix B or Bhatia and Szegö, 1970, Chapter V), we obtain (asymptotic) stability.

If S is a compact metric space, then the set of limit points of any trajectory is non-empty. If additionally f is Lipschitz continuous, Theorem 2 implies that the semi-flow is weakly continuous and from standard results in dynamic systems theory (see e.g. Bhatia and Szegö, 1970, Chapter II) the set of limit points is compact and connected. Moreover, average payoff is constant on this set. Hence, each ω-limit point of a trajectory is stationary and, by Proposition 1, a Nash equilibrium. □

We use the above general result to show local or global asymptotic stability of an equilibrium. We will demonstrate this for two classes of games: Games with an equilibrium that satisfies evolutionary robustness and games with a unique Nash equilibrium.

Proposition 2. *Let S be compact and f Lipschitz continuous. If a doubly symmetric game has a unique Nash equilibrium P^*, then P^* is globally asymptotically stable under BNN.*

Proof. From Lemma 1, we get that P^* is asymptotically stable. From the proof, one sees immediately that $\Lambda(Q) := E(P^*, P^*) - E(Q, Q)$ is a global Lyapunov function, and global convergence follows. □

Proposition 3. *Consider a doubly symmetric game with Lipschitz continuous payoff function f and compact metric strategy space S. If P^* is evolutionarily robust, then P^* is asymptotically stable with respect to BNN.*

Proof. By definition of evolutionary robustness, we have for Q close to P^*,

$$\begin{aligned}\Lambda(Q) &= E(P^*, P^*) - E(Q, Q)\\ &= E(P^*, P^*) - E(P^*, Q) + E(P^*, Q) - E(Q, Q)\\ &\geqslant E(P^*, P^*) - E(P^*, Q)\\ &= E(P^*, P^*) - E(Q, P^*) \geqslant 0,\end{aligned}$$

where the last equality follow from symmetry of f, and the last inequality from the fact that every evolutionary robust population is a symmetric Nash equilibrium. Note that the first inequality above becomes strict unless $Q = P^*$. The result then follows from Lemma 1. □

For replicator dynamics and finite (doubly symmetric) games, ESS is equivalent to asymptotic stability. So one might conjecture that evolutionary robustness is equivalent to asymptotic stability for BNN in the current setting. However, this is not the case as the linear-quadratic case introduced in Example 2 shows. We are now ready to completely characterize the dynamic behavior of BNN dynamics for this case.

Proposition 4. *Let $S = [-A, B] \subset \mathbb{R}$ be an interval around 0 and $f(x, y) = -x^2 + axy - y^2$ be a linear-quadratic game. For all parameters $a \in \mathbb{R}$, $(0, 0)$ is a strict Nash equilibrium.*

- *For $a < 2$, 0 is a strict maximizer of mean payoff and BNN dynamics globally converge to δ_0.*
- *For $a > 2$, δ_0 is unstable. There are two other symmetric strict Nash equilibria at the boundary of S, $(-A, -A)$ and (B, B). δ_{-A} and δ_B are strict local maximizers of mean payoff and hence asymptotically stable under BNN.*

- *For $a = 2$, there is a continuum of pure strategy Nash equilibria (x, x) for all $x \in S$ and BNN dynamics converge to this set of Nash equilibria.*

In particular, δ_0 is asymptotically stable if and only if it is CSS.

Proof. The proposition follows from Lemma 1 once we note that (cf. (8))

$$E(P, P) = aP_1^2 - 2P_2 = (a-2)P_1^2 - \operatorname{Var}(P),$$

where $\operatorname{Var}(\cdot)$ denotes the variance. Thus, maximizing mean payoff implies that $\operatorname{Var}(P) = 0$. For $a < 2$, δ_0 strictly maximizes mean payoff. For $a > 2$, both δ_{-A} and δ_B are local maximizers. To see instability of δ_0, note that $E(\delta_0, \delta_0) = 0$ and $E(\delta_x, \delta_x) = (a-2)x^2 > 0$ for $x \neq 0$. Furthermore, average payoff $E(P(t), P(t))$ is strictly increasing if one starts close to δ_0. Lyapunov's instability theorem then yields that δ_0 is unstable. For $a = 0$, $E(P, P) = -\operatorname{Var}(P)$. Hence, all pure strategies and only those maximize mean payoff. By Lyapunov's stability theorem, the set of all pure strategies is globally asymptotically stable. □

Note that for $1 < a < 2$, 0 is a CSS but δ_0 is not evolutionary robust (see Eqs. (13) and (15)). Therefore, evolutionary robustness is sometimes too strong a condition and is not necessary for asymptotic stability. For linear-quadratic games, Proposition 4 shows that CSS is necessary and sufficient for asymptotic stability.[6]

The next proposition shows that CSS is a necessary condition for general symmetric payoff functions.

Proposition 5. *Let S be an interval in $\mathbb{R}$, with x^* in the interior of S, and let f be twice continuously differentiable and symmetric. If δ_{x^*} is asymptotically stable with respect to BNN, then x^* satisfies the CSS condition* (13)*.*

Proof. By (11), every asymptotically stable state δ_{x^*} must correspond to a local maximum of mean payoff $E(P, P)$. In particular, x^* must be a maximum of $f(x, x)$. The necessary second order condition for x^* to be a maximum is $f_{xx}(x^*, x^*) + 2f_{xy}(x^*, x^*) + f_{yy}(x^*, x^*) \leqslant 0$, which reduces due to symmetry of f to the CSS condition (13). □

Given the insight from Example 2 and the previous proposition one might hope that for general payoff functions, CSS is also sufficient for a homogeneous population $P^* = \delta_{u^*}$ to be asymptotically stable with respect to BNN. Since this result holds for quadratic payoff function, one may further conjecture that the general result can be proven by using a second order Taylor approximation of the payoff function. The following example shows that this is unfortunately not the case in general.

Example 3. Let $S = [-1, 1]$ and $f(x, y) = 10x^4 - x^2 - xy$.

