

THE UNIVERSITY of York

Discussion Papers in Economics

No. 1999/37

Coordination and Equilibrium Selection in Mean Defined Supermodular Games Under Payoff Monotonic Selection Dynamics

by

Huw David Dixon and Ernesto Somma

Department of Economics and Related Studies University of York Heslington York, YO10 5DD

Coordination and equilibrium selection in mean-defined supermodular games under payoff monotonic selection dynamics

Huw D. Dixon*

Ernesto Somma

Department of Economics
and Related Studies,
University of York,
York YO1 5DD UK
Tel: (44) 1904 433788
${ m email:hdd1@york.ac.uk}$

Dipartimento di Scienze Economiche Università degli Studi di Bari Via C. Rosalba 70124 Bari Tel (39) 80 5049100 email: e.somma@dse.uniba.it

December 1999

Abstract

This paper introduces the class of mean defined supermodular games. The characteristic feature of these games is that, given an order on the strategy sets of the players, the payoff to each player depends on his own strategy and the average of the population play. We characterise the set of the Nash equilibria and analyse their dynamic properties under payoff monotonic selection dynamics. Weak Nash equilibria, both in pure and mixed strategies, are unstable. The only asymptotically stable equilibria of the game are symmetric strict equilibria where each player uses the same strategy. We show that the strategies that do not survive the process of iterated deletion of strictly dominated strategies vanish in the long run. As a corollary to this latter result, we show that if the game is dominance solvable then the dynamics converges from any initial interior state.

Keywords: Strategic complementarities, supermodular games, bounded rationality, replicator dynamics, coordination failures.

^{*}We would like to thank the ESRC for their financial support (Grant number R000 23 6179) and the EU TMR program.

1 Introduction

In recent years strategic complementarities have received growing attention in the economic as well the game theoretic literature. The definition of strategic complementarity is due to Bulow et al. (1985) and identifies situations in which for each player the marginal benefit to an increase in his action is increasing in the level of other players' actions. A number of interactions among individuals or firms share this characteristic.

In macroeconomics, the vast literature on *coordination failures* (Diamond, 1982; Howitt, 1985; Kiyotaki, 1988; Cooper and John, 1988, among the others) focuses on the presence of strategic complementarities to explain the inability of macroeconomic systems to achieve the Pareto-dominant equilibrium, in the presence of multiple equilibria. The complementarity between the actions independently undertaken by the agents is not accounted for and the system can get stuck at an inefficient equilibrium which is Paretodominated by another equilibrium. More recently Matsusaka and Sbordone (1995) found evidence that, in the presence of strategic complementarities, if consumers are pessimistic about the state of the economy, there can be a slowdown in output even if their beliefs are not based on economic fundamentals. Redding (1996) explores the macroeconomic consequences of strategic complementarities between investment in human capital and R&D. He shows that for certain parametrisations multiple Pareto-rankable equilibria may occur. In this case a governmental policy coordinating agents expectations may be welfare improving. Support for macroeconomic models with complementarities is further provided by Cooper and Haltinwanger (1996), whereas the role of strategic complementarities in shaping the business cycle is discussed in Cooper (1998). In the IO literature Bertrand oligopoly with differentiated products exhibits strategic complementarities: when competitors increase their prices, the marginal profits from the own price increase rise. Network externalities can be modelled as strategic complementarities, as more firms, for example, adopt the same technological standard or use the same telecommunication network, the marginal benefit to other firms doing the same increases (Katz and Shapiro, 1986; Farrell and Saloner, 1986). The relevance of strategic complementarities in modern manufacturing and retail firms is emphasised in Bagwell and Ramey (1994); Milgrom and Roberts (1990a, 1995). Matsuyama (1995) presents a review of models of monopolistic competition used in macroeconomics, international economics as well as growth and development theory with different kinds of complementarities that can generate multiplier effects, business cycles, underdevelopment traps etc.

The crucial characteristic of the models with strategic complementari-

ties discussed above is the possible presence of Pareto-rankable equilibria. The existence of Pareto-rankable equilibria, especially in macroeconomic or general equilibrium models, generates mixed reactions. According to many authors this feature is what makes such models interesting creating the possibility for coordination failures. According to others coordination failures will rarely occur in practice. Because of the large gains to be made from coordination, the agents will eventually find ways of coordinating their actions. Moreover, according to these critics, coordination failures are logically inconsistent with equilibrium analysis. The validity of equilibrium analysis rests on the assumption that agents are able to coordinate, in some unspecified way, their expectations and actions so that a specific equilibrium is played. If agents are able to coordinate their expectations, they should be able to coordinate a change in their expectation so that a Pareto-dominant equilibrium will be played.

One obvious answer to the first objection is that in many cases the number of individuals involved and the number of activities to be coordinated is so large that even in the long run coordination may fail to be achieved notwithstanding the potential gains. Turning to the second objection, as Matsuyama(1995,p. 724) argues:

One possible response is that coordinating expectations is much easier than coordinating changes in expectations. The former can be achieved historically through conventions, customs, cultural beliefs, ideologies, or other processes of learning..... These arguments, of course, have to be represented in an explicit dynamic setting..

In this paper we present a family of games encompassing most of the models with complementarities referred above and we analyse the issues of coordination, equilibrium selection and stability of the equilibria within an evolutionary framework using a general class of selection dynamics known as payoff monotonic dynamics.

Strategic complementarity finds a more general and formal mathematical representation in terms of supermodularity of the objective function and restrictions on the structure of the action space. All the models referred above can infact be represented as supermodular games. These were introduced by Topkis (1979) and later developed by Vives (1990); Milgrom and Roberts (1990b); Milgrom and Shannon (1994). We introduce a family of such games that we call *mean-defined supermodular games*. The characteristic feature of this class of games is that the payoff to each player depends on his own strategy and the average of the population play. This feature restricts the realm of application of mean-defined supermodular games to the cases where an average strategy is meaningful. This restriction however is not so severe; supermodular games have in fact (partially) ordered strategy sets which requires a quantitative dimension that makes the computation of average strategy possible. In addition, the type of interaction we are interested in often involves large numbers of players interacting simultaneously and the assumption that the payoffs are determined by the player's strategy and the average of the population play makes the dynamics of the model analytically tractable.

The paper is organised as follows. Section 2 summarises definitions, properties and results of supermodular games that are relevant for the subsequent analysis. In section 3 the class of mean-defined supermodular games is introduced and the set of the Nash equilibria is characterised. As a first result we show that the property of monotonicity of best replies applies to this class of games. Furthermore, we show that the greatest and least pure strategy Nash equilibria are also the greatest and the least element of the set of pure strategies that survive the process of iterated deletion of strictly dominated strategies. This same result is presented in Milgrom and Roberts (1990b) but needs to be proved for this class of games as the payoff to the players depends on a statistic of the strategy profile rather than the profile itself. Section 4 studies the issues of coordination, equilibrium selection and dynamical properties of the equilibria for this class of games under payoff monotonic selection dynamics. The main conclusion is that weak Nash equilibria, both in pure and mixed strategies, are unstable. The only asymptotically stable equilibria of the game are symmetric strict equilibria where each player uses the same strategy. Finally, we extend the results of Samuelson and Zhang (1992) concerning the long run survival of serially dominated strategies. We show that the strategies that do not survive the process of iterated deletion of strictly dominated strategies vanish in the long run. As a corollary to this latter result, we show that if the game is dominance solvable then the dynamics converges from any initial interior state. Section 5 presents an application of the analysis carried out in the previous sections to a model of decentralised trading resembling the "coconut economy" of Diamond (1982) using a specific type of payoff monotone dynamics known as replicator dynamics.

2 Supermodular games

Supermodular games are games in which, given a partial order on the strategy set of each player, the marginal returns from an increase in one's strategy are increasing in the strategy played by the rivals. In many applications, where the strategy space is one-dimensional this is simply the result of strategic complementarity between players' strategies. More generally, if the strategy sets of players are multidimensional, supermodularity is an assumption of complementarity among the components of each player's strategy that ensure that these components move in the same direction when the rivals' strategies change.

Let K be the finite or infinite set of players. The strategy set for each player $k \in K$ is denoted as I_k and it is assumed to be completely ordered with generic element *i*. If the strategy set is finite, the cardinality of I_k is denoted by m_k . The strategy space is $\Sigma = \times_{k \in K} I_k$ whose elements $\mathbf{s} = (i_k, i_{-k}) \in \Sigma$ define a strategy profile; the payoff function is given by $\Pi : \Sigma \mapsto \Re$.

Definition 2.1. A game G is supermodular (strictly supermodular) if the following conditions are satisfied for each $k \in K$:

- (C1) I_k is a compact subset of \Re ;
- (C2) Π is upper semi-continuous in i_k (for fixed i_{-k});
- (C3) Π has increasing (strictly increasing) differences ¹, i.e. $\Pi(i_k, i_{-k}) \Pi(i'_k, i_{-k}) \geq \Pi(i_k, \hat{i}_{-k}) \Pi(i'_k, \hat{i}_{-k})$ for all $(i_k, i'_k) \in I_k$ and $(i_{-k}, \hat{i}_{-k}) \in I_{-k}$ such that $i_k > i'_k$ and $i_{-k} > \hat{i}_{-k}$

We study games with strategy sets in \Re ; in more general cases, condition (C1) requires that the strategy set for each player is a complete lattice. The real line is a lattice and every compact subset of it is a complete lattice.

