Project Number 001907

DELIS

Dynamically Evolving, Large-scale Information Systems

Integrated Project

Member of the FET Proactive Initiative **Complex Systems**

Deliverable D4.3.1

An evaluation of the currently known ”evolutionary game theory” approaches

Start date of the project:	January 2004
Duration:	48 months
Project Coordinator:	Prof. Dr. math. Friedhelm Meyer auf der Heide Heinz Nixdorf Institute, University of Paderborn, Germany

Due date of deliverable:	June 2004
Actual submission date:	October 2004
Dissemination level:	PU – public

Work Package 4.3:	An evaluation of the currently known "evolutionary game theory" approaches
Participants:	Department of Computer Science, University of Bologna (UniBO), Italy Computer Science Institute (CTI), Patras, Greece University of Cyprus (UCY)

Authors of deliverable:	Ozalp Babaoglu (babaoglu@cs.unibo.it) David Hales (dave@davidhales.com) Edoardo Mollona (emollona@cs.unibo.it) Paul Spirakis (spirakis@cti.gr) Spyros Kontogiannis (kontog@cti.gr) Panagiota Panaggopoulou Marios Mavronicolas (mavronic@cs.ucy.ac.cy)

Preface

This report comprises the complete D4.3.1 deliverable as specified for work-package WP4.3 in Subject SP4 of the DELIS (Dynamically Evolving Large-scale Information Systems) Integrated Project.

The essential goal of the DELIS project is to understand, predict, engineer and control large evolving information systems. These are often arranged in the form of dynamic networks (for example: the Internet or Peer-2-Peer applications running on-top of such physical infrastructure).

WP4.3 brings together a team of specialists in Game Theory, Evolutionary Economics and Organisational Science and Agent-Based Modelling.

The purpose of this report is to outline the relevant work and state-of-the-art in each field (including limitations) with application to DELIS-like systems and then to outline integrative on-going problems, ideas and projects that bring these techniques together to advance the state-of-the-art.

In the following chapters we overview techniques and ideas from classical game theory, evolutionary game theory, individual and agent-based computational modelling, and organisation science.

We indicate their limitations and propose on-going work to integrate techniques to overcome some of those limitations. In this context we identify some ambitious (yet realistic) long-term goals.

In this early (6 month) deliverable the participants in WP4.3 report on planned and on-going collaborative and supportive projects and work.

The next scheduled deliverable for WP4.3 is at the 18 month stage.

This report can therefor be viewed as an outline of the on-going work that will take place within WP4.3 over the next 12 months (though it should not read as exhaustive or strict).

Contents

Chapter 1

The game-theoretic approach

1.1 Classical Game Theory

Game theory can be defined as the study of mathematical models of conflict and cooperation between intelligent rational decision-makers. Game theory provides general mathematical techniques for analyzing situations in which two or more individuals make decisions that will influence one another's welfare.

The object of study in game theory is the *game*, which is a formal model of an interactive situation involving two or more individuals. The individuals involved in a game are called the *players*. Game theory generally makes two basic assumptions about players: they are rational and they are intelligent. A player is *rational* if he always chooses an action which gives the outcome he most prefers, given what he expects his opponents to do. In game theory, building on the fundamental results of decision theory, we assume that each player's objective is to maximize the expected value of his own payoff, which is measured in some *utility* scale. A player is *intelligent* if he knows everything that we know about the game and he can make any inferences about the situation that we can make.

1.1.1 History of Game Theory

Some game-theoretic ideas can be traced to the 18th century, but the major development of the theory began in the 1920s with the work of the mathematician Emile Borel and the polymath John von Neumann. A decisive event in the development of the theory was the publication in 1944 of the book *Theory of Games and Economic Behavior* by John von Neumann and Oskar Morgenstern. This book provided much of the basic terminology and problem setup that is still in use today.

In 1950, John Nash demonstrated that finite games always have an equilibrium point, at which all players choose actions which are best for them given their opponents' choices. This central concept of noncooperative game theory has been a focal point of analysis since then. In the 1950s game-theoretic models began to be used in economic theory and political science, and psychologists began studying how human subjects behave in experimental games. Since the 1970s, it has driven a revolution in economic theory. Additionally, it has found applications in sociology and psychology, and established links with evolution and biology. Subsequently, game theoretic methods have come to dominate microeconomic theory and are used also in many other fields of economics and a wide range of other social and behavioural sciences.

Game theory received special attention in 1994 with the awarding of the Nobel prize in economics to John Nash, John Harsanyi and Reinhard Selten.

At the end of the 1990s, a high-profile application of game theory has been the design of auctions. Prominent game theorists have been involved in the design of auctions for allocating rights to the use of bands of the electromagnetic spectrum to the mobile telecommunications industry.

1.1.2 Game-Theoretic Models

Games can be described formally at various levels of detail. Once we define the set of players, we may distinguish between between two types of models: those in which the sets of possible actions of *individual* players are primitive and those in which the sets of possible joint actions of *groups* of players is primitive. Models of the first type are referred to as **noncooperative**, while those of the second type are referred to as **cooperative** or **coalitional**.

A **strategic (or, normal form)** game is a model of a situation in which each player chooses his plan of action once and for all, and all players' decisions are considered to be made simultaneously (that is, when choosing a plan of action each player is not informed of the plan of action chosen by any other player). By contrast, the model of an **extensive form** game specifies the possible orders of events; each player can consider his plan of action not only at the beginning of the game but also whenever he has to make a decision.

A game with **perfect information** models a situation in which the participants are fully informed about each others' moves, while in a game with **imperfect information** the participants may be imperfectly informed.

1.1.3 Strategic-Form Games

Of the different forms for games that were introduced in the preceding subsection, the simplest conceptually is the strategic (or normal form) game.

