WORKING PAPERS IN ECONOMICS

No. 08/04

BURKHARD HEHENKAMP AND ODDVAR M. KAARBØE

EQUILIBRIUM SELECTION IN SUPERMODULAR GAMES WITH MEAN PAYOFF TECHNOLOGIES

Department of Economics

UNIVERSITY OF BERGEN

Equilibrium selection in supermodular games with mean payoff technologies*

Burkhard Hehenkamp† and Oddvar M. Kaarbøe‡

November 10, 2004

Abstract

We examine an evolutionary model of equilibrium selection, where all individuals interact with each other, recurrently playing a strictly supermodular game. Individuals play (myopic) best responses to the current population profile, occasionally they pick an arbitrary strategy at random. To address the robustness of equilibrium selection in this simultaneous play scenario, we investigate whether different best-response approximations can lead to different long run equilibria.

Keywords: equilibrium selection, supermodular games, simultaneous play, best-response approximation

JEL-Class.-No.: C72, C73

1. Introduction

The theory of supermodular games provides a framework for the analysis of systems marked by strategic complementarities. First introduced by Topkis (1979) and further explored by Milgrom and Roberts (1990) and Vives (1990), it includes models of oligopolistic competition, macroeconomic coordination failure, Bertrand price competition, bank runs, R&D competition, or Becker's (1990) model of individual consumers' demand for restaurant seats or theater tickets. Supermodular games are characterized by the following properties: (i) each player's action set is partially ordered; (ii) marginal returns to increasing one's action rise with increases in the other players' action; and (iii), in the case of multidimensional actions, marginal returns to any single component of the player's action rise with increases in his other components. As a result, supermodular

*We thank Birgitte Sloth and Xander Tieman for comments.

†Department of Economics, University of Dortmund, D-44221 Dortmund, Germany. *Corresponding Author.* Email: b.hehenkamp@wiso.uni-dortmund.de. Phone: ++49-231-755-3215. Fax: ++49-231-755-3027.

‡Programme for Health Economics (HEB), Department of Economics, University of Bergen. Fosswinckelsg. 6, N-5007 Bergen. Norway. Email: oddvar.kaarboe@econ.uib.no
URL http://www.econ.uib.no/stab/oddvar.html

games exhibit monotonically increasing best-response functions, whenever actions can be completely ordered.

In the presence of increasing best-response functions, the possibility of multiple equilibria arises, creating an equilibrium selection problem: without any plausible equilibrium selection at hand, these types of games lack predictive power in that it remains unclear how players co-ordinate on equilibrium play. Among other things, Kandori and Rob (1995) address this objection (KR henceforth). Applying tools from stochastic evolutionary game theory (Kandori et al, 1993; Young, 1993), they show how the geometry of best-response functions helps in singling out a generically unique long run equilibrium.

Even though the context of random matching is the one most widely studied in evolutionary game theory, many models of economic interaction are better characterized by *simultaneous play*.[1] This interaction scenario refers to situations where the entire population plays the game under consideration simultaneously. Each player's objective function depends on his own action and on some summary statistics of other (or all) players' behavior. This statistics is typically taken to be some average (arithmetic, geometric, harmonic etc.) of the current action profile. For instance, in models of monopolistic competition individual profit depends on own price and the geometric mean of other players' prices (cf. Blanchard and Kiyotaki, 1987); in production externality models an individual's productivity (or production) depends on the average production level in the economy (cf. Cooper and Haltiwanger, 1996); in coordination games with simultaneous play a player's best response is to exactly match the average action of other players.

While KR mainly focus on the random matching scenario, they also provide an example of linear payoff functions that can be interpreted in both a random matching and in a simultaneous play scenario. Since the analyses for both interpretations coincide, the authors conclude "that random matching is not an essential part of [their] model." Investigating class coordination games with simultaneous play, Robles (1997) develops results that fully characterize the set of long run equilibria. He confirms KR's conclusion in that he extends their approach to the simultaneous play scenario. Robles allows for general linear summary statistics, which are taken to summarize the current play of the population. Since this summary statistics can in general assume more values than pure actions exist, he has to transform the summary statistics into the set of pure actions. To this end, he picks a specific rounding function. For the salient case, where the summary statistics is taken to be the mean, it turns out that the set of long run equilibria is bounded away from the extreme strategies. It always includes the Pareto efficient Nash equilibrium.

Recently, Robles' (1997) finding has been challenged by Hansen and Kaarbøe (2002, HK henceforth). Considering class coordination games with mean payoff technologies, they show that equilibrium selection is sensitive to how the summary statistics is transformed into pure actions. More specifically, choosing the appropriate rounding function, any limit state of the mutation-free process can be the unique long run equilibrium of

[1] Schelling (1973) introduced 'simultaneous play' into economics. In biology, Maynard Smith (1982) coined the 'playing the field' model. Crawford (1991, 1997) also argues for introducing genuine simultaneous interaction into the evolutionary literature.

the stochastic evolutionary process. Put differently, equilibrium selection depends on the shape of best-response functions off-side the grid of pure actions.

The main question we address here is whether this sensitivity of equilibrium selection to the rounding function carries over to supermodular games with simultaneous play. To this end, we provide a sufficient condition under which the answer is negative. Notice that we can rephrase the problem whether rounding matters as a problem of approximation. We start from a symmetric strictly supermodular 'original' game, where both the individual action and the mean of others' actions are defined on a continuous action space. To facilitate comparison with the papers by HK, KR, and Robles, we represent the symmetric original game by its best-response function. Subsequently, we discretize the individual action set so that, off-side the grid of pure actions, best-responses have to be 'rounded' to pure actions. To capture any reasonable way of rounding, we introduce a class of feasible approximating best-response functions. The problem of approximation is then, whether, for sufficiently fine discretizations of the individual action set, equilibrium selection does depend on feasible best-response approximations.

2. The model

We consider a population of players that, simultaneously and recurrently, play the same symmetric strictly supermodular stage game. Payoffs depend on own action choice and on the mean of other players' actions. From time to time, players adjust their behavior adopting (myopic) best replies against the summary statistic. Rarely, players pick an arbitrary action at random.

2.1. The original game

Let $\mathcal{N} := \{1, \ldots, N\}$ denote the set of N players, which we index by $n \in \mathcal{N}$. Similarly, let $\mathcal{M} := [0, M] \subseteq \mathbb{R}_+$ represent the set of actions, $m \in \mathcal{M}$. The payoff of any player $n \in \mathcal{N}$ depends on his own action $m_n \in \mathcal{M}$ and the mean of *other* players' actions

$$\mu(m_{-n}) = \frac{1}{N-1} \sum_{n' \in \mathcal{N} \setminus \{n\}} m_{n'}.$$

Payoffs are symmetric in players' identity, i.e.

$$\pi_n = \pi(m_n, \mu(m_{-n})),$$

for some $\pi : \mathcal{M}^2 \to \mathbb{R}_+$ and all players $n \in \mathcal{N}$. Moreover, we assume payoffs to be continuous and *strictly supermodular*, the latter of which says that, for all $m' < m''$, the difference $\pi(m'', \mu) - \pi(m', \mu)$ is strictly increasing in μ. Denote this stage game by $\Gamma := (\mathcal{N}, \mathcal{M}^2, \pi)$.

Alternatively, we could take the mean with respect to *all* players' actions. For some *given* supermodular game, the specification of μ might affect the set of equilibria and, similarly, the set of equilibria selected by the evolutionary process. However, notice that our findings only make use of the assumption of supermodularity but not of the particular specification of μ. Therefore, they remain valid for *both* specifications of μ.

One might think of Γ as a game that arises from some "larger" game $\widehat{\Gamma} = (\mathcal{N}, \widehat{\mathcal{M}^2}, \widehat{\pi})$ after all iteratively dominated actions have been removed. In this case, the following assumptions imposed upon the larger game $\widehat{\Gamma}$ enable one to shrink down the set of actions: (i) $\widehat{\mathcal{M}}$ is a path-connected complete lattice, e.g. a closed interval $[\underline{M}, \overline{M}]$; (ii) $\widehat{\pi}(m_n, \mu)$ is upper semi-continuous in m_n, for any given μ, and continuous in μ, for any given m_n; (iii) $\widehat{\pi}(m_n, \mu)$ is bounded; and (iv) $\widehat{\pi}(m_n, \mu)$ is supermodular in m_n and has increasing differences in m_n and μ. Notice that the game specified above, $\Gamma = (\mathcal{N}, \mathcal{M}^2, \pi)$, satisfies these assumptions. Given assumptions (i)-(iv), we could then apply Theorem 5 in Milgrom and Roberts (1990, p. 1265). It states that there exist a largest and a smallest serially undominated (pure) action, each of which corresponds to a symmetric Nash equilibrium (NE) in pure actions. Accordingly, we think of $m = 0$ as the smallest NE action and $m = M$ as the largest NE action after all strictly dominated actions have iteratively been removed from the 'larger' game $\widehat{\Gamma}$.

