TI 1999-096/1
Tinbergen Institute Discussion Paper

Equilibrium Selection in Games with Macroeconomic Complementarities

Oddvar M. Kaarbøe
Alexander F. Tieman

Tinbergen Institute
The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Keizersgracht 482
1017 EG Amsterdam
The Netherlands
Tel.: +31.(0)20.5513500
Fax: +31.(0)20.5513555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31.(0)10.4088900
Fax: +31.(0)10.4089031

Most TI discussion papers can be downloaded at
http://www.tinbergen.nl

Equilibrium Selection in Games with Macroeconomic Complementarities*

Oddvar M. Kaarbøe† and Alexander F. Tieman‡

December 20, 1999

Abstract

We apply the stochastic evolutionary approach of equilibrium selection to macroeconomic models in which a complementarity at the macro level is present. These models often exhibit multiple Pareto-ranked Nash equilibria, and the best response-correspondence of an individual increases with a measure of the aggregate state of the economy. Our main theoretical result shows how the equilibrium that is singled out by the evolutionary dynamics is directly related to the underlying externality that creates the multiplicity problem in the underlying macroeconomic stage game. We also provide clarifying examples from the macroeconomic literature.

JEL classification: C63, C72, C73, E19, L16

*The authors thank Gerard van der Laan, Terje Lensberg and Paul Frijters for comments.

†Corresponding author. University of Bergen, Dept. Economics, Fosswinckelsg. 6, 5007 Bergen, Norway. Ph: +47-55589226, Fax: +47-55589210, E-mail: oddvar.kaarboe@econ.uib.no

‡Tinbergen Institute and Free University, Dept. Econometrics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. Ph: +31-20-4446022, Fax: +31-20-4446020, E-mail: XTieman@econ.vu.nl. URL http://www.econ.vu.nl/medewerkers/xtieman

1. Introduction

The purpose of this paper is to apply the stochastic evolutionary approach of equilibrium selection proposed by Blume (1993), Kandori, Mailath, and Rob (1993), and Young (1993) to macroeconomic models of coordination failure. In macroeconomic models of coordination failure, individuals' decisions are interdependent. Without the assumption of perfect coordination (through the Walrasian auctioneer) formalization of this interdependency has lead to models with multiple Nash equilibria. Often these equilibria are Pareto-ranked and the models are capable of explaining why the economy might become stuck in a socially inefficient equilibrium. The main result of this paper is on evolutionary selection between these Pareto-ranked equilibria.

The game theoretical framework for analysis of the existence of multiple Pareto-ranked Nash equilibria was given by Cooper and John (1988). They showed that multiple equilibria can arise in a (static) game when strategic complementarities are present, and spillovers exist between agents at the level of payoffs. Cooper and John (1988) also divided the literature in three groups depending on whether agents are interconnected through the production process (e.g. in Bryant (1983)), through an exchange arrangement (e.g. in Diamond (1982)), or whether agents are in multisector models of imperfect competition where agents are specializing in what they produce and interconnected in what they consume (e.g. in Hart (1982) and Weitzman (1982)).

While the typical paper in this literature does give conditions for the economy to exhibit multiple Pareto-ranked equilibria, it contains no analysis of which set of equilibria is most likely to be observed.[1] This is perhaps surprising since some models are capable of generating a continuum of equilibria that have markedly different properties.[2] Therefore pinning down which equilibrium is selected is vital if any predictive power is to come from coordination failure models. In particular, the necessity of equilibrium selection is important for sensitivity analysis, and for investigating if the presence of coordination failure implies a call for more active (government) interventions.

In this paper we address these issues by embedding the main features of the different macroeconomic models of coordination failure into the class of strict supermodular games played simultaneously by the entire population. In this class of games, introduced by Topkis (1979), and further explored by Milgrom and Roberts (1990) and Vives (1990), the best-response correspondences are increasing, so the players' strategies are 'strategic complements'. When the best-response of an individual player increases with some measure of the aggregate state of the economy, a macroeconomic complementarity is present. In the coordination failure literature this measure is often taken to be some average of the current strategy profile. For example, in models of monopolistic competition (e.g. Blanchard and Kiyotaki (1987)), the marginal payoffs of one's own price depend positively on the (geometric) average of the other agents' prices, and in production externality models (e.g. Cooper and Haltiwanger (1996)) the marginal payoff of an individual's effort (or production) depends positively on the average level of effort exerted in the economy. In this article we present such examples. Thus we think of agents'

[1] One exception is Cooper (1994) who proposes a history-dependent selection criterium.

[2] Bryant (1983) models an input game that generates a continuum of equilibria. Cooper and John (1988) give examples which show that both the search externality model (Diamond (1982)) and the multisector model of imperfect competition (Hart (1982)) are capable of generating a continuum of equilibria.

strategic variables as being e.g. prices, quantities or effort levels. One contribution of this paper is to show how evolutionary selection (learning) and mutations (mistakes and experimentation) interact to pin down which equilibrium is selected in these macroeconomic models of coordination failure if the players repeatedly play a strict supermodular stage game with a summary statistic of the population state. More specifically, we show how the equilibrium that is most likely to be observed (in the long run) directly relates to the underlying externality that creates the multiplicity problem in the first place.

To establish this result, we expand the evolutionary literature on equilibrium selection in strict supermodular games outlined in Kandori and Rob (1995) (hereafter KR) in two directions. First we follow a recent paper by Kaarbøe (1999) and analyze equilibrium selection in the class of strict supermodular games when the interaction takes place on the population-wide level. In the model, an agent plays against a summary statistic of the population state, as e.g. the average price in a setting where we have price competition. Thus we abandon the context of random pairing commonly used in evolutionary literature, and focus on an interaction structure that fits the framework of macroeconomics.[3,4] Secondly, we follow the traditional macroeconomic approach and let (almost-)rational players exist in the population. These players form almost-rational expectations about next period's state, and play a best-response to these expectations. We define almost-rational expectations as expectations which are correct in the absence

[3]Schelling (1973) introduced simultaneous play in economics. In biology this type of interaction is called "playing the field" (Maynard Smith (1982)). Robles (1997) and Hansen and Kaarbøe (1997) study coordination games with a simultaneous play interaction structure. Crawford (1991, 1997) also argues for introducing genuine simultaneous interaction into the evolutionary literature.

[4]KR give one example of simultaneous play under the assumption that individual payoffs are linear in the opponents' strategies. Hence, many economic models do not fit into their framework.

of mutations, i.e. when no mutations occur between the choosing of actions and the realization of payoffs. We call these players dynamic optimizers. In addition, boundedly rational players coexist with the dynamic optimizers. These players are called myopic optimizers since they play a best-reply to the population state they observe, i.e. they form adaptive expectations. By using this framework we provide an evolutionary analysis in which selection among the equilibria is independent of the presence of boundedly rational players, while the time spent outside the set of equilibria is dependent of myopic optimizers being present. The insight stemming from the interaction between the almost-rational and the boundedly rational type is that in equilibrium, boundedly rational agents do better than almost-rational agents if almost-rational agents have to pay a positive per period cost for performing the (complicated) calculations or for gathering the relevant information associated with the forming of almost-rational expectations. Thus, in equilibrium, selection works against this more sophisticated type, see also e.g. Conlisk (1980), Droste and Tuinstra (1998) and Droste, Hommes, and Tuinstra (1999).

The structure of this paper is as follows. Section 2 provides the general model. In section 3 we present the theoretical analysis, while section 4 provides three macroeconomic examples that fit the proposed structure, namely a production externality model, Bertrand competition with demand externalities, and a model of search and matching. Section 5 states some concluding remarks.

2. The Model

We consider a finite population $\mathcal{N} := \{1, 2, ..., N\}$ consisting of N players. In each period $t = 1, 2, \ldots$ all players interact in a market structure, i.e. they all 'play the field'. Each player's payoff is determined by his own action and by the value of a summary statistic $\sigma(\cdot)$ of the actions chosen by all players in the population.

2.1. The Stage Game

The stage game we posit is a symmetric strict supermodular game, in which the role of the column player is played by the summary statistic $\sigma(\cdot)$. The individual player thus chooses a row. The set $\mathcal{M}(\gamma)$ of pure actions in a strict supermodular stage game is by definition partially ordered. Here we follow KR and assume that $\mathcal{M}(\gamma)$ contains a finite number of $\gamma + 1$ actions, and is completely ordered from low (action 0) to high (action M), i.e. $\mathcal{M}(\gamma) := \{0, \frac{M}{\gamma}, 2\frac{M}{\gamma}, ..., M\}$ for an arbitrary $\gamma \in \mathbb{N}^+$. The only condition required is that all the Nash equilibria in the game played on the continuous action space $[0, \infty)$ belong to this grid. This indicates that we see the action space $\mathcal{M}(\gamma)$ as a discrete approximation of the continuous action space $[0, \infty)$.