Definition 1.1.1 *A* **strategic game** $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ *consists of*

- *a finite set N of players*

- *for each player $i \in N$ a nonempty set C_i of allowable actions available to player i. C_i is usually called the action set of player i.*

- *for each player $i \in N$ a utility function $u_i : C \to \mathbb{R}$.*

Here $C \equiv \times_{i \in N} C_i$ denotes the set of all possible combinations of strategies (usually called **pure strategies profiles** or **configurations**) that may be chosen by all the players.

Given any strategic game $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$, a **mixed strategy** for any player i is a probability distribution over C_i. We let $\Delta(C_i)$ denote the set of all possible mixed strategies for player i. To emphasize the distinction from mixed strategies, the actions in C_i are called **pure strategies**.

A **mixed strategies profile** is any vector that specifies one mixed strategy for each player, so the set of all mixed strategy profiles is $\times_{i \in N} \Delta(C_i)$. That is, $\mathbf{p} \equiv \left(p_1, \ldots, p_{|N|}\right)$ is a mixed strategy profile in $\times_{i \in N} \Delta(C_i)$ iff, for each player i and each pure strategy c_i in C_i, $p_i(c_i) \in \mathbb{R}_{\geq 0}$, is the probability that player i would choose c_i under $\mathbf{p}$, and so $\sum_{c_i \in C_i} p_i(c_i) = 1$, $\forall i \in N$.

Although we may be uncertain about what profile of pure strategies will be chosen by players when Γ is played, Bayesian decision theory guarantees that there exists some probability distribution over the set of pure strategy profiles $C = \times_{i \in N} C_i$ that quantitatively expresses our beliefs about the players' strategy choices. Furthermore, as we assume that all the players choose their strategies simultaneously, our beliefs about the game should correspond to some mixed strategies profile $\mathbf{p} \in \times_{i \in N} \Delta(C_i)$.

For any mixed strategies profile $\mathbf{p}$, let $u_i(\mathbf{p})$ denote the expected utility that player i would get when the players independently choose their pure strategies according to $\mathbf{p}$:

$$u_i(\mathbf{p}) = \sum_{c \in C} P(\mathbf{p}, c) \cdot u_i(c), \quad \forall i \in N \,,$$

where, $P(\mathbf{p}, c) \equiv \prod_{j \in N} p_j(c_j)$ is the occurrence probability of any configuration $c \in C$ wrt $\mathbf{p}$.

For any player $i \in N$ and any mixed strategy $\tau_i \in \Delta(C_i)$, we let $\mathbf{p}^{-i} \oplus \tau_i$ denote the mixed strategies profile in which the i-component is τ_i and all other components are as in $\mathbf{p}$. Similarly, for any $c_i \in C_i$, we let $\mathbf{p}^{-i} \oplus c_i$ denote the mixed strategies profile in which the i-component is the probability distribution e_i which assigns probability equal to 1 to the strategy $c_i \in C_i$ and probability equal to 0 to each other strategy in C_i, and all other components are as in $\mathbf{p}$. Finally, we denote by $C^{-i} \equiv \times_{j \neq i} C_i$ the cartesian product of the action sets of all players but for i.

1.1.4 Solution Concepts for Strategic Games

Nash Equilibrium

The most commonly used solution concept in game theory is that of Nash equilibrium. This notion captures a *steady state* of the play of a strategic game in which each player holds the correct expectation about the other players' behaviour and acts rationally. A Nash equilibrium recommends a strategy to each player, so that no player can improve upon its own strategy by moving *unilaterally* to some other strategy, ie, provided that the other players adopt the strategy indicated by the equilibrium. Since the other players are also rational, it is reasonable for each player to expect its opponents to follow the recommendation as well.

Definition 1.1.2 *A pure strategies profile $c \in C$ is a* **Pure Nash Equilibrium (PNE)** *of the strategic game $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ iff*

$$u_i(c) \geq u_i(c^{-i} \oplus d_i), \quad \forall i \in N, \quad \forall d_i \in C_i \,.$$

Definition 1.1.3 *A mixed strategies profile* $\mathbf{p}$ *is a* **Nash Equilibrium (NE)** *of the strategic game $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ iff*

$$u_i(\mathbf{p}) \geq u_i(\mathbf{p}^{-i} \oplus \tau_i), \quad \forall i \in N, \quad \forall \tau_i \in \Delta(C_i) \,,$$

or equivalently, iff

$$\forall i \in N,\ \forall c_i \in C_i, \quad p_i(c_i) > 0 \Rightarrow c_i \in \arg\max_{d_i \in C_i} u_i(\mathbf{p}^{-i} \oplus d_i) \equiv BR_i(\mathbf{p}) \,.$$

The set of actions $BR_i(\mathbf{p}) \subseteq C_i$ is called the **best replies** *set for player $i \in N$ wrt $\mathbf{p}$. A NE $\mathbf{p}$ for which $\forall i \in N$, $\{p_i\} = BR_i(\mathbf{p})$ is a* **strict Nash Equilibrium***, and clearly has to be a PNE.*

The set of PNE of a strategic game is a subset of its set of Nash equilibria. There exist games for which the set of PNE is empty. There are also games for which the set of Nash equilibria is empty. However, every game in which each player has finitely many strategies has at least one Nash equilibrium, as the general existence theorem of John Nash shows:

Theorem 1 (Nash, 1951) *Given any finite game Γ in strategic form, there exists at least one equilibrium in $\times_{i \in N} \Delta(C_i)$.*

Correlated Equilibrium

In a correlated equilibrium, someone tosses a coin and gives each player partial information about the result, in the form of a recommendation of how to play. An equilibrium exists in that no player has an incentive to disobey the recommendation. Formally:

Definition 1.1.4 *A* **correlated equilibrium** *p of the strategic game $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ is a distribution on the joint strategy set $\times_{i\in N} C_i$ such that $\forall i, \forall t \neq t' \in C_i$*

$$\sum_{c^{-i}\in C^{-i}} p(c^{-i} \oplus t) \cdot u_i(c^{-i} \oplus t) \geq \sum_{c^{-i}\in C^{-i}} p(c^{-i} \oplus t) \cdot u_i(c^{-i} \oplus t') \,.$$

That is, for each recommended strategy t for player i, this player has no incentive to deviate to any other strategy $t' \neq t$.