Let $g(x, y) = -x^2 - xy$ be the quadratic approximation (through a second order Taylor approximation) at $x = y = 0$. Since $g(x, y)$ satisfies condition (16), Theorem 3 below implies that

[6] Since (13) does not apply for the two boundary equilibria, one needs to check the CSS definition (12) directly. For $a > 2$, δ_{-A} and δ_B are CSS.

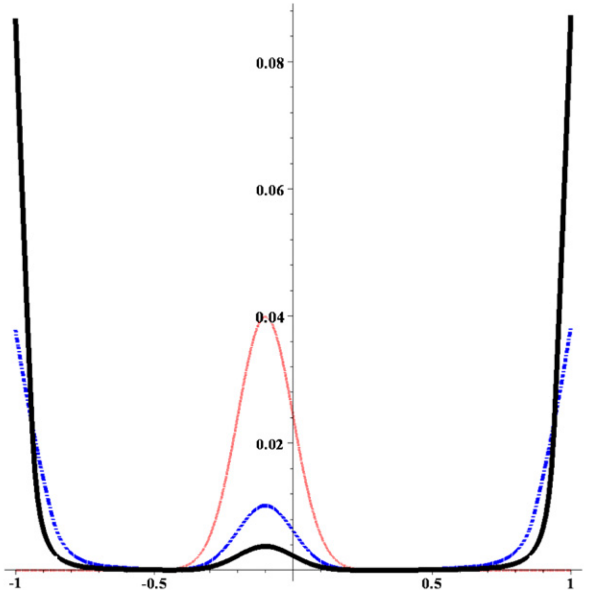

Fig. 2. Simulation of time-discretized BNN dynamics with payoff function $f(x, y) = 10x^4 - x^2 - xy$ on strategy space $\{-1, -0.99, \ldots, 0.99, 1\}$. Initial distribution (light gray) is truncated normal. The dashed line shows simulation after 10 iterations, the solid black line after 50 iterations.

δ_0 is globally asymptotically stable for the payoff function $g(x, y)$. However, for the actual payoff function f, the BNN dynamics converges to $Q = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$ from any initial $P(0) \neq \delta_0$. This follows from Proposition 2 or Theorem 3 since Q is the unique Nash equilibrium of this negative semidefinite game. Fig. 2 shows a typical simulation of time-discretized BNN dynamics where the initial population is a discretized, truncated normal distribution whose mean can be arbitrarily close to 0. Clearly, the dynamics diverges to $Q = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$.

Note that $(0, 0)$ is not a Nash equilibrium for the payoff function $f(x, y)$, and hence not even a stationary point under BNN. However it is a local strict Nash equilibrium and satisfies condition (13) for CSS. Hence, δ_0 is stable for the replicator dynamics w.r.t. initial distributions with support close to 0 and attracts such initials whose support is an interval containing 0, see Cressman and Hofbauer (2005), Cressman et al. (2006).

The class of doubly symmetric games may appear restrictive (and it is) but there are many games that can easily be transformed into a such game as the following example shows.

Example 4. Consider a Bertrand game with heterogeneous goods. Two firms simultaneously set prices p_1 and p_2. The demand function for firm i given by $1 - p_i + \gamma p_{-i}$ with $0 < \gamma < 2$. Firm 1's payoff function is then

$$f(p_1, p_2) = p_1(1 - p_1 + \gamma p_2).$$

By adding $p_2 - p_2^2$ to the payoff function, this game can be transformed into a doubly symmetric one.[7] Thus, by Proposition 2 BNN dynamics globally converge (also for the original game) to the unique symmetric Nash equilibrium $p_1^* = p_2^* = \frac{1}{2-\gamma}$.

[7] Note again that $\sigma(x, P)$ is not affected by adding terms that depend only on y. Thus, the dynamics do not change.

7. Negative semidefinite games

In this section we will consider games with an expected payoff function that is negative semi-definite in the sense that for all $P, Q \in \Delta$

$$E(P-Q, P-Q) \leqslant 0 \tag{16}$$

(see e.g. Hofbauer, 2000 for the corresponding property in finite games; in Hofbauer and Sandholm, 2007 these are called 'stable games'). Note first that linear quadratic games like $f(x,y) = -x^2 + axy$ satisfy condition (16) if and only if $a \leqslant 0$ as one can easily check. Furthermore, every symmetric zero-sum game satisfies condition (16). By definition of a symmetric zero-sum game, $f(x,y) + f(y,x) = 0$ for all $x, y \in S$. This implies that $E(P,Q) + E(Q,P) = 0$, and in particular $E(P,P) = 0$. Therefore, $E(P-Q, P-Q) = 0$. Further examples for negative semi-definite games include contests (see Example 5 below) and the War of Attrition (see Example 6). Finally, it is well known (see e.g. Hofbauer and Sigmund, 1988, p. 122) that all finite games with an interior ESS satisfy (16) with strict inequality for all $P \neq Q$.

Lemma 2. (1) *Under condition* (16)*, the set of Nash equilibria is convex.*

(2) *If either there exists a strict Nash equilibrium or condition* (16) *holds with strict inequality* (*and at least one NE exists*)*, there is a unique Nash equilibrium, which is, furthermore, globally evolutionarily robust.*

Proof. (1) Suppose P^* and Q^* are Nash equilibria. By condition (16) we have that

$$E(P^*,P^*) + E(Q^*,Q^*) \leqslant E(P^*,Q^*) + E(Q^*,P^*) \leqslant E(Q^*,Q^*) + E(Q^*,P^*),$$

which implies that $E(P^*,P^*) = E(Q^*,P^*)$. Thus, any convex combination $P_\lambda = \lambda P^* + (1-\lambda)Q^*$ is also a best reply against P^*, and similarly against Q^*. Since for all Q

$$\begin{aligned} E(P_\lambda, P_\lambda) &= \lambda E(P_\lambda, P^*) + (1-\lambda)E(P_\lambda, Q^*) \\ &\geqslant \lambda E(Q, P^*) + (1-\lambda)E(Q, Q^*) \\ &= E(Q, P_\lambda), \end{aligned}$$

P_λ is also a Nash equilibrium which proves that the set of Nash equilibria is convex.

