Adv. Complex Systems (1999) **1**, 329–363

Generalized Replicator Dynamics as a Model of Specialization and Diversity in Societies

Vivek S. Borkar[†]
Department of Computer Science and Automation,
Indian Institute of Science, Bangalore 560 012, India

Sanjay Jain[‡]
Centre for Theoretical Studies,
Indian Institute of Science, Bangalore 560 012, India

Govindan Rangarajan[§]
Department of Mathematics and Centre for Theoretical Studies,
Indian Institute of Science, Bangalore 560 012, India

ABSTRACT. We consider a generalization of replicator dynamics as a non-cooperative evolutionary game-theoretic model of a community of N agents. All agents update their individual mixed strategy profiles to increase their total payoff from the rest of the community. The properties of attractors in this dynamics are studied. Evidence is presented that under certain conditions the typical attractors of the system are corners of state space where each agent has specialized to a pure strategy, and/or the community exhibits diversity, i.e, all strategies are represented in the final states. The model suggests that new pure strategies whose payoff matrix elements satisfy suitable inequalities with respect to the existing ones can destabilize existing attractors if N is sufficiently large, and be regarded as innovations that enhance the diversity of the community.
KEYWORDS: Generalized replicator dynamics, evolutionary game theory, interacting agents, diversity, specialization.

[†] E-mail: *borkar@csa.iisc.ernet.in*
[‡] Also at Jawaharlal Nehru Centre for Advanced Scientific Research, Bangalore, and Associate Member of ASICTP, Trieste. E-mail: *jain@cts.iisc.ernet.in*
[§] Also at Jawaharlal Nehru Centre for Advanced Scientific Research, Bangalore. E-mail: *rangaraj@math.iisc.ernet.in*

1. Introduction

Replicator dynamics [Hofbauer and Sigmund, 1988; Weibull, 1995, Samuelson, 1996; Vega-Redondo, 1996; and references therein] is a standard model in evolutionary biology describing the dynamics of growth and decay of a number of species under selection pressure. Formally, consider s species in an interactive environment, each genetically programmed to a particular behaviour with which it is identified. Tag the possible behaviours as 'strategies' belonging to the set $S = \{1, 2, \ldots, s\}$ in the sense of game theory. Let $A = [[a_{ij}]]$, $1 \leq i, j \leq s$, denote the payoff matrix with the following interpretation: the i-th species (playing strategy i) receives a payoff ($\approx$ contribution to its fitness) of a_{ij} on interaction with species j. The rate of increase of the population $y_i(t)$ of species i at time t is assumed to be proportional to $y_i(t)$ and the birth rate minus the death rate. The latter is assumed constant, the former a constant modulated by an additive term equal to the average (current) payoff to i. Letting $x_i(t) = y_i(t)/(\sum_j y_j(t))$ denote the population share of i, simple algebra (see, e.g., Weibull, 1995) leads to replicator dynamics

$$\dot{x}_i(t) = x_i(t)[\sum_j a_{ij}x_j(t) - \sum_{k,j} x_k(t)a_{kj}x_j(t)]. \tag{1.1}$$

The description makes the underlying biological formalism clear. A species is not allowed to change its strategy or 'genetic traits'. That is, 'ontogenetic adaptation' whereby individuals adapt to environment in their lifetime, is not permitted. Eq. (1.1) stands for 'phylogenetic adaptation' by the species as a whole in response to birth-death mechanism modulated by selection pressure. This makes eminent biological sense, because ontogenetic adaptation is not inherited, phylogenetic is.