A Nash equilibrium is simply a correlated equilibrium where p is a product of *independent* distributions for each player. So Nash's theorem implies that a correlated equilibrium exists in every finite game.

1.1.5 Complexity Issues

Although Nash's Theorem guarantees the existence of at least one Nash equilibrium in every finite game, the computation of Nash equilibria has been long observed to be difficult and it is one of the most important algorithmic problems for which no polynomial-time algorithms are known.

While the question of how complex it is to construct a Nash equilibrium remains wide open, important concrete advances have been made in determining the complexity of related questions. For example, 2-person zero-sum games (ie, when in each pure strategy profile the utilities of the two players sum up to zero) can be solved using linear programming in polynomial time [59]. In [21] it was shown that the problems of (1) determining whether a Nash equilibrium with certain properties exists and (2) counting the number of Nash equilibria are $\mathcal{NP}-$hard.

1.2 Nash on the Network

A key problem in network management is that of routing traffic in order to optimize network performance. Given the traffic to be routed, we wish to find a routing that optimizes certain performance parameters, such as the average travel time (ie, *total latency* of the whole traffic), the maximum delay experienced over all users (ie, *maximum delay* among the users) or the maximum delay experienced over all edges (ie, *maximum congestion* occurring in the network). Solutions to these kind of problems can be extremely hard since the latency experienced by each edge usually depends on the amount of traffic traveling through that edge. Moreover, in large-scale or evolving networks, such as the Internet, there is no authority possible to employ a centralized traffic management. Besides the lack of central regulation even cooperation of the users among themselves may be impossible due to the fact that the users may not even know each other. A natural assumption in the absence of central regulation and coordination is to assume that network users behave selfishly and aim at optimizing their own individual welfare. In order to understand the mechanisms in such non-cooperative network systems, it is of great importance to investigate the selfish behaviour of users and their influence on the performance of the entire network.

Since each user seeks to determine the shipping of its own traffic over the network, different users may have to optimize completely different and even conflicting measures of performance. A natural framework in which to study such multi-objective optimization problems in a non-cooperative network is (non-cooperative) game theory. We can view network users as independent agents participating in a non-cooperative game and expect the routes chosen by users to form a Nash equilibrium in the sense of classical game theory. The theory of Nash equilibria [68] provides us with an important mathematical tool in analyzing the behaviour of selfish users in non-cooperative networks: a Nash equilibrium is a state of the system such that no user can decrease its individual cost by unilaterally changing its strategy. Users selfishly choose their private strategies, which in our environment correspond to paths (or probability distributions over paths) from their sources to their destinations. When routing their traffics according to the strategies chosen, the users will experience an expected latency caused by the traffics of all users sharing edges. Each user tries to minimize its private cost, expressed in terms of its expected individual latency. This in general contradicts the goal of optimizing the *social cost* which measures the global performance of the whole network in terms of eg the average delay experienced by the users or the maximum delay experienced over all users.

Many interesting problems arise in such environments. A first one is the problem of analyzing the degradation of the global performance due to the selfish behaviour of its users, often termed as the **coordination ration** or the **price of anarchy** [73]. A second problem deals with the existence and tractability of *pure* Nash equilibria in non-cooperative networks. A third problem is the Nashification problem, i.e. the problem of converting any given non-equilibrium routing into a Nash equilibrium without increasing the social cost.

1.2.1 The Price of Anarchy

Koutsoupias and Papadimitriou [55] introduced the question of *how far can a Nash equilibrium, obtained by selfish behaviour, be from the optimal routing?* They only considered networks with a single pair of source and sink, with m parallel links between the two, and linear latency functions for each edge. In their proposed model, they assume that n agents have each an amount of traffic w_i, $i = 1, \ldots, n$ to send from the origin to the destination. The set of pure strategies for agent i is therefore $\{1, \ldots, m\}$, and a mixed strategy is a distribution on this set.

Systems of parallel links, albeit simple, represent an appropriate model for several, diverse networking problems. Consider, for example, broadband networks where bandwidth is preallocated to different virtual paths that do not interfere; thus, these paths result effectively in a system of parallel links between source/destination pairs. As a second example, consider a multimedia network with several servers that are shared by the network customers; each customer distributes its applications among the servers, while competing with the other customers on the common available resources. Modeling each server as a link, the parallel links model fits well such a framework.

As mentioned above, [55] initiated the study of the price of anarchy and showed the following results for networks consisting of m parallel links. It was shown that for two identical links the price of anarchy is exactly $\frac{3}{2}$. For two links (not necessarily identical, that is, with possibly different speeds) the price of anarchy is $\phi = \frac{1+\sqrt{5}}{2}$. For m identical links the price of anarchy is $\Omega(\frac{\log m}{\log\log m})$ and it is at most $3 + \sqrt{4m \ln m}$. Finally, it was shown that the price of anarchy for any number of tasks and m (not necessarily identical) links is $O(\sqrt{\frac{s_1}{s_m}\sum_{j=1}^{m}\frac{s_j}{s_m}}\sqrt{\log m})$, where s_j is the speed of link j, and $s_1 \geq s_2 \geq \cdots \geq s_m$.