Denote the set of *pure NE actions* by $\mathcal{M}^{NE}$, i.e.

$$\mathcal{M}^{NE} = \{m^* \in \mathcal{M} \,|\, \pi(m^*, m^*) \geq \pi(m, m^*) \,\forall m \in \mathcal{M}\}.$$

We assume the set of NE actions to be finite, $K := \#\mathcal{M}^{NE} < \infty$. Without loss of generality, we index the NE actions such that $\mathcal{M}^{NE} = \{m_1^*, \ldots, m_K^*\}$ with $m_1^* < \ldots < m_K^*$, $m_1^* = 0$ and $m_K^* = M$.

The following proposition collects two well-known properties of supermodular games.

Proposition 1. *Let $\mathcal{BR}(\cdot)$ represent the best-response correspondence of the original game. Then we have (i) $(m_1, \ldots, m_N) \in \mathcal{M}^{NE} \Longrightarrow m_n = m_1$, for all $n \in \mathcal{N}$; and (ii) $m \in \mathcal{BR}(\mu), m' \in \mathcal{BR}(\mu')$, and $\mu < \mu'$ imply $m \leq m'$.*

Proof. Claim (ii) is contained in Milgrom and Roberts (1990). To establish (i), assume that $(m_1, \ldots, m_N)$ represents an asymmetric Nash equilibrium, i.e. $m_n < m_{n'}$ for some $n \neq n'$,

$$\pi(m_n, \mu(m_{-n})) \geq \pi(m_{n'}, \mu(m_{-n})), \tag{2.1}$$

and

$$\pi(m_{n'}, \mu(m_{-n'})) \geq \pi(m_n, \mu(m_{-n'})). \tag{2.2}$$

By supermodularity, $\pi(m_{n'}, \mu) - \pi(m_n, \mu)$ is strictly increasing in μ. Since $m_n < m_{n'}$ implies $\mu(m_{-n'}) < \mu(m_{-n})$, it follows that

$$\pi(m_{n'}, \mu(m_{-n'})) - \pi(m_n, \mu(m_{-n'})) < \pi(m_{n'}, \mu(m_{-n})) - \pi(m_n, \mu(m_{-n})).$$

However, because of (2.1) and (2.2) the LHS is positive and the RHS is negative, respectively, which yields a contradiction. Thus, claim (i) holds true. ■

Part (i) says that no asymmetric NE in pure actions exists. According to part (ii), strict supermodularity implies weak monotonicity of the best-response correspondence.

The discretized game

Applying the standard framework of stochastic evolutionary game theory, which is based on finite state space Markov chains, we discretize the original game. To this end, we consider a discretized action space with equidistant actions, $\mathcal{M}_\delta = \{0, \delta, 2\delta, \ldots, M\}$, where $M = L\delta$ for some $L \in \mathbb{N}$ and $\delta > 0$. Since we intend to compare the robustness of NE actions in the original game with the robustness of NE actions in the approximated game, it is warranted to assume that all NE actions of the original game belong to the discretized action space, i.e. $\mathcal{M}^{NE} \subseteq \mathcal{M}_\delta$.

Notice that the mean can assume more values than pure actions exist in $\mathcal{M}_\delta$. Taking this into account, we denote the discretized game by $\Gamma_\delta = (\mathcal{N}, \mathcal{M}_\delta \times \mathcal{M}, \pi)$. To indicate that payoffs coincide in the original and in the discretized game, we use the same label 'π' to denote the pay-off function although, mathematically, these are different functions. Observe also that the discretized game Γ_δ inherits the property of strict supermodularity from the original game Γ.

Let $\mathcal{M}_\delta^{NE}$ denote the set of pure NE actions of Γ_δ. Obviously, we have $\mathcal{M}^{NE} \subseteq \mathcal{M}_\delta^{NE}$. However, in order to compare the robustness of NE actions in the original and in the approximated game, respectively, it is also warranted to assume that no additional equilibrium arises in the discretized game, i.e. $\mathcal{M}^{NE} \supseteq \mathcal{M}_\delta^{NE}$. Combining both inclusions, we impose the following assumption on the discretized game:

Assumption A.

$$\mathcal{M}^{NE} = \mathcal{M}_\delta^{NE}. \tag{2.3}$$

The set of grid sizes, $\delta > 0$, satisfying (2.3) is denoted by $\mathcal{D} := \{\delta > 0 : M_\delta^{NE} = M^{NE}\}$. Observe that $\mathcal{D} \neq \emptyset$, since we assumed the number of Nash equilibria to be finite, $K < \infty$.

Assuming (2.3) is not as innocent as it might appear at first sight, as, for some games, *every* discretization of the action set results in existence of an additional NE action. For instance, discretizing the action set of the standard Bertrand oligopoly game results in the additional NE action where players charge the lowest price strictly above marginal cost.[2] Therefore, it is important to notice that, for the class of supermodular games with mean payoff technologies, which is under consideration here, this is not the case. Accordingly, the class of games covered by our results is not further restricted by Assumption A.

Imposing Assumption A, the first result of Proposition 1 carries over to the discretized game Γ_δ. As for the second, we first have to define best respones in terms of the discretized game.

Best-response approximation

To facilitate comparison with the papers by Hansen and Kaarbøe (2002), Kandori and Rob (1995), and Robles (1997), we take best-response functions as our starting-point. We impose the following assumption:

[2]See e.g. Alós-Ferrer (1999).

Assumption B. The best-response correspondence of the original game is single-valued and continuous as a function. The best-response function is denoted $BR(\cdot)$, where $BR : \mathcal{M} \to \mathcal{M}$ and $\mu \mapsto BR(\mu)$.

Observe that strict supermodularity immediately implies the best-response function to be strictly increasing. It then follows that all Nash equilibria in pure actions are strict.

Assumption B essentially incorporates the "continuity" assumption as originally proposed by Kandori and Rob (1995). To see this, let us first restate their assumption: whenever two pure actions, m and m', represent best responses to mixed actions α and α', respectively, i.e. $m \in \mathcal{BR}(\alpha)$ and $m' \in \mathcal{BR}(\alpha')$, then any intermediate action, $m < m'' < m'$, should represent a *unique* best response to some convex combination of α and α', i.e. there should exist a $\lambda \in (0,1)$ such that $\mathcal{BR}(\lambda\alpha + (1-\lambda)\alpha') = \{m''\}$. Imposing this assumption to hold true for the discretized game Γ_δ and *any* discretization $\delta \in \mathcal{D}$, would imply that the best-response correspondence is single-valued, which in fact means it is a function.

Turning towards the best-response approximation of the discretized games Γ_δ, $\delta \in \mathcal{D}$, one has to restrict the class of feasible best-response functions, $b : \mathcal{M} \to \mathcal{M}_\delta$. First of all, we have already mentioned earlier that the mean can take more values than pure actions exist in $\mathcal{M}_\delta$. Therefore, we have to specify which values the best-response function can assume for $\mu \in \mathcal{M} \setminus \mathcal{M}_\delta$. Second, the best-response function should reflect that the discretized game Γ_δ is supermodular. Third, representing the discretized game, $b(\cdot)$ should approximate the best-response function $BR(\cdot)$ of the original game. Fourth and finally, the sets of Nash equilibria corresponding to $b(\cdot)$ and $BR(\cdot)$ should coincide. Only then, we have the same candidates for equilibrium selection.

Accordingly, if $b : \mathcal{M} \to \mathcal{M}_\delta$ denotes a best-response approximation of the original best-response function $BR(\cdot)$, we assume the following:

(i) $b(\cdot)$ is (weakly) increasing,

(ii) $\forall m \in \mathcal{M}_\delta : BR^{-1}(m) \subset b^{-1}(m)$, and

(iii) $\forall m \in \mathcal{M}_\delta : BR(m) = m \Longleftrightarrow b(m) = m$.