More recently, however, the model has crossed the boundaries of biology. It has been taken up by economists as a model of economic learning. The motivation comes from the problem of equilibrium selection in game theory. Having concurred that Nash equilibrium is the most reasonable equilibrium concept for noncooperative games, one is confronted with the problem of choosing a 'natural' one from the overabundance thereof in a typical scenario. Many refinements were proposed, culminating in an elaborate theory of equilibrium selection [Harsanyi and Selten, 1988]. A recent departure from this trend has been to leave the static equilibrium framework in favour of dynamic models of disequilibrium and look for their asymptotic equilibria. Replicator dynamics and its variants have been a prominent front runner in this activity [see e.g. Mailath, 1992; Fudenberg and Levine, 1995]. The attractive feature of this model of evolutionary economics is that it builds in bounded rationality (agents are not perfect optimizers with perfect knowledge, infinite memory and computational ability, but make simple decisions based on immediate payoffs) and inertia (the adaptation takes place in incremental steps). Though originally meant to model 'phylogenetic learning' in biology, economists have successfully recovered it as a limiting case of models of individual, or ontogenetic learning [Borgers and Sarin, 1993; Ritzberger and Weibull, 1995] thus deriving 'macrobehaviour' from 'micromotives' [Schelling, 1978]. Prompted by this, there have been alternative models of evolutionary learning in economics [Young, 1993; Ellison and Fudenberg, 1995, for example].

Another motivation for considering dynamic models comes from the study of complex adaptive systems. A major question in this area is how do complex organizations emerge spontaneously in certain systems (e.g., chemical, biological and social systems). Implicit in the question is a notion of time evolution. One is interested the mechanisms whereby, starting from initial conditions where there is no organizational complexity, the system can end up in a state where an organizational character is discernible in which specific, nonrandom interaction between components of the system resulting in an overall coherence has emerged. In this context generalizations of replicator dynamics have been considered in the literature (see Stadler, Fontana and Miller, 1993, and references therein).

In this paper we discuss another generalization [Borkar, Jain and Rangarajan, 1997], motivated by somewhat different concerns than earlier works. Mathematically, our model is simply a multipopulation variant of replicator dynamics [Weibull, 1994, Chapter 5]. Here, N distinct populations, each with its strategy profile, evolve under selection pressure. Let $p_i^\alpha(t)$ denote the fraction of α-th population that plays strategy i. The dynamics then is

$$\dot{p}_i^\alpha(t) = p_i^\alpha(t)\left[\sum_{\beta\neq\alpha}\sum_j a_{ij}p_j^\beta(t) - \sum_{\beta\neq\alpha}\sum_{k,j} p_k^\alpha(t)a_{kj}p_j^\beta(t)\right],\ 1 \leq \alpha \leq N,\ 1 \leq i \leq s. \tag{1.2}$$

Our interpretation, however, is different. We view Eq. (1.2) as a model of ontogenetic learning wherein the superscript α, $1 \leq \alpha \leq N$, stands for α-th agent, $p_i^\alpha(t)$ being the probability with which she plays strategy i at time t. Thus $p^\alpha(t) \equiv (p_1^\alpha(t), p_2^\alpha(t), \ldots, p_s^\alpha(t))^T$ is her 'mixed strategy' at time t and $\sum_{\beta\neq\alpha}\sum_{k,j} p_k^\alpha(t)a_{kj}p_j^\beta(t)$ represents the payoff received by α from the rest of the community at time t. Eq. (1.2) then shows how agents adapt their strategies in response to payoffs received. The interpretation of the r.h.s. of Eq. (1.2) is that each agent increases the weight of those pure strategies that do better than her current (mixed) strategy in the environment of the current strategies of other agents, and decreases the weight of those that do worse, in proportion to the difference. The community as a whole evolves as all the individuals simultaneously make their updates. To emphasize this shift in paradigm from populations to individuals, we dub Eq. (1.2) 'generalized replicator dynamics' or GRD, rather than 'multipopulation replicator dynamics'. For contrast, Eq. (1.1) will be called 'pure replicator dynamics' or PRD.

Not only our interpretation, but the focus of our study is also different from that of Weibull [1994, Chapter 5]. Specifically, our primary aim will be to identify and analyse two key phenomena, which we dub 'specialization' and 'diversification'. By 'specialization', we refer to the situation in which every agent plays a pure strategy, i.e., for each α, the probability vector $p^\alpha(t)$ is concentrated at a single point of the strategy space. In other words, while having the choice of choosing from the set of mixed strategies, thereby apportioning her resources (energy, time etc) on more than one activity, she chooses to put them all into a single activity. The kind of situations we are interested in are those in which this phenomenon goes hand in hand with another, that of 'diversification'. We say that diversification occurs in a community if for each possible strategy choice, there is at least one agent at any time who opts for it with positive probability.