(2) Let P^* be a Nash equilibrium and $Q \neq P^*$. Then

$$E(P^*-Q, Q) = E(P^*-Q, Q-P^*) + E(P^*-Q, P^*).$$

The first term is nonnegative by condition (16) and the second term is nonnegative by definition of a Nash equilibrium. For a strict Nash equilibrium, the second term is strictly positive. If (16) holds with strict inequality, the first term is strictly positive. In either case, $E(P^*-Q, Q) > 0$, that is, P^* is globally evolutionarily robust. This in turn implies that there is no other Nash equilibrium because $E(Q,Q) < E(P^*,Q)$ for all $Q \neq P^*$. □

We may now proceed to study the global stability properties of Nash equilibria in negative semidefinite games.

Theorem 3. *For negative semidefinite games* (16)*, define the function*

$$H(P) = \frac{1}{2}\int_S \sigma_+(x,P)^2 \mu(\mathrm{d}x).$$

The following statements hold true:

(1) H *is nonnegative and decreases to* 0 *along every solution of BNN.*
(2) *If* S *is a compact metric space,* f *is continuous, and* μ *a measure with full support, then every trajectory of BNN converges to the set of Nash equilibria.*
(3) *In particular, every strict Nash equilibrium and every equilibrium that satisfies* (16) *with strict inequality is globally asymptotically stable.*

Proof. Let us first determine the gradient of $\sigma(x, P)$ with respect to P at some point Q. We have

$$\nabla\sigma(x, P)(Q) = E(\delta_x, Q) - E(P, Q) - E(Q, P).$$

From this, we obtain via the chain rule

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}H(P) &= \int_S \sigma_+(x, P)\nabla\sigma(x, P)(\dot{P})\mu(\mathrm{d}x) \\ &= \int_S \sigma_+(x, P)\big[E(\delta_x, \dot{P}) - E(P, \dot{P}) - E(\dot{P}, P)\big]\mu(\mathrm{d}x) \\ &= \Sigma(P)\big(E\big(R^P, \dot{P}\big) - E(P, \dot{P}) - E(\dot{P}, P)\big),\end{aligned}$$

where we have used the definition of the relative excess measure R^P. By definition of the dynamics $\dot{P}_t$, we proceed to

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}H(P) &= \Sigma(P)^2\big(E\big(R^P, R^P - P\big) - E\big(P, R^P - P\big) - E\big(R^P - P, P\big)\big) \\ &= \Sigma(P)^2\big(E\big(R^P - P, R^P - P\big) - E\big(R^P - P, P\big)\big).\end{aligned}$$

The first term in parentheses, $E(R^P - P, R^P - P) \leqslant 0$ by Assumption (16), and the second term, $E(R^P - P, P) \geqslant 0$ by definition of the relative excess measure (see (4)). Thus, we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t}H(P) \leqslant 0$$

and the inequality is strict whenever P is not a stationary point.

If S is a compact metric space and f Lipschitz continuous, then Δ is compact (in the weak topology) and BNN generates a weakly continuous semi-flow by Theorem 2. Hence, the ω-limit set (in the weak topology) of a trajectory $P(t)$ is non-empty and contained in the set of $P \in \Delta$ for which $\frac{\mathrm{d}}{\mathrm{d}t}H(P) = 0$, which is the set of Nash equilibria by Proposition 1. Thus, the BNN dynamics converge to the convex set of Nash equilibria. In particular, every strict Nash equilibrium and every equilibrium that satisfies (16) with strict inequality, is a unique equilibrium and, therefore, globally asymptotically stable.

For general measure spaces S, we proceed similar to Brown and von Neumann (1950) and Nikaido (1959). The above expression implies

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}H(P) &\leqslant -\Sigma(P)^2 E\big(R^P - P, P\big) \\ &= -\Sigma(P)^2\bigg(\int E(\delta_x, P)\frac{\sigma_+(x, P)}{\Sigma(P)}\mu(\mathrm{d}x) - E(P, P)\bigg)\end{aligned}$$

$$= -\Sigma(P)^2 \int \left[E(\delta_x, P) - E(P, P)\right] \frac{\sigma_+(x, P)}{\Sigma(P)} \mu(\mathrm{d}x)$$

$$= -\Sigma(P)^2 \int \frac{\sigma_+(x, P)^2}{\Sigma(P)} \mu(\mathrm{d}x) = -2\Sigma(P)H(P). \tag{17}$$

Since f is bounded, $\sigma_+(x, P) \leqslant 2\|f\| =: \frac{1}{c}$. Hence,

$$\Sigma(P) = \int_S \sigma_+(x, P)\mu(\mathrm{d}x) \geqslant c \int_S \sigma_+(x, P)^2 \mu(\mathrm{d}x) = cH(P).$$

Therefore, (17) implies the differential inequality

$$\frac{\mathrm{d}}{\mathrm{d}t} H(P) \leqslant -2cH(P)^2,$$

which integrates to $H(P(t)) \leqslant \frac{H(P(0))}{1+2cH(P(0))t}$. Hence $H(P(t)) \to 0$ as claimed. □

Example 5. Contests. Let $S = [a, b]$ for some numbers $a < b$. Two players exert an effort level of x and y in S, respectively, to obtain a prize worth $K > 0$. The probability that player 1 wins the prize is $p(x, y)$, and the probability that player 2 wins is $p(y, x) = 1 - p(x, y)$. Costs are given by some cost function $c(x)$. The payoff function is thus $f(x, y) = Kp(x, y) - c(x)$. An example is the all-pay auction for which $p(x, y) = 1$ if $x > y$, $p(x, y) = 1/2$ if $x = y$, and $p(x, y) = 0$ else.