Condition (i) preserves the supermodular structure, condition (ii) incorporates the idea of approximation, and condition (iii) imbeds Assumption A. Notice that $b(\cdot)$ can no longer be strictly increasing, for the mean can assume more values than actions exist in $\mathcal{M}_\delta$. Since $\mathcal{M}_\delta$ is finite, it follows that $b(\cdot)$ is piece-wise constant.

Condition (ii) requires that, for any action of the discretized game, $m \in \mathcal{M}_\delta$, the following should hold true: whenever m is optimal against mean μ in the original game, i.e. $\mu = BR^{-1}(m)$, then m remains optimal against μ under the best-response approximation $b(\cdot)$. Notice that $BR^{-1}(\cdot)$ denotes the inverse function to $BR(\cdot)$, while $b^{-1}(m)$ represents the inverse image of m under $b : \mathcal{M} \to \mathcal{M}_\delta$, i.e. $b^{-1}(m) = \{\mu \in \mathcal{M} : b(m) = \mu\}$.

Finally, condition (iii) ensures that, for any $\delta > 0$, any best-response representation of the discretized game, $b \in \mathcal{B}_\delta$, satisfies Assumption A, i.e. $\mathcal{M}_\delta^{NE}(b) = \mathcal{M}^{NE}$, where $M_\delta^{NE}(b)$ denotes the set of pure NE actions under best-response approximation $b(\cdot)$.

For any $\delta \in \mathcal{D}$, we denote the set of all feasible approximated best-response functions by

$$\mathcal{B}_\delta := \{b : \mathcal{M} \to \mathcal{M}_\delta \ : \ b \text{ satisfies properties (i)-(iii)}\}.$$

2.2. The evolutionary process

We now turn towards the dynamic specification of the evolutionary process. To this end, we take a bird's-eye view and look at the population as a whole rather than at the behavior of single players.

Given any action profile at time t, $(m_1^t, m_2^t, \ldots, m_N^t)$, a state of the evolutionary process is a frequency distribution of all actions used at that time, i.e. $s^t = (s_0^t, s_\delta^t, \ldots, s_M^t)$ is such that $s_m^t = \#\{n \in \mathcal{N} : m_n^t = m\}$ represents the number of players employing action m, for all $m \in \mathcal{M}_\delta$. Obviously, we must have $\sum_{m\in\mathcal{M}_\delta} s_m^t = N$ so that the state space is given by

$$\mathcal{S}_\delta = \left\{ s \in \{0, 1, \ldots, N\}^{L+1} : \sum_{m\in\mathcal{M}_\delta} s_m = N \right\}.$$

Let s^m denote the state where all players employ action m.

To determine the payoff earned by any action $m \in \mathcal{M}_\delta$, given the current state is $s^t \in \mathcal{S}_\delta$, we have to remove this player from the population profile s^t. Let s_{-m}^t be the according population profile of *other* players' actions, i.e. $s_{-m}^t = (s_0^t, \ldots, s_{m-\delta}^t, s_m^t - 1, s_{m+\delta}^t, \ldots, s_M^t)$, where action $m \in \mathcal{M}_\delta$ has to be actually played by at least one player, i.e. $s_m^t \geq 1$. Given s_{-m}^t, we write with slight abuse of notation

$$\mu(s_{-m}^t) = \frac{1}{N-1}\left[\sum_{m'\neq m} m' \cdot s_{m'}^t + m \cdot (s_m^t - 1)\right],$$

for m such that $s_m^t \geq 1$. Obviously, we have $\mu(s_{-m}^t) = \mu(m_{-n}^t)$ if (and only if) player n uses action m. Then the payoff earned from action $m \in \mathcal{M}_\delta$ can be derived as

$$\pi_m^t = \pi(m, \mu(s_{-m}^t)).$$

Adaptation dynamics

For $\delta > 0$, fix $b \in \mathcal{B}_\delta$ arbitrarily. With some probability $\zeta \in (0,1)$, each player revises his action, adopting a (myopic) best response against other players' mean action,

$$m_n^{t+1} \in b(\mu(s_{-m}^t)),$$

for all players $n \in \mathcal{N}$. With complementary probability $1-\zeta$, the player sticks with his previous action, i.e. $m^{t+1} = m^t$. For simplicity, we assume the revision probability, $\zeta \in (0,1)$, to be independent across players. (Allowing $\zeta \in (0,1)$ to vary across players would not change our results.)

Mutation dynamics

Once players have adapted their behavior, their choice is subject to random shocks, which might be considered as mistakes or mutations. With probability $\xi > 0$, each player picks an arbitrary action from the action set $\mathcal{M}_\delta$. Similar to the above, the mutation probability $\xi > 0$ is assumed to be independent across players.

The overall process

The overall process (involving adaptation plus mutation) constitutes a discrete-time Markov process with finite state space $\mathcal{S}_\delta$. The dynamics of adaptation and mutation fully characterize the transition matrix of the process, which we denote $P(\xi, \zeta, \delta) = (p_{ss'}(\xi, \zeta, \delta))_{s,s' \in \mathcal{S}_\delta}$. Here, the transition probability $p_{ss'}(\xi, \zeta, \delta)$ represents the probability of the dynamic system to make a transition from $s \in \mathcal{S}_\delta$ to $s' \in \mathcal{S}_\delta$ within one period — given the mutation probability ξ, the adaptation probability ζ, and the grid size δ. The mutation-free or *pure adaptation process* corresponds to $P(0, \zeta, \delta)$.

The presence of mutations implies that every transition from one state to another has positive probability. Therefore, the Markov process is aperiodic, (positive) recurrent and irreducible for any $\xi > 0$. A standard result in the theory of Markov chains then guarantees existence and uniqueness of an invariant distribution $\phi(\xi, \zeta, \delta)$. This distribution describes the long-run frequencies with which every state is observed (with probability one) along any sample path. Accordingly, if $\Delta(\mathcal{S}_\delta)$ denotes the set of probability measures defined on $\mathcal{S}_\delta$, then we have $\phi(\xi, \zeta, \delta) \in \Delta(\mathcal{S}_\delta)$. As a characteristic feature of these types of processes, the stationary distribution does not depend on the starting-point. Therefore, we can assume the process to start in an arbitrary state $s^0 \in \mathcal{S}_\delta$.

To capture the idea of rare mutations, we look at the behavior of the process as the rate of mutation becomes small or, formally, at the *limit distribution*

$$\phi^*(\zeta, \delta) := \lim_{\xi \to 0} \phi(\xi, \zeta, \delta).$$

Based upon arguments in Freidlin and Wentzell (1984), Young (1993) has shown this limit to exist. Following Young, we call any state $s \in \mathcal{S}_\delta$ that has strictly positive probability under $\phi^*(\zeta, \delta)$, i.e. with $\phi^*_s(\zeta, \delta) > 0$, stochastically stable.

3. Approximation

In this section we show that equilibrium selection does not depend on best-response approximation provided that discretization of the action set is sufficiently fine (Theorem 3.1). To this end, we first establish that approximation of the original best-response function by functions from $\mathcal{B}_\delta$ is uniform (Lemma 1). Subsequently, we build on this 'vertical' approximation result to establish a 'depth' approximation result (Lemma 2), relating the 'depth' of the original best-response function to that of the best-response approximations from $\mathcal{B}_\delta$. As already illustrated by Kandori and Rob (1995), it is the 'depth' rather than the 'size' of basins of attraction that determines equilibrium selection for strictly supermodular games. Since Lemma 2 only applies to interior NE actions, we continue with analyzing the case of boundary NE actions (Lemma 3). Our main result, Theorem 3.1, concludes this section.

We start with determining the limit sets of the pure adaptation process. Only these limit sets represent candidates for stochastic stability.