We assume monotonicity of the summary statistic, i.e. when there are two states[5] s and s' with $s \succ s'$, where $\succ$ refers to first order stochastic dominance[6], $\sigma(s) > \sigma(s')$. Furthermore, we assume that the summary statistic $\sigma(\cdot)$ only takes values on the action

[5] The exact definition of a state follows in section 2.2.

[6] A (frequency or probability) distribution or density $f(\cdot)$ on $\mathcal{M}(\gamma)$ with cumulative distribution $F(a) = \sum_{k \leq a} f(k)$ first order stochastically dominates a distribution $g(\cdot)$ on $\mathcal{M}(\gamma)$ with cumulative distribution $G(a) = \sum_{k \leq a} g(k)$ when for all $a \in \mathcal{M}(\gamma)$, it holds that $F(a) \leq G(a)$, with strict inequality for at least one a.

grid, i.e. $\sigma(s) \in \mathcal{M}(\gamma)$, $\forall s$, and that the summary statistic takes the value m when the state is such that all players play action $m \in \mathcal{M}(\gamma)$. Examples of a summary statistic meeting these conditions are the arithmetic, harmonic and geometric average of the actions in the population. This implies that we model the summary statistic as a piecewise constant function from the state space to the action space. [7]

We also follow KR's assumption on 'continuity' of the best-responses on the grid, i.e. we assume that if $br(\sigma(s)) = m$ and $br(\sigma(s'')) = m''$ and $m < m' < m''$, then there exists an $\alpha \in (0,1)$ such that $br(\alpha\sigma(s) + (1-\alpha)\sigma(s'')) = m'$, where $br(\sigma(s))$ denotes the set of best-responses to the summary static $\sigma(s)$ of state s. Since we have continuity of the summary statistic, we can then also find a $\beta \in (0,1)$ such that $br(\sigma(\beta s + (1-\beta)s'')) = m'$.

We denote a player's payoff when she plays action $m \in \mathcal{M}(\gamma)$ against the summary statistic $\sigma(\cdot)$ by $u(m, \sigma(\cdot))$. We can now define supermodularity for a stage game in which an agent plays against a summary statistic.

Definition 1. *The stage game is a strict supermodular game if for any pair of strategies* $0 \leq m < m' \leq M$ *the payoff differences* $u(m', \sigma(\cdot)) - u(m, \sigma(\cdot))$ *are strictly increasing in* $\sigma(\cdot)$.

In this setup, we define the equilibrium concept as follows

[7] In Kaarbøe (1999) it is shown that modelling 'playing the field' when the grid of the action space is arbitrarily fine (i.e. when $\gamma \to \infty$), the difficulties which arise from piecewise constant best responses in coordination games (see Robles (1997) and Hansen and Kaarbøe (1997)) do not arise in a supermodular stage game.

Definition 2. *A state s is an equilibrium state if under s each player n, $n = 1, \ldots, N$, plays an action in $br(\sigma(s))$.*

The ordering of the actions and the supermodularity of the stage game ensures that, in case the stage game exhibits multiple stable equilibria, these equilibria are rankable in the Pareto sense (see also section 3).

2.2. The Players' Types and the State Space

At every time $t = 1, 2, \ldots$ each player in the population uses a particular update rule, which we refer to as a player's type. An update rule specifies how a player updates the action she plays in the stage game. We consider the update rules 'myopic optimization' (μ) and 'dynamic optimization' (δ) and thus we accommodate two types of players. Players that update according to μ are called myopic optimizers, while players that update according to δ are called dynamic optimizers.

In every period t each player is characterized by a pair (m, i) identifying the action m she currently plays and her type i. For every period t, the state $s(t) = (s_0^\mu(t), s_0^\delta(t), ..., s_M^\mu(t), s_M^\delta(t))$ is a vector whose m-th element, $s_m^i(t)$ $i = \mu, \delta$, represents the number of type i-players using action $m \in \mathcal{M}(\gamma)$ at time t. Thus, the state space is given by $\mathcal{S} = \{1, 2, \ldots, N\}^{2(\gamma+1)}$, where for every $s \in \mathcal{S}$, we have that $\sum_{m \in \mathcal{M}(\gamma)} (s_m^\mu + s_m^\delta) = N$. We take account of the state $s(t)$ at the start of period t. The total number of players playing action m at time t is $s_m(t) = s_m^\mu(t) + s_m^\delta(t)$, the number of myopic optimizers in the population at time t is $N_t^\mu = \sum_{m \in \mathcal{M}(\gamma)} s_m^\mu(t)$ and the number of dynamic optimizers in the population at time t is $N_t^\delta = \sum_{m \in \mathcal{M}(\gamma)} s_m^\delta(t) = N - N_t^\mu$.

Note that $s(t)$ is a frequency distribution and thus the concept of first order stochastic dominance between states is well defined.

As mentioned before, at each time t, all players play the stage game simultaneously and 'against' the same summary statistic of the entire population state. Therefore, the average payoff of an action $m \in \mathcal{M}$ at time t, $\overline{u}_m(t)$, is equal to the realized payoff of action m, $u(m, \sigma(s(t)))$, where $s(t)$ is the state at the time the stage game is played (see also Figure 2.1, p. 13).

For completeness we label the action that player $n \in \mathcal{N}$ is playing as $a_n \in \mathcal{M}(\gamma)$ and we define the indicator functions I_n^t for $t = 0, 1, \ldots,$ and $n \in \mathcal{N}$ as

$$I_n^t = \begin{cases} 1, & \text{if player } n \text{ is of type } \delta \text{ at the beginning of period } t, \\ 0, & \text{otherwise.} \end{cases} .$$

A player that is of type i, $i = \mu, \delta$, makes a positive cost at each period respectively to calculate a best-response, or to infer next-period behavior of the other players, and then calculate a best-response to that. We assume that the dynamic optimizers have at least as high per-period costs that the myopic optimizers, both because of the more complicated calculation dynamic optimizers have to make, and because they need to gather more information to perform those calculations. We normalize the costs of myopic optimizers to zero and pose a non-negative per-period cost $c \geq 0$ for dynamic optimizers.

2.3. The Mutation-Free Dynamics

The mutation-free dynamics operates on two levels. At the first level, at each time t, all players get the possibility to revise their action. They do so according to a particular update rule. At the second level, at each time t, some (possibly all or none) players get the possibility to switch update rules, i.e. to update their types. We now specify both processes in detail.

2.3.1. Updating Actions

At every $t = 1, 2, ...$ and before play is conducted each player gets the opportunity to update his action. The player chooses the new action as dictated by her type.

Assumption A (on myopic optimizers). Myopic optimizers observe the structure of the stage game and, at each time t, the summary statistic $\sigma(s(t))$. Subsequently they play a myopic pure best-response to $\sigma(s(t))$ in the stage game at time $t + 1$. In case of ties, there is a positive probability of choosing each of the best actions.

Assumption B (on dynamic optimizers). Dynamic optimizers observe the stage game and, at each time t, the state $s(t)$, the number of players that switch type due to type adjustment, and the birth & death process. Based on this information they form expectations which are correct when no mutations occur, i.e. they correctly predict the state $\widetilde{s}(t)$ after all players have chosen an action, but before mutations occur. They consequently choose a pure best-reply to these expectations. In case of ties, there is a positive probability of choosing each of the best actions.

Dynamic optimizers form almost-rational expectations, in the sense that their beliefs incorporate all relevant information in the model, except the information regarding mutations.[8] Thus, at time t, dynamic optimizers correctly predict the actions all players want to play at time $t + 1$, but do not take into account that mutations might alter the actual actions (and thus the payoffs) played in the stage game (see also Figure 2.1, p. 13).

Now we are able to define the update rules in terms of the state space. At time t, a myopic optimizer switches to action $m^* \in \arg\max_{m \in \mathcal{M}(\gamma)} u(m, \sigma(s(t)))$, while a dynamic optimizer switches to action $m^* \in \arg\max_{m \in \mathcal{M}(\gamma)} u(m, \sigma(\widetilde{s}(t, m)))$, where $\widetilde{s}(t, m)$ denotes the state in which all myopic optimizers play a best reply to $\sigma(s(t))$ and the dynamic optimizers play action m.[9] This implies that dynamic optimizers explicitly take account of the number of myopic and dynamic optimizers playing each action in the population and that they infer which action these other players will play in the stage game. Thus, dynamic optimizers' action m^* follows from calculating a fixed point of the correspondence that takes the behavior of the myopic optimizers as given, and reflects the optimal reaction to the current state by all dynamic optimizers. For both types of players it holds that when m^* is not unique, they pick an arbitrary action from their set of argmax-es.