Mavronicolas and Spirakis [62] greatly extended some of the bounds above and showed the following results in the so-called *fully-mixed model*, which is a special class of Nash equilibria in which each user assigns *non-zero* probability to each and every link. They showed that for m identical links in the fully-mixed Nash equilibrium the price of anarchy is $\Theta(\frac{\log m}{\log\log m})$. For m (not necessarily identical) links and n identical weights in the fully-mixed Nash equilibrium, if $m \leq n$, then the price of anarchy is $\Theta(\frac{\log n}{\log\log n})$.

Following this line of work, Czumaj and Vöking [25] gave an exact description of the price of anarchy depending on the number of links and the ratio of the speed of the fastest link over the speed of the slowest link. The main result of [25] is that the price of anarchy on m parallel links of possibly different speeds is $\Theta(\frac{\log m}{\log\log\log m})$.

[24] presented a thorough study of the case of general, monotone delay functions on parallel links, with emphasis on delay functions from queuing theory. Unlike the case of linear cost functions, they showed that the price of anarchy for non-linear delay functions in general is far worse and often even unbounded.

[33] studied the case of a network with a distinguished source – destination pair (s,t), where each $s-t$ path has length exactly l and each node lies on an $s-t$ path. It was shown that the price of anarchy when the number of resources (links) is m and resource delays equal to their loads is at most $8e(\frac{\log m}{\log\log m} + 1)$.

In [77] the price of anarchy in a multi-commodity network game among infinitely many users each of negligible demand was studied. The social cost in this case is expressed by the total delay payed by the whole flow in the system. For linear resource (link) delays, the price of anarchy is at most $\frac{4}{3}$.

For general, continuous, non-decreasing resource delay functions, the total delay of any Nash flow is at most equal to the total delay of an optimal flow for double flow demands.

In the subsequent paper [78] it was proved that for the same setting, it is actually the class of allowable latency functions and not the specific topology of a network that determines the price of anarchy in the corresponding game.

For recent treatments see [60] [35], [36], [34].

1.2.2 Pure Nash Equilibria and Nashification

A natural question arising in game theory is what kind of games possess a pure Nash equilibrium and under which circumstances can such an equilibrium be found in polynomial time. Rosenthal [76] introduced a class of games, called **congestion games**, in which each player chooses a particular combination of resources (called an **action**) out of a set of allowable actions for it, constructed from a basic set of primary resources for all the players. The **delay** associated with each primary resource is a non-decreasing function of the number of players who choose it, and the total delay received by a player is the sum of the delays associated with the primary resources he chooses. Each game in this class possesses at least one PNE. This result follows from the existence of a real-valued function (an *exact potential* [67]) over the set of pure strategies profiles with the property that the gain of a player unilaterally shifting to a new strategy is equal to the corresponding increment of the potential function.

Formally, a congestion game consists of a set N of n players, a finite set E of resources, and the **action set** (ie, the pure strategies available) of player i is $S_i \subseteq 2^E \setminus \emptyset$. We are also given non-decreasing **delay functions** $d_e : \{0, 1, \ldots, n\} \mapsto \mathbb{R}_{\geq 0}$ for all the primary resources $e \in E$. The players' payoffs in this game are computed as follows. Let $s = (s_1, \ldots, s_n)$ be a configuration of the players (ie, a pure strategies profile) and let $f_s(e) = |\{i : e \in s_i\}|$. Then $u_i(s) = -\sum_{e \in s_i} d_e(f_s(e))$.

A congestion game in which all users have the same payoff function and the same action set is called **symmetric**. In a **network congestion game** the action sets S_i are presented as paths in a network. We are given a directed network $G = (V, E)$, with the edges playing the role of resources, two nodes $s_i, t_i \in V$ for each player i and a delay function over the set of edges. The subset of E available as the action set of player i is the set of all paths from s_i to t_i. Then we refer to a **multicommodity network congestion game**. If all origin-destination pairs of the users coincide with a unique pair (s, t) we have a **single-commodity network congestion game** and then all users share the same action set. In a **weighted congestion game** we allow the users to have different demands for service, and thus affect the resource delay functions in a different way, depending on their own weights.

As already mentioned, the class of (unweighted) congestion games is guaranteed to have at least one PNE. In [30] it is proved that a PNE for any unweighted single-commodity network congestion game can be constructed in polynomial time, no matter what resource delay functions are consider (as long as they are non-decreasing with loads). On the other hand, it is shown that even for a symmetric congestion game or an unweighted mlticommodity network congestion game, it is PLS-complete to find a PNE, though it certainly exists.

For the special case of single-commodity (parallel-edges) network congestion games where users have varying demands, it was shown in [32] there is always a PNE which can be constructed in polynomial time. It was also shown that it is $\mathcal{NP}-$hard to construct the best or the worst PNE.

[64] deals with the problem of weighted parallel-edges congestion games with user-specific costs: each allowable action of a user consists of a single resource and each user has its own private cost function for each resource. It is shown that all such games involving only two users, or only two possible actions for all the users, or equal delay functions, always possess a PNE. On the other hand, it is shown that even a single-commodity, 3-user, 3-actions, weighted parallel-edges congestion game may not possess a PNE (using 3-wise linear resource delays).

[31] deals with the problem of *Nashification*, ie, how to convert any given non-equilibrium routing into a PNE without increasing the social cost. An efficient algorithm for the Nashification problem allows to compute a Nash equilibrium with low social cost by first computing an appropriate non-equilibrium routing with known algorithms for the scheduling problem and then converting this routing into a Nash equilibrium. One way to nashify a configuration, is to perform a sequence of greedy selfish steps. A greedy selfish step is a user's change of its current pure strategy to its best pure strategy with respect to the current strategies of all other users. [31] presents an algorithm which nashifies any pure routing in a weighted parallel-edges congestion game with resource delays identical to their loads, by performing such selfish greedy steps. However, this process may take exponential time even if the edges have identical capacities.