Proposition 2. *Fix the grid size $\delta \in \mathcal{D}$, fix the revision probability $\zeta \in (0,1)$, and let $b \in \mathcal{B}_\delta$ be arbitrary. Then each limit set is a singleton. The set of the corresponding limit states, $\mathcal{S}^0$, is in one-to-one correspondence with the set of NE actions, i.e.*

$$\mathcal{S}^0 = \{s^m \in \mathcal{S}_\delta : m \in \mathcal{M}^{NE}\}.$$

Proof. The claim can be established along the lines of Theorem 2 in Kandori and Rob, 1995, p. 400. ■

Recall that, for any fixed $\delta \in \mathcal{D}$, we have $\mathcal{M}^{NE} = \mathcal{M}^{NE}_\delta(b)$ for any $b \in \mathcal{B}_\delta$ by definition of $\mathcal{B}_\delta$. Then Proposition 2 says the following. First, all limit sets consist of single states. This implies that the process of pure adaptation does not display cycling behavior or drift. Second, only monomorphic states, where all players use the same action, represent candidates for stochastic stability. Third, the corresponding action must represent a NE action. Fourth, the set of limit states, $\mathcal{S}^0$, does not depend on revision probability $\zeta \in (0,1)$. And finally, the set of limit states neither depends on the best-response approximation $b \in \mathcal{B}_\delta$.

While the set of limit states does not depend on the best-response approximation $b \in \mathcal{B}_\delta$, the limit distribution in general will do so. To capture this dependence, let us denote the set of NE actions corresponding to stochastically stable states by $\mathcal{M}^*(b)$, i.e.

$$\mathcal{M}^*(b) := \{m \in \mathcal{M}^{NE} : \phi^*_{s^m}(\delta) > 0\}.$$

Determining the set of stochastically stable states, we rely on mutation-cost analysis such as first introduced by Kandori et al (1993) and Young (1993). For each $m \in \mathcal{M}^{NE}$, define a m-tree T as a directed spanning tree on $\mathcal{M}^{NE}$ such that, for every $m' \in \mathcal{M}^{NE} \backslash \{m\}$, the collection T contains a path $P_T(m') = \{(m', m_1), (m_1, m_2), \ldots, (m_k, m)\}$ leading from m' to m, i.e. m is the root of T. We denote the set of all m-trees by $\mathcal{T}_m$. Furthermore, for any best-response approximation $b \in \mathcal{B}_\delta$, let $C^b(m', m'')$ be the minimum number of mutations required by the overall process to make a transition from m' to m'', where $m', m'' \in \mathcal{M}^{NE}$. Finally, we introduce the stochastic potential $C^b(m)$ of any NE action $m \in \mathcal{M}^{NE}$ by setting $C^b(m) := \min_{T \in \mathcal{T}_m} \sum_{(m', m'') \in T} C^b(m', m'')$. Then, Lemma 1 in Young (1993) allows us to characterize stochastic stability: $m \in \mathcal{M}^*(b)$ if and only if $C^b(m) \leq C^b(m')$ for all $m' \in \mathcal{M}^{NE}$.

The following Lemma states that approximation of the original best-response function $BR(\cdot)$ by best-response functions from $\mathcal{B}_\delta$ is uniform.

Lemma 1. *(Uniform Approximation of the Original Best Response)* $\forall \varepsilon > 0 : \exists \overline{\delta} > 0 : \forall \delta \in (0, \overline{\delta}) : \forall b \in \mathcal{B}_\delta : \quad \sup_{\mu \in \mathcal{M}} |BR(\mu) - b(\mu)| < \varepsilon$

Proof. We show that setting $\overline{\delta}(\varepsilon) := \varepsilon$ does the job. To this end, fix $\varepsilon > 0$. Let $BR^{-1}(\mathcal{M}_\delta)$ denote the set of mean actions $\mu \in \mathcal{M}$ such that some action $m \in \mathcal{M}_\delta$

represents a best response to μ, i.e. $BR^{-1}(\mathcal{M}_\delta) = \{\mu \in \mathcal{M} : BR(\mu) = m$ for some $m \in \mathcal{M}_\delta\}$.

First, if $\widehat{\mu} \in BR^{-1}(\mathcal{M}_\delta)$ then $BR(\widehat{\mu}) = b(\widehat{\mu}) = m$ for some $m \in \mathcal{M}_\delta$ and all $b \in \mathcal{B}_\delta$ by property (ii) in the definition of $\mathcal{B}_\delta$. Hence, $BR(\widehat{\mu}) - b(\widehat{\mu}) = 0$.

Second, for $\widehat{\mu} \in \mathcal{M} \setminus BR^{-1}(\mathcal{M}_\delta)$, set $\underline{\mu} := \max\{\mu \in BR^{-1}(\mathcal{M}_\delta) : \mu < \widehat{\mu}\}$ and $\overline{\mu} := \min\{\mu \in BR^{-1}(\mathcal{M}_\delta) : \mu > \widehat{\mu}\}$. These numbers are well defined because of $\{0, M\} \subset BR^{-1}(\mathcal{M}_\delta)$. By selection of $\widehat{\mu}$ and construction of $\underline{\mu}$ and $\overline{\mu}$, we have $\underline{\mu} < \widehat{\mu} < \overline{\mu}$ and $BR(\overline{\mu}) = BR(\underline{\mu}) + \delta$. Moreover, it follows that $BR(\underline{\mu}) = b(\underline{\mu})$ and $BR(\overline{\mu}) = b(\overline{\mu})$, for any $b \in \mathcal{B}_\delta$, which in turn implies $BR(\overline{\mu}) - BR(\underline{\mu}) = b(\overline{\mu}) - b(\underline{\mu}) = \delta$. Since $BR(\cdot)$ is strictly increasing and $b \in \mathcal{B}_\delta$ is weakly increasing, we obtain $BR(\underline{\mu}) < BR(\widehat{\mu}) < BR(\overline{\mu}) = BR(\underline{\mu}) + \delta$ and $b(\underline{\mu}) \leq b(\widehat{\mu}) \leq b(\overline{\mu}) = b(\underline{\mu}) + \delta$, respectively. Combining the two chains of inequalities, we obtain $-\delta < BR(\widehat{\mu}) - b(\widehat{\mu}) < \delta$.

Thus, setting $\overline{\delta}(\varepsilon) := \varepsilon$, we have $\sup_{\mu \in \mathcal{M}} |BR(\mu) - b(\mu)| < \delta < \varepsilon$, for all $\delta \in (0, \overline{\delta})$ and all $b \in \mathcal{B}_\delta$. ■

In fact, we have established the following corollary:

Corollary 1. *Fix $\delta \in \mathcal{D}$. For all $b \in \mathcal{B}_\delta$, we have:* $\sup_{\mu \in \mathcal{M}} |BR(\mu) - b(\mu)| < \delta$.

We continue with introducing further notation preparing the mutation-cost analysis. Choose $\delta \in \mathcal{D}$ arbitrarily. By Proposition 2, only NE actions represent candidates for stochastic stability. Let $m', m'' \in \mathcal{M}^{NE}$ be two adjacent NE actions such that $m' < m''$ and define *upward* and *downward depth* of the original best-response function *between* m' and m'' as

$$\begin{aligned} D^{BR}(m', m'') \ &: \ = \min\{a \in [0,1] : BR(aM + (1-a)\mu) \geq \mu \text{ for all } \mu \in [m', m'']\} \qquad \text{and} \\ D^{BR}(m'', m') \ &: \ = \min\{a \in [0,1] : BR((1-a)\mu) \leq \mu \text{ for all } \mu \in [m', m'']\}, \qquad (3.1) \end{aligned}$$

respectively. Notice that both minima exist, because $\widetilde{BR}_z : [0,1] \to \mathcal{M}$, $a \longmapsto \widetilde{BR}_z(a) := BR(az + (1-a)\mu)$, is a continuous function on a compact intervall for each $z = 0, M$ and all $\mu \in [m', m'']$ and because $\widetilde{BR}_M(1) = BR(M) = M \geq \mu$ and $\widetilde{BR}_0(1) = BR(0) = 0 \leq \mu$ holds true for all $\mu \in [m', m''] \subset \mathcal{M}$. Moreover, observe that $\widetilde{BR}_M(\cdot)$ is strictly increasing, while $\widetilde{BR}_0(\cdot)$ is strictly decreasing.