Neither type of player will ever play a strictly dominated action. In the symmetric stage game, it holds that if action $m \in \mathcal{M}(\gamma)$ is strictly dominated for the row player, column m of the payoff matrix of the stage game is also strictly dominated. Because

[8] This is done because we believe that from the point of the players a key feature of mutations is their unpredictability, in the sense that not even the probability distribution of the mutations is known to the players.

[9] In case the $\arg\max$ is not unique, the state $\widetilde{s}(t, \cdot)$ should be denoted as a function of the vector containing the action of each dynamic optimizer in the population.

of these two observations, we can restrict attention to symmetric strict supermodular stage games from which all strictly dominated strategies have been iteratively removed. A well-know result on supermodular games (see e.g. KR, Proposition 6) then yields that

Proposition 1. *After all strictly dominated strategies have been iteratively removed from the game, the smallest and largest of the remaining strategies are pure strategy Nash Equilibria (NE for short).*

In section 3 we will establish the link between these Nash equilibria of the stage game and the equilibrium concept as given in Definition 2 for games in which all players play the field.

2.3.2. Updating Types

After the stage game has been played, average realized payoffs among dynamic optimizers and myopic optimizers become common knowledge. Explicit account is taken of the costs $c \geq 0$ that dynamic optimizers make. We label the realized average payoffs $\overline{u}_t^\delta$ and $\overline{u}_t^\mu$ respectively. Hence,

$$\begin{aligned} \overline{u}_t^\mu &= \frac{1}{N_t^\mu} \sum_{n \in \mathcal{N}} u\left(a_n, \sigma\left(s\left(t\right)\right)\right)\left(1 - I_n^t\right), \\ \overline{u}_t^\delta &= \frac{1}{N_t^\delta} \sum_{n \in \mathcal{N}} \left[u\left(a_n, \sigma\left(s\left(t\right)\right)\right) - c\right] I_n^t, \end{aligned}$$

We posit that with probability $\theta \in [0, 1]$ each player receives the opportunity of revising her type. Thus, at each time t, each player takes an independent draw from a Bernoulli trial. With probability $\theta \in [0, 1]$ this draw produces the outcome 'learn', and the player chooses a new type, in the way described below. With the complementary probability $1 - \theta$, the draw produces the outcome 'do not learn' and the player stays with her current type. The possibility to update types is called a learning draw. In the case a player receives the opportunity of revising her type, she changes type if and only if average payoff to the other type in the last period is strictly higher than the average payoff her own type received in the same period, i.e. a player of type i, $i = \mu, \delta$, switches type iff $\overline{u}_t^i < \overline{u}_t^j$, $j = \mu, \delta$, $j \neq i$. When a player changes type we assume that she starts to play the action prescribed by her new type.

Note that we explicitly admit for $\theta = 0$, i.e. no player ever updates his type, and $\theta = 1$, i.e. updating types is always prompt and no inertia is displayed. Furthermore, note that although we let players switch type based on the perception of average payoffs per type, this is not an essential feature of our model. Other type switching rules, like e.g. switching to the other type if the average payoff of that type is higher than the realized payoff of the player, do not alter the basic results. In fact, such a rule comes down to the same behavior when best-responses are unique and thus all players of the same type play the same action.

2.4. Mutation Dynamics

At both levels we allow the above described mutation-free dynamics to be slightly perturbed by deviations. We refer to these deviations as a birth & death process on the type level and as mutations on the action level. Before we specify both processes in details, we present a graphic overview of the sequence of events during a time period t in figure

2.1.

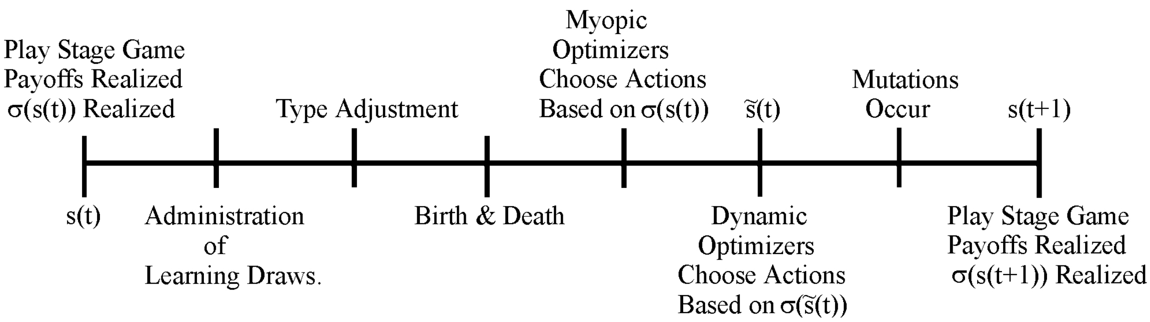

Figure 2.1: The sequence of events during a time period.

2.4.1. Birth & Death at the Type Level

The birth & death process lets to it that some randomly chosen players 'die' and replaces them by newborn players. This event takes place after the stochastic type adjustment according to the mutation-free dynamics has taken place. We assume that at each time t there is a positive probability $2\kappa \geq 0$ that a player dies. Setting $\kappa = 0$ means that no birth & death takes place. When $\kappa > 0$ and a player dies, she is replaced by a newborn player, who is of either type with equal probability $\frac{1}{2}$. This formulation boils down to each player having a probability $\kappa \geq 0$ of being replaced by a new player of the other type.

Formally, let the random variables $V_t\left(\widehat{N}_t^{\mu}, \kappa\right)$ and $W_t\left(\widehat{N}_t^{\delta}, \kappa\right)$ denote the number of myopic optimizers and dynamic optimizers respectively that switch type at time t, where $\widehat{N}_t^i$, is the number of type i players in the population at time t after players were able to revise their type according to the mutation-free type adjustment dynamics. Note that $V_t\left(\widehat{N}_t^{\mu}, \kappa\right)$ and $W_t\left(\widehat{N}_t^{\delta}, \kappa\right)$ both have a binomial distribution $Bin\left(n, p\right)$ with parameters

$n = \widehat{N}_t^{\mu}$ and $n = \widehat{N}_t^{\delta}$ respectively and $p = \kappa$. Now, we have that

$$N_{t+1}^{\mu} = \widehat{N}_t^{\mu} - V_t\left(\widehat{N}_t^{\mu}, \kappa\right) + W_t\left(\widehat{N}_t^{\delta}, \kappa\right),$$

and

$$N_{t+1}^{\delta} = \widehat{N}_t^{\delta} - W_t\left(\widehat{N}_t^{\delta}, \kappa\right) + V_t\left(\widehat{N}_t^{\mu}, \kappa\right).$$

2.4.2. Mutations at the Action Level

At each time t, each player is subject to some common and independent (across players and time) probability of implementing an action other than the one described by her type. We label such a deviation from the prescribed action a 'mutation'. Mutations can be thought of as representing the making of a mistake in the implementation or calculation phase, as suggested by Kandori, Mailath, and Rob (1993) and Young (1993) and implemented in e.g. Van Damme and Weibull (1998), where more costly mistakes are assumed to occur less frequently.[10]

When a mutation occurs, the player chooses any action in a purely arbitrarily manner. The mutations happen after all players have calculated which action they want to play but prior to play being conducted.

[10] Alternatively, a mutation can be thought of as a boundedly rational player experimenting, because she thinks her action will not be a best-response to the state she plays against (which is correct when the system is out of equilibrium). However, such a boundedly rational player is unable to calculate this best-response, since she cannot form correct expectations on next period's state. Therefore, some experimentation can be of use. In this context it seems natural to assume that myopic optimizers experiment, while dynamic optimizers do not, since they already make on average correct predictions about the state they will be playing against and experimentation thus is of no use to them. This would result in a higher combined rate of experimentation for myopic optimizers than for dynamic optimizers, like in e.g. Kaarbøe and Tieman (1999). However, in the current setting such a generalization does not alter the basic results of the model and we will therefore not implement it in the model.

Assumption C (on mutations). At every time t, each player makes a mistake or experiments in the implementation of her action with some common and independent probability $\varepsilon > 0$. In that case she plays an action $m \in \mathcal{M}(\gamma)$ with positive probability on each $m \in \mathcal{M}(\gamma)$, i.e. the probability distribution of mutations has full support. Furthermore, this distribution is fixed over time and independent of ε.