As it is well known, (in strictly) supermodular games:

- (i) No asymmetric NE in pure strategies exist;
- (ii) If *i* is a best response to *i'* and *î* is a best response to *i*, where i > i', then $i \ge i$;
- (iii) The set of pure-strategy Nash equilibria is non empty and possesses greatest and least equilibria $\underline{s}, \overline{s};$
- (iv) For a generically selected supermodular game all pure strategy Nash equilibria are strict;

¹In general, supermodularity is a property stronger than (C3), however, the two concepts coincides when the strategy set in singledimensional and/or the function Π is defined on a product of ordered sets.

(v) Let α and α' be two probability distributions defined over the strategy set I_k Assume $\alpha \succ \alpha'$, where \succ refers to first-order stochastic dominance. Then we have:

 $\min BR\left(\alpha'\right) \ge \max BR(\alpha).$

Properties (i) and (ii) are standard results, (iv) and (v) are due to Kandori and Rob (1995, Prop. 6 and 7 respectively). Property (v) extends the monotonicity property (ii) to comparison across mixed strategies when these are partially ordered according to *first-order stochastic-dominance*.

Property (v) is based on the fact that the expected value of any increasing function under probability distribution α' is no smaller than its expected value under a stochastically dominated probability distribution α .

3 Mean-defined supermodular games

The objects of our study are symmetric N-person games in "ordered normal form" in which the payoff to a player is determined by his strategy and the average of the population play. K denotes the set of players, $\Sigma = \times_{k \in K} I_k$ is the strategy space. The strategy sets can be continuous or discrete, we assume that they have least upper bound and greatest lower bound in the set. Given a strategy profile **s**, the average strategy is simply given by $\mu(\mathbf{s}) = \sum_{k \in K} i_k / |K|$. Finally the payoff function maps a player's strategy and the current value of the statistic to his payoff: $\Pi : \{i_k, \mu(\mathbf{s})\} \mapsto \Re$ and it is assumed to satisfy condition (C2) and strictly (C3). The game , = $\{K, \Sigma, \mu(\mathbf{s})\}$ is then the mathematical object we call mean-defined supermodular game. Due to the particular features of the game properties (i) to (v) need not to hold. Some of them will be holding in general whereas others will depend on additional assumptions about the payoff function.

Proposition 3.1 (Monotonicity of best replies). Let , = $\{K, \Sigma, \mu(\mathbf{s})\}$ be a mean-defined supermodular game. If *i* is a best response to μ and $\tilde{\imath}$ is a best response to $\tilde{\mu}$, where $\tilde{\mu} > \mu$, then $\tilde{\imath} > i$.

Proof. Assume the contrary; i.e. let $i > \tilde{i}$. By construction $\Pi(i, \mu) - \Pi(\tilde{i}, \mu) \ge 0$. Strictly increasing differences will then imply that $\Pi(i, \tilde{\mu}) - \Pi(\tilde{i}, \tilde{\mu}) > 0$. A contradiction.

Next we show that the set of pure strategy Nash equilibria is non empty. The largest and smallest elements of the set are the largest and smallest strategies surviving the iterative elimination of strictly dominated strategies. This is a standard property of N-player supermodular games, (see Milgrom and Roberts, 1990b). This need to be proved here as the payoff to player k depends on a summary statistic of the strategy profile **s** rather than the profile itself. Let $H \subseteq I_k$ be any subset of the entire strategy set of player k. If all players are restricted to H then $\mu \in [\underline{i}^H, \overline{\imath}^H]$ the smallest and largest elements of H. Let $R^H = [\underline{\mu}, \overline{\mu}]$ define the possible range of μ given H. A pure strategy i for player k is said to be strictly dominated by another pure strategy $\hat{\imath}$ if it is the case that for all $\mu \quad \Pi(i, \mu) < \Pi(\hat{\imath}, \mu)$. Given any subset H of I_k , we define the set of undominated responses to R^H by:

$$U(R^{H}) = \{ i \in I_{k} \mid (\forall \tilde{i} \in I_{k}) (\exists \mu \in R^{H}) \mid \Pi(i,\mu) \ge \Pi(\tilde{i},\mu) \}$$

Let $\overline{U}(\mathbb{R}^H)$ define the interval $[\inf(U(\mathbb{R}^H)), \sup(U(\mathbb{R}^H))]$ We use U to represent the process of iterative elimination of strictly dominated strategies as follows. Define $I_k^0 = I_k$ the full strategy set for player k, for $\tau \geq 1$ define $I_k^{\tau} = U^{\tau-1}(I_k)$. A strategy i is serially undominated if $i \in U(I_k^{\tau})$ for all τ . U is a monotone nondecreasing function, that is if $H' \subset H$ then $U(H') \subset U(H)$.

Proposition 3.2. Let , = $\{K, \Sigma, \mu(\mathbf{s})\}$ be a mean-defined supermodular game. The set of pure strategy Nash equilibria is non-empty and it contains the largest and smallest serially undominated strategies $\overline{\imath}$ and \underline{i} .

Proof. First we need to show that the smallest and the largest best replies to $\underline{\mu}$ and $\overline{\mu}$ respectively are the smallest and largest undominated responses to R^H . Consider any subset H of the strategy set of player k. The average $\mu \in [\underline{\mu}^H, \overline{\mu}^H]$. Let $\beta(\underline{\mu}^H)$ be the set of best replies to the smallest element in R^H and similarly for $\beta(\overline{\mu}^H)$. All strategies $\hat{\iota} < \underline{\beta}(\underline{\mu}^H)$, where $\underline{\beta}(\underline{\mu}^H)$ is the smallest element in the set of best replies, are strictly dominated by this latter; infact for each player k we know that increasing differences imply that for any $H \subseteq I_k$

$$\Pi(\underline{\beta}(\underline{\mu}^{H}),\underline{\mu}^{H}) - \Pi(\hat{\imath},\underline{\mu}^{H}) < \Pi(\underline{\beta}(\underline{\mu}^{H}),\mu) - \Pi(\hat{\imath},\mu)$$

for all $\hat{i} < \underline{\beta}(\underline{\mu}^H)$ and $\mu > \underline{\mu}^H$. The left hand side of the inequality is positive by construction and so will be the right hand side for all μ . Similarly all strategies $\tilde{i} > \overline{\beta}(\overline{\mu}^H)$ are strictly dominated by this latter. We conclude that $\overline{U}(R^H) = [\underline{\beta}(\underline{\mu}^H), \overline{\beta}(\overline{\mu}^H)]$. Define the largest and smallest elements in the strategy set I_k^0 as $\bar{\imath}^0$ and \underline{i}^0 . For $r \ge 1$ let $\bar{\imath}^r = \overline{\beta}(\bar{\imath}^{r-1})$ and $\underline{i}^r = \underline{\beta}(\underline{i}^{r-1})$. We now show that $U^r(I_k) \subset [\underline{i}^r, \bar{\imath}^r]$. This is true for r = 0; assume it is true for $r \le k$. Then

$$U^{k+1}(I_k) \subset U([\underline{i}^k, \overline{\imath}^k]) \subset [\underline{i}^{k+1}, \overline{\imath}^{k+1}]$$

the first inclusion follows from the observation made earlier that U in nondecreasing, and the second follows from the first part of the proof. It follows that $\{\underline{i}^r\}$ is nondecreasing and $\{\overline{\imath}^r\}$ is nonincreasing. These sequences have limits \underline{i} and $\overline{\imath}$ respectively. Finally we show that these are best replies to themselves. We have proved that for all $r \ge 0$, $\underline{\beta}(\underline{\mu}^r) \ge \underline{\mu}^r$ and $\overline{\beta}(\overline{\mu}^r) \le \overline{\mu}^r$. Suppose \underline{i} is not a best reply to itself, then there exists some $i \in [\underline{i}, \overline{\imath}]$ such that

$$\Pi(i,\underline{i}) - \Pi(\underline{i},\underline{i}) > 0$$

Increasing differences imply that

$$\Pi(i,\hat{i}) - \Pi(i,\hat{i}) > 0 \ \forall \hat{i}$$

and hence \underline{i} is strictly dominated by i, which contradicts \underline{i} being the limit of the sequence. The same argument works for \overline{i} .

It is clear from the proof of Proposition 3.2 that the largest and smallest serially undominated strategies are symmetric Nash equilibria. Asymmetric pure strategy equilibria are also possible as well as mixed strategy equilibria. Since players are *playing the field*, the mixed strategy equilibria require that positive probabilities are attached only to the best replies to the expected average μ . For this reason, as the next proposition shows, mixed strategy equilibria and asymmetric pure strategy equilibria are intimately linked. Let $\sigma = (\sigma_k, \sigma_{-k})$ be a mixed strategy profile with support $C(\sigma)$ and denote the expected average as $\mu(\sigma) = \frac{1}{|K|} \sum_k \sigma_k I_k$. A strategy profile $\sigma = (\sigma_k, \sigma_{-k})$ is a Nash equilibrium if for all players $k \in K$ $C(\sigma_k) \subseteq \beta(\mu(\sigma))$. In equilibrium therefore $\Pi(\sigma_k, \mu(\sigma)) = \Pi(i_k, \mu(\sigma))$ for all $i_k \in C(\sigma_k)$.