That is, no available strategy is completely discarded by the community as a whole. Note that specialization is not a prerequisite for diversification. We do not require that agents in a diversified society use pure strategies in order for it to be considered as such. In fact, we do not even insist that they necessarily play distinct strategies. The society would still be deemed as being diversified if all of them played the same mixed strategy that puts positive probability mass on every strategy option. We do, however, look for situations where both specialization and diversification occur at the same time: while each agent puts all her eggs in one basket, no basket is left empty. Needless to say, we expect such behaviour only when the payoff matrix A satisfies additional constraints. For example, if a strategy choice is strictly dominated by every other, one would indeed expect the society as a whole to opt out of it. Thus a significant part of the effort will be directed towards precisely nailing down what those conditions on A might be.

The motivation for such an exercise is almost self-evident. We wish to capture qualitatively a phenomenon that recurs across a variety of evolutionary paradigms. If one looks at the evolution of human society from a primitive hunter-gatherer society to a modern town, one sees progressive specialization of essential tasks among individuals while no essential task is being completely discarded by the society as a whole. If one looks at the evolution of species on earth, one sees genetic strands diversifying to fill in specific ecological niches till their common ancestry is no longer obvious. In economics, the evolution from old manufacturing and trading patterns to a modern corporation is marked by extreme specialization and decentralization of tasks, while the corporation as a whole concurrently does a multitude of tasks that go into the meeting of its final objective.

This is just by the way of motivation. We delegate the issues of interpretation of the mathematical model to the discussion at the end of the paper and return to the mathematical formalism per se.

The quantity $\bar{x}_i \equiv (1/N)\sum_{\alpha=1}^{N} p_i^\alpha$ is the population average of the probability that the strategy i is being played by the entire collection of agents and can be interpreted as the relative weight of the i^{th} strategy in the population. One might wonder whether this averaged out quantity satisfies the PRD equation when the individual p_i^α satisfy GRD. The answer is no even in the limit of large number of agents, as will be discussed later. Nevertheless, we find that in the large N limit GRD contains a fair amount of "memory" about some of the PRD structures, in particular the interior equilibrium point of PRD, as we shall note later. The important point being made here, however, is that although our concept of diversification bears some similarity with concepts such as persistence and permanence for the PRD [Hofbauer and Sigmund, 1988], it is not subsumed or implied by these.

This paper is organized as follows: Section 2 reviews some known properties of PRD and GRD attractors, sets the notation and outlines the similarities and differences between PRD and GRD. This is followed by further results on the non-corner attractors of GRD. Section 3 identifies the conditions for the existence of maximal diversity in GRD, or preservation of all strategies in the population. Here we note the importance of large N in promoting diversity in GRD. Explicit sufficient conditions for $s = 2$ and 3 are given, as well as partial results for higher s. Section 4 studies corner solutions in GRD which correspond to specialization in the community. Some theorems concerning the coexistence of specialization and diversity in the community are proved. Finally in section 5 we summarize our results from the previous sections, discuss their possible implications concerning specialization, diversity and innovation in economic communities, and mention some future directions.

2. NON-CORNER ATTRACTORS OF GRD

2.1. NON-CORNER EQUILIBRIA OF PRD

In this subsection we recall some well known facts about PRD (see Hofbauer and Sigmund, 1988, chapters 13, 19) that will be useful later and set some notation. Let J denote the simplex of s-dimensional probability vectors, $J = \{\mathbf{x} = (x_1, \cdots, x_s)^T \in \mathbf{R}^s | x_i \geq 0,\ \sum_{i=1}^{s} x_i = 1\}$. J is the full configuration space of PRD dynamics, and is invariant under it, i.e., an s-dimensional probability vector remains so under evolution given by (1.1). Let L denote a subset of S with at most $s-1$ elements. Then the set $F_L = \{\mathbf{x} \in J | x_k = 0, k \in L\}$ is called a face of J. In particular if $L = \{k\}$, the corresponding face $F_{\{k\}}$ (denoted F_k for brevity) is such that on this face the k^{th} strategy is not represented in the population; the full diversity of available strategies is lost. All the faces F_k together constitute the boundary of J. The relative boundary of each F_k in turn is a union of the 'subfaces' $F_{kj} \equiv F_{\{k,j\}} = F_k \cap F_j$, $\partial F_k = \cup_{j \neq k} F_{kj}$, and so on. Finally, if L has exactly $s-1$ elements, say $L = S - \{i\}$, the corresponding face F_L is a singleton, and will be called the 'corner' C_i of J; only strategy i survives in the population at C_i. All faces, subfaces, corners are invariant under PRD; if $x_k = 0$ at some time, it remains so thereafter. It follows that corners are trivially equilibrium points.