The composition of the adjustment processes generates a discrete-time Markov-process over the finite state space $\mathcal{S}$, whose transition matrix is denoted by $P(\varepsilon, \kappa) = (p_{ss'}(\varepsilon, \kappa))$. An element $p_{ss'}(\varepsilon, \kappa)$ represents the transition probability of moving to state s' at time $t + 1$ conditional on being in state s at time t. The dynamics without mutations at the action level and without birth & death corresponds to $P(0, 0)$.

The occurrence of mutations and birth & death implies that every transition has positive probability. It is a standard result that such Markov chains have a unique stationary (invariant) probability distribution. Let $\phi_\kappa(\varepsilon)$ denote the unique invariant distribution of $P(\varepsilon, \kappa)$ for each $\varepsilon > 0$ and fixed $\kappa > 0$. Our aim is to characterize the unique invariant distribution $\phi^*_\kappa := \lim_{\varepsilon \to 0} \phi_\kappa(\varepsilon)$. Using arguments in Freidlin and Wentzell (1984), Young (1993) has shown that this limit exists. The states that have strict positive measure under ϕ^*_κ are called long run equilibria.

3. Theoretical Results

In this section we first state some useful results on the best-response structure and the equilibria in supermodular games in which all players play the field. Then, in Proposition 6, we characterize the limit sets of the Markov process, when no mutations are present and there is no birth & death and no type switching. That is, we look at what happens when we have a population consisting of a fixed number of myopic optimizers and a fixed number of dynamic optimizers. The proposition says that the limit sets correspond one-to-one with the set of pure strategy Nash equilibria of the stage game. It also argues that the presence of myopic optimizers makes it possible that time is spent out-of-equilibrium. Second, we admit for stochastic type switching and birth & death, but still set the probability of mutation to $\varepsilon = 0$. This leads to the result in Proposition 7 that still the limit sets correspond one-to-one with the set of pure strategy Nash equilibria of the stage game. Furthermore, we establish that, in the presence of positive costs $c > 0$ for dynamic optimizers, once in equilibrium, the fraction of dynamic optimizers will become small. Third, we present the main result of this section in Proposition 8. Here, we have both mutations and stochastic type switching and birth & death present in the model. This main result is on selection among the multiple Pareto rankable equilibria and states that the equilibrium with the deepest basin of attraction will be selected. The proposition is followed by Corollary 1, which establishes the same result for the special case in which only dynamic optimizers are present in the model.

We first state a proposition on monotonicity of best responses over the set of states. A similar proposition was proven for random matching among agents in Kandori and Rob (1995) and for the case of playing the field in Kaarbøe (1999). Let $br(\sigma(s))$ denote the set of best responses to $\sigma(s)$.

Proposition 2. *Let s and s' be two states such that $s \succ s'$. Then* $\min br(\sigma(s)) \geq \max br(\sigma(s'))$.

Proof.

The proof is by contradiction. Take an arbitrary $m \in br(\sigma(s))$ and $m' \in br(\sigma(s'))$ and suppose $m < m'$. Then, the strict supermodular structure of the stage game yields

$$u(m, \sigma(s')) - u(m', \sigma(s')) > u(m, \sigma(s)) - u(m', \sigma(s)).$$

The nature of best responses yields $u(m, \sigma(s)) - u(m', \sigma(s)) \geq 0$ and $u(m, \sigma(s')) - u(m', \sigma(s')) \leq 0$. Thus, we have

$$u(m, \sigma(s')) - u(m', \sigma(s')) \leq 0 \leq u(m, \sigma(s)) - u(m', \sigma(s)),$$

a contradiction with the above.

Since the above holds for any $m \in br(\sigma(s))$ and $m' \in br(\sigma(s'))$, it also holds that $\min br(\sigma(s)) \geq \max br(\sigma(s'))$. ¤

Second, we stow that a state to which there is more than one best-response is strongly unstable in the sense that it is not even a stationary point of the best-response dynamics. A similar result was proven for the case of random matching by KR (Proposition 8) and for the case of playing the field by Kaarbøe (1999). By $\# A$ we denote the number of elements in the set A.

Proposition 3. *A state s for which $\# br(\sigma(s)) > 1$ cannot be a stationary point of the best-response dynamics.*

Proof.

Consider a population with only myopic optimizers and an arbitrary state s for which $\# br(\sigma(s)) > 1$. According to assumption A on the updating of myopic optimizers, at the next stage, all myopic optimizers randomize over all the $\# br(\sigma(s))$ elements in $br(\sigma(s))$. Thus, with positive probability, the state s is left. Therefore a state s for which $\# br(\sigma(s)) > 1$ is not a stationary point of the best-response dynamic. ¤

Third, we prove that states which mimic mixed actions are also strongly unstable in the sense that they are not even stationary points of the best-response dynamic.

Proposition 4. *A state s in which multiple actions are played (i.e. a polymorphic or mixed state s) cannot be a stationary point of the best-response dynamics.*

Proof.

Consider a population with only myopic optimizers. Furthermore, consider an arbitrary mixed state s, in which myopic optimizer i plays action a_i and myopic optimizer j plays action $a_j \neq a_i$. If not both $a_i \in br(\sigma(s))$ and $a_j \in br(\sigma(s))$, either i or j (or both) will change his action, since they both get the possibility to revise their action at each time t, and the state is thus not stationary under the best-response dynamics. Thus, in order for the mixed state s to be stationary, all actions played under s must be elements of the set $br(\sigma(s))$. But then Proposition 3 states that s is not a stationary point of the best-response dynamic. ¤

Obviously, when a polymorphic state s is unstable under best-response dynamics, it will surely not be stable in dynamics where some players play best-responses to last period's state and other players are dynamic optimizers or in dynamics where all players are dynamic optimizers (see also the proof of Proposition 6 below), since dynamic optimizers play a fixed point of the best-response correspondence. Therefore, the results in Propositions 3 and 4 should not be interpreted as results on populations consisting only of myopic optimizers, but have a much broader scope and are applicable to any population consisting of an arbitrary non-negative number of myopic optimizers and an arbitrary non-negative number of dynamic optimizers.

From the above propositions it follows directly that an equilibrium state s under which multiple actions are played is strongly unstable in the sense specified above. For this reason we only focus on equilibria in which all players play the same action. This leads to the following proposition.

Proposition 5. *No stable asymmetric equilibria in pure actions exist. The set of stable equilibrium action pairs $\mathcal{M}^*(\gamma)$ is a subset of the main diagonal of the stage game. The set S^* of stable equilibria contains only strict equilibria.*

Proof.

Consider a state s. If multiple actions are played under s, it is strongly unstable in the sense of Proposition 4. Thus it cannot be a stable equilibrium. If s is a stable equilibrium, all players play the same action, say action $m \in \mathcal{M}(\gamma)$. By assumption, this leads to a summary statistic $\sigma(s) = m$. Thus, when stable equilibria exist, they lie on the main diagonal of the stage game. From Proposition 1 we have that at least one equilibrium exists.

The fact that all stable equilibria are strict is a straightforward corollary from Proposition 3, which establishes that the set of best-responses to a stable equilibrium state s can only be a singleton set. ¤

Thus we have established the link between symmetric Nash equilibria in the stage game and (stable) symmetric equilibria in the playing the field setting. We now define the set of stable (symmetric) equilibria as $S^* = \{s^* | \exists\, m^* \in \mathcal{M}(\gamma)$ such that $s_{m^*} = s^{\mu}_{m^*} + s^{\delta}_{m^*} = N$ and $m^* = br(\sigma(s^*))\}$ and the set of stable equilibrium action pairs as $\mathcal{M}^*(\gamma) := \{(m^*, \sigma(s^*)) \mid s^* \in S^*, m^*, \sigma(s^*) \in \mathcal{M}(\gamma)$ and $m^* = br(\sigma(s^*))\}$. In an equilibrium with $(m^*, \sigma(s^*)) \in \mathcal{M}^*(\gamma)$, all players get the same payoff from the stage game $u(m, \sigma(s))$. Therefore, Pareto ranking stable equilibria comes down to comparing scalar payoffs.

We now state the convergence result for the model without any randomness, i.e. without mutation, stochastic type switching and without birth & death.[11]

Proposition 6. *Consider the model without any randomness and with both a fixed non-negative number of myopic optimizers, and a fixed non-negative number of dynamic optimizers present in the population, i.e. the model with $\varepsilon = 0$, $\theta = 0$ and $\kappa = 0$. Then, the limit sets of the model correspond one-to-one with the collection of pure strategy (strict) NE.*

Proof.