The upward (downward) depth between adjacent actions $m' < m''$ characterizes the minimum share of players that have to mutate to the highest (lowest) action M (0) in order to move the dynamics into the basin of attraction of NE action m'' (m'). Accordingly, the numbers

$$C^{BR}(m', m'') := \max\{\lceil N \cdot D^{BR}(m', m'') \rceil, 1\} \qquad (3.2)$$

characterize the minimum number of mutating players required by the overall process to make a transition from m' to m'', for any $m', m'' \in \mathcal{M}_\delta^{NE}$ (where, for any $x \in \mathbb{R}_+$, the function $\lceil x \rceil$ denotes the smallest integer number such that $x \leq \lceil x \rceil$). If $D^{BR}(m', m'') = 0$ then one mutation is required to make a transition from m' to m'', since $\{s^{m'}\}$ is a limit state, i.e. $C^{BR}(m', m'') = 1$ for all $N \in \mathbb{N}$.

Figure 3.1 illustrates the upward and downward depth of the original best-response function between adjacent NE actions $m' < m'' < m'''$, respectively. We have $D^{BR}(m', m'') = a_1$, $D^{BR}(m'', m''') = 0$, $D^{BR}(m''', m'') = a_2$, and $D^{BR}(m'', m') = 0$.

Include Figure 3.1 about here.

Finally, let $D^b(m', m'')$ and $C^b(m', m'')$ represent the corresponding notions where some best-response approximation $b \in \mathcal{B}_\delta$ replaces the original best-response function $BR(\cdot)$. Since $b(\cdot)$ is a step-function and hence discontinuous at some points, it is no longer clear that the minimum really exists. However, the set

$$\{a \in [0,1] : b(aM + (1-a)\mu) \geq \mu \text{ for all } \mu \in [m', m'']\}$$

is non-empty, since it always contains $a = 1$. Therefore, it suffices to replace the minimum operator by the infimum operator. We set

$$\begin{aligned} D^b(m', m'') \quad &: \quad = \inf\{a \in [0,1] : b(aM + (1-a)\mu) \geq \mu \ \forall \mu \in [m', m'']\} \qquad \text{and} \\ C^b(m', m'') \quad &: \quad = \max\{\lceil N \cdot D^b(m', m'') \rceil, 1\}. \end{aligned}$$

Similar notation applies to the case of downward depth.

The following two lemmas characterize the approximation behavior of best-response step-functions from $\mathcal{B}_\delta$ in terms of both upward and downward depth. At first, Lemma 2 covers the case of interior NE actions, $m', m'' \in \mathcal{M}^{NE} \setminus \{0, M\}$. It states that, for sufficiently close approximations $b \in \mathcal{B}_\delta$, both upward and downward depth of the respective approximating function become arbitrarily close to upward and downward depth of the original best-response function $BR(\cdot)$, respectively. Subsequently, Lemma 3 deals with the case of boundary NE actions, $m' = 0$ or $m'' = M$.

Lemma 2. *(Uniform Approximation of Depth) Consider adjacent NE actions $m', m'' \in \mathcal{M}^{NE} \setminus \{0, M\}$ such that $m' \neq m''$. Then, for any $\varepsilon > 0$, there exists some $\overline{\delta} > 0$ such that, for all $\delta \in \mathcal{D} \cap (0, \overline{\delta})$, we have*

$$\left| D^{BR}(m', m'') - D^b(m', m'') \right| < \varepsilon \qquad \textit{for all } b \in \mathcal{B}_\delta.$$

Proof. We establish the claim for the case of upward depth. The opposite case can be dealt with similarly. Accordingly, fix $m', m'' \in \mathcal{M}^{NE} \setminus \{0, M\}$ such that $m' < m''$ and $m \in \mathcal{M}^{NE}$ implies either $m \leq m'$ or $m \geq m''$. Moreover, choose $\delta' > 0$ sufficiently small such that $\max\{m', m''\} < M - \delta'$ and $\min\{m', m''\} > \delta'$.

First, by Lemma 1, we have that

$$BR(\mu) - \delta < b(\mu) < BR(\mu) + \delta \qquad \text{for all } \mu \in \mathcal{M} \tag{3.3}$$

and all $b \in \mathcal{B}_\delta$, such that $\delta \in (0, \delta') \cap \mathcal{D}$.

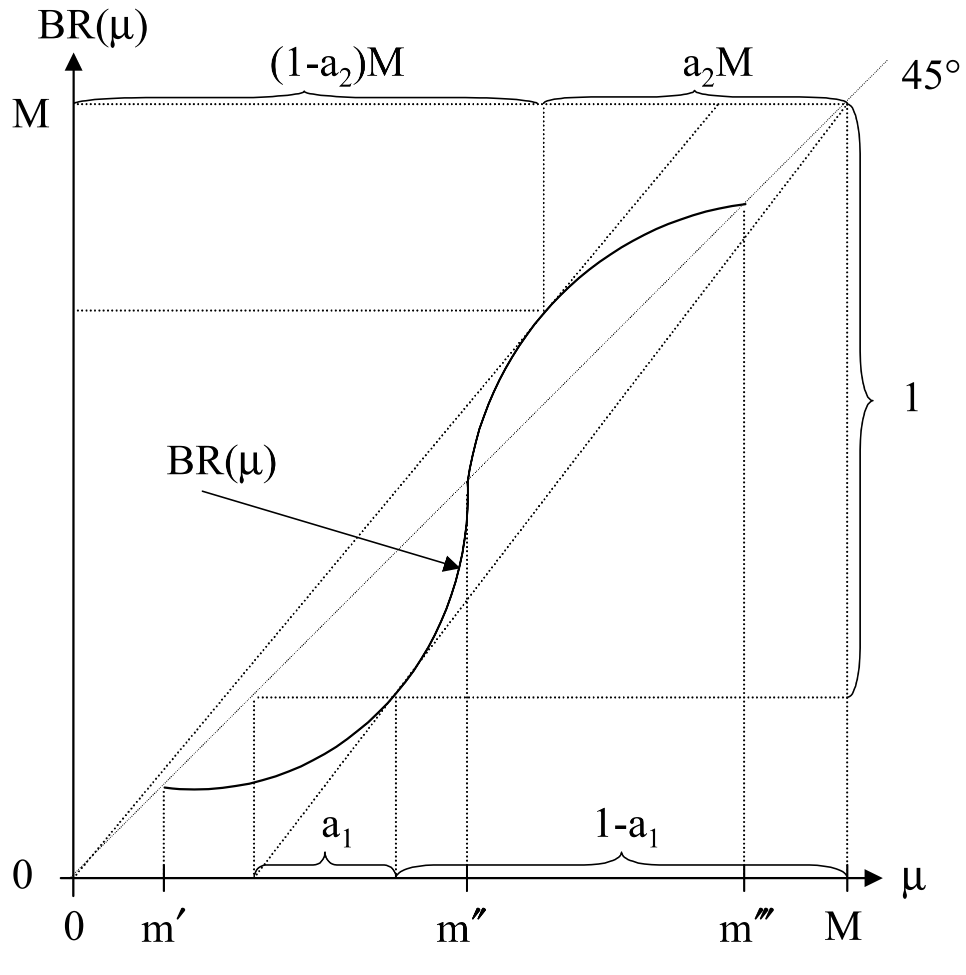

Figure 3.1: Upward and downward depth of the original best-response function

Second, we extend the notion of $D^{BR}(m', m'')$ to $D^{BR+\delta}(m', m'')$ and $D^{BR-\delta}(m', m'')$. I.e., we set

$$D^{BR+\eta}(m', m'') := \min\{1, \inf\{a \in [0,1] : BR(aM + (1-a)\mu) + \eta \geq \mu \; \forall \; \mu \in [m', m'']\}\},$$

for $\eta \in \{-\delta, \delta\}$. We set $\inf \emptyset := \infty$ in order to take into account that, for $\eta = \delta > 0$, the set

$$\{a \in [0,1] : BR(aM + (1-a)\mu) + \eta \geq \mu \; \forall \; \mu \in [m', m'']\}$$

might be empty. From (3.3), it follows that

$$D^{BR+\delta}(m', m'') \leq D^b(m', m'') \leq D^{BR-\delta}(m', m'') \qquad \forall b \in \mathcal{B}_\delta \; \forall \delta \in \mathcal{D} \cap (0, \delta'). \tag{3.4}$$

Third, we show that

$$D^{BR+\eta}(m', m'') = \max_{\mu \in [m', m'']} \frac{BR^{-1}(\mu - \eta) - \mu}{M - \mu}, \qquad \text{for all } \eta \in [-\delta, \delta] \tag{3.5}$$

and all $\delta \in \mathcal{D} \cap (0, \delta')$. To see this, we transform the definition of $D^{BR+\eta}(m', m'')$:

$$\begin{aligned} D^{BR+\eta}(m', m'') &= \min\{a \in [0,1] : BR(aM + (1-a)\mu) + \eta \geq \mu \text{ for all } \mu \in [m', m'']\} \\ &= \min\left\{a \in [0,1] : a \geq \frac{BR^{-1}(\mu - \eta) - \mu}{M - \mu} \text{ for all } \mu \in [m', m'']\right\} \\ &= \min\left\{a \in [0,1] : a \geq \max_{\mu \in [m', m'']} \frac{BR^{-1}(\mu - \eta) - \mu}{M - \mu}\right\} \\ &= \max_{\mu \in [m', m'']} \frac{BR^{-1}(\mu - \eta) - \mu}{M - \mu}. \end{aligned}$$

The second equality holds true because $\mu - \eta \in (0, M)$ holds true for all $\mu \in [m', m'']$ by selection of δ'.