Proposition 3.3. A mixed strategy profile $\sigma = (\sigma_k, \sigma_{-k})$ is a Nash equilibrium of , $= \{K, \Sigma, \mu(\mathbf{s})\}$ iff there exists a pure strategy Nash equilibrium of the game \mathbf{s} such that $C(\sigma) = C(\mathbf{s})$.

Proof. We first prove the *if* part. If **s** is a Nash equilibrium of the game $C(\mathbf{s}) \subseteq \beta(\mu(\mathbf{s}))$. Any randomization on $C(\mathbf{s})$ such that the expected average is equal to $\mu(\mathbf{s})$ is then a mixed strategy equilibrium. Necessity is proved by contradiction. Let σ be a mixed strategy Nash equilibrium and assume that there is no $\mathbf{s} \mid C(\mathbf{s}) = C(\sigma)$ that is a pure strategy asymmetric Nash equilibrium. This implies that there is no $\mu \in [\underline{i}, \overline{i}] \mid C(\sigma) \subset \beta(\mu)$, where \underline{i} and \overline{i} are the smallest and largest pure strategies in $C(\sigma)$. Therefore there are $i \in C(\sigma) \notin \beta(\mu)$. An obvious contradiction.

4 The evolutionary analysis

In this section we study the issues of stability and selection of the equilibria of mean-defined supermodular games under a specific class of dynamics

known as *payoff monotonic dynamics*. This type of dynamics may be used to represent at an aggregate level a learning mechanism in an economy populated by a large, but finite, population of boundedly rational agents who play repeatedly the same game. Specifically we envisage an economy where the entire population plays at each time t the same mean-defined supermodular game. The assumption of large populations is necessary for a *strategic* and a technical reason. Large populations allow us to disregard the possibility for any single player to alter future play of his opponents. Technically, large populations are needed for the smoothness of the vector field associated with the dynamics.

To represent a model of social evolution we need to make the strategy set of each player finite. In the evolutionary metaphor strategies represent phenotypes, i.e. agents type characterised by the use of a particular strategy. Moreover we assume that the game is symmetric. For this reason we will not be interested in which strategy is played by each individual player but rather we will focus on the fraction of the population of players playing each strategy. Finally we assume that players play pure strategies only. Adopting the same notation of previous sections, the finite strategy set of the representative agent is I, with cardinality m; Σ denotes the strategy space. The configuration of strategies in the economy is summarised by the state vector **x** in \Re^m whose element x_i denotes the population share associated with strategy i.² All elements of \mathbf{x} are nonnegative and add up to one. The set of all such vectors define the simplex Δ . For each state vector we define the support $C(\mathbf{x}) = \{i \in I : x_i > 0\}$ as the set of existing strategies in \mathbf{x} . Let $\mu(\mathbf{x}) = \mathbf{x}' I$ denote the average strategy induced by \mathbf{x} . This definition of the average strategy is equivalent to the one adopted previously $\mu(\mathbf{s})$. To each strategy profile s corresponds a unique state vector \mathbf{x} ; since the identity of players does not matter $\mu(\mathbf{x})$ is easier to work with. The payoff to a player posting strategy i in state x is denoted as $\Pi(i, \mu(\mathbf{x}))$.

The dynamics is defined over the mixed-strategy simplex Δ in terms of the growth rates of the populations shares associated with each pure strategy $i \in I$ of the game. The dynamic process takes the following general form

$$\dot{x}_i = g_i(\mathbf{x})x_i \tag{1}$$

The function g_i indicates the growth rate of the population share attached to strategy *i*. Recall the following definition:

Definition 4.1. A regular growth-rate function g is *payoff monotonic* if, for

 $^{^{2}}$ Being the game symmetric we drop the subscript from the strategy sets of players. Henceforth subscripts refer to strategies.

all $x \in \Delta$

$$\Pi(i, \mu(\mathbf{x})) > \Pi(\tilde{i}, \mu(\mathbf{x})) \iff g_i(\mathbf{x}) > g_{\tilde{i}}(\mathbf{x})$$

Regularity of the growth function ensures that the dynamics is well behaved in the sense of inducing a unique solution to the system through any initial condition, a solution that never leaves the simplex Δ . According to the Picard-Lindelof theorem a solution to the system $\dot{\mathbf{x}} = \mathbf{g}(\mathbf{x})\mathbf{x}$ exists and is unique if the vector field of the system is Lipschitz continuous. The population state remains in the simplex Δ if the weighted sum of the growth rates $\sum_{i=1}^{m} g_i(\mathbf{x})x_i$ is constantly equal to zero. We assume that the conditions ensuring regularity are satisfied.

We are interested in identifying equilibrium population profiles. As it is clear from the definition, monomorphic population states in which only one strategy is played are stationary points of the dynamics. In other words the dynamic process (1) comes to rest whenever all strategies in the support of \mathbf{x} earn the same payoff. For the sake of predictions one is generally interested in equilibria that are stable. Two concepts of stability are employed: i) Lyapunov stability and ii) asymptotic stability. Broadly speaking, a population state $\mathbf{x} \in \Delta$ is Lyapunov stable if all solutions that start sufficiently close to \mathbf{x} stay close; a state $\mathbf{x} \in \Delta$ is asymptotically stable if it is Lyapunov stable and if trajectories starting sufficiently close to \mathbf{x} eventually approach \mathbf{x} as $t \to \infty$.

One can derive implications from payoff monotonicity and payoff positivity to aggregate Nash equilibrium behaviour. In particular it has been shown that (Nachbar, 1990; Bomze, 1986)

- (a) If $\mathbf{x} \in int(\Delta)$ is stationary in (1), then \mathbf{x} is a Nash equilibrium of the stage game.
- (b) If $\mathbf{x} \in \Delta$ is Lyapunov stable in (1), then \mathbf{x} is a Nash equilibrium of the stage game.
- (c) If $\mathbf{x} \in \Delta$ is the limit to some interior solution to (1), then \mathbf{x} is a Nash equilibrium of the stage game.
- (d) If $\mathbf{x} \in \Delta$ is a strict Nash equilibrium then \mathbf{x} is asymptotically stable.

By virtue of these results limit states as well as Lyapunov and/or asymptotically stable states are to be found in the set of Nash equilibria of the stage game. As the following proposition shows, instability characterises all weak Nash equilibria, both symmetric and asymmetric. Let $\xi(t, \mathbf{x}^{\circ})$ define the solution through a point $\mathbf{x}^{\circ} \in \Delta$ to the system (1).

Proposition 4.1. Let \mathbf{x} and \mathbf{z} be respectively an asymmetric and a symmetric weak Nash equilibrium of , = { $K, \Sigma, \mu(\mathbf{x})$ }. Then under any regular payoff monotonic dynamics both \mathbf{x} and \mathbf{z} are not Lyapunov stable.

Proof. A state $\mathbf{x} \in \Delta$ is Lyapunov stable if every neighbourhood B of \mathbf{x} contains a neighbourhood B° of \mathbf{x} such that the flow $\xi(t, \mathbf{x}^{\circ}) \in B$ for all $\mathbf{x}^{\circ} \in B^{\circ} \cap \Delta$ and $t \geq 0$. In other words all forward orbits from $B^{\circ} \cap \Delta$ are contained in B:

$$\gamma^{+}(\mathbf{x}^{\circ}) = \{\mathbf{x} \in \Delta \mid \mathbf{x} = \xi(t, \mathbf{x}^{\circ}) \text{ for some } t \ge 0\} \subset B \; \forall \mathbf{x}^{\circ} \in B^{\circ} \cap \Delta$$

We first show that asymmetric equilibria are unstable. Let $\mu(\mathbf{x})$ be the average effort associated to \mathbf{x} . Consider a forward orbit $\gamma^+(\mathbf{y})$ emanating from a state \mathbf{y} in a small neighbourhood of \mathbf{x} with $C(\mathbf{x}) = C(\mathbf{y})$, i.e. the two states have the same support. In particular assume that $y_i > x_i \ \forall i \in C(\mathbf{x})$ such that $i < \mu(\mathbf{x})$ and let $\mu(\mathbf{y}) < \mu(\mathbf{x})$. By assumption in \mathbf{x}

$$\Pi(\tilde{\imath}, \mu(\mathbf{x})) - \Pi(i, \mu(\mathbf{x})) = 0$$

for all $i, i \in C(\mathbf{x})$. Strictly increasing differences imply that

$$\Pi(\tilde{\imath},\mu(\mathbf{y})) - \Pi(i,\mu(\mathbf{y})) < 0$$

for all $i < \tilde{i} \in C(\mathbf{y})$. Payoff monotonicity of the dynamics ensures then that the growth rates of each strategy obey the order:

$$g_i(\mathbf{y}) > g_{\tilde{i}}(\mathbf{y})$$
 for all $i < \tilde{i} \in C(\mathbf{y})$