A necessary condition for a state to be a limit set, is to be a stationary point of the best-response correspondence. From Proposition 4 we see that a mixed state cannot be

[11] Krishna (1991) presents a similar global convergence result in supermodular games with multiple equilibria, but considers a different adjustment process, namely fictitious play dynamics.

a stationary point of the best-response correspondence. Thus it cannot be a limit set. Since a non-NE pure state is not a stationary point of the best-response correspondence either, it can also not be a limit set.

Every pure strategy NE corresponds to a monomorphic state s which is a best-reply to itself. In the non-generic case that the NE corresponding to s is not strict, $\#br(\sigma(s)) > 1$ and Proposition 3 states that s cannot be a limit set. When all players play a certain strict NE action m, the myopic optimizers will play m next period also. Next period, the dynamic optimizers will play an action which is a stationary point of their best-reply correspondence, given that all myopic optimizers play m. The current NE action m is one such fixed point. Consider any other point which involves all dynamic optimizers playing one NE (say m') and all myopic optimizers playing m. In that case it must hold that the payoff from the stage game to the dynamic optimizers (without accounting for the costs $c \geq 0$) is at least as high as the payoff from the stage game to the myopic optimizers. Thus, the myopic optimizers will change their action next period, and, as already established above, such a state is therefore not a stationary point of the best-response correspondence and thus it is not stable and cannot be a limit set.

Since generically all NE are strict, a possible path that leads to an arbitrary strict NE in finite time follows easily. A strict NE corresponding to state s, has a basin of attraction of positive size, i.e. there exist states from which playing best-responses will lead the state to the NE state s. Furthermore, each state in $\mathcal{S}$ lies in at least one such basin of attraction, i.e. there are no states from which no best-response path leads to a NE state. Given that this is true for best-response dynamics, it is also true for a setting in which some players play best-response to last period's state and others have almost-rational expectations, as the dynamic optimizers in this model. The presence of dynamic optimizers only speeds up the time it takes to reach the NE state s, since they play a best-response to the state in which the myopic optimizers have already taken a best-response and in which all dynamic optimizers are also taking best responses. The above also holds for a setting with only dynamic optimizers present. In that case the movement to the NE state will be prompt and thus no time will be spent out of equilibrium. ¤

Since the proposition looks at the model without type switching, the costs dynamic optimizers make every period to form predictions play no role in the above. The result establishes that without stochasticity, only pure strategy NE are limit sets. Thus, these states are the only possible candidates for stationary states in the model with stochasticity, and selection of stable states boils down to selection among the pure strategy NE.

When we admit stochastic type switching and birth & death, but still set the probability of mutation to $\varepsilon = 0$, we get a result similar to the one above, with the additional feature that dynamic optimizers are selected against once the system is in a limit set.

Proposition 7. *Consider the model without mutations, i.e.* $\varepsilon = 0$*. Then, the limit sets of the model correspond one-to-one with the collection of pure strategy NE. For* $c > 0$ *and* $\theta > 0$*, once a limit set is reached, the fraction of dynamic optimizers in the population becomes small and will remain small. It may reach zero, but, with probability* 1*, it will not remain zero forever.*

Proof.

The last claim is established first: The birth & death process sees to it that no type can become extinct forever, i.e. when a type disappears from the population, it will be reintroduced in finite time with probability 1. So, most of the time, both the number of myopic optimizers and the number of dynamic optimizers in the population will be positive. From Proposition 6 it is then clear that the limit sets of the model correspond one-to-one with the set of pure strategy NE. Once a limit set is reached, the system cannot leave it, since the probability of a mutation is zero, and since players that change type (through the birth & death process or through type switching) will still play the same action after having changed type. In a limit set, all players play the same (NE) action m and hence get the same payoff from the stage game. However, when $c > 0$, the dynamic optimizers still incur a positive cost c. Thus, at a limit set, $\overline{u}_t^{\delta} < \overline{u}_t^{\mu}$ and consequently all players that get the possibility to revise their type will choose to become myopic optimizers. This drives the fraction of dynamic optimizers in the population down when $\theta > 0$. With positive probability, the fraction of dynamic optimizers reaches zero this way. From the above, it is then clear that with probability 1 dynamic optimizers will be reintroduced in finite time. Hence, most of the time, the fraction of dynamics optimizers in the population will be small. ¤

Dynamic optimization is more costly than myopic optimization. Therefore, in a situation in which both types realize the same payoffs from the stage game, myopic optimizers have an advantage over the more sophisticated dynamic optimizers. In the presence of a positive costs for forming correct expectations, it is only out of equilibrium that dynamic optimization may yield an advantage over myopic optimization.

From the above, we have seen that without the presence of mutations, the dynamics take the system to a limit set, where it stays forever after. Which limit sets the system will be in depends crucially on the initial state $s(0)$. Thus the system without mutation is non-ergodic or, in other words, it displays path dependency. Adding a small probability of mutations transforms the system into an ergodic system, which has a unique invariant distribution, i.e. which does not display path dependency. Our main theoretical result characterizes the limit of this unique invariant distribution when the mutation rate is taken to zero in the limit. This limit invariant distribution puts all probability weight on a subset of the set of all limit sets. Thus this method selects between limit sets. A crucial role in the selection between limit sets is played by the depth of the basins of attraction of the different limit sets. The basin of attraction of a limit set l, denoted by $B(l)$ is the set of all states that will converge to l under the dynamics without mutations. The right (left) basin of attraction of a limit set l, denoted by $B^R(l)$ ($B^L(l)$) is the set of all states $\{s|\sigma(s) > \sigma(l)\}$ ($\{s|\sigma(s) < \sigma(l)\}$) that will converge to l under the dynamics without mutations. We can now define the depth of the basin of attraction at a state s.

Definition 3. *The depth of the basin of attraction at state s is $\min_{\hat{s}\in br(\sigma(s))} |\sigma(s) - \hat{s}|$, i.e. the absolute value of the minimum difference between the summary statistic $\sigma(s)$ and the best-reponse correspondence $br(\sigma(s))$ at s.*

We use this definition to define the depth of the basin of attraction of a limit set l.

Definition 4. *The depth $D(l)$ of the basin of attraction of limit set l is $\max_{s\in B(l)} \min_{\hat{s}\in br(\sigma(s))} |\sigma(s) - \hat{s}|$, i.e. the maximal depth over all states in the basin* *of attraction of limit set l. The depth $D^R(l)$ of the right basin of attraction of limit set l is* $\max_{\{s\in B(l)|\sigma(s)>\sigma(l)\}} \min_{\hat{s}\in br(\sigma(s))} |\sigma(s)-\hat{s}|$. *The depth $D^L(l)$ of the left basin of attraction of limit set l is* $\max_{\{s\in B(l)|\sigma(s)<\sigma(l)\}} \min_{\hat{s}\in br(\sigma(s))} |\sigma(s)-\hat{s}|$.

For completeness, we define $D^L(0) = D^R(M) = 0$.

Note that in a graphical representation of individual best-responses as a function of the summary statistic (see e.g. Figure 4.1, p. 32), the depth of the basin of attraction is precisely the distance between the best-response correspondence and the 45^0 line. The depth of the right (left) basin of attraction of a limit set l is a measure for the number of mutations necessary to reach the higher (lower) limit set adjacent to l.

We also define the depth of the (right and left) basin of attraction relative to the depths of other basins of attraction. Consider a set E of adjacent limit sets. We say that limit set $\tilde{l} \in E$ has the deepest (right or left) basin of attraction relative to (the other elements of) E when $\tilde{l}$ has the deepest (right or left) basin of attraction in the model restricted to the actions played under E.

Proposition 8 establishes the set of long run equilibria in the presence of mutations and states that, generically, in the long run the system spends most of the time in a unique NE when the probability of mutation is small. This will be the NE selected by an iterative procedure based on the relative depths of basins of attraction. Before stating the proposition, we first give a description of this iterative procedure.

We start the procedure by constructing the set of stable limit sets $E^0 = \{l | l \text{ is a stable limit set}\}$. A limit set l is stable when it has multiple elements in its basin of attraction $B(l)$. We now look for the (strictly) deepest right and left basin of attraction among all elements of E^0. Label these limit sets l^R and l^L. We now consider four distinct possibilities.

1. When $l^L \leq l^R$, we construct the set $E^1 = \left\{l \in E^0 | l^L \leq l \leq l^R\right\}$.

2. When $l^L > l^R$ and $D^L\left(l^L\right) > D^R\left(l^R\right)$, we construct the set $E^1 = \left\{l \in E^0 | l^L \leq l \leq M\right\}$.

3. When $l^L > l^R$ and $D^L\left(l^L\right) < D^R\left(l^R\right)$, we construct the set $E^1 = \left\{l \in E^0 | 0 \leq l \leq l^R\right\}$.