Fourth, since the right hand side of (3.5) is continuous and strictly decreasing in η, so is the left hand side. Hence, the outer expressions of relation (3.4) converge as we take the limit $\delta \to 0$,

$$\begin{aligned} \lim_{\delta \to 0} D^{BR+\delta}(m', m'') &= D^{BR}(m', m'') \qquad \text{and} \\ \lim_{\delta \to 0} D^{BR-\delta}(m', m'') &= D^{BR}(m', m''). \end{aligned}$$

Finally, this in turn implies that also the nested expression in (3.4) converges, i.e., for any $\varepsilon > 0$, there exists some $\overline{\delta} > 0$ such that, for any $\delta \in \mathcal{D} \cap (0, \min\{\overline{\delta}, \delta'\})$ and any $b \in \mathcal{B}_\delta$, we have

$$\left| D^{BR}(m', m'') - D^b(m', m'') \right| < \varepsilon.$$

■

Figure 3.2 illustrates the intuition underlying the proof of Lemma 2. It depicts the original best-response function $BR(\cdot)$ as well as its shifted graph $BR(\cdot) - \delta$. By the corollary to Lemma 1, the shifted best-response $BR(\cdot) - \delta$ provides a lower bound on all feasible best-response approximations $b \in \mathcal{B}_\delta$, i.e. $BR(\cdot) - \delta < b(\cdot)$, for any given grid size $\delta \in \mathcal{D}$. Turning towards the upward depth between m and m', the *lower* bound $BR(\cdot) - \delta$ on $b(\cdot)$ transforms into an *upper* bound on upward depth, i.e. $D^b(m', m'') \leq D^{BR-\delta}(m', m'')$. Similarly, the *upper* bound $BR(\cdot) + \delta$ translates into a *lower* bound on upward depth, i.e. $D^{BR+\delta}(m', m'') \leq D^b(m', m'')$. It is then evident that, as $\delta \to 0$, we have $D^{BR+\eta}(m', m'') = D^{BR}(m', m'')$ for $\eta \in \{-\delta, \delta\}$ and hence approximation of upward depth is uniform.

Include Figure 3.2 about here.

According to Lemmas 1 and 2, the family of best-response functions, $b \in \mathcal{B}_\delta$, not only approximates the original best-response function $BR(\cdot)$ vertically, which is ensured by the assumptions underlying $\mathcal{B}_\delta$. For interior NE actions, approximation is also in terms of upward and downward depth of the best-response function.

This finding is by no means trivial, since (i) it was not clear that approximation by functions in $\mathcal{B}_\delta$ is uniform and (ii) that we could overcome the problem that best-reponse step-functions are not invertible. The latter was key, since we needed to draw inferences from approximation in the range $\mathcal{M}_\delta$ to approximation in the domain $\mathcal{M}$.

Turning towards the case of boundary Nash equilbrium actions, Lemma 3 shows that approximation in terms of depth is arbitrarily close, provided the equilibrium action under consideration is stable with respect to the original best-response function $BR(\cdot)$. If it is unstable with respect to $BR(\cdot)$, then approximation in terms of depth is no longer arbitrarily close, but the equilibrium action remains unstable under any best-response approximation $b \in \mathcal{B}_\delta$, for any $\delta \in \mathcal{D}$ sufficiently small.

Lemma 3. *Let $m^*_{K-1} > 0$ and $m^*_K = M$ be adjacent NE actions.*
(i) Suppose $m^*_K = M$ is stable, i.e. $BR(\mu) > \mu$ for all $\mu \in (m^*_{K-1}, m^*_K)$. Then, for any $\varepsilon > 0$, there exists some $\overline{\delta} > 0$ such that for all $\delta \in \mathcal{D} \cap (0, \overline{\delta})$ and all $b \in \mathcal{B}_\delta$ we have

$$\left| D^{BR}(m', m'') - D^b(m', m'') \right| < \varepsilon,$$

where $m', m'' \in \{m^*_{K-1}, m^*_K\}, m' \neq m''$.
(ii) Suppose M is unstable, i.e. $BR(\mu) < \mu$ for all $\mu \in (m^*_{K-1}, m^*_K)$. Then, for any $\varepsilon > 0$, there exists some $\overline{\delta} > 0$ such that for all $\delta \in \mathcal{D} \cap (0, \overline{\delta})$ and all $b \in \mathcal{B}_\delta$ we have

$$\begin{aligned} D^{BR}(m^*_{K-1}, M) - \varepsilon &\leq D^b(m^*_{K-1}, M) \leq 1 \qquad \text{and} \\ 0 &\leq D^b(M, m^*_{K-1}) < \varepsilon. \end{aligned} \tag{3.6}$$

A similar result obtains for the case of $m^*_1 = 0$ and $m^*_2 < M$.