We claim that contests are negative semidefinite. To see this, note that because of $p(x, y) + p(y, x) = 1$, we have

$$E(P, Q) + E(Q, P) = K - \int c(x)P(\mathrm{d}x) - \int c(y)Q(\mathrm{d}y).$$

It follows that $E(P, P) = \frac{K}{2} - \int c(x)P(\mathrm{d}x)$, and this implies

$$E(P - Q, P - Q) = E(P, P) + E(Q, Q) - E(P, Q) - E(Q, P) = 0.$$

Now assume that $p(\cdot, \cdot)$ and $c(\cdot)$ are continuously differentiable, $p(\cdot, \cdot)$ is strictly concave in x, and $c(\cdot)$ is convex. If there exists $\bar{x} \in (a, b)$ with $K \frac{\partial}{\partial x} p(\bar{x}, \bar{x}) = c'(\bar{x})$, then $(\bar{x}, \bar{x})$ is a strict Nash equilibrium (which is then unique by Lemma 2). By Theorem 3, $\delta_{\bar{x}}$ is globally asymptotically stable under BNN dynamics.

Finally, we will demonstrate how our techniques are useful even when applied to games with discontinuous payoff function, like the war of attrition.

Example 6. War of attrition. Consider two players fighting for a prize worth V to both players. A strategy is to choose a length of time $x \in [0, M]$ for which one is prepared to stay in the race. Fighting is costly. The payoffs are given as follows

$$f(x, y) = \begin{cases} V - y & \text{if } x > y, \\ \frac{V}{2} - x & \text{if } x = y, \\ -x & \text{if } x < y, \end{cases}$$

that is, a player gets the prize if he stays longer in the race than his rival but has to share if they stay equally long.

We assume that $M > V/2$. Otherwise waiting until the end is always profitable. Bishop and Cannings (1978) show that there is a unique Nash equilibrium, which has the following equilibrium distribution $P^\star$ with $t^\star = M - V/2$,

$$P^\star([0,x]) = \begin{cases} 1 - \mathrm{e}^{-x/V} & \text{if } x \leqslant t^\star, \\ 1 - \mathrm{e}^{-t^\star/V} & \text{if } t^\star < x < M, \\ 1 & \text{if } x = M. \end{cases}$$

Bishop and Cannings (1978) show that $P^\star$ is an ESS. They also show (Bishop and Cannings, 1978, p. 118) that

$$E(P-Q, P-Q) = -\int_0^M \big(P([s,M]) - Q([s,M])\big)^2 \,\mathrm{d}s. \tag{18}$$

In particular, the war of attrition is a negative semidefinite game. As the payoff function f is not continuous, we cannot apply the second part of the above theorem. Nevertheless, its conclusion still holds true provided the measure μ that defines the excess measure in the definition of BNN dynamics (see Eq. (2)) puts some weight on the point M. The intuition for this assumption is as follows: the Nash equilibrium has a mass point on M, but strategies close to M are not being played in equilibrium. If μ is the Lebesgue measure, the excess measure has a density with respect to the Lebesgue measure that is zero close to M, and strictly positive at M. However, a single value of the density does not contribute to the distribution, and thus, the excess measure puts no weight on M if μ is the Lebesgue measure. Consequently, there is in general no hope that BNN generates some mass on or around M if one uses the Lebesgue measure. Therefore, we assume that μ puts some small mass on M. The following proposition shows that this is sufficient for convergence.

Proposition 6. *Assume that $\mu = \mathrm{d}x + \varepsilon\delta_M$ for some (small) $\varepsilon > 0$. In the War of Attrition, every trajectory of BNN converges to the unique Nash equilibrium.*

Proof. Without loss of generality, we set $V = 1$ in the proof. Consider the Lyapunov function $H(P)$ as in Theorem 3, where we take $\mu = \mathrm{d}x + \varepsilon\delta_M$, the sum of the Lebesgue measure on $[0, M]$ and a point mass on M. The proof of $H(P(t)) \to_{t\to\infty} 0$ does not use continuity of f.

We show next that $H(P)$ is lower semi-continuous in the weak topology in the sense that $H(P) \leqslant \liminf H(P^n)$ if (P^n) converges in the weak topology to P. By symmetry, we have

$$\begin{aligned} 1 &= \iint 1_{\{x<y\}} P(\mathrm{d}x)P(\mathrm{d}y) + \iint 1_{\{y<x\}} P(\mathrm{d}x)P(\mathrm{d}y) \\ &\quad + \iint 1_{\{x=y\}} P(\mathrm{d}x)P(\mathrm{d}y) \\ &= 2\iint 1_{\{x<y\}} P(\mathrm{d}x)P(\mathrm{d}y) + \iint 1_{\{x=y\}} P(\mathrm{d}x)P(\mathrm{d}y). \end{aligned}$$

It follows that average payoff can be written as

$$\begin{aligned} E(P,P) &= \iint 1_{\{x<y\}} P(\mathrm{d}x)P(\mathrm{d}y) + \frac{1}{2}\iint 1_{\{x=y\}} P(\mathrm{d}x)P(\mathrm{d}y) \\ &\quad - \iint \min(x,y) P(\mathrm{d}x)P(\mathrm{d}y) \\ &= \frac{1}{2} - \iint \min(x,y) P(\mathrm{d}x)P(\mathrm{d}y). \end{aligned}$$