4. Finally, when $l^L > l^R$ and $D^L\left(l^L\right) = D^R\left(l^R\right)$, we construct the set $E^1 = E^{1,1} \cup E^{1,2}$, with $E^{1,1} = \left\{l \in E^0 | 0 \leq l \leq l^R\right\}$ and $E^{1,2} = \left\{l \in E^0 | l^L \leq l \leq M\right\}$.

Now we apply the same procedure sketched above relative to the set E^1. In case E^1 consists of only one component we find the left and right basins of attraction with maximal depth in E^1, and select a subset E^2 of E^1 in the way outlined above. When $E^1 = E^{1,1} \cup E^{1,2}$, we apply the procedure to the sets $E^{1,1}$ and $E^{1,2}$ separately and define the union of the resulting sets as E^2. We stop the iterative procedure when no further selection can be accomplished, i.e. when $E^{n+1} = E^n$, $n \in \mathbf{N}$, which is when a either the set E^n, $n \in \mathbf{N}$, is reduced to a singleton set or when the set consists of multiple limit sets with $D^R(\widetilde{l}) = D^L(l)$ for $\widetilde{l} = \max\left\{\bar{l} \in E^n | \bar{l} < l\right\}$, for all $l \in E^n \setminus \min\left\{\overset{*}{l} \in E^n\right\}$, $n \in \mathbf{N}$. In this case the basins of attractions have the same depth. The states which are in the set at which this iterative procedure terminates are called states with the iterative deepest right and left basins of attraction. We now state that these are precisely the long run equilibria of the model.

Proposition 8. *The set of long run equilibria consists of the states with the iterative relative deepest right and left basins of attraction. For a generic supermodular stage* *game, there will be only one long run equilibrium when the action set is going to approximate a continuous state space, i.e. when $\gamma \to \infty$ and subsequently the population size increases, i.e. $N \to \infty$.*

Proof.

Above, we have shown that the best-responses are monotonically increasing (Proposition 2), that polymorphic states cannot be stable (Proposition 4), and that the limit sets of the model without mutations correspond one-to-one with the set of pure strategy Nash equilibria. Thus, we have shown that all conditions for Theorems 4, 5, and 6 in KR are satisfied. Applying KR's Theorems 4, 5, and 6 leads to the conclusion that the iterative procedure sketched above yields the long run equilibria of the model, as follows.

Take an arbitrary element $e \in E^1$ and an arbitrary limit set $l \notin E^1$. Then, it takes more mutations to get from e to l than vice versa. Thus, when $\varepsilon \to 0$, the move from l to e is arbitrarily much more likely to be observed than the move from e to l. This establishes that when $\varepsilon \to 0$, a fraction 1 of the time is spent in (a subset of) E^1, i.e. that the set of long run equilibria is a subset of E^1. It is possible that at some elements of E^1, the system only spends a fraction 0 of the time, i.e. that some limit sets $e' \in E^1$ are not long run equilibria. This is the case when a subset $\widetilde{E}^1 \subset E^1$ is reached from $E^1 \backslash \widetilde{E}^1$ with a number of mutations that is strictly less than the number of mutations it takes to reach $E^1 \backslash \widetilde{E}^1$ from $\widetilde{E}^1$. A criterion for selecting such a subset is thus the depths of the right and left basin of attraction, which results in the observation that $\widetilde{E}^1 = E^2$.

In the iterative procedure the focus is on the set E^n in which the system will spend a fraction 1 of the long run time. At each step, the procedure establishes which of the states in E^n will be left more easily than others, thus establishing a strict subset $E^{n+1} \subset E^n$ on which the system will spend a fraction 1 of the long run time. When the procedure yields no further selection, i.e. when $E^{n+1} = E^n$, the system spends a strictly positive fraction of long run time in each element of E^n. In other words, E^n is the set of long run equilibria.

The second statement in the proposition is straightforward: Consider a supermodular stage game with randomly drawn payoffs from $\mathbf{R}$. When we take $\gamma \to \infty$, the continuity of best-responses sees to it that the (piecewise constant) best-response correspondence approximates a smooth curve arbitrarily close. Since the grid of the state space gets ever finer with increasing N, i.e. ever more mixed strategy sets can be mimicked by states $\widetilde{s} \in S$, the probability that two NE states s and s' are such that it takes the same number of mutations to reach s from s', as it takes to reach s' from s is decreasing in the population size, when the best-reply correspondence is (arbitrarily close to) a smooth curve. ¤

In a model where only two stable equilibria are present, a direct consequence of this proposition is that the equilibrium with the deepest basin of attraction is selected as the unique long run equilibrium. A result similar to Proposition 8 holds for a setting in which all players in the population form (almost-) rational expectations.

Corollary 1. *Consider the model with only dynamic optimizers present and no birth & death. Then the set of long run equilibria consists of the states with the iterative relative deepest basin of attraction. For a generic supermodular stage game, there will be only one long run equilibrium when the action set is going to approximate a continuous state* *space, i.e. when* $\gamma \longrightarrow \infty$ *and subsequently the population size increases, i.e.* $N \rightarrow \infty$.

Proof.

The proof of Proposition 8 does not rely on a mixture of types of players in the population. Therefore, all statements follow also for a population consisting only of dynamic optimizers. ¤

4. Examples

In this section we consider three specific examples of games exhibiting macroeconomic complementarities. For simplicity we present a framework with a continuous action space, i.e. we have taken the limit $\gamma \longrightarrow \infty$ in the theoretical model above. Thus the action space is embedded into a continuous action space $\widetilde{\mathcal{M}} := [0, M]$.

The specific functional form we choose for the complementarities is not important. The only requirement is that the functional form generates best-responses that are monotonically increasing in the summary statistic of the actions chosen by the individual players. In all of the examples we pose the following general functional form of the complementarities

$$\eta\left(\sigma\left(s\right)\right) = \frac{\exp(\sigma\left(s\right))}{z + \frac{1}{10}\exp(\sigma\left(s\right))}.$$

In the examples, this function yields the best-response structure of a supermodular game and it is sufficiently simple to illustrate the possible results. The parameter $z \in [0, \infty)$ indicates the strength of the complementarity at the macro level. Note that the lower the value of z is, the stronger is the complementarity. The factor $\frac{1}{10}$ is a scaling parameter which is of no importance for the results.

4.1. A Model of Effort with Production Externalities

Consider a specific economy where the macroeconomic complementarity is present through the production process (see e.g. Cooper and Haltiwanger (1996)). Furthermore, assume that the population size N is large, ensuring that an individual's action does not influence the summary statistic. Specifically, at each time $t = 1, 2, ...$ each agent n, which might be either of type μ or δ, has effort $e_t^n \in \widetilde{\mathcal{M}}$ as input and produces output $y_t^n = e_t^n \cdot \nu\left(\overline{E}_t\right)$, with $\overline{E}_t = \frac{1}{N}\sum_{n=1}^{N} e_t^n$. Thus, the output one unit of effort creates, depends on the average production in the economy through a function $\nu(\cdot)$ of the average effort. We assume $\nu(\cdot)$ to be increasing. The agents each consume their own production. The technology is supplemented with the preferences of an agent over consumption and effort given by

$$\begin{aligned} U^{\mu}(y_t^n, e_t^n) &= y_t^n - \frac{1}{2}(e_t^n)^2 = e_t^n \cdot \nu\left(\overline{E}_t\right) - \frac{1}{2}(e_t^n)^2 \equiv \overline{U}^{\mu}(e_t^n, \overline{E}_t) \\ U^{\delta}(y_t^n, e_t^n) &= y_t^n - \frac{1}{2}(e_t^n)^2 - c = e_t^n \cdot \nu\left(\overline{E}_t\right) - \frac{1}{2}(e_t^n)^2 - c \equiv \overline{U}^{\delta}(e_t^n, \overline{E}_t), \end{aligned}$$

where $c \geq 0$ is the per-period calculation-cost dynamic optimizers incur. Note that this game is supermodular since $\frac{\partial^2 \overline{U}}{\partial e_t^n \partial \overline{E}_t} = \nu'\left(\overline{E}_t\right) > 0$.

At time t myopic optimizers have observed $\overline{E}_{t-1}$. Their best guess is that the summary statistic will remain unchanged between time $t-1$ and t, .i.e. $\mathbf{E}\overline{E}_t = \overline{E}_{t-1}$, where $\mathbf{E}$ is the expectaion operator. They play a best-reply to these expectations. Maximizing w.r.t. e_t^n yields $e_t^n = \nu\left(\mathbf{E}\overline{E}_t\right) = \nu\left(\overline{E}_{t-1}\right)$. Thus, since all myopic optimizers update their action each period, they all play $e_t^n = \nu\left(\overline{E}_{t-1}\right)$.