The regularity of the dynamics ensures that $C(\mathbf{y})$ is invariant for all t. We show that $\mathbf{y}' = \xi(t, \mathbf{y}) \prec \mathbf{y} \ \forall t > 0$ where \prec refers to first-order stochastic dominance; $\mathbf{y}' \prec \mathbf{y}$ implies that the cumulative distribution of \mathbf{y}' lies everywhere above the cumulative distribution of \mathbf{y} . Suppose that this was not the case, i.e. let

$$\sum_{k \le i} y'_k \le \sum_{k \le i} y_k \text{ for some } i \in C(\mathbf{y})$$

then

$$\left(1 - \sum_{k \le i} y'_k\right) \ge \left(1 - \sum_{k \le i} y_k\right) = \sum_{k > i} y'_k \ge \sum_{k > i} y_k$$

which contradicts strictly increasing differences and payoff monotonicity. The expected value of any increasing function \mathbf{y} is larger than the expected value of the same function under \mathbf{y}' . Thus the average effort $\mu(\mathbf{y}') < \mu(\mathbf{y})$. As $t \to \infty$ the average effort reduces monotonically and $\xi(t, \mathbf{y}) \to \mathbf{y}^*$ whose support $C(\mathbf{y}^*)$ is a singleton. Hence \mathbf{x} is not Lyapunov stable for all t. Next we show that weak symmetric equilibria are unstable. In a symmetric weak equilibrium all players adopt the same strategy i^* which is a weak best reply to itself. There exists some other strategy \tilde{i} such that $\Pi(i^*, i^*) - \Pi(\tilde{i}, i^*) = 0$. Consider a state $\tilde{\mathbf{z}}$ with support over \tilde{i} and i^* and let ϵ be the population share of \tilde{i} . Assume that $\tilde{i} < i^*$; clearly $\mu(\tilde{\mathbf{z}}) < \mu(\mathbf{z})$ and by the usual argument of strictly increasing differences $\Pi(i^*, \mu(\tilde{\mathbf{z}})) - \Pi(\tilde{i}, \mu(\tilde{\mathbf{z}})) < 0$ The dynamics will then lead the system away from \mathbf{z} towards an equilibrium state where only \tilde{i} is played. *Mutatis mutandis* the same applies to the case where $\tilde{i} > i^*$.

At this stage we are not able to rule out the presence of limit cycles for the general case. However if the game is dominance solvable all strategies but the equilibrium one are iteratively strictly dominated. A well known result, due to Samuelson and Zhang (1992), holds that if a pure strategy does not survive the iterative elimination of strategies that are strictly dominated by another pure strategy, then the strategy does not survive under a payoff monotonic selection dynamics. This result was originally stated for multipopulation games with random matching (see Theorem 1 in Samuelson and Zhang, 1992, p.371). Here we extend their theorem to mean-defined supermodular games. The result need to be proved here as their proof expressly uses the linearity of the expected payoffs in the population shares and this is a property that our model lacks.

Theorem 4.1. Suppose that strategy *i* does not survive the process of elimination of iteratively strictly dominated strategies of , = { $K, \Sigma, \mu(\mathbf{x})$ }. Then for every payoff monotonic dynamics given an interior initial state \mathbf{x}^0 we have

$$\lim_{t \to \infty} \xi_i(t, \mathbf{x}^0) = 0$$

Proof. Let $\mu(\mathbf{x}^0, t)$ be the average strategy at time t given an interior initial state \mathbf{x}^0 . Clearly $\mu(\mathbf{x}^0, t) \in [\underline{i}, \overline{\imath}]$ where \underline{i} and $\overline{\imath}$ are the smallest and the largest strategies in I respectively. Define D^k as the set of strategies that do not survive the kth round of elimination of strictly dominated strategies. Let \underline{D}^k (\overline{D}^k) be the set of strategies in D^k such that $i^k < \underline{i}^{k+1}$ $(i^k > \overline{\imath}^{k+1})$ where \underline{i}^{k+1} $(\overline{\imath}^{k+1})$ is the smallest (largest) strategy to survive round k+1 of deletion. For k = 0, all strategies in \underline{D}^0 are strictly dominated by \underline{i}^1 which

implies that

$$\frac{d}{dt} \left[\frac{\xi_{i^0}(t, \mathbf{x}^0)}{\xi_{\underline{i}^1}(t, \mathbf{x}^0)} \right] = \left[g_{i^0}(\mathbf{x}) - g_{\underline{i}^1}(\mathbf{x}) \right] \frac{\xi_{i^0}(t, \mathbf{x}^0)}{\xi_{\underline{i}^1}(t, \mathbf{x}^0)} < 0$$

for all $t \ge 0$ and all $i^0 \in \underline{D}^0$.

By continuity of g, there exists a $\varepsilon > 0$ such that $[g_{i^0}(\mathbf{x}) - g_{\underline{i}^1}(\mathbf{x})] < -\varepsilon \quad \forall t \ge 0$ hence

$$\frac{d}{dt}\frac{\xi_{i^0}(t,\mathbf{x}^0)}{\xi_{\underline{i}^1}(t,\mathbf{x}^0)} < -\varepsilon \ \frac{\xi_{i^0}(t,\mathbf{x}^0)}{\xi_{\underline{i}^1}(t,\mathbf{x}^0)}$$

and therefore

$$\frac{\xi_{i^0}(t, \mathbf{x}^0)}{\xi_{\underline{i}^1}(t, \mathbf{x}^0)} < \frac{x_i^0}{x_{\underline{i}^1}^0} \exp(-\varepsilon t) \quad \forall t \ge 0$$

Being $\xi_{\underline{i}^1}(t, \mathbf{x}^0) < 1 \quad \forall t$, it follows that $\xi_{i^0}(t, \mathbf{x}^0) \to 0$. This shows that there exists a $\eta > 0$ and time T^0 such that

$$\xi_{i^0}(t, \mathbf{x}^0) < \eta \quad \forall i^0 \in \underline{D}^0 \text{ and all } t \ge T^0$$

The same sort of argument applies to strategies in \overline{D}^0 . Let τ^0 be the time such that $\xi_{i^0}(t, \mathbf{x}^0) < \eta \quad \forall i^0 \in \overline{D}^0$ and all $t \ge \tau^0$.

Without loss of generality assume that $T^0 > \tau^0$. For all $t > T^0$ the population shares attached to the strategies in D^0 are small enough for us to consider the reduced game with strategy set $I^1 = \{i : i \notin D^0\}$ for each player. For $t \ge T^0$ the average strategy $\mu(\mathbf{x}^0, t)$ is contained in a small neighbourhood of the interval $[\underline{i}^1, \overline{\imath}^1]$. The same process will take place for all strategies in D^1 ; the share of the population playing those strategies will become negligible at some finite time T^1 and τ^1 . Since the process of iteration ends in a finite number of rounds in finite games, only a finite number of iterations of the argument are required and we conclude that there is a finite time T^* after which the shares of all serially dominated strategies converge to zero.

A payoff monotonic selection dynamics removes all serially dominated strategies irrespective of whether the dynamics converges or not. If the game is dominance solvable then the set of serially undominated strategies is a singleton and the dynamics must converge. We state this formally in the next corollary.

Corollary 4.2. Consider , = $\{K, \Sigma, \mu(\mathbf{x})\}$. If the process of elimination of iteratively strictly dominated strategies yields one strategy for each player then the system converges from any interior initial state under any regular payoff monotonic selection dynamics.

5 An application: The coconut economy

In this final section we apply the analysis developed above to a model of search inspired to the *coconut economy* of Diamond (1982). The model exhibits, under opportune parametrisations, multiple Pareto-ranked equilibria and coordination failures. Whereas in a Walrasian economy these cannot occur, once the perfectly competitive paradigm is abandoned, the causes of such phenomena are numerous. In the Diamond (1982) coconut model the auctioneer is replaced by a random matching mechanism; in Howitt (1985) the cost of transactions are made dependent on the level of economic activity; Kiyotaki (1988) analyses a monopolistically competitive market with a multiplicity of equilibria; Cooper and John (1988) discuss the conditions for coordination failures in models of imperfect competition.

Consider an economy populated by a large number K of individuals. Each individual starts with one unit of a perishable good in each period ("coconuts"), and there may be some coconuts lying around loose for anyone to pick up. Each agent wishes to consume one and only one coconut per period, and there is a taboo against eating your own coconuts ³. To consume a coconut he/she must either trade with another agent, or pick up a loose coconut. Individuals are scattered around the economy and a search must be undertaken in order to find a trading partner: search is costly and the probability of agent k trading is a function of his search effort i_k and the average search effort of the rest of the population. The model is closely related to Diamond (1982), except that here individuals do not choose the level of production, but only the search intensity. We will analyse first the constituent search model with continuous strategy sets so that effort levels icome from a compact, convex subset of the real line, $\mathcal{I} \subset \Re_+$.

Once individuals have chosen their search effort the actual search takes place. Should two agents meet, the gross utility from trade is L. The matching function P gives the probability of individual i meeting someone to trade with as a function of individual and collective search effort:

$$Pr[\text{trade}] = P(i, \ \mu) \tag{2}$$

where

$$\mu = \frac{1}{K} \sum_{k=1}^{K} i_k \tag{3}$$

³" but there is a taboo against eating (coco-)nuts one has picked oneself", (Diamond, 1982, page 893).

is the average effort, and i_k is the individual effort. As it is standard, we assume that P is increasing in both i and μ ($P_i > 0$ and $P_{\mu} > 0$), and most importantly that the cross derivative $P_{i\mu} > 0$. Furthermore, we assume $P(0,\mu) = 0$ and $P_{ii} \leq 0$, and that the matching function is the same for all agents k. The expected payoff from search to the representative agent is (dropping the k subscript):

$$\Pi(i,\mu) = L \cdot P(i,\mu) - 1/2i^2 + k \tag{4}$$

Where L is the gross benefit from trade, and k > 0 is a normalisation chosen to ensure that all payoffs are non-negative. We assume that agents are risk neutral, and choose search effort to maximize (4).