In [33] it is proved that even for a weighted single-commodity network congestion game with resource delays being either linear or 2-wise linear functions of the loads, there may be no pure Nash equilibrium. It is also shown that there exist weighted single-commodity network congestion games which admit no exact potential function, even when the resource delays are identical to their loads. Nevertheless, it is proved that for the case of a weighted multicommodity network congestion game with resource delays proportional to their loads (ie, $\forall e \in E, d_e(x) = a_e \cdot x : \ a_e \geq 0$ is a constant), at least one PNE exists and can be computed in pseudo-polynomial time.

1.3 Challenges: Adaptive and Dynamic Systems

As already stated, classical (ie, static) game theory mostly deals with completely rational individuals who are engaged in a given interaction (a "game") with other individuals (co-players) and have to adopt strategies (for selecting among a set of allowable actions) that maximize their own (exogenously fixed) payoff function, given the strategies of their co-players. Of course the value of each player's payoff is dependent on the other players' strategies for choosing their own actions.

The initial goal of game theorists was to determine the major principles of rational behaviour, by means of thought experiments involving fictitious players who were assumed to know such a theory and to know that their equally fictitious co-players would use it. At the same time, it was expected that rational behaviour would prove to be optimal against irrational behaviour. It turned out that this was too much.

When the "trembling hand" dogma came into the scene things changed dramatically. Now a perfect strategy should take into account that the co-players, instead of taking the perfect strategic decisions, occasionally do the wrong thing. The question is how often this "occasional" behaviour appears. From allowing an infinitesimal margin of error to assuming that the faculties of the players are limited (implying bounded rationality) it takes only a small step. But, as long as the players are no longer constrained to be rational (due to their limited faculties), the can learn, adapt and **evolve**.

Thus, the new major task of game theory became the description of the dynamics of model games defined by strategies, payoffs *and* adaptation mechanisms. Evolutionary game theory deals with an entire (typically large) population where all individuals are "destined" to adopt some fixed strategy (usually considered as a **behavioural type**). Strategies with higher payoff than the average payoff given the current state of the population (ie, types with larger than average fitness) are expected to spread within the population (by learning, copying successful strategies, inheriting strategies, or even by infection). The frequencies of the types in the population thus change according to their payoffs, which in turn depend on the frequencies of the other types in the population. Observe that in this case it is the different types of behaviour (and not the individuals) that are the actual players participating in a kind of a repeated game. The subject of evolutionary game theory is exactly the study of the *dynamics* of this feedback loop. A very good presentation of evolutionary dynamics in strategic games, is in [51]. For a thorough study of evolutionary dynamics in extensive form games, the reader is referred to [22].

Numerous paradigms for modeling individual choice in a large population have been proposed in the literature. For example, if each player chooses its own strategy so as to optimize its own payoff ("one against all others" scenario) given the current population state (ie, other players' strategies), then the aggregate behaviour is described by the **best-response dynamics** [61]. If each time an arbitrary user changes its strategy for any other strategy of a strictly better (but not necessarily the best) payoff, then the aggregate behaviour is described by the **better-response dynamics** or Nash dynamics [76]. In case that pairs of players are chosen at random and these players engage in a bimatrix game ("one against one" scenario) whose payoff matrix determines according to some rule the gains of the types adopted by these two players, then we refer to **imitation dynamics**, the most popular version of which is the **replicator dynamics** [81].

The proposed dynamical systems for describing evolutionary games, despite their appealing and highly intuitive definitions, have some strong weaknesses. For example, the best/better-response dynamics admit multiple solution trajectories from a single initial point, whose terminal (rest) points vary significantly in their aggregate performance. On the other hand, the replicator dynamics admits trajectories, whose rest points are not necessarily Nash equilibria (eg, when all the users of a population choose exactly the same strategy, even if this strategy is strictly dominated, there is nothing to imitate – this is a general weakness of the imitation dynamics).

A typical way out of such situations is to introduce small amounts of noise to the underlying models. For example, we may add (small) payoff perturbations to the optimization dynamic models, or we may add occasional arbitrary behaviour to the model of replication. Of course, such modifications create new difficulties: for the optimization models the rest points may now be slightly away from the Nash equilibria, while the replication models may admit some oscillation phenomena.

Since evolutionary game theory is a dynamical process (hopefully) ending up in some equilibrium which may also demonstrate robustness against invasion or infection (stability), we consider it as a *self-organization process* in a very large population of entities that adopt some strategies representing either computer programs trying to prevail in the market, or competing network protocols, or viruses trying to spread all over the Internet. Thus the prime concern of an algorist is to determine the principal rules defining the (global or local) convergence to stable states, to propose computationally efficient algorithms for constructing such stable states, to be able to compare the trajectories of two phenomenically different laws of motion, or even to describe how the underlying infrastructure (eg, the network infrastructure in which a virus might spread) is involved in the evolution of a population. We comment here that our perspective has to do with *strategic evolution*, ie, the evolution in time of the vector of the frequencies of the behavioural types (ie, possible strategies) in the whole population, and not with traditional evolution (under the biological viewpoint) where the sub-populations actually evolve in size.