Dynamic optimizers correctly predict what the myopic optimizers will do, as they form almost-rational expectations, which take this into account. Thus, dynamic optimizers to determine their effort level e_t^n, they solve the following problem

$$\widetilde{\overline{E}}_t = \frac{1}{N}\sum_{n=1}^{N} e_t^n \left(1 - I_n^{t+1}\right) + \frac{1}{N}\sum_{n=1}^{N} e_t^n I_n^{t+1} = \frac{N_{t+1}^{\mu}}{N} \cdot \nu\left(\overline{E}_{t-1}\right) + \frac{N - N_{t+1}^{\mu}}{N} \cdot \nu\left(\widetilde{\overline{E}}_t\right).$$

w.r.t. $\widetilde{\overline{E}}_t$. [12] Note that I_n^{t+1} is just the indicator function of player n being of type δ after type switching and birth & death took place at time t, since the type of a player will not change between birth & death at time t and the beginning of period $t+1$. This explains the use of N_{t+1}^{μ} in the equation. With $\widetilde{\overline{E}}_t$ determined by solving the above equation, the dynamic optimizers play $\nu\left(\widetilde{\overline{E}}_t\right)$.

Whenever $\nu\left(\widetilde{\overline{E}}_t\right) \neq \nu\left(\overline{E}_{t-1}\right)$ the dynamic optimizers obviously get a payoff from the stage game that is at least as high as the myopic optimizers' payoff, since the dynamic optimizers play a best response to the actual state $\widetilde{\overline{E}}_t$. It depends on the payoffs and on the cost $c \geq 0$ dynamic optimizers incur, whether they actually have a higher net utility $U^{\delta}\left(y_t^n, e_t^n\right)$ (that is utility of consumption and effort minus costs).

To illustrate the results, we assume complementarities in production of the form

$$\nu\left(\overline{E}_t\right) = \frac{\exp(\overline{E}_t)}{z + \frac{1}{10}\exp(\overline{E}_t)},$$

where $z \in [0, \infty)$ is a parameter that indicates the strength of the externality.

The next figure shows how a typical individual best response function depends on

[12]Note that $\phi(\overline{E}_{t-1})$ is known at time t. The problem can be solved numerically.

the average effort in the economy.

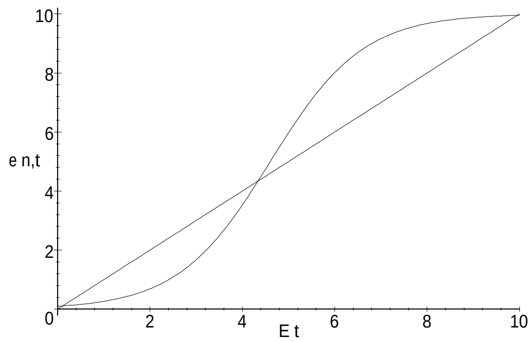

Figure 4.1: The best-response function for $z = 10$.

This economy can exhibit either one or two stable equilibria[13], depending on the strength of the production externalities, i.e. on the parameter $z \in [0, \infty)$. To analyze the relationship between z and the intersection point of the curve $e = \nu\left(\overline{E}\right)$ and the line $e = \overline{E}$ (i.e. the equilibria in the economy) we calculate

$$\frac{\exp(\overline{E})}{z + (\frac{1}{10})\exp(\overline{E})} = \overline{E}.$$

Solving this equation (numerically) yields that for $z \in [0, 2.5]$, the economy exhibits a unique Nash equilibrium (near) 10; for $z > 91$ a unique Nash equilibrium (near) 0 is established. For values of the externality-parameter z between 2.5 and 91 the economy exhibits multiple stable equilibria, one near 10 and one near 0. The Nash equilibrium where all agents choose to put 10 units of effort into the production process, is the Pareto efficient equilibrium. When the economy exhibits multiple equilibria, i.e. when $2.5 < z < 91$, we know from Proposition 8 that the equilibrium state with the deepest

[13]This economy also exhibits one unstable equilibrium. We do not focus on the unstable equilibria, since, per definition, they can be upset by one mutation.

basin of attraction will be the long run equilibrium of the model. We now calculate the depth of the basins of attraction as a function of the externality-parameter z. The depth of the basin of attraction at point $\overline{E}$ is defined by the absolute value

$$g_z\left(\overline{E}\right) = \left|\overline{E} - \frac{\exp(\overline{E})}{z + \frac{1}{10}\exp(\overline{E})}\right|.$$

We solve for the (arg)maxima of this expression. This yields

$$\overline{E}_1 = \ln 10\left(4 + \sqrt{15}\right)z = \ln 10\left(4 + \sqrt{15}\right) + \ln z \approx 4.366 + \ln z,$$

$$\overline{E}_2 = \ln 10\left(4 - \sqrt{15}\right)z = \ln 10\left(4 - \sqrt{15}\right) + \ln z \approx .23915 + \ln z.$$

The corresponding maximal values of $g_z(.)$ are

$$g_z\left(\overline{E}_1\right) = g_z\left(\ln 10\left(4 + \sqrt{15}\right)z\right) \approx |\ln z - 4.507|$$

$$g_z\left(\overline{E}_2\right) = g_z\left(\ln 10\left(4 - \sqrt{15}\right)z\right) \approx |\ln z - .88787|$$

Note that for all $z \in \mathbf{R}_+$, it is the case that $\overline{E}_1 > \overline{E}_2$ and thus that $g_z\left(\overline{E}_1\right) \approx |\ln z - 4.507| := h(z)$ is the depth of the basin of attraction of the equilibrium near 10 and $g_z\left(\overline{E}_2\right) \approx |\ln z - .88787| := k(z)$ is the depth of the basin of the equilibrium near 0. To see which equilibria is selected for a given externality-parameter z, we plot the depth-functions $k(z)$ (thick) and $h(z)$

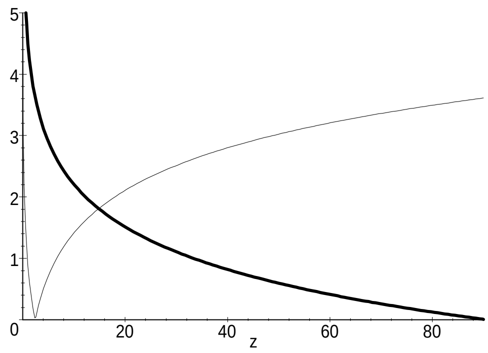

Figure 4.2: The depths of the basins of attraction as functions of z.

We conclude that given a multiplicity problem, for strong externalities ($z < 14.84$), the Pareto-efficient equilibrium is the long run equilibrium, and for weak externalities ($z > 14.84$), the Pareto-inferior equilibrium is selected in the long run.

We end this example by drawing the best response function for two different values of the externality parameter z. The first figure shows an economy with a strong externality ($z = 8$). Hence the Pareto-efficient equilibrium is the long run equilibrium, while in the second figure, the externality is weak ($z = 30$), and the Pareto-inferior equilibrium is selected.

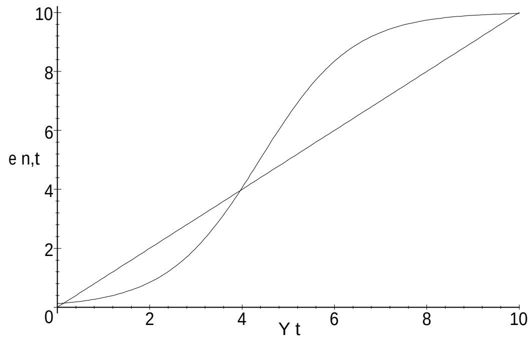

The best-response function for $z = 8$.

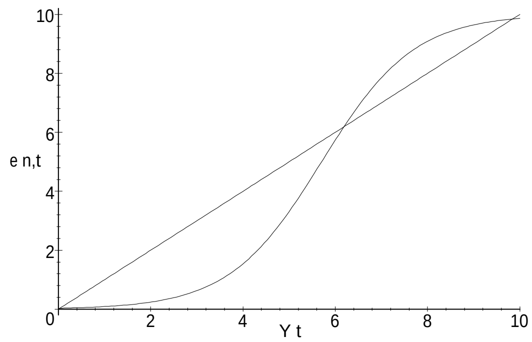

The best-response function for $z = 30$.