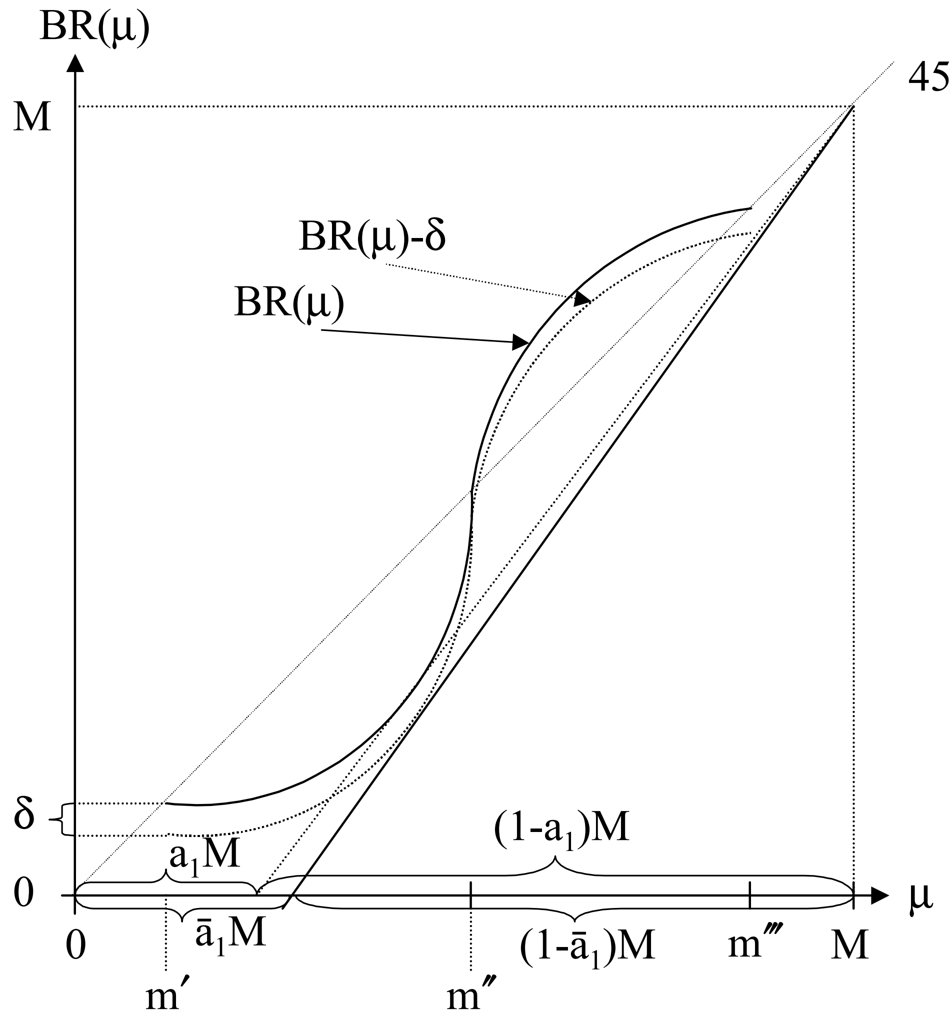

Figure 3.2: Uniform approximation of upward depth for interior NE actions

Proof. Claims (i) and (ii) can be established using techniques similar to those in Lemma 2. To prove claim (i) one has to replace $BR(\mu)-\delta$ by $\underline{b}(\mu) := \max\{BR(\mu)-\delta, \mu\}$ for μ close to M, and to see that $\underline{b}(\mu) \leq b(\mu)$ for all $b \in \mathcal{B}_\delta$. As to the first part of claim (ii), note that $b(\mu) = M-\delta$ for all μ close to M (so that $b(\mu) = M$ implies $\mu = M$) would be compatible with $b \in \mathcal{B}_\delta$. Therefore, the only uniform upper bound on $D^b(m^*_{K-1}, M)$ is 1. Similarly, regarding the second part of claim (ii), observe that $b(\mu) = M$ for all μ close to M is also compatible with $b \in \mathcal{B}_\delta$. Accordingly, though we have $D^{BR}(M, m^*_{K-1}) = 0$, it can be that $D^b(M, m^*_{K-1}) > 0$. ■

Figure 3.3 illustrates part (ii) of Lemma 3. Set $\underline{\mu} := BR^{-1}(M-\delta)$. For any fixed $\delta \in \mathcal{D}$ and arbitrary $\widehat{\mu} \in (\underline{\mu}, M]$, the following shape of best-response approximations $b_{\widehat{\mu}}$ is compatible with $b_{\widehat{\mu}} \in \mathcal{B}_\delta$:

$$b_{\widehat{\mu}}(\mu) = \begin{cases} M & \text{for } \mu \in [\widehat{\mu}, M] \text{ and} \\ M-\delta & \text{for } \mu \in [\underline{\mu}, \widehat{\mu}). \end{cases}$$

Altering $\widehat{\mu}$ allows to shift upward and downward depth arbitrarily within the ranges given in Lemma 3(ii).

Include Figure 3.3 about here.

As a corollary to part (ii) of Lemma 3, we note that actions that are unstable under $BR(\cdot)$ remain unstable under any best-response approximation $b \in \mathcal{B}_\delta$ provided that discretization is sufficiently fine. According to inequality (3.6), there is a uniform lower bound on $D^b(m^*_{K-1}, M)$ holding for all $b \in \mathcal{B}_\delta$. This uniform lower bound converges to $D^{BR}(m^*_{K-1}, M)$ as $\varepsilon > 0$ goes to zero. When M is unstable, we have $D^{BR}(m^*_{K-1}, M) > 0$ so that, for $\varepsilon > 0$ sufficiently small, we even have $D^{BR}(m^*_{K-1}, M) - \varepsilon > \varepsilon$. Since the left hand side of the latter is the uniform lower bound on all $b \in \mathcal{B}_\delta$, it follows that $D^b(M, m^*_{K-1}) < \varepsilon \leq D^{BR}(m^*_{K-1}, M) - \varepsilon \leq D^b(m^*_{K-1}, M)$. Thus, M is unstable with respect to all $b \in \mathcal{B}_\delta$ provided that $\varepsilon > 0$ is chosen sufficiently small.

We continue introducing notation required to state our main theorem. Let $D^{BR}(m) := \min_{T \in \mathcal{T}_m} \sum_{(m',m'') \in T} D^{BR}(m', m'')$, where, for non-adjacent NE actions $m_i < m_j$, we set $D^{BR}(m_i, m_j) = \max_{i<k\leq j} D^{BR}(m_{k-1}, m_k)$ and $D^{BR}(m_j, m_i) = \max_{i<k\leq j} D^{BR}(m_k, m_{k-1})$. Similarly, define $C^{BR}(m) := \min_{T \in \mathcal{T}_m} \sum_{(m',m'') \in T} C^{BR}(m', m'')$ and, for non-adjacent NE actions $m_i < m_j$, $C^{BR}(m_i, m_j) = \max_{i<k\leq j} C^{BR}(m_{k-1}, m_k)$ and $C^{BR}(m_j, m_i) = \max_{i<k\leq j} C^{BR}(m_k, m_{k-1})$ in accordance with Theorem 5 in Kandori and Rob (1995). We call any NE action $m \in \mathcal{M}^{NE}$ *mutation-dominant* if and only if $D^{BR}(m) < D^{BR}(m')$ for all $m' \in \mathcal{M}^{NE} \setminus \{m\}$.

The observation in Lemmas 2 and 3 has a direct consequence stated in Theorem 3.1 below: For any strictly supermodular game Γ with a mutation-dominant NE action, equilibrium selection does not depend on the best-response approximation $b \in \mathcal{B}_\delta$, provided that discretization is sufficiently fine (δ sufficiently small). Moreover, it is the mutation-dominant NE action that is selected. I.e. the selected equilibrium action can be characterized in terms of the original game.

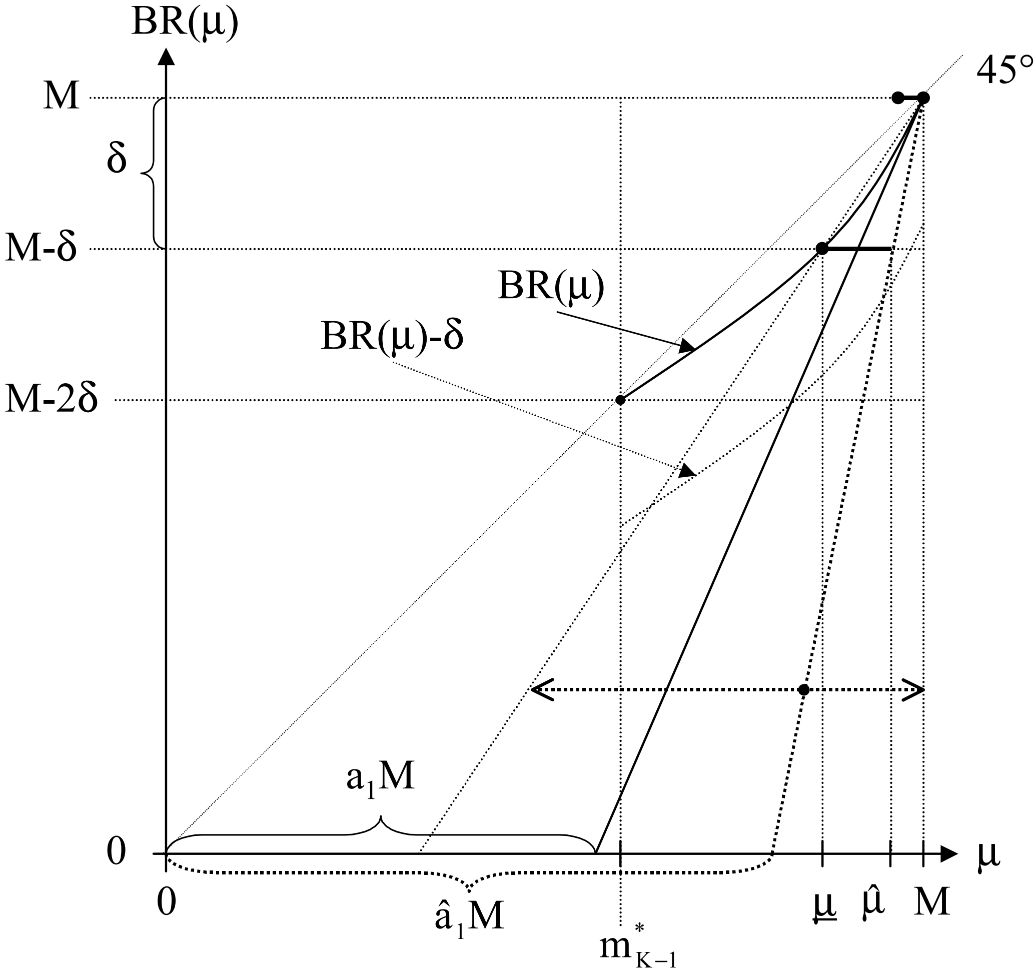

Figure 3.3: The case of boundary NE actions

Theorem 3.1. *Consider any strictly supermodular game that has a mutation-dominant NE action $m^* \in \mathcal{M}^{NE}$. Then there is some $\overline{\delta} > 0$ such that, for all $\delta \in \mathcal{D} \cap (0, \overline{\delta})$ and all $b \in \mathcal{B}_\delta$, we have $\mathcal{M}^*(b) = \{s^{m^*}\}$, i.e. s^{m^*} is the unique stochastically stable state.*

Proof. We prove the claim for the case where both boundary actions, $m = 0, M$, are stable so that we can apply part (i) of Lemma 3. In case of unstable boundary actions the proof has to be adapted, fallling back on part (ii) of the Lemma.