1.3.1 Evolutionary Stable Strategies

Consider a large population of players that encounter randomly chosen opponents. We assume that there is a fixed set $R = \{r_i\}_{i\in[N]}$ of N primitive actions of each player, as well as a set of n possible strategies (or **behavioural types**) $P = \{\mathbf{p}_j\}_{j\in[n]} \subseteq \Delta(R) \equiv \left\{\mathbf{z} \in [0,1]^N : \sum_{i\in[n]} z_i = 1\right\}$, which are nothing more than mixed strategies on R. If all players use the same *strict* NE $\mathbf{p} \in P$, then each individual deviating (unilaterally) from this strategy will be definitely penalized. That is, dissident behaviour is not spread. On the other hand, we cannot assume that every non-strict NE is a proof against invasion by a minority, unless they encounter an *evolutionary stable* population. To see this, consider an arbitrary non-strict NE $\mathbf{p} \in P$. For this equilibrium we know that at least one player $i \in N$ has a best-replies set that contains more than one strategies. All these strategies seem equally profitable for i. Nevertheless, it may be the case that some of them leads to an unstable state that will cause then other players to defect to more profitable strategies for them, a.s.o.

An evolutionary equilibrium is a variant of the notion of Nash equilibrium, designed to model situations in which the players' actions are determined by forces of evolution. Consider the case where the members of a single population of organisms (ie, all individuals may adopt types of behaviour from the same set) interact with each other in a *pairwise* fashion. In each encounter each organism has a type from the set P. The organisms do not consciously choose behavioural types; rather, they either *inherit* them from their forebears or these are assigned to them by *mutation.*

A type of behaviour in the population is said to be an **evolutionary stable strategy** if whenever all members of the populations adopt it, no dissident type can invade the population under the influence of natural selection. So suppose for example that, when two individuals of some given species meet, they play the game $\Gamma = (\{1,2\}, (R,R), (u_1,u_2))$ that is *symmetric*, in the sense that the contending individuals have the same action set R, and

$$u_1(a,b) = u_2(b,a) = U[a,b], \quad \forall a,b \in R\,,$$

for some $N \times N-$dimensional payoff matrix U. Here, an individual's utility represents some contribution to the reproductive fitness of the corresponding type adopted by this individual. The following is a formal definition of an evolutionary stable strategy:

Definition 1.3.1 *In a large population in which behavioural types evolve as above according to a symmetric two-player game* Γ *whose payoff matrix represents the change in fitness of the different primitive actions from R, a behavioural type* $\mathbf{p} \in P$ *is called an* **evolutionary stable strategy (ESS)** *iff* $\forall \mathbf{q} \in P \setminus \{\mathbf{p}\},\ \exists \bar{\varepsilon}(\mathbf{q}) > 0 : \forall \delta \in (0, \bar{\varepsilon}(\mathbf{q})),$

$$
\begin{aligned}
u_1(\mathbf{q}, \delta\mathbf{q} + (1-\delta)\mathbf{p}) &= \mathbf{q}^T U[\delta\mathbf{q} + (1-\delta)\mathbf{p}] \\
&< \mathbf{p}^T U[\delta\mathbf{q} + (1-\delta)\mathbf{p}] \\
&= u_1(\mathbf{p}, \delta\mathbf{q} + (1-\delta)\mathbf{p})
\end{aligned}
$$

where, $\bar{\varepsilon}(\mathbf{q})$ *is called the* **invasion barrier** *wrt to the invading type* $\mathbf{q}$.

The above definition states that the equilibrium strategy $\mathbf{p}$ should be strictly better than any alternative strategy $\mathbf{q}$ when the frequency of $\mathbf{q}$ in the population is sufficiently small (ie, there is a small positive probability of meeting an individual who is using the alternative strategy $\mathbf{q}$). It is well known that not all symmetric two-player games possess an evolutionary stable strategy. Indeed, even populations whose evolution is described by a 3×3 symmetric game may not possess any ESS.

In the above, for simplicity, we gave the definition of an evolutionary stable strategy in the context of a symmetric two-person game, but ESS is defined for asymmetric games as well, or even in scenarios where the evolution of the population is determined by broader conflicts of individuals, rather than pairwise encounters of randomly chosen opponents.

An equivalent definition of the ESS (which is a straightforward corollary of definition 1.3.1) is the following: A behavioural type $\mathbf{p} \in P$ is an ESS iff $\forall \mathbf{q} \in P \setminus \{\mathbf{p}\}$ the following two conditions hold:

(I) **[Equilibrium Condition]** $\mathbf{q}^T U\mathbf{p} \leq \mathbf{p}^T U\mathbf{p}$

(II) **[Stability Condition]** $\mathbf{q}^T U\mathbf{p} = \mathbf{p}^T U\mathbf{p} \Rightarrow \mathbf{q}^T U\mathbf{q} < \mathbf{p}^T U\mathbf{q}$

Condition (I) demands that $\mathbf{p}$ is a Nash equilibrium for the bimatrix game determining the law of motion (no invader $\mathbf{q} \in P$ can do better than the resident $\mathbf{p} \in P$, against the resident), while condition (II) states that, in case that an invader does equally well with the resident against the resident, then the resident must be strictly better than the invader, against the invader. The following theorem provides a simple (yet *inefficient*, due to the possibly too many types of behaviour) test whether a specific strategy $\mathbf{p}$ is an ESS:

Theorem 2 (Hofbauer & Sigmund, Theorem 6.4.1 [51]) *The strategy* $\mathbf{p}$ *is an ESS iff*

$$\forall \mathbf{q} \neq \mathbf{p},\ \mathbf{p}^T U\mathbf{q} > \mathbf{q}^T U\mathbf{q}$$

in some neighbourhood of $\mathbf{p}$.