We are interested in the interior of J, i.e., $J^\circ = J - \cup_{k=1}^{s} F_k$. (The superscript $^\circ$ for a set will always denote its relative interior.) For trajectories that remain in J°, all strategies are always present, since $0 < x_i < 1$ for all i. A point $\mathbf{x} \in J^\circ$ that is a stationary point of (1.1) will be called an interior equilibrium point (IEP) of PRD. $\mathbf{x}$ is an IEP of PRD only if the $s+1$-dimensional vector $\mathbf{X} = (x_0, \mathbf{x})$ satisfies

$$B\mathbf{X} = \mathbf{E_0}, \tag{2.1}$$

where B is the $s+1$-dimensional matrix

$$B = \begin{pmatrix} 0 & 1 & 1 & \cdots & 1 \\ -1 & & & & \\ -1 & & A & & \\ \vdots & & & & \\ -1 & & & & \end{pmatrix} \tag{2.2}$$

and $\mathbf{E_0}$ is the $s+1$-dimensional unit vector $(1, 0, 0, \cdots, 0)^T$. (In (2.1) $\mathbf{X}$ is represented as an $s+1$-dimensional column vector). To see this, note that the 0^{th} component of (2.1) is just the normalization condition $\sum_{i=1}^{s} x_i = 1$, and the i^{th} component is the equation $\sum_{j=1}^{s} a_{ij} x_j - x_0 = 0$. The latter statement means that the payoff for all strategies is x_0, independent of strategy, which, together with the normalization condition is equivalent to the vanishing of the bracket on the r.h.s. of (1.1). The solution to (2.1) is

$$x_i = B^{-1}{}_{i0} = u_i/\det B, \quad u_i \equiv \text{Cofactor of } B_{0i}, \tag{2.3}$$

$$x_0 = B^{-1}{}_{00} = \det A/\det B, \tag{2.4}$$

provided $\det B \neq 0$. Note from the structure of B that

$$\det B = \sum_{i=1}^{s} u_i. \tag{2.5}$$

x_i being probabilities must be greater than or equal to zero, and for an interior point, must be positive. This is guaranteed from (2.3) and (2.5) if

A1: $u_i \neq 0 \;\; \forall \;\; i$, and all u_i have the same sign.

A1 is a necessary and sufficient condition (on the payoff matrix A) for PRD to have an isolated IEP, which, if it exists, is unique and given by (2.3). Sufficiency is evident from the above. To see that **A1** is also necessary, observe that $\det B = 0$ would imply that the solution of the linear equation (2.1) is not isolated, since we could always add to **X** a vector proportional to a zero eigenvector of B.

Dynamics within F_k can be studied using the submatrix $A^{(k,k)}$ of the full payoff matrix A ($M^{(p,q)}$ denotes the matrix obtained from a matrix M by deleting its p^{th} row and q^{th} column; we will follow the convention for labelling the indices that if the indices of M go over a set T, then the row indices of $M^{(p,q)}$ will go over the set $T - \{p\}$, and the column indices over the set $T - \{q\}$). Denote $u_i^{(k)} \equiv$ Cofactor of$(B^{(k,k)})_{0i}$, $i \in S - \{k\}$, and define the following condition:

A1$_k$: $u_i^{(k)} \neq 0 \;\; \forall \;\; i \neq k$, and all $u_i^{(k)}$ have the same sign.