We can therefore define the best-response function for the representative agent (dropping the subscript) $\beta(\mu) : \mu \mapsto i$:

$$\beta(\mu) = \arg \max_{i \in \mathcal{I}} \Pi(i, \mu)$$
(5)

It is a fundamental assumption of search models that the best response is increasing in μ : this is ensured by the assumption that $P_{i\mu} > 0$.

Whilst we are interested in looking at matching functions P which yield multiple equilibria, it is useful to consider at first properties of matching functions that usually yield unique equilibria. For example, a common assumption about the matching function is that it is homogeneous of degree 0 in (i, μ) , so that it can be written as a function \hat{P} of the relative search intensity i/μ , yielding the payoff function:

$$\Pi = L \cdot \hat{P}(\frac{i}{\mu}) - 1/2i^2 + k$$

Where \hat{P} is strictly increasing in $\frac{i}{\mu}$ when $\mu > 0$. Let us consider the game $G = \{K, \mathcal{I}, \mu(\mathbf{s})\}$. All equilibria are symmetric, so that $i = \mu$. Let \hat{i} be the set of Nash-equilibria: then it is easy to show that there is at most one $\hat{i} \in \hat{i}$ where $\hat{i} > 0$.

For the search technology to be represented by \hat{P} , it requires that the matching probabilities are scale independent, so that the levels of i and μ do not matter. This is a strong assumption. In this paper, we assume that the matching function is a logistic function of the following form for each agent k:

$$P(i_k,\mu) = \frac{i_k^{\alpha}}{c + \exp(a - b \ \mu)} \qquad \forall k \in (1,\dots,K) \text{ and } i_k,\mu \in \mathcal{I}$$
 (6)

where $\alpha \in (0, 2), c > 0, b > 0, c > 0$ with the joint restriction on (α, a, c, b) that $P(i, \mu) \in [0, 1]$. This function has the property that the marginal effect





Figure 1: The matching function

Figure 2: The matching function

of individual search on P given μ is constant if $\alpha = 1$, decreasing if $\alpha < 1$ and increasing if $\alpha > 1$. However, the effect of μ on P is sigmoidal. Note that if $\mu = 0$ then P(i,0) > 0 for i > 0. This reflects the fact that there is still a probability of finding a coconut even if no one else is searching, due to the possibility of picking up *loose* coconuts or finding a non searching agent. The function (6) is depicted in figure 1 and 2 with the parameter values $\alpha = 0.9, a = 4, b = 2$ and c = 3.

In traditional search models, there are usually two sides to the market (for example firms and workers): trades depend on inputs from both sides of the market (see for example Pissarides (1990) Chapter 4). In our context there is only one sort of agent: total possible trades are fixed ($P(\text{trade}) \in [0, 1]$). This follows the Diamond (1982) setup. In the standard model the issue of returns to scale in the matching process is crucial for the existence of multiple equilibria. In general, increasing returns to scale are required for the game to exhibit multiple equilibria. In our setting this is not the case, although the matching function presents varying return to scale under opportune parametrisation, the presence of multiple equilibria does not hinge upon increasing returns to search. ⁴. In particular the nature of the returns depends on the value of the parameter α . When $\alpha < 1$ the matching function initially has increasing returns for low levels of search and then decreasing,

⁴This contrasts the findings of Diamond (1984).

whereas for $\alpha \geq 1$ we obtain initially decreasing and then increasing returns to search.

The payoff function is (dropping the k subscript):

$$\Pi = \frac{Li^{\alpha}}{c + \exp(a - b\mu)} - 1/2i^2 + k$$
(7)

This function is strictly concave in i, and strictly supermodular. It can give rise to multiple equilibria. From (7) the best response function is ⁵

$$\beta(\mu) = \left[\frac{\alpha L}{c + \exp(a - b\mu)}\right]^{\frac{1}{2-\alpha}}$$
(8)

We depict the best-response function in figure 3 for the same parameter values used in figure 1. In this case there are 3 pure strategy equilibria; these are Pareto-ranked, so that A is the worst and C the best. A is a low level equilibrium where everyone is undertaking a low level of search and so the consumption of the population is low: at C there is high search and high consumption.

It is easy to verify that payoff function (7) has strictly increasing differences in its two arguments i and μ^{6} and hence the game is strictly supermodular.

5.1 The equilibrium analysis of the search game

In this section we briefly describe the set of Nash equilibria of the search game for both the continuous and discrete version. As we have seen in the first part of the paper, for the subsequent evolutionary analysis we need to consider finite strategy sets for each player. Using the same notation adopted in section 2 Let I and; $\Sigma = \times_{k \in K} I_k$ denote the strategy set and the strategy space respectively for the discrete case. We write the discrete search game formally as $G = \{K, \Sigma, \mu(\mathbf{s})\}$ where $\Pi : \{I \times \mu\} \mapsto \Re$ is given by (7). The pure strategy Nash equilibria of the game fall into two categories. Symmetric equilibria where each player adopts the same strategy that is then equal to the population average strategy, and asymmetric equilibria where different players adopts different strategies. Let $\beta(\mu(\mathbf{s}))$ be the set of best replies to the population average effort $\mu(\mathbf{s})$. A strategy profile \mathbf{s} defines a Nash

⁵In the derivation of the best response we have set the derivative $\frac{\delta \mu}{\delta i} = 0$. Due to the assumption of a large population the effect of a change in the level of effort exerted by one player on the average population effort is negligible.

⁶Note that this is equivalent to the condition that $\Pi(i, \mu) - \Pi(i, \mu')$ is nondecreasing in i for all $\mu \ge \mu'$.



Figure 3: The best reply function

equilibrium if each player is taking best responses to $\mu(\mathbf{s})$, where $\mu(\mathbf{s})$ is the population average effort implied by s. Consider the pair $(s, \mu(s))$ and let $C(\mathbf{s})$ be the support of \mathbf{s} . Obviously $\mu(\mathbf{s})$ need not belong to $C(\mathbf{s})$. Consider first the case where $\mu(\mathbf{s}) \in C(\mathbf{s})$, then $(\mathbf{s}, \mu(\mathbf{s}))$ is a Nash equilibrium only if $\mu(\mathbf{s}) \in \beta(\mu(\mathbf{s}))$. From strict concavity of the payoff function in i it follows that $\beta(\mu(\mathbf{s})) = \{\mu(\mathbf{s})\}$, i.e. the set of best replies is a singleton. Each player is playing the same strategy $i = \mu(\mathbf{s})$ and the pair $(\mathbf{s}, \mu(\mathbf{s}))$ is a symmetric strict Nash equilibrium where all players exert the same level of effort. Suppose instead that $\mu(\mathbf{s}) \notin C(\mathbf{s})$. Again $(\mathbf{s}, \mu(\mathbf{s}))$ is a Nash equilibrium only if $i \in \mathcal{I}$ $\beta(\mu(\mathbf{s})) \forall i \in C(\mathbf{s})$ and this implies that $(\mathbf{s}, \mu(\mathbf{s}))$ is a weak Nash equilibrium. From strict concavity of the payoff function it is immediately clear that $\mu(s)$ must not belong to I either, and that there are at most two best replies to any given $\mu(\mathbf{s})$. These are asymmetric equilibria, different players exert different levels of effort that are optimal given the population average. Moreover there is a whole set of strategy profiles compatible with $\mu(\mathbf{s})$ and the requirement that $i \in \beta(\mu(\mathbf{s})) \ \forall \ i \in C(\mathbf{s})$. Because of the symmetry of the game the identity of the players does not matter and we do not need to consider such strategy profiles separately.

In the continuous version of the game, all pure strategy NE are symmetric because of strict concavity of the payoff function in *i* given μ . Only when the strategy sets are discretised, asymmetric equilibria can occur. The obser-

vation that for any given μ the payoff function (7) is strictly concave makes it easier to compute the asymmetric Nash equilibria of the game. These infact can only occur *near* the symmetric equilibria of the continuous game. For $\alpha = 1$ the payoff function is also symmetric around its maximum for any given μ^{7} . More precisely two consecutive strategies $(i < \hat{i}) \in I$ can form an asymmetric NE iff $\mu, \tilde{\beta}(\mu) \notin C(\mathbf{s})$ and i, \hat{i} are such that:

$$\left|\tilde{\beta}(\mu) - i\right| = \left|\tilde{\beta}(\mu) - \hat{i}\right| = \min(\left|\tilde{\beta}(\mu) - i'\right|) = \delta/2 \quad \forall \ i' \in I$$

where δ defines is the distance between i and \hat{i} , $\tilde{\beta}(\mu)$ is the best reply to μ defined over the continuous strategy set \mathcal{I} . The symmetry of the payoff function coupled with strict concavity implies infact that strategies equally apart from $\tilde{\beta}(\mu)$ earn the same payoff against μ . The condition reported above implies that $\tilde{\beta}(\mu) = (i+\hat{i})/2 = i+\delta/2$ hence whenever $\left|\tilde{\beta}(\mu) - \mu\right| > \delta/2$ we cannot observe an asymmetric NE since this will imply that $\hat{i} > i > \mu$ an obvious contradiction. The range of values that need to be considered to check for the existence of asymmetric equilibria is centred around the strict NE of the continuous game. The size of this interval is rather small and is a decreasing function of δ . In particular the interval is given by the following:

$$\delta \left[\frac{L \ b \ \exp(a - bi^*)}{\left[c + \exp(a - bi^*)\right]^2} - 1 \right]^{-1}$$

where i^* are the NE of the continuous game.