For example, consider a large population of individuals that demonstrate any combination of a *hawkish* or a *dovish* behaviour. A hawk is aggressive and fights against its opponents until either the opponent retreats or they are both injured. On the contrary, a dove is willing to compromise (when

	...Hawk and gets	...Dove and gets
Hawk meets...	$\frac{G-C}{2}$	G
Dove meets...	0	$\frac{G}{2}$

Table 1.1: The payoff matrix of the Hawk-Dove evolutionary game. The total gain from each duel is G, while the injury deficit (between hawks) is $C > G$. When two hawks contend, they share the loss in fitness they cause to their type, while a hawk against a dove gets all the gain for its own type. Finally, two contending doves share the gain peacefully, without harming each other.

the opponent is another dove) or retreat immediately (when it faces a hawk that never compromises). The matrix determining the change in fitness by pairwise encounters is the following: This game is easily shown to have a unique ESS $\hat{\mathbf{p}} = \left(\hat{p}_h = \frac{G}{C}, \hat{p}_d = 1 - \frac{G}{C}\right)$, since $\forall \mathbf{p} = (p_h, 1-p_h) \neq \hat{\mathbf{p}},\ \hat{\mathbf{p}}^T U \mathbf{p} - \mathbf{p}^T U \mathbf{p} = \frac{(G - C \cdot p_h)^2}{2C} > 0$.

Population Games

Up to this point we assumed that the success of a behavioural type depends on the outcome of *pairwise* encounters with randomly chosen opponents. This need not always be the case. Indeed, in many examples the success of a type does not always depend only on the type of the opponent, but also on the frequencies of all types in the whole population. The ESS theory can be extended to such cases as well. Suppose that the payoff functions u_i depend on the vector $\mathbf{m} \in \Delta(R)$ of **expected frequencies** of the primitive actions $r_i \in R$ in the population. The expected frequency of each of the primitive actions is just the *expected* number of individuals adopting the specific action (given the exact frequencies vector $\mathbf{x} \in \Delta(P)$ of the behavioural types in the whole population) over the total population size. That is, $\forall i \in [N],\ m_i = \sum_{j \in [n]} p_j(i) \cdot x_j = (\mathbf{x}^T P)_i$, where $P = [\mathbf{p_1}, \mathbf{p_2}, \ldots, \mathbf{p_n}]^T$ also denotes the $n \times N$ matrix each row of whom is a different behavioural type. The strategy mix of the subpopulations $\mathbf{m} \in \Delta(R)$ can also be seen as the **average behavioural type**. Since a $\mathbf{p_j}-$strategist adopts each primitive action $r_i \in R$ with probability $p_j(i)$, his payoff will be given by

$$\sum_{i \in [N]} p_j(i) \cdot u_i(\mathbf{m}) = \mathbf{p_j} \cdot \mathbf{u}(\mathbf{m}),$$

where $\mathbf{u}(\mathbf{m}) \equiv (u_1(\mathbf{m}), \ldots, u_N(\mathbf{m}))$. In accordance with the definition of an ESS given in theorem 2, we define the notion of a **local ESS**:

Definition 1.3.2 $\mathbf{p} \in P$ *is a* **local ESS** *if* $\mathbf{p} \cdot \mathbf{u}(\mathbf{q}) > \mathbf{q} \cdot \mathbf{u}(\mathbf{q})$ *holds for all* $\mathbf{q} \neq \mathbf{p}$ *in some neighbourhood of* $\mathbf{p}$ *in* P.

It is quite interesting that now many ESS may coexist in the interior of $\Delta(R)$, while there could only be at most one ESS when the individuals were involved only in pairwise encounters.

1.3.2 The Replicator Dynamics.

The evolutionary stable strategies are only one of the many alternatives to describe stability in evolutionary games. Indeed, there have been proposed many other notions of stability, differing in their implicit underlying dynamics. In certain situations, the underlying dynamics can be modeled by a differential equation on the simplex $\Delta(R) = \{\mathbf{z} \in [0,1]^N : \sum_{i\in[N]} z_i = 1\}$. The dynamic process, whose outcome an ESS is supposed to be, is the so called replicator dynamics, a system of deterministic differential equations. Unlike the case of an evolutionary stable strategy, each individual is now associated with a *pure* strategy (ie, it chooses with certainty a specific primitive action to take). The counterpart of a mixed strategy in this model is a polymorphic population in which each primitive action is adopted by different individuals in the population. Thus, the actual players of the game are the primitive actions and not the individuals in the population.

Suppose that we split time in periods of length $\tau \in (0,1]$ and in each period, a fraction τ of the whole population gives birth, with each $i-$strategist (ie, an individual that adopts the primitive action r_i) giving birth to $f_i(\mathbf{x})$ offspring who adopt the same action. The function $f_i : \Delta(R) \mapsto \mathbb{R}_{\geq 0}$ is the fitness function of the primitive action r_i. The expected number of $i-$strategists at time $t+\tau$ is given by $N_i(t+\tau) = N_i(t) + \tau N_i(t) f_i(\mathbf{x}(t))$, while the total population size is $N(t+\tau) = \sum_{i\in R}[N_i(t) + \tau N_i(t) f_i(\mathbf{x}(t))] = N(t) + \tau \sum_{i\in[n]} N_i(t) f_i(\mathbf{x}(t))$. The change in frequencies of the primitive actions in the whole population can then be described as follows: $\forall i \in [n]$,

$$\begin{aligned}
x_i(t+\tau) \equiv \frac{N_i(t+\tau)}{N(t+\tau)} &= \frac{N_i(t)\cdot[1+\tau f_i(\mathbf{x}(t))]}{N(t) + \tau\sum_{j\in[n]} N_j(t) f_j(\mathbf{x}(t))} \\
&= \frac{x_i(t)\cdot[1+\tau f_i(\mathbf{x}(t))]}{1 + \tau\sum_{j\in[n]} x_j(t) f_j(\mathbf{x}(t))} \\
&= \frac{x_i(t)\cdot[1+\tau f_i(\mathbf{x}(t))]}{1+\tau \bar{f}(\mathbf{x}(t))} \qquad \Rightarrow \\
\frac{x_i(t+\tau) - x_i(t)}{\tau} &= \frac{x_i(t)\cdot[f_i(\mathbf{x}(t))] - \bar{f}(\mathbf{x}(t))}{1+\tau \bar{f}(\mathbf{x}(t))}
\end{aligned}$$