It is clear from the above that if **A1**$_k$ is satisfied, then the relative interior of F_k, $F_k^\circ = F_k - \partial F_k$ has an equilibrium point denoted $\mathbf{x}^{(k)} = (x_1^{(k)}, \ldots, x_s^{(k)})$ with

$$x_i^{(k)} = \frac{u_i^{(k)}}{\det B^{(k,k)}}, \quad i \neq k, \quad x_k^{(k)} = 0,. \tag{2.6}$$

with the payoff to all agents at this equilibrium point given by

$$x_0^{(k)} = \frac{\det A^{(k,k)}}{\det B^{(k,k)}}. \tag{2.7}$$

A1$_k$ is sufficient but not a necessary condition for F_k° to have an equilibrium point. An equilibrium point can also arise if $u_k = 0$ and all other u_i are non-vanishing and have the same sign. However the latter equilibrium point is non-generic, i.e., it happens only when the components of A satisfy an equation ($u_k = 0$). This situation is structurally unstable: it can be destroyed by a small perturbation of A. Thus generically, F_k° has at most one equilibrium point (one, given by (2.6) if **A1**$_k$ holds, and none if it is violated). The generic equilibrium point is characterized by equal payoffs to all strategies except k, which is just ignored, and the non-generic one by the fact that even the k^{th} strategy receives the same payoff as all other strategies. An analogous situation holds for the equilibrium points in the interior of subfaces F_{kj}, etc.

For all trajectories that do not converge to the boundary of J, the barycentre of every invariant probability measure supported in the interior is the IEP of PRD. To see this, define $\bar{x}_i = \lim_{T\to\infty} \frac{1}{T}\int_0^T dt\, x_i(t)$ and $c = \lim_{T\to\infty} \frac{1}{T}\int_0^T dt \sum_{k,j=1}^{s} x_k(t) a_{ij} x_j(t)$ along some common convergent subsequence. From (1.1) it follows that $\dot{x}_i/x_i = (Ax)_i - x^T Ax$. Integrating this w.r.t. t from zero to T, dividing by T and taking the limit along appropriate subsequence gives

$$\lim_{T\to\infty} \frac{1}{T}[\ln x_i(T) - \ln x_i(0)] = (A\bar{x})_i - c. \tag{2.8}$$

Since by assumption x_i is bounded away from zero, the l.h.s. vanishes. Therefore $(A\bar{x})_i$ is independent of i, which, together with the normalization condition on $\bar{x}$ means that $\bar{x}_i$ must coincide with the IEP (2.3). Similarly, for generic A, for all trajectories in F_k that do not converge to its relative boundary, the barycentre of every invariant probability measure supported in F_k° is given by (2.6).

In the remainder of this section we will show that above results for PRD have appropriate generalizations in GRD, and also discuss some departures.

2.2. GRD PRELIMINARIES

2.2.1. NOTATION, DEFINITION OF SPECIALIZATION AND DIVERSITY

The configuration space of GRD will be denoted $J^N = \Pi_{\alpha=1}^N J^{(\alpha)}$ where $J^{(\alpha)}$ is a copy of J for the α^{th} agent. A point of J^N will be denoted $p = (\mathbf{p}^1, \mathbf{p}^2, \dots, \mathbf{p}^N)$, where $\mathbf{p}^\alpha = (p_1^\alpha, p_2^\alpha, \dots, p_s^\alpha) \in J^{(\alpha)}$. For many purposes it will be convenient to think of J^N to be embedded in $\mathbf{R}^{Ns} = \Pi_{\alpha=1}^N \mathbf{R}^s(\alpha)$ where $\mathbf{R}^s(\alpha)$ is a copy of $\mathbf{R}^s$ for the α^{th} agent and $J^\alpha \subset \mathbf{R}^s(\alpha)$. Then we can write $p = (p_1, p_2, \dots, p_{Ns}) \equiv (p_b)$, $p_b \geq 0$. b runs over the index set $S_N = \{1, 2, \dots, Ns\} = \cup_{\alpha=1}^N S(\alpha)$, where the set $S(\alpha) = \{(\alpha-1)s+1, (\alpha-1)s+2, \dots, \alpha s\}$ will be called the α^{th} block (of indices). $\hat{e}_b$ will denote the unit vectors along the axes of R^{Ns}, the d^{th} component of $\hat{e}^b$ is $(\hat{e}^b)_d = \delta_d^b$. Thus $J^N = \{p \in \mathbf{R}^{Ns} |\ p_b \geq 0,\ \sum_{b\in S(\alpha)} p_b = 1\ \ \forall\ \ \alpha = 1, 2, \dots, N\}$, and $J^{(\alpha)} = \{p \in \mathbf{R}^{Ns} |\ p_a \geq 0\ \forall a,\ p_b = 0,\ b \notin S(\alpha),\ \sum_{b\in S(\alpha)} p_b = 1\}$. The notation $b = (\alpha, i)$ will denote that b is the i^{th} element in the α^{th} block.