As $\min(x, y)$ is continuous in (x, y), $E(P, P)$ is continuous in the weak topology. For points x with $P(\{x\}) = 0$, $E(\delta_x, P) = P([0, x)) - \int \min(x, y)P(\mathrm{d}y)$. By the Portmanteau Theorem, $P \mapsto P([0, x))$ is continuous at P in the weak topology for these x. The function $x \mapsto \int \min(x, y)P(\mathrm{d}y)$ is continuous in the weak topology because the integrand is continuous. We conclude that $\sigma_+(x, P)$ is continuous at P in the weak topology for all x with $P(\{x\}) = 0$. Now let $P^n \to P$ in the weak topology. Then $\lim \sigma_+(x, P^n) = \sigma_+(x, P)$ for all points x with $P(\{x\}) = 0$. As the set of points x with $P(\{x\}) = 0$ has full Lebesgue measure, and the payoff function is bounded, we get by dominated convergence that $\lim \frac{1}{2} \int_S \sigma_+(x, P^n)^2\, \mathrm{d}x = \frac{1}{2} \int_S \sigma_+(x, P)^2\, \mathrm{d}x$. Hence, the first part of H is continuous in the weak topology. Now consider $\sigma_+(M, P) = 1 - \frac{1}{2}P(\{M\}) - \int xP(\mathrm{d}x)$. By the Portmanteau Theorem, $P(\{M\}) \geqslant \limsup P^n(\{M\})$. Therefore, $\sigma_+(M, P) \leqslant \liminf \sigma_+(M, P^n)$. This finally establishes $H(P) \leqslant \liminf H(P^n)$.

As H is lower semi-continuous in the weak topology, we conclude that every limit point P^0 of BNN dynamics satisfies $H(P^0) = 0$. It follows that $\sigma(x, P^0) \leqslant 0$ for μ-almost all $x \in [0, M]$. It remains to be shown that this implies $\sigma(x, P^0) \leqslant 0$ for all x. As μ has a point mass on M, we have $\sigma(M, P^0) \leqslant 0$. Now consider some $x < M$. There exists a sequence (x^n) that converges to x from the right and satisfies $P^0(\{x^n\}) = 0$ as well as

$$\sigma\left(x^n, P^0\right) \leqslant 0 \quad \text{for all } n.$$

It follows that

$$\begin{aligned} E\left(P^0, P^0\right) &\geqslant \lim E\left(\delta_{x^n}, P^0\right) = \lim P^0\left(\left[0, x^n\right)\right) - \int \min\left(x^n, y\right)P^0\left(\mathrm{d}y\right) \\ &= P^0\left([0, x]\right) - \int \min(x, y)P^0(\mathrm{d}y) \\ &\geqslant P^0\left([0, x)\right) + \frac{1}{2}P^0\left(\{x\}\right) - \int \min(x, y)P^0(\mathrm{d}y) \\ &= E\left(\delta_x, P^0\right). \end{aligned}$$

This establishes $\sigma(x, P^0) \leqslant 0$ for all $x < M$. □

8. Discussion

To put the contribution of the current paper into perspective it is useful to compare the BNN dynamics and its properties to its two main “competitors” when it comes to describing human learning behavior in games.

Under the *replicator dynamics*, the population share of a pure strategy grows at the rate of its excess payoff (the current payoff minus the current average payoff of the population). Replicator dynamics emerge if one interprets payoff as number of offspring and uses a purely selective Darwinian model (Taylor and Jonker, 1978). It is also possible to give an economic microfoundation based on imitation (Schlag, 1998). The replicator dynamics are the prototype of *regular selection dynamics* where selection pressure operates on existing strategies. Regular selection dynamics do not allow the appearance (or reappearance) of extinct strategies.

In contrast, both, BNN and best response dynamics, are *innovative* dynamics, that is, new strategies can enter the population. *Best response dynamics* (Gilboa and Matsui, 1991; Hofbauer, 1994) move in the direction of the current best reply. They emerge when (new) individuals in the population behave myopically rational in the sense that they exactly know the current population distribution and play a best response to this distribution. Since small changes in the population state can change the best response to it, best response dynamics are not *continuous*, in contrast to the other two considered dynamics.[8]

BNN dynamics can be viewed as the prototype of innovative, better-reply dynamics that lie in between the replicator dynamics, which arguably require too little rationality to do justice to human behavior, and the best response dynamics, which may require too much of it. With BNN dynamics, all strategies with a positive excess payoff increase proportionally to their current excess payoff. Strategies with negative current excess payoff decrease at the same rate. In some sense, BNN dynamics require less rationality than best response dynamics as they just require that some better than average strategy be chosen by new members (see also the microfoundation in Appendix B). Moreover, the probability of choosing such a strategy is proportional to the excess payoff, which is close in spirit to the replicator dynamics.

An important property that all three dynamics share is that of being *myopic adjustment dynamics* (Swinkels, 1993), i.e. they tend to move in the direction of better replies. As a consequence, asymptotically stable states of all dynamics satisfy strong rationality postulates. In this sense, all dynamics tend to support rational predictions of equilibrium play. On the other hand, Berger and Hofbauer (2006) have recently given an example of a game in which a strictly dominated strategy does not die out under BNN. For replicator dynamics, Samuelson and Zhang (1992) have shown that strictly dominated strategies always become extinct (as they do, obviously, under the best response dynamics).

We now consider rest points. For regular selection dynamics, all pure strategies are rest points. For replicator dynamics, a rest point P is a Nash equilibrium in a game that has only the support of P as pure strategies. Under BNN dynamics, the set of rest points coincides with the set of Nash equilibria. As we show in our paper, in games with a continuum strategy set, the latter statement holds when the payoff function is continuous. The property of Nash stationarity is also shared by the best response dynamics.

As far as equilibrium predictions of the three dynamics are concerned, it is best to consider specific examples that highlight the similarities and differences that can generally occur.