5.2 The evolution of search behaviour

Next we describe the evolutionary model of search. We assume that at the aggregate level in the economy operates a selection mechanism which rewards those strategies that on average have proved to be more profitable. This process is analytically described by the replicator dynamics.

The replicator dynamics is a special case of payoff monotonic dynamics and it assumes that the subpopulation playing a particular strategy grows in proportion to the difference between the payoff it secures and the average population payoff. We think that the choice of the replicator as a description of the dynamic adjustment toward the equilibrium is justified by its relative flexibility. This type of dynamics can in fact result from different assumptions about individual behaviour ⁸. The economy-wide learning mechanism

⁷Note that $\frac{d^2\Pi}{di^2} = -1$.

⁸We overlook the details which can be found in (Weibull, 1995, chapter 4) and Schlag (1998) among the others.

is represented by the following continuous time replicator dynamics:

$$\dot{x}_i = g_i(\mathbf{x})x_i \tag{9}$$

where $g_i(\mathbf{x}) = [\Pi(i, \mu(\mathbf{s})) - \sum x_i \Pi(i, \mu(\mathbf{s}))]$ and

$$\sum x_i \Pi(i, \mu(\mathbf{s})) = \frac{L\mathbf{x}' I^{\alpha}}{c + \exp\left(a - b\mu\right)} - 1/2\mathbf{x}' I^2 + k \tag{10}$$

is the population average payoff. Again this formulation differs from the standard replicator dynamics since here the growth rate is not linear in the population shares x_i and we need to check that the dynamic system given in (9) has a Lipschitz continuous vector field. That this is the case is proved by noting that the vector field has continuous partial derivative with respect to **x**. It is easily verified that the solution remains in the simplex if the sum of growth rates is equal to zero, something which visual inspection of expressions (9) and (10) proves to be the case.

As discussed above, the Nash equilibria of $G = \{K, \Sigma, \mu(\mathbf{s})\}$ falls in two categories, symmetric and asymmetric equilibria. Symmetric Nash equilibria are strict and therefore represent asymptotically stable states of the dynamics. Asymmetric equilibria are weak and do not share this stability property as they do not satisfy the weaker concept of Lyapunov stability. Let $\xi(t, \mathbf{x}^{\circ})$ define the solution through a point $\mathbf{x}^{\circ} \in \Delta$ to the system (9).

Proposition 5.1. Let $\mathbf{x} = (x_i, x_i)$ be an asymmetric Nash equilibrium of $G = \{K, \Sigma, \mu(\mathbf{s})\}$. Then \mathbf{x} is not Lyapunov stable.

Proof. See the proof of Proposition (4.1)

Instability of the asymmetric Nash equilibria does not imply that these cannot be limit points to some interior solution. However an (unmodeled) shock can lead the population away from these Nash equilibrium states.

Robust predictions require asymptotic rather than Lyapunov stability. Asymptotic stability is a robust property in the sense that small perturbations to the vector field of the dynamics do not destroy it. If we allow for unmodeled drift and/or for the possibility that the model do not capture features that are likely to perturb the system then we should rely on asymptotic stability. This latter holds for the symmetric equilibria of $G = \{K, \Sigma, \mu(\mathbf{s})\}$ only. For this reason we characterise for the game at hand a set-valued generalisation of asymptotically stable state due to Ritzberger and Weibull (1995).

Consider a subset H of the strategy set I of the game and define a surviving set $H \subset I$ as follows:

Definition 5.1. $H \subset I$ is a long run surviving set if its simplex $\Delta(H)$ is asymptotically stable and H does not contain a non empty subset L for which $\Delta(L)$ is asymptotically stable.

The above definition can be found in Weibull (1995, p.118)⁹ and requires the surviving set H to be minimal with respect to the property that if initially all pure strategies not in H are present in sufficiently small proportions, then they will vanish over time. Proposition 4.10 in Weibull (1995, p.149) gives a necessary condition for the asymptotic stability of H under any payoff positive dynamics. In the next proposition we show that the same necessary condition holds for our game. Using Weibull's notation let

$$\alpha^0(H) = \{ i \in I : \Pi(i, \hat{i}) > \Pi(\hat{i}, \hat{i}) \text{ for some } \hat{i} \in H \}$$

Proposition 5.2. Consider the game $G = \{K, \Sigma, \mu(\mathbf{s})\}$. If $\Delta(H)$ is Lyapunov stable under the replicator dynamics (9) then $\alpha^0(H) \subset H$.

Proof. The proof of is similar in structure to that of Proposition 4.10 in Weibull (1995) the only difference being that in our case we have to allow for the dependence of the payoffs on the average strategy. Suppose that $\alpha^0(H) \not\subset H$. Then there exists some $i \notin H$ and some $\hat{i} \in H$ such that

$$\Pi(i,\hat{i}) > \Pi(\hat{i},\hat{i})$$

Recall that the second argument in the profit function $\Pi(\cdot)$ represents the average effort. Consider a subset $M = \{i, \hat{i}\} \subset I$. A solution orbit starting in $\Delta(M)$ remains in the simplex forever. Consider now a state $\mathbf{y} \in \Delta(M)$ arbitrarily close to a state \mathbf{x} which put weight $1 \text{ to} = \hat{i} \in \Delta(H)$. The average effort $\mu(\mathbf{y}) \simeq \mu(\mathbf{x}) = \hat{i}$ and so the continuity of the growth rates $g'_i s$ ensures that $g_i < 0, g_i > 0$. The flow $\xi_i(t, \mathbf{y}) > y_i$ for any t > 0, thus moving the average nearer to i. Increasing differences ensure that the solution orbit will forever move away from $\Delta(H)$ and hence $\Delta(H)$ is not Lyapunov stable. \Box

Then $\alpha^0(H) \subset H$ requires that H contains the strictly better replies to all strategies in H. We show that if this necessary condition is met, then the set H does not respect the definition (5.1) as it contains a proper subset spanning an asymptotically stable face of the simplex Δ .

Proposition 5.3. The only long run survivor sets of $G = \{K, \Sigma, \mu(\mathbf{s})\}$ are singletons and correspond to the strict Nash equilibria of the game.

 $^{^{9}{\}rm This}$ definition was first introduced in Ritzberger and Weibull (1995) within a more complex multipopulation model.

Proof. Consider a set $H \subset I$ and let \underline{i} and $\overline{\imath}$ be the smallest and largest elements in H. H is a long run survivor set only if, from Proposition (5.2), $\alpha^0(H) \subset H$. This requires the reversal of the sign of the difference $\beta(i) - i$. This difference must be positive for \underline{i} , i.e. $\beta(\underline{i}) > \underline{i}$, and nonpositive for $\overline{\imath}$, i.e. $\beta(\overline{\imath}) < \overline{\imath}$. This reversal does not violate the monotonicity of best responses, *iff* at least one $i \in H$ is a symmetric Nash equilibrium, which in turn makes the set H not minimal with the respect to the property of being asymptotically stable.

We are not able to rule out the possibility of limit cycles for all parametrisation of the payoff function (4). However we can prove that if the game $G = \{K, \Sigma, \mu(\mathbf{s})\}$ exhibits a unique pure strategy NE and given any interior initial state the dynamics will converge to the unique NE. The proof is given in Corollary (4.2).

5.3 The relevance of initial conditions

In the analysis carried out so far we have been able to characterise the equilibria of the dynamics and their stability properties. We also ruled out the existence of limit cycles in the special case of dominance solvability of the stage game. In all other cases the game typically has a number of asymmetric (weak) and symmetric (strict) Nash equilibria. The asymmetric equilibria are unstable whereas the strict equilibria are asymptotically stable.

One last issue that so far defies analytical characterisation is the effect that the initial distribution has on the resulting equilibrium.

We are not able to pinpoint which initial conditions will prompt a specific equilibrium. We can safely dismiss unstable equilibria as these will never be observed unless the initial state of the system happens to be at the unstable equilibrium. The question then really is: what can be said about the asymptotically stable equilibria? Each of these has a basin of attraction but a characterisation is very difficult to obtain. The addition of a *drift* term to the replicator equation (9) will not alter the asymptotic properties of the asymptotically stable states and therefore it not useful for equilibrium selection.

The issue of equilibrium selection is of particular interest when the game exhibits multiple strict Nash equilibria and especially when, as in our model, these are Pareto-rankable. The types of dynamics analysed in this work are deterministic and therefore cannot select one strict equilibrium from the set of strict Nash equilibria. With deterministic dynamics we can only perform what some authors (Binmore et al., 1995; Binmore and Samuelson, 1997)

refer to as *long-run* analysis as opposed to the *ultralong-run* analysis. The long run refers to a time span sufficient for the sample path of the dynamics to reach an equilibrium in the vicinity of which it will spend a long time. The ultralong run, instead refers to "the length of time required for mutations and other rare events to occur with sufficient frequency to make a stationary distribution relevant." Samuelson (1998).