We now define a face of J^N. Let L be any subset of S_N such that it contains at most $s-1$ indices from every block. Then the set $F_L = \{p \in J^N | p_b = 0,\ b \in L\}$ will be called a face of J^N. A face $F_{L'}$ will be called a 'subface' of F_L if L is a proper subset of L'. Clearly if $L = \emptyset$, $F_L = J^N$. If L is of order $N(s-1)$, i.e., it has exactly $s-1$ elements from each block, then F_L is a singleton and will be called a corner of J^N. If for some α, $\mathbf{p}^\alpha$ goes to a corner C_i of $J^{(\alpha)}$ ($p_j^\alpha = \delta_{ij}$), this of course means that the agent α has opted for the pure strategy i, and we say that α has 'specialized' in strategy i. At a corner of J^N, every agent has specialized to some strategy or the other and we say that the community is 'fully specialized'. J^N is invariant under GRD, as the L^1 norm and positivity of every $\mathbf{p}^\alpha$ is preserved under (1.2). Every face of J^N is invariant under GRD, as $p_i^\alpha(t) = 0\ \forall\ t$ if it is zero for some t. Therefore corners are trivially equilibrium points. This does not, however, mean that they are stable. If corners of J^N are the only asymptotically stable attractors of the dynamics defined by (1.2), we will say that the latter exhibits ***specialization***. For then, the community will always end up being fully specialized, starting from (a small perturbation of) any initial condition.

If $L = \{k, s+k, 2s+k, \ldots, (N-1)s+k\}$ for some $k \in S$ the corresponding face F_L is one where every agent has opted out of the k^{th} strategy. Such faces will be particularly important for us and we will denote them as F_k. They are analogues of the F_k defined in the previous subsection, where a particular strategy k becomes extinct from the population. At a face F_k the full diversity of strategies is lost. Conversely the complement of the set $\cup_{k=1}^{s} F_k$ in J^N corresponds to the situation where every strategy is represented in the community and we denote this set by J_D^N, with the subscript D denoting the fact that the full 'diversity' of strategies is maintained. Note that we use the word 'diversity' not to signify the variation between individual agents, but as a measure of the size of set of strategies being pursued. Indeed we can have no variation but full diversity if all agents pursue the same mixed strategy: for all α, $\mathbf{p}^\alpha = \mathbf{c} \in J^\circ$. When $\mathbf{p}^\alpha$ is independent of α, the community is completely 'homogeneous' since all agents are doing the same thing. Also we remark that the community can be fully specialized and diversified at the same time: each agent chooses a pure strategy and every strategy is chosen by some agent or the other.

In order to maintain the full diversity of strategies in the society, no strategy k should be dynamically opted out of by all agents. That is, no trajectory that starts in the interior of J^N should end up in any of the faces F_k. Better still, any trajectory in a face F_k should become unstable with respect perturbations that take it into the interior of J^N. The basic idea is that while under the dynamics (1.2) a trajectory in F_k will always remain in it, in practice in a real society, agents will always be subject to influences that draw their attention to strategies that do not exist in the population, but are valid strategies within the capabilities of the agents. Such influences can be thought of as perturbing the society away from the face F_k. If F_k is unstable to small perturbations, then this perturbation will grow. If all faces F_k are unstable, no strategy can in practice become extinct from the population. We will say that the dynamics (1.2) exhibits ***diversity*** if every trajectory in each of the faces F_k of J^N becomes unstable at some time or the other with respect to small perturbations that deform it away from F_k. This notion of diversity is analogous to the notions of permanence [Schuster, Sigmund and Wolff, 1979], persistence [Butler, Freedman and Waltman, 1986, and references therein], etc., which have been defined in PRD to study the extinction of strategies (see Hofbauer and Sigmund, 1988 for definitions).