For payoff functions of the form $f(x,y) = -x^2 + axy$ we understand completely the dynamics of replicator, BNN and best response dynamics. These games have a strict (symmetric) Nash equilibrium at 0. Note that $a/2$ is the slope of the best reply function in these games. The Nash equilibrium 0 is *evolutionary robust* (ER) if $a < 1$. Asymptotic stability for replicator dynamics is a subtle issue, see Cressman (2005), Cressman and Hofbauer (2005), Cressman et al. (2006), Oechssler and Riedel (2002). Replicator dynamics cannot converge to a point mass on 0 if 0 is not in the support of the initial distribution. A meaningful definition of asymptotic stability thus restricts to initial populations that have 0 in their support. Cressman and Hofbauer (2005), building on Oechssler and Riedel (2002), show that the point mass on 0 is asymptotically stable for replicator dynamics (in this restricted sense) if $a < 2$. Eshel and Sansone (2003) have shown that this condition is also necessary for asymptotic stability. In summary, asymptotic stability for the replicator dynamics is equivalent to evolutionary robustness. In this paper, we show that the point mass on 0 is asymptotically stable for BNN if and only if 0 is CSS or continuously stable (Eshel, 1983), which is the case if $a < 2$. The (smooth) best response dynamics are here

[8] In general, the best response dynamics are not continuous, even for continuous payoff functions. However, under suitable regularity and smoothness conditions plus quasi-concavity of the payoff function, best replies are unique and depend smoothly on the strategy of the opponent. In these classes of games, the best response dynamics are also continuous.

Table 1
Properties of the different dynamics

Property	BNN	Replicator	Best response
regular selection	no	yes	no
innovative	yes	no	yes
continuous	yes	yes	no
rest points = Nash equilibria	yes*	no	yes
dominated strategies die out	no	yes	yes
quadratic games:			
asymptotically stable points =	CSS ($a < 2$)	ER ($a < 1$)	CSS ($a < 2$)

* For continuous payoff functions.

$\dot{x} = BR(x) - x = (a/2 - 1)x$. Again, 0 is asymptotically stable if and only if $a < 2$, i.e. when it is CSS.

We collect the above discussion in Table 1.

Thus, although the different dynamics share many features, there are some straightforward testable differences that could be checked in a simple laboratory experiment. For example, a testable difference between BNN and replicator dynamics arises in the following game. Consider the quadratic game $f(x, y) = -x^2 + 3/2xy$ and take the strategy set to be $S = [0, 1]$. Under BNN, players' behavior converges to 0 (according to Proposition 4 above), whereas it converges to 1 under replicator dynamics if initially enough players use the strategy 1 and no player uses strategies in $[1/2, 1)$. An interesting experiment would thus be to study the convergence or nonconvergence of behavior for this game in the lab.

Acknowledgments

We would like to thank two referees of this journal seminar audiences at Berkeley, Stanford, UC Santa Barbara, Madison, Zürich, UCL, at the Econometric Society World Congress in London, and in particular Vince Crawford, Ken Judd, and Bill Sandholm for comments and discussions. We acknowledge financial support from the Austrian Science fund (P 15281), UCL's Centre for Economic Learning and Social Evolution (ELSE), and the SFB 504 by the Deutsche Forschungsgemeinschaft.

Appendix A

Proof of Theorem 1. The strategy for proving the theorem is the following. Denote by

$$F(Q) := \int_{\cdot} \sigma_+(x, Q)\mu(\mathrm{d}x) - Q(\cdot) \int_S \sigma_+(x, Q)\mu(\mathrm{d}x)$$

the right-hand side of the BNN-dynamics. Since F is neither bounded nor globally Lipschitz continuous on $\mathcal{M}^e$, we construct in the following lemma an auxiliary function $\tilde{F}$ which has these properties and coincides with F on Δ (see also Bomze, 1991). In particular, we show that $\tilde{F}$ satisfies a global Lipschitz condition

$$\exists K > 0 \quad \text{s.t.} \quad \forall \mu, \nu \in \mathcal{M}^e, \quad \left\| \tilde{F}(\mu) - \tilde{F}(\nu) \right\| \leqslant K \|\mu - \nu\|,$$

where $\|\cdot\|$ denotes the variational norm on $\mathcal{M}^e(S,\mathcal{A})$. The variational norm is given by

$$\|\mu\| = \sup_g \left|\int g \, \mathrm{d}\mu\right|,$$

where the sup is taken over all measurable functions $g : S \to \mathbb{R}$ bounded by 1, $\sup_{s\in S} |g(s)| \leqslant 1$. Endowed with the variational norm, $\mathcal{M}^e$ is a Banach space (see Shiryaev, 1995).

Standard arguments (see e.g. Zeidler, 1986, Corollary 3.9) then imply that the ordinary differential equation

$$\dot{Q}(t) = \tilde{F}\big(Q(t)\big), \quad Q(0) = P$$

has a unique solution $(Q(t))$. Finally, since $\dot{Q}(t)(S) = 0$, $Q(t)$ never leaves Δ, which implies that $(Q(t))$ also solves differential equation (2) on Δ.

Lemma 3. *Suppose f is bounded, then there exists a bounded, Lipschitz continuous function $\tilde{F} : \mathcal{M}^e \to \mathcal{M}^e$, which coincides with F on Δ,*

$$\tilde{F}(P) = F(P), \quad \forall P \in \Delta.$$

Proof. We define $\tilde{F}$ as

$$\tilde{F}(Q) = \big(2 - \|Q\|\big)_+ F(Q).$$

$\tilde{F}$ is zero for $\|Q\| \geqslant 2$. It is bounded and coincides with F on Δ because probability measures have norm 1. It remains to show that $Q \mapsto F(Q)$ is Lipschitz for $\|Q\| \leqslant 2$.