With deterministic dynamics, the issue of equilibrium selection can only be addressed somewhat informally by restricting the attention to stable equilibria since the dynamics is carried away from unstable equilibria by (generally unmodeled) shocks or mutations. Any further analysis of equilibrium selection requires the modeling of an explicitly and truly stochastic dynamic and the study of the properties of its sample path as time goes to infinity and either the population size grows to infinity and/or the mutation rate goes to zero. The adoption of a truly stochastic dynamic allows for the analysis of robustness of the equilibria against sequences of small shocks or simultaneous small shocks that together form a big perturbation that can move the path away from stable equilibria. No matter how small or infrequent, these shocks alter the nature of the dynamic process. Instead of being dependent on the initial population state the process may become ergodic and have an asymptotic stationary distribution that is history independent. This type of ultralong run analysis was pioneered by Foster and Young (1990); Fudenberg and Harris (1992); Young (1993); Kandori et al. (1993); Samuelson (1994); Binmore et al. (1995); Cabrales (1996); Binmore and Samuelson (1997) and develops along two main different lines. The papers by Foster and Young (1990); Fudenberg and Harris (1992); Cabrales (1996) study continuous time stochastic systems based on the replicator dynamics where the stochastic term is represented by a Wiener process. Foster and Young (1990) consider a single-population replicator dynamic and add a Wiener process with no cross-variance and a state-dependent variance. Their system of stochastic differential equations takes the general form:

$$dx_{i,t} = x_{i,t}g_i(\mathbf{x}_t) + \sigma(i|x_i)dW_t(i)$$

where $\sigma(i|x_i)$ is the variance function and W is a standard Wiener process. They compute the limit of the long run distributions, letting the variances of the Wiener process go to zero and they show that in 2 X 2 coordination games the distributions converge to the Pareto-efficient and risk-dominant equilibrium.

The approach followed by Fudenberg and Harris (1992) differs from the one just described in that a stochastic term is added directly to the payoffs to each strategy i and then the equations for the evolution of population shares

are derived. The advantage in using payoff shocks is that of being consistent with a nonnegligible level of noise in models with a continuum of agents. As the population grows large, i.i.d. shocks to individual agents tend to become deterministic and some form of correlation between shocks is necessary to explain aggregate noise with a continuum of agents. ¹⁰ Fudenberg and Harris (1992) then analyse 2 X 2 games and show that if the game has 2 strict Nash equilibria, the system is not ergodic and converges with probability 1 to one of the two equilibria with relative probabilities depending on the initial conditions. By adding a flow of deterministic mutations, the system becomes ergodic as it can never reach the boundary and they show that the limit of the ergodic distributions, computed by taking both the payoff variances and the deterministic mutation flow to zero, converges to the risk-dominant strategy. ¹¹ This result, however, does not carry over to N-player games as shown in Cabrales (1996).

Alternatively, the papers by Young (1993); Kandori et al. (1993); Samuelson (1994); Binmore et al. (1995); Binmore and Samuelson (1997) study discrete time, autonomous, finite-population stochastic adjustment models. The analysis carried out in these papers relies heavily on the results in the theory of Markov chains. Common to the abovementioned papers is the procedure with which the limiting ultralong run distributions are derived. The procedure typically entails 4 steps:

- 1. Specification of the state space. This is generally given by the number of agents in each player population playing each strategy and can include also some information about actions played in a number of previous periods;
- 2. Specification of the deterministic dynamics (best-response dynamics or other type of dynamics such as payoff monotone dynamics etc.) and derivation of the Markov transition matrix;
- 3. Introduction of a noise term which affects the Markov transition matrix (some technical assumptions are needed to make the Markov matrix ergodic which ensures that the process has a unique invariant distribution);

¹⁰Other authors, notably Binmore et al. (1995), employ a stochastic differential equation to approximate the limit of the ultralong run distribution of discrete-time finite population model. By taking limits appropriately, they derive a stochastic differential equation which is then used to compute the ultralong run limit of the system as the noise goes to zero and not to model a system with nonnegligible noise level.

¹¹In 2 X 2 games the risk-dominant equilibrium need not be the Pareto-dominant equilibrium unless the game is a coordination game.

4. Derivation of the limiting distribution as the noise goes to zero.

The main result in Young (1993); Kandori et al. (1993) is that in 2 X 2 games, the risk-dominant equilibrium is selected as the unique steady state.

In general games the risk-dominant equilibria may fail to be Pareto efficient and the quite robust conclusion from the analysis of stochastic adjustment models is that selection dynamics tend to select equilibria that are relatively resistant to mutations (risk-dominant equilibria) and this may conflict with the criterion of Pareto efficiency.

This result, however, is partly questioned by Binmore et al. (1995) who show that, incorporating sources of noise intrinsic to the selection process¹² in 2X2 games, the limiting distribution may select the pay-off dominant equilibrium.

The authors also derive a link between a continuous time version of the replicator dynamics and the long-run behaviour of the Markov process they study by showing that if the population is sufficiently large and the length of each time period is sufficiently short then the sample path of their model can be approximated arbitrarily closely by a solution of the replicator dynamics which incorporates a *deterministic* noise term. However, in the ultralong run the replicator is not a good approximation and the ultralong run stationary distributions are derived form the original Markov process adopting the methods of Freidlin and Wentzell (1984).

The ultralong run analysis of the models presented is well beyond the scope of this work and would require a new specification of the model. However, we think that the effect that the initial distribution has on the resulting equilibrium is an important issue since the equilibria are Pareto-rankable. To address this issue further we resort to simulations.

In our stylised economy the level of activity is summarised by the average search effort put forth by the population of agents. It is the level of activity that drives the dynamics. The higher the average effort the higher being the marginal payoff to an increase in the individual search effort.

The questions we ask are: Which initial conditions will prompt the economy to reach the Pareto dominant equilibrium? Will the level of activity in the economy fluctuate along the evolutionary path or will it monotonically reach its equilibrium value?

From an analytical point of view the fact that the average effort will increase or decrease along the evolution path to equilibrium depends not only

¹²This form of noise differs in a fundamental way from the one considered in Kandori et al. (1993). In this latter, the only source of randomness is represented by mutations. In Binmore et al. (1995) even in the absence of mutations, the selection process is noisy since agents need not always adjust their strategy towards the current best reply.

on its initial value but on the details of the distribution of the population of players across strategies as well. What emerges from the simulations is that for a wide set of initial distributions the system will converge to the equilibrium with an average effort above (below) the initial average if the best reply to it is above (below) it. Moreover the time profile of the average effort converges monotonically to its equilibrium value. For some initial distributions this does not happen. Two effects interacts in a complex way to affect the time pattern of the level of activity in the economy. In each period the strategies closer to the best reply to the current average enjoy the highest growth rates and the agents tend to group around the best reply. The distance between this latter and the current average though needs not to decrease monotonically since the average effort is determined not by the growth rates but by the growth of the strategies in absolute terms. The new average results infact from the combined effect produced by agents increasing their search effort and agents reducing theirs as the population groups around the current best reply. The exact details of the distribution are brought to bear and hence it is very difficult to derive analytical sufficient conditions for monotonicity of the average population effort in the economy to hold. Nevertheless, for sufficiently uniform initial distributions, i.e. distributions where the population is not clustered around few strategies, the outcome of the dynamics can be predicted comparing the initial average effort with the best response to it. If the difference between the two is positive, then the system converges toward the higher equilibrium, whereas if it is negative, the economy converges to the low level equilibrium. The simulations strongly corroborate this intuition; we generated infact over one hundred random initial distributions and constantly obtained this result. The next picture shows the evolution of the population average effort for randomly generated initial distributions.

[Fig. 6.4 here]

In cases a), the difference between the initial average and the best reply is positive and the system converges to the high level equilibrium; in cases b) the difference is negative and the low level equilibrium is attained.

To test further our conjecture we ran four hundred additional simulations. To improve our control over the characteristics of the initial states we adopted the following procedure to generate the different \mathbf{x}^0 s. We start by noting that a population state is formally equivalent to the probability mass function (PMF) of a discrete distribution defined over a finite domain. A very broad and flexible class of such distributions is that of the generalised power series



27

Figure 4: The evolution of the population average effort for random initial states

distributions. The PMF can be written in the form:

$$Pr[X = x] = \frac{a_x \theta^x}{\sum_{x=0}^{\infty} a_x \theta^x}, \quad x = 0, 1, \dots, \ \theta > 0$$

By choosing θ and a_x appropriately we can generate all sorts of initial distributions with different average and variance, unimodal or multimodal. The results obtained seem to reinforce our conjecture. Whenever the initial average was below (above) the smallest (largest) of the NE, the system converged to this latter. Figure 5 shows nine initial states generated as generalised power series distributions and figure 6 shows the evolution of the average search effort. Again in cases a), the difference between the initial average and the best reply is positive and the system converges to the high level equilibrium whereas in cases b) the difference is negative and the low level equilibrium is attained.