It is evident that the boundary of J^N is given by the union of all faces in which exactly one of the p_b's is zero, $\partial J^N = \cup_{b=1}^{Ns} F_{\{b\}}$. In general the relative boundary of F_L is given by $\partial F_L = \cup_{b \in \bar{L}} \{p \in F_L | p_b = 0\}$, where $\bar{L}$ is the complement of L in S_N. The relative interior of a face F_L is defined as $F_L^\circ = F_L - \partial F_L$. Note that unlike in PRD here the boundary of J^N not the same as $\cup_{k=1}^{N} F_k$; the latter is a subset of ∂J^N. Similarly $\cup_{j \neq k} F_{kj}$ (where $F_{kj} = F_k \cap F_j$) is a subset of the boundary of F_k.

In general, $p \in J^N$ is an equilibrium point for Eq. (1.2) if for every $b = (\alpha, i) \in S_N$, either $p_b = 0$, or,

$$\sum_{\beta \neq \alpha} \sum_{j} a_{ij} p_j^\beta = \sum_{\beta \neq \alpha} \sum_{k,j} p_k^\alpha a_{kj} p_j^\beta. \tag{2.9}$$

Note that the right hand side of the above equation is independent of i. Combining this with the fact that $(p_1^\alpha, p_2^\alpha, \cdots, p_s^\alpha)$ is a probability vector, one can replace Eq. (2.9) by: $\sum_{\beta \neq \alpha} \sum_j a_{ij} p_j^\beta$ is the same for all i for which $p_i^\alpha > 0$. At an equilibrium point an agent receives the same payoff for all the pure strategies she has not opted out of.

2.2.2. Review of some known results for GRD

We now recall some known facts about equilibria (and other attractors) of GRD from Chapter 5, Weibull (1995). (See this reference for relevant definitions.)

Theorem 2.1 If strategy i is either dominated or iteratively dominated, $p_i^\alpha(t) \to 0 \ \forall \alpha$.

Theorem 2.2 All interior equilibria and Liapunov stable equilibria of Eq. (1.2) are Nash equilibria. Also, if an interior trajectory of Eq. (1.2) converges to a point, the same must be a Nash equilibrium.

Theorem 2.3 Any compact set in the relative interior of a face cannot be asymptotically stable.

Note that strict Nash equilibria are perforce pure strategy Nash equilibria and therefore correspond to corner solutions which we identify as 'specialization' in behavioural sense.

2.2.3. Differences between PRD and GRD

Note that (1.1) is of the form $\dot{x}_i = x_i(f_i(\mathbf{x}) - \bar{f}(\mathbf{x}))$, with $f_i(\mathbf{x}) = \sum_{i \in S} a_{ij} x_j$ and $\bar{f}(\mathbf{x}) = \mathbf{x}^T \mathbf{f}(\mathbf{x})$. The GRD equation (1.2) does not have quite this form. Let $\bar{A}$ denote the $Ns \times Ns$ matrix given by

$$\begin{pmatrix} 0 & A & A & \cdots & A \\ A & 0 & A & \cdots & A \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ A & A & \cdots & \cdots & 0 \end{pmatrix} \tag{2.10}$$

where 0 is a $s \times s$ zero matrix. Then indeed (1.2) can be written as $\dot{p}_b = p_b(f_b(p) - \bar{f}(p))$, with $f_b(p) = \sum_{c \in S_N} \bar{A}_{bc} p_c$. But now $\bar{f}(p)$ is no longer given by $p^T f(p)$, but by $p^{(\alpha_b)^T} f(p)$, where $p^{(\alpha_b)}$ is the projection of p onto the block α_b to which b belongs ($p_c^{(\alpha_b)} = p_c$ if c is in the same block as b and zero otherwise).

Second, note that in GRD every agent plays against other agents only, not against herself. That is, if agent α were to play pure strategy i, her payoff,