The estimates

$$\big|E(\delta_x, Q)\big| \leqslant \|f\|_\infty \|Q\|, \qquad \big|E(P, Q)\big| \leqslant \|f\|_\infty \|P\| \|Q\| \tag{19}$$

imply that for each $x \in S$, the functions $Q \mapsto \sigma(x, Q)$ and hence also $Q \mapsto \sigma_+(x, Q)$ are Lipschitz (for $\|Q\| \leqslant 2$) with a Lipschitz constant L independent of x. Then the map $Q \mapsto \hat{F}(Q)$ with $\hat{F}(Q)(A) = \int_A \sigma_+(x, Q)\mu(\mathrm{d}x)$ from $\mathcal{M}^e$ into itself is Lipschitz with Lipschitz constant $L\mu(S)$. In particular, also $Q \mapsto \Sigma(Q) : \mathcal{M}^e \to \mathbb{R}$ is Lipschitz. Hence $F(Q)$ is Lipschitz in Q. □

Proof of Theorem 2. In the following, we will use the metric $\|\cdot\|_{BL}$ on Δ which metrizes the weak topology (cf. Shiryaev, 1995, p. 352). Endowed with the BL-norm, $\mathcal{M}^e$ is a Banach space. For a Lipschitz continuous, bounded function $g : S \to \mathbb{R}$, let

$$\|g\|_{BL} := \sup_{x\in S} \big|g(x)\big| + \sup_{x\neq y} \frac{|g(x) - g(y)|}{d(x, y)}. \tag{20}$$

Abusing notation slightly, we define the dual norm $\|\cdot\|_{BL}$ on $\mathcal{M}^e(S,\mathcal{A})$ via

$$\|Q\|_{BL} = \sup\left\{\int g \, \mathrm{d}Q\right\}, \tag{21}$$

where the supremum is taken over all bounded, Lipschitz continuous functions g with $\|g\|_{BL} \leqslant 1$.

We prove below that we have

$$\big|\sigma(x, P) - \sigma(x, Q)\big| \leqslant L\|P - Q\|_{BL} \tag{22}$$

for some constant $L > 0$ and all populations P, Q and all strategies x. The same Lipschitz estimate holds true when we pass to the positive part, so we have

$$\left|\sigma_+(x,P) - \sigma_+(x,Q)\right| \leqslant L\|P-Q\|_{BL}. \tag{23}$$

This implies that the right-hand side of (2) is Lipschitz in the norm (21). The claim follows then from Gronwall's Lemma (see e.g. Zeidler, 1986, Propositions 3.10 and 3.11), which implies that

$$\left\|P(t) - Q(t)\right\|_{BL} \leqslant \mathrm{e}^{Lt}\left\|P(0) - Q(0)\right\|_{BL} \tag{24}$$

(for a possibly different L) and hence continuous dependence of solutions on initial conditions for finite time.

It remains to prove the claim (22). By boundedness and Lipschitz continuity of f, there exist constants $L_0, L_1 > 0$ such that for all strategies x, y, x', y'

$$\left|f(x,y)\right| \leqslant L_0, \tag{25}$$

$$\left|f(x,y) - f(x',y')\right| \leqslant L_1 \max\left\{\mathrm{d}(x,x'), \mathrm{d}(y,y')\right\}. \tag{26}$$

For $x = x'$, the latter inequality yields

$$\left|f(x,y) - f(x,y')\right| \leqslant L_1 d(y,y'). \tag{27}$$

Let R be a population. Define the function

$$g(y) = \int_S f(x,y) R(\mathrm{d}x).$$

Since R is a probability measure, (25) carries over to g,

$$\left|g(y)\right| \leqslant L_0,$$

and so does (27),

$$\left|g(y) - g(y')\right| \leqslant L_1 d(y,y').$$

Hence, g is a bounded and Lipschitz continuous function with $\|g\|_{BL} \leqslant L_0 + L_1$. We thus obtain

$$\begin{aligned}\left|E(R, P-Q)\right| &= \left|\int_S g(y)(P-Q)(\mathrm{d}y)\right| \\ &\leqslant \|g\|_{BL}\|P-Q\|_{BL} \\ &\leqslant (L_0 + L_1)\|P-Q\|_{BL}.\end{aligned} \tag{28}$$

By a symmetric argument, we also have

$$\left|E(P-Q, R)\right| \leqslant (L_0 + L_1)\|P-Q\|_{BL}. \tag{29}$$

Now, to prove our claim, note that

$$\begin{aligned}\left|\sigma(x,P) - \sigma(x,Q)\right| &\leqslant \left|E(\delta_x, P-Q)\right| + \left|E(Q,Q) - E(P,P)\right| \\ &\leqslant \left|E(\delta_x, P-Q)\right| + \left|E(Q-P, Q)\right| + \left|E(P, Q-P)\right|.\end{aligned}$$

Applying (28) for $R = \delta_x$ and $R = P$, as well as (29) for $R = Q$, we finally obtain

$$\left|\sigma(x,P) - \sigma(x,Q)\right| \leqslant 3(L_0 + L_1)\|P-Q\|_{BL},$$

and the proof is complete. □

Appendix B

Although the BNN dynamics have been mentioned as learning process by a number of authors (see e.g. Skyrms, 1990 or Swinkels, 1993), probably the most convincing microfoundation has recently been provided by Sandholm (2005, 2006). It runs as follows.

In a large population of players each player receives an opportunity to revise his strategy in each round according to a Poisson process with arrival rate R. When given the opportunity, a player switches from his current strategy s to a new strategy s' with conditional switching rate $\rho_{ss'}$. Benaim and Weibull (2003) and Sandholm (2003) show that the resulting stochastic process is well approximated by the deterministic continuous time process

$$\dot{p}_s = \sum_{s' \in S} p_{s'} \rho_{s's} - p_s \sum_{s' \in S} \rho_{ss'}, \tag{30}$$

where p_s is the share of the population playing pure strategy s. The first term describes the "inflow" to strategy s and the second term represents the "outflow" from s.

Our favorite interpretation of this revision opportunity for a player is that a player (the "parent") "dies," or leaves the population for other reasons, and is replaced by his "child," who imitates the parent's current strategy without really knowing why he should follow this strategy, i.e. without knowing its payoff. (Arguably, this is how education works all too often.)

What should the switching rate ρ_{ij} depend on? In contrast to evolutionary or imitation dynamics which allow only to play strategies which are already in use in the population, BNN dynamics are based on the assumption that unused strategies can be sampled. In particular, a player randomly samples one strategy, all strategies being equally likely to be sampled (this can be generalized. In fact, in Eq. (2) in the paper we do just that by allowing for a general measure $\mu(\mathrm{d}x)$ for sampling new strategies).

The player then compares the expected payoff of the sampled strategy $u(e_s, p)$ to the current average payoff of all strategies in the population $u(p, p)$ and switches only if the $u(e_s, p) > u(p, p)$. Furthermore, if the expected gain is positive, the player switches with a probability that is proportional to the expected gain in payoffs $u(e_s, p) - u(p, p)$. Plugging those switching probabilities into (30) yields

$$\dot{p}_s = \sum_{s' \in S} p_{s'} \bigl[u(e_s, p) - u(p, p)\bigr]_+ - p_s \sum_{s' \in S} \bigl[u(e_{s'}, p) - u(p, p)\bigr]_+,$$

which reduces to the discrete BNN dynamics

$$\dot{p}_s = \bigl[u(e_s, p) - u(p, p)\bigr]_+ - p_s \sum_{s' \in S} \bigl[u(e_{s'}, p) - u(p, p)\bigr]_+. \tag{31}$$

The continuous strategy version of (31) is given by Eq. (2) in the paper.

References

Benaim, M., Weibull, J.W., 2003. Deterministic approximation of stochastic evolution in games. Econometrica 71, 873–903.

Berger, U., Hofbauer, J., 2000. The Nash dynamics. Mimeo, University of Vienna.

Berger, U., Hofbauer, J., 2006. Irrational behavior in the Brown–von Neumann–Nash dynamics. Games Econ. Behav. 56, 1–6.

Bhatia, N.P., Szegö, G.P., 1970. Stability Theory of Dynamical Systems. Springer, Berlin–Heidelberg–New York.

Billingsley, P., 1968. Convergence of Probability Measures. Wiley, New York.

Bishop, D., Cannings, C., 1978. A generalized war of attrition. J. Theoret. Biol. 70, 85–124.
Bomze, I., 1991. Cross entropy minimization in uninvadable states of complex populations. J. Math. Biol. 30, 73–87.
Borel, E., 1921. La théorie du jeu les équations intégrales à noyau symétrique. Comptes Rendus de l'Académie 173, 1304–1308. English translation by Savage, L.: The theory of play and integral equations with skew symmetric kernels. Econometrica 21 (1953) 97–100.
Brown, G., von Neumann, J., 1950. Solutions of games by differential equations. In: Kuhn, H., Tucker, A. (Eds.), Contributions to the Theory of Games. In: Ann. of Math. Stud., vol. 24. Princeton Univ. Press, Princeton.
Cressman, R., 2005. Stability of the replicator equation with continuous strategy space. Math. Soc. Sci. 50, 127–147.
Cressman, R., Hofbauer, J., 2005. Measure dynamics on a one-dimensional continuous trait space: Theoretical foundations for adaptive dynamics. Theoret. Pop. Biol. 67, 47–59.
Cressman, R., Hofbauer, J., Riedel, F., 2006. Stability of the replicator equation for a single-species with a multi-dimensional continuous trait space. J. Theoret. Biol. 239, 273–288.
Eshel, I., 1983. Evolutionary and continuous stability. J. Theoret. Biol. 103, 99–111.
Eshel, I., Motro, U., 1981. Kin selection and strong evolutionary stability of mutual help. Theoret. Pop. Biol. 19, 420–433.
Eshel, I., Sansone, E., 2003. Evolutionary and dynamic stability in continuous population games. J. Math. Biol. 46, 445–459.
Friedman, D., 1991. Evolutionary games in economics. Econometrica 59, 637–666.
Gilboa, I., Matsui, A., 1991. Social stability and equilibrium. Econometrica 59, 859–867.
Hofbauer, J., 1994. Stability for the best response dynamics. Mimeo, University of Vienna.
Hofbauer, J., 2000. From Nash and Brown to Maynard Smith: Equilibria, dynamics and ESS. Selection 1, 81–88.
Hofbauer, J., Sandholm, W., 2007. Evolution in games with randomly disturbed payoffs. J. Econ. Theory 132, 47–69.
Hofbauer, J., Sigmund, K., 1988. The Theory of Evolution and Dynamical Systems. Cambridge Univ. Press.
Laslier, J.-F., Picard, N., 2002. Distributive politics and electoral competition. J. Econ. Theory 103, 106–130.
Maynard Smith, J., 1974. The theory of games and the evolution of animal conflicts. J. Theor. Biol. 47, 209–222.
Nash, J., 1951. Non-cooperative games. Ann. of Math. 54, 286–295.
Nikaido, H., 1959. Stability of equilibrium by the Brown von Neumann differential equation. Econometrica 27, 654–671.
Oechssler, J., Riedel, F., 2001. Evolutionary dynamics on infinite strategy spaces. Econ. Theory 17, 141–162.
Oechssler, J., Riedel, F., 2002. On the dynamic foundation of evolutionary stability in continuous models. J. Econ. Theory 107, 223–252.
Samuelson, L., Zhang, J., 1992. Evolutionary stability in asymmetric games. J. Econ. Theory 57, 363–391.
Sandholm, W.H., 2003. Evolution and equilibrium under inexact information. Games Econ. Behav. 44, 343–378.
Sandholm, W.H., 2005. Excess payoff dynamics and other well-behaved evolutionary dynamics. J. Econ. Theory 124, 149–170.
Sandholm, W.H., 2006. Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. Mimeo, University of Wisconsin.
Schlag, K., 1998. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. J. Econ. Theory 78, 130–156.
Shiryaev, A., 1995. Probability, second ed.. Springer-Verlag, New York.
Skyrms, B., 1990. The Dynamics of Rational Deliberation. Harvard Univ. Press, Cambridge, MA.
Swinkels, J., 1993. Adjustment dynamics and rational play in games. Games Econ. Behav. 5, 455–484.
Taylor, P., Jonker, L., 1978. Evolutionary stable strategies and game dynamics. Math. Bioscience 40, 145–156.
Weibull, J., 1996. The mass action interpretation. J. Econ. Theory 69, 165–172.
Zeidler, E., 1986. Nonlinear Functional Analysis and its Applications, vol. I. Springer-Verlag, New York.