where $\bar{f}(\mathbf{x}) \equiv \sum_{j\in[n]} x_j \cdot f_j(\mathbf{x}) = \mathbf{x}^T \cdot \mathbf{f}(\mathbf{x})$ is the **average fitness** in the population and $\mathbf{f}(\mathbf{x}) \equiv (f_1(\mathbf{x}), \ldots, f_n(\mathbf{x}))$ is the **fitness vector**. By allowing $\tau \to 0$ we get from this discrete time model the following replicator equation:

$$\dot{x}_i(t) = \frac{d}{dt} x_i(t) = x_i(t)[f_i(\mathbf{x}(t)) - \bar{f}(\mathbf{x}(t))] \tag{1.1}$$

Observe that the **rate of increase** $\frac{\dot{x}_i}{x_i}$ of each behavioural type $r_i \in R$ may be seen as a *measure of evolutionary success*. In that case, it is clear that the replicator equation is the natural dynamics reflecting the evolution of the frequencies vector in the population over time.

Replicator Dynamics vs stability

Having fixed the dynamics for the description of our evolutionary system, the next question is what one can learn from the replicator equation wrt stability and equilibrium notions. A crucial observation that also holds for a more general family of dynamics that include the replicator dynamics, called the **imitation dynamics**, is that the support of the population remains always the same: no initially zero-frequency type may ever be used by any individual, and no initially non-zero-frequency type may ever become extinct. On the other hand, there may be types whose frequency will become very small, albeit they never reach zero. Indeed, this is one of the most crucial drawbacks of this family of dynamics, since it does not allow the adoption of new behavioural types that were not initially used in the population. We say that the simplex of *used* types is **invariant** under the replicator dynamics. As an extreme example, observe that any population that adopts exclusively any specific type is definitely a rest point[1] of the replicator dynamics, since there is actually nothing different from this type to imitate. The following theorem gathers some folklore propositions demonstrating the close connection of the rest points of the replicator dynamics and the notions of NE and ESS:

Theorem 3 (Samuelson 1997 [79]) *Consider the replicator equation 1.1 and the corresponding symmetric two-person game $\Gamma = (\{1,2\}, \{R,R\}, U)$ determined by the changes in fitness of the different types.*

1. $(\mathbf{x},\mathbf{x}) \in \Delta(R)$ *is a NE* $\Rightarrow$ $\mathbf{x}$ *is a rest point. The opposite need not always be true.*
2. $\mathbf{x}$ *is a stable point* $\Rightarrow$ $(\mathbf{x},\mathbf{x})$ *is a NE. The inverse is not true in general.*
3. $\mathbf{x}$ *is asymptotically stable* $\Rightarrow$ $(\mathbf{x},\mathbf{x})$ *is a perfect, isolated equilibrium. The inverse is not true.*
4. $\mathbf{x}$ *is ESS* $\Rightarrow$ $\mathbf{x}$ *is asymptotically stable. The inverse is not true.*
5. *In a 2-player, 2-strategies game,* $\mathbf{x}$ *is asymptotically stable* $\Leftrightarrow$ $\mathbf{x}$ *is an ESS.*

[1]A **rest point** of a system of differential equations $\left\{\frac{dx_i}{dt} = f(\mathbf{x})\right\}_{i\in[n]}$ is a point $\mathbf{z} \in \Delta(R) : f(\mathbf{z}) = 0$. A rest point $\mathbf{z}$ is called **stable** if for any neighborhood V of $\mathbf{z}$, there is a neighborhood $U \subseteq V$ s.t. $\mathbf{x}(0) \in U \Rightarrow \forall t \geq 0,\ \mathbf{x}(t) \in V$. A stable point $\mathbf{z}$ is called **asymptotically stable** or **attractor** if there is some neighborhood U of $\mathbf{z}$ s.t. $\mathbf{x}(0) \in U \Rightarrow \lim_{t\to\infty} \mathbf{x}(t) = \mathbf{z}$.

1.3.3 Stochastic Models.

The dynamics models based on systems of differential equations is not the only way to describe the behaviour of an evolutionary game. For example, Boylan [12, 13] who shows that a choice model may give behaviour that differs markedly from the behaviour described by a system of differential equations. The crucial assumption in Boylan's model is the finite population size. Each individual is then considered to play exactly one game with a single opponent in each period, where all the possible matchings of opponents are equally likely to appear. After each such match the contending individuals die and are replaced by two new individuals, whose behavioural types are determined by the payoff matrix of the game. Notice that by construction of the model, the population size always remains the same. It is only the frequencies of the behavioural types that alter. We thus have a finite stochastic process to study. Boylan uses either a deterministic, discrete analogue of the replicator dynamics, or derives a continuous replicator dynamics that is given by assuming that the population size tends to infinity and the length of each period tends to zero. In both cases it is proved that these dynamic systems are good approximations of the considered stochastic process.

Another approach is the Kandori-Mailath-Rob (KMR) model [54]. The model also uses a symmetric normal-form game played by a single population of N individuals, each of whom is characterized by a pure strategy. In each period, these individuals are matched in pairs to play the game an *infinite* number of times. At each iteration, each individual is equally likely to be matched with any of the $N-1$ other individuals in the population. Consequently, there is no randomness in the average payoffs received in each period, and in each period an individual (playing a fixed pure strategy) gets