We also employed a binomial distribution to generate initial states. Mean and variance are controlled by a single parameter. The shape of the distributions obtainable is unimodal and in over one hundred simulations, allowing for different values of the parameter and for different strategy sets we obtained the same result.



29

Figure 5: Nine initial bimodal distributions



30

Figure 6: The evolution of the population average effort for the initial bimodal distributions

6 Conclusions

In this paper we have introduced the class of mean-defined supermodular games and analysed the issues of selection and dynamic stability of the equilibria of this type of games under a general class of evolutionary selection dynamics known as *payoff monotonic dynamics*.

Two features characterise this class of games: the payoff function is supermodular and the payoff to each player depends on his own strategy and the average of the population play.

Supermodularity of the payoff function implies the existence of strategic complementarities between players' strategies. These in turn may led to the existence of multiple Pareto-rankable equilibria and eventually to *coordina-tion failures*.

The assumption that the payoff of each player depends on a summary statistic of the population strategy profile and not directly on the profile itself makes this class of games particularly suitable for evolutionary analysis of N-player games where the usual random matching mechanism cannot be justified. Evolutionary models generally assume large populations of players that are randomly matched to play a two-person game. This scenario cannot be usefully employed to describe situations in which the entire population of players interacts simultaneously in a N-player game. Models of monopolistic competition, network games as well as search models of the type presented in the paper exemplify such instances.

As for the analysis of the Nash equilibria of the mean-defined supermodular games, we have shown that these games share two fundamental properties of supermodular games. Namely the property of monotonicity of best replies and the property according to which the greatest and least pure strategy Nash equilibria are also the greatest and the least element of the set of pure strategies that survive the process of iterated deletion of strictly dominated strategies. This same result is presented in Milgrom and Roberts (1990b) but needs to be proved for this class of games as the payoff to the players depends on a statistic of the strategy profile rather than the profile itself.

As for the issues of coordination, equilibrium selection and dynamical properties of the equilibria for this class of games under payoff monotonic selection dynamics, the main conclusion is that weak Nash equilibria, both in pure and mixed strategies, are unstable. The only asymptotically stable equilibria of the game are symmetric strict equilibria where each player uses the same strategy. Finally, we extend the results of Samuelson and Zhang (1992) concerning the long run survival of serially dominated strategies. We show that the strategies that do not survive the process of iterated deletion of strictly dominated strategies vanish in the long run. As a corollary to this latter result, we show that if the game is dominance solvable then the dynamics converges from any initial interior state.

In the second part of the paper we present an application of the analysis carried out in the previous sections to a model of decentralised trading resembling the "coconut economy" of Diamond (1982) using a specific type of payoff monotone dynamics known as replicator dynamics. Agents are endowed each period with one unit of perishable good (coconut) that they have to trade with another agent before consumption. Agents are scattered around and a search must be undertaken to find a trading partner. The return to individual search is increasing in the population average search effort and the level of activity in the economy coincides with total consumption which in turn is increasing in the average level of search activity. This type of economy may be characterised by *coordination failures* in that the system may fail to achieve the Pareto-dominant equilibrium. In our model the potential causes for such failures are the presence of a random matching mechanism replacing the auctioneer and the bounded rationality of the agents. The dynamic analysis of such economies is interesting in that it provides, at least partially, an answer to the two often cited criticisms to this literature on *coordination* failures about the static nature of the models and the lack of a selection mechanism that helps to forecast which equilibrium will be selected. In the choice of the dynamic process we opted for a particular process known as replicator dynamics. If the game is dominance solvable, the dynamics converges to the pure strategy Nash equilibrium from any given interior initial distribution. This result is important in that it rules out the possibility of limit cycles where the economy perpetually fluctuates along states characterised by changing levels of consumption. The weak Nash equilibria are unstable. This in turn has relevant consequences for the sake of predictions. If we concede that the deterministic model employed is in fact missing some aspects of the problem under analysis, robust predictions call for stability or better asymptotic stability. In addition we show that the only long run survivor sets (in the sense of Ritzberger and Weibull (1995)) are singletons and correspond to the strict Nash equilibria of the game. In the presence of multiple equilibria the dynamics exhibits strong path dependence. For a large number of initial distributions the convergence towards a high or low equilibrium can be predicted just by looking at the difference between the initial average population search effort and the best reply to it. Whenever this is positive, the system converges towards the Pareto-dominated equilibrium characterised by a low level of consumption; when the best reply is larger than the initial average the system converges to the Pareto-dominant equilibrium. This result does not carry over to all possible initial distributions and for these it is not possible to forecast the final equilibrium.

We believe that the framework adopted in this paper is particularly suited to analyse the dynamics of this type of economies with agents interacting repetitively and in an uncoordinated fashion. The model presented is simple and more fundamental questions may be addressed allowing for additional complications. In a possible extension we might introduce a spatial dimension to the model so that agents can choose to realise a directed search and return to the place where they had previously encountered a trading partner. Trade will cluster in certain locations and a process of endogenous formation of a market will then emerge.

References

- Bagwell, K. and Ramey, G. (1994). Coordination economies, advertising and search behavior in retail markets. *American Economic Review*, 84:498–517.
- Binmore, K. and Samuelson, L. (1997). Muddling through: Noisy equilibrium selection. Journal of Economic Theory, 74:235–265.
- Binmore, K., Samuelson, L., and Vaughan, R. (1995). Musical chairs: Modelling noisy evolution. Games and Economic Behavior, 11:1–35.
- Bomze, I. (1986). Non-cooperative two-person games in biology: A classification. International Journal of Game Theory, 15:31–57.
- Bulow, J., Geanakoplos, J., and Klemperer, P. (1985). Multimarket oligopoly: strategic substitutes and complements. *Journal of Political Economy*, 93:425–39.
- Cabrales, A. (1996). Stochastic replicator dynamics. Mimeo, Mimeo Universitat Pompeu Fabra.
- Cooper, R. (1998). Business cycles: Theory, evidence and policy implications. Scandinavian Journal of Economics, 100:213–237.
- Cooper, R. and Haltinwanger, J. (1996). Evidence on macroeconomic complementarities. *Review of Economics and Statistics*, 78:78–93.
- Cooper, R. and John, A. (1988). Coordinating coordination failures in keynesian models. *Quarterly Journal of Economics*, 103:441–463.
- Diamond, P. A. (1982). Aggregate demand management in search equilibrium. Journal of Political Economy, 90:881–894.
- Diamond, P. A. (1984). A search equilibrium approach to the foundations of macroeconomics. MIT Press, Cambridge, Mass.
- Farrell, J. and Saloner, G. (1986). Installed base and compatibility: Innovation, product preannouncements, and predations. *American Economic Review*, 76:940–955.
- Foster, D. and Young, P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38:219–232.
- Freidlin, M. and Wentzell, A. (1984). Random Perturbations of Dynamical Systems. Springer-Verlag, New York.

- Fudenberg, D. and Harris, C. (1992). Evolutionary dynamics with aggregate shocks. Journal of Economic Theory, 57:420-441.
- Howitt, P. (1985). Transaction costs in the theory of unemployment. American Economic Review, 75:88–101.
- Kandori, M., Mailath, G., and Rob, R. (1993). Learning, mutation and long run equilibria in games. *Econometrica*, 61:29–56.
- Katz, M. and Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94:822–841.
- Kiyotaki, N. (1988). Implications of multiple expectational equilibria under monopolistic competition. Quarterly Journal of Economics, 103:695–714.
- Matsusaka, J. G. and Sbordone, A. M. (1995). Consumer confidence and economic fluctuations. *Economic Enquiry*, 33:296–318.
- Matsuyama, K. (1995). Complementarities and cumulative processes in models of monopolistic competition. *Journal of Economic Literature*, 33:701– 729.
- Milgrom, P. and Roberts, J. (1990a). The economics of modern manufacturing: Technology, strategy and organization. *American Economic Review*, 80:511–528.
- Milgrom, P. and Roberts, J. (1990b). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58:1255–1277.
- Milgrom, P. and Roberts, J. (1995). Complementarities and fit: Strategy, structure and organizational change in manufacturing. *Journal of Account*ing Economics, 19:179–208.
- Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econo*metrica, 62:157–180.
- Nachbar, J. (1990). Evolutionary selection dynamics in games: convergence and limit properties. *International Journal of Game Theory*, 19:59–89.
- Pissarides, A. C. (1990). Equilibrium unemployment theory. Blackwell, Oxford.
- Redding, S. (1996). The low-skill, low-quality trap: Strategic complementarities between human capital and r & d. *Economic Journal*, 106:458–470.

- Ritzberger, K. and Weibull, J. W. (1995). Evolutionary selection in normalform games. *Econometrica*, 63:1371–1399.
- Samuelson, L. (1994). Stochastic stability in games with alternative best replies. *Journal of Economic Theory*, 64:35–65.
- Samuelson, L. (1998). Evolutionary Games and Equilibrium Selection. MIT Press, Cambridge Mass.
- Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory*, 57:363–391.
- Schlag, K. (1998). Why imitate and if so how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*.
- Topkis, D. M. (1979). Equilibrium points in nonzero-sum N-person submodular games. SIAM Journal of Control and optimization, 17:773–787.
- Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal* of Mathematical Economy, 19:305–321.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge, Mass.
- Young, P. (1993). Evolution of conventions. *Econometrica*, 61:57–84.