4.2. Bertrand Competition with Demand Externalities

Consider an economy in which a large number N of producers, labelled $n = 1, \ldots, N$, compete through Bertrand competition. Each producer n produces a unique good n. Demand $D_n(.,.)$ for good n at producer n is a function of the price p_n of good n and an index of prices P of all goods on the market, i.e.

$$D_n(p_n, P) = a - b_n p_n + \nu(P),$$

with $a > 0$ and $b_i > 0$. From this demand function we see that the goods are substitutes. The degree of substitutability is determined by the form of the (non-negative, increasing) demand externality function $\nu(P)$.

We assume that agents have no fixed costs and that they derive utility from profit only, thus

$$U_n(p_n, P) = \pi_n(p_n, P) = (p_n - k_n) D_n(p_n, P),$$

where $k_n > 0$ is the marginal cost of production for player n. This game is supermodular, since $\frac{\partial^2 U_n}{\partial p_n \partial P} = \nu'(P) > 0$. Furthermore, competition of the individual agent is against a summary statistic of the actions of all agents in the population. Thus, the example fits the framework of our model.

In calculating best-responses, each agent takes P as given. The first order condition yields

$$p_n = \frac{1}{2b_n}\nu(P) + \frac{a}{2b_n} + \frac{1}{2}k_n.$$

Thus,we have best-responses of the form $p_n(P) = C_1 + C_2\nu(P)$, with $C_1 > 0$ and $C_2 > 0$ constants. When the demand externalities are of the form

$$\nu(P) = \frac{\exp(P)}{z + \frac{1}{10}\exp(P)},$$

with z a parameter indicating the degree of substitutability between the different products, the best-response function $p_n(P)$ is just a linear transformation of the best-response function $e^n = \nu\left(\overline{E}\right)$ in section 4.1. Since such a transformation does not affect the structure of the model, the analysis of the former section applies to this example.

4.3. A Model of Search and Matching

Consider a simple search model in the spirit of Diamond (1982) (see also e.g. Fudenberg and Tirole (1991)) with a large number N of players, labelled $n = 1, \ldots, N$. Denote the intensity of search of agent n by $s_n \geq 0$. The utility of an agent is determined by the probability of finding a trading partner, i.e. being matched to some other agent, by the gain from trade $\alpha > 0$ when a partner is found and by the costs $k(s_n)$ of searching at intensity s_n. The probability for agent n of finding a trading partner depends on his own search intensity and an increasing function of the average search intensity in the market $\nu(S)$, with $S = \frac{1}{N}\sum_{n=1}^{N} s_n$. Thus, the utility function looks like

$$U_n(s_n, S) = \alpha s_n \nu(S) - k(s_n).$$

This example exhibits a supermodular game, since $\frac{\partial U_n}{\partial s_n \partial S} = \alpha\nu'(S) > 0$. Furthermore, each agent competes against a summary statistic of the actions of all agents in the entire population. Thus, this example fits the framework of our model.

In calculating best-responses, each agent takes S as given. The first order condition yields $\alpha\nu(S) - k'(s_n) = 0$. When we further assume costs to be quadratic, i.e. $k(s_n) = \beta s_n^2$, with $\beta > 0$, we get that $s_n(S) = \frac{\alpha}{2\beta}\nu(S)$. When we take $\nu(S) = \frac{\exp(S)}{z + \frac{1}{10}\exp(S)}$, the best-response function $s_n(S)$ is again a linear transformation of the best-response function in section 4.1. Thus, the previous analysis applies.

5. Concluding Remarks

The main points of this paper have been to show how the stochastic evolutionary approach of equilibrium selection applies to macroeconomic models of coordination failure, and how the equilibrium which is singled out by the dynamic is directly related to the underlying externality that creates the multiplicity problem in the underlying stage game. It is important to note that the unique equilibrium selection result is obtained in the limit as the probability of mutations goes to zero. We believe however, that this is an analytical construction which is just a benchmark case for what happens in the real economy, where the mutation rate will not be zero. Therefore, the uniquely selected equilibrium is better interpreted as the equilibrium that will be observed the largest fraction of the time. The economy can also remain at other equilibria for a non-negligible fraction of time. One might interpret the movement of the system between several equilibria as a rather crude form of business cycle.

To obtain these results, we have expanded the evolutionary literature on equilibrium selection in two directions. First, we abandon the random pairing assumption and analyze equilibrium selection in a game where players interact in a market structure. Secondly, we allow players to posses different degrees of sophistication and provide an evolutionary framework in which selection among the different types take place in the class of strict supermodular games. We believe that our result that in equilibrium (selection works against the more sophisticated agents, when sophistication comes with a cost) provides a possible rationale for modelling not only rational agents, but also boundedly rational agents in macroeconomic models.

References

Blanchard, O. and N. Kiyotaki (1987). Monopolistic competition and the effects of aggregate demand. *American Economic Review* 77, 647–666.

Blume, L. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–424.

Bryant, J. (1983). A simple rational expecations Keynes-type model. *Quarterly Journal of Economics* 97, 525–529.

Conlisk, J. (1980). Costly optimizers versus cheap imitators. *Journal of Economic Behavior and Organization* 1, 275–293.

Cooper, R. (1994). Equilibrium selection in imperfectly competitive economies with multiple equilibria. *Economic Journal* 104, 1106–1122.

Cooper, R. and J. Haltiwanger (1996). Evidence on macroeconomic complementarities. *Review of Economics and Statistics*, 78–93.

Cooper, R. and A. John (1988). Coordinating coordination failures in Keynesian models. *Quarterly Journal of Economics* 103, 441–463.

Crawford, V. (1991). An "evolutionary" interpretation of Van Huyck, Battalio and Beil's experimental results of coordination. *Games and Economic Behavior* 3, 25–59.

Crawford, V. (1997). *Theory and Experiments in the Analysis of Strategic Interaction*, Volume 1, pp. 206–242. New York: Cambridge University Press.

Diamond, P. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy 90*, 881–894.

Droste, E., C. Hommes, and J. Tuinstra (1999). Endogenous fluctuations under evolutionary pressure in Cournot competition. Working paper, CeNDEF, University of Amsterdam.

Droste, E. and J. Tuinstra (1998). Evolutionary selection of behavioral rules in a Cournot model: a local bifurcation analysis. Discussion paper no. 9886, CentER, Tilburg University.

Freidlin, M. and A. Wentzell (1984). *Random Perturbations of Dynamical Systems*. New York: Springer Verlag.

Fudenberg, D. and J. Tirole (1991). *Game Theory*. Cambridge, Massachusetts: The MIT Press.

Hansen, P. and O. Kaarbøe (1997). On equilibrium selection mechanisms in coordination games with simultaneous play. Working paper, no. 11, Department of Economics, University of Bergen, Bergen, Norway.

Hart, O. (1982). A model of imperfect competition with Keynesian features. *Quarterly Journal of Economics 97*, 109–138.

Kaarbøe, O. (1999). Equilibrium selection in supermodular games with simultaneous play. Mimeo, Dept. of Economics, University of Bergen, Bergen, Norway.

Kaarbøe, O. and A. Tieman (1999). Equilibrium selection under different learning modes in supermodular games. Discussion paper, nr. TI 99-061/1, Tinbergen Institute, Amsterdam, The Netherlands.

Kandori, M., G. Mailath, and R. Rob (1993). Learning, mutation and long run equilibria in games. Econometrica 61, 29–56.

Kandori, M. and R. Rob (1995). Evolution of equilibria in the long run: A general theory and applications. Journal of Economic Theory 65, 383–414.

Krishna, V. (1991). Learning in games with strategic complementarities. Mimeo, Harvard Business School.

Maynard Smith, J. (1982). Evolution and the Theory of Games. Cambridge: Cambridge university Press.

Milgrom, P. and J. Roberts (1990). Rationalizability, learning, and equilibrium in games with strategic complementaries. Econometrica 58, 1255–1277.

Robles, J. (1997). Evolution and long run equilibria in coordination games with summary statistic payoff technologies. Journal of Economic Theory 75, 180–193.

Schelling, T. (1973). Hockey helmets, concealed weapons, and daylight saving: A study of binary choices with externalities. Journal of Conflict Resolution 17, 381–428.

Topkis, D. (1979). Equilibrium points in nonzero-sum n-person submodular games. SIAM Journal of Control and Optimization 17, 773–787.

Van Damme, E. and J. Weibull (1998). Evolution with mutations driven by control costs. Discussion paper, nr. 9894, CentER, Tilburg University.

Vives, X. (1990). Nash equilibrium with strategic complementarities. Journal of Mathematical Economics 19, 305–321.

Weitzman, M. (1982). Increasing returns and the foundation of unemployment theory. *Economic Journal* 92, 787–804.

Young, H. (1993). The evolution of conventions. *Econometrica* 61, 57–84.