Let $m^* \in \mathcal{M}^{NE}$ be such that

$$D^{BR}(m^*) = \min_{T \in \mathcal{T}_{m^*}} \sum_{(m',m'') \in T} D^{BR}(m', m'') < \min_{T \in \mathcal{T}_m} \sum_{(m',m'') \in T} D^{BR}(m', m'') = D^{BR}(m),$$

for all $m \in \mathcal{M}^{NE} \setminus \{m^*\}$. Pick a minimum cost m^*-tree $T^* \in \mathcal{T}_{m^*}$ so that $D^{BR}(m^*) = \sum_{(m',m'') \in T^*} D^{BR}(m', m'')$ and, for some arbitrary other NE action $m \neq m^*$, an arbitrary m-tree $T \in \mathcal{T}_m$. Choose $0 < \varepsilon < D^{BR}(m) - D^{BR}(m^*)$. By Lemmas 2 and 3, there exists $\overline{\delta} > 0$ such that for all $\delta \in \mathcal{D} \cap (0, \overline{\delta})$ and all $b \in \mathcal{B}_\delta$, we have $\left| D^{BR}(m', m'') - D^b(m', m'') \right| < \varepsilon/(2K-2)$, for arbitrary adjacent NE actions $m' \neq m''$. (Recall that $K < \infty$ denotes the number of NE actions). Notice that by definition of $D^{BR}(m', m'')$ and by Theorem 5 in Kandori and Rob (1995), the depth between non-adjacent NE actions can be determined focusing on adjacent NE actions.

For $T^* \in \mathcal{T}_{m^*}$, it then follows that

$$\begin{aligned} D^b(m^*) &= \sum_{(m',m'') \in T^*} D^b(m', m'') < \sum_{(m',m'') \in T^*} \left[D^{BR}(m', m'') + \frac{\varepsilon}{2(k-1)} \right] \\ &= D^{BR}(m^*) + \frac{\varepsilon}{2}, \end{aligned}$$

whereas, for $T \in \mathcal{T}_m$, we have

$$\begin{aligned} D^b(m) &= \sum_{(m',m'') \in T} D^b(m', m'') > \sum_{(m',m'') \in T} \left[D^{BR}(m', m'') - \frac{\varepsilon}{2(k-1)} \right] \\ &= D^{BR}(m^*) - \frac{\varepsilon}{2}. \end{aligned}$$

Combining both inequalities, we thus obtain

$$D^b(m) - D^b(m^*) \geq D^{BR}(m) - D^{BR}(m^*) - \varepsilon > 0, \tag{3.7}$$

for any $b \in \mathcal{B}_\delta$, where the last inequality holds true by selection of ε.

To establish unique stochastic stability of m^*, it suffices to show that $C^b(m) - C^b(m^*) > 0$, for N sufficiently large. However, recalling the definitions $C^b(m) := \min_{T \in \mathcal{T}_m} \sum_{(m',m'') \in T} C^b(m', m'')$ and $C^b(m', m'') := \max\{\lceil N \cdot D^b(m', m'') \rceil, 1\}$, the claim directly follows from 3.7. ■

4. Conclusions

The main insight of this article has been that, for strictly supermodular games with simultaneous play, rounding of the original best-response function to pure actions has no important impact on the long run equilibrium. Put differently, best-response approximation does not affect equilibrium selection provided that discretization is sufficiently fine.

Our analysis has been warranted, since Hansen and Kaarbøe (2002) have shown for coordination games that equilibrium selection strongly depends on best-response approximations. Notice that their result remains valid even for very fine discretizations of the individual action set.

The difference between Hansen and Kaarbøe (2002) and the present article is the following. At first, for strictly supermodular games exhibiting a finite number of Nash equilibria, the best-response function of the original game does not coincide with the 45°-line. Therefore, the basins of attraction of each Nash equilibrium must display a strictly positive depth (recall from Kandori and Rob, 1995, that the relative size of these depths determines equilibrium selection). For sufficiently fine discretizations of the individual action grid, it thus follows that also the depth of basins of attraction under the approximating best-response function is strictly positive. Our core result then was to establish that the depths corresponding to different feasible best-response approximations actually come arbitrarily close to each other provided that discretization is sufficiently fine.

For coordination games this is different. As shown by Hansen and Kaarbøe, rounding essentially *determines* the depth of the basins of attractions. Since for sufficiently fine discretizations all depths of basins of attraction become arbitrarily close to zero, it remains always possible to affect the relative depth of these basins by constructing appropriate best-response approximations. This is what Hansen and Kaarbøe exploit to generate their results.

References

[1] C. Alos-Ferrer, The discretization of continuum strategy spaces, working paper, 1999.

[2] G. Becker, A note on restaurant pricing and other examples of social influences on price, J. of Polit. Econ. 99 (1990), 1109-1116.

[3] O. Blanchard, N. Kiyotaki, Monopolistic competition and the effects of aggregate demand, Amer. Econ. Rev. 77 (1987), 647-666.

[4] R. Cooper, J. Haltiwanger, Evidence on macroeconomic complementarities, Rev. of Econ. and Stat. 78 (1996), 78-93.

[5] V. Crawford, An "evolutionary" interpretation of Van Huyck, Battalio and Beil's experimental results of coordination, Games Econ. Behav. 3 (1991), 25-59.

[6] V. Crawford, Theory and experiments in the analysis of strategic interaction, in: D. Kreps, K. Wallis (Eds.), Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress, Econometric Society Monographs No. 27, Cambridge University Press, Cambridge, U.K., 1997, 206-242.

[7] M. Freidlin, A. Wentzell, Random Perturbations of Dynamical Systems, Springer Verlag, New York, 1984.

[8] P. Hansen, O. Kaarbøe, On equilibrium selection mechanisms in coordination games with simultaneous play, Econ. Theory 20 (2002), 793-807.

[9] M. Kandori, G. Mailath, R. Rob, Learning, mutation, and lon run equilibria in games, Econometrica 61 (1993), 29-56.

[10] M. Kandori, R.Rob, Evolution of equilibria in the long run: A general theory and applications, J. of Econ. Theory 65 (1995), 383-415.

[11] J. Maynard Smith, Evolution and the Theory of Games, Cambridge University Press, Cambridge, UK, 1982.

[12] J. Milgrom, P. Roberts, Rationalizability, learning, and equilibrium in games with strategic complementarities, Econometrica 58, 1255-1277.

[13] J. Robles, Evolution and long run equilibria in coordination games with summary statistic payoff technologies, J. Econ. Theory 75 (1997), 180-193.

[14] T. Schelling, Hockey helmets, concealed weapons, and daylight saving: A study of binary choices with externalities, J. of Confl. Resol. 17 (1973), 381-428.

[15] D. M. Topkis, Supermodularity and Complementarity, Princeton University Press, Princeton, NJ, 1998.

[16] X. Vives, Nash equilibrium with strategic complementarities, J. of Math. Econ. 19 (1990), 305-321.

[17] P. Young, The evolution of conventions, Econometrica 61 (1993), 57-84.

Department of Economics
University of Bergen
Fosswinckels gate 6
N-5007 Bergen, Norway
Phone: +47 55 58 92 00
Telefax: +47 55 58 92 10
http://www.svf.uib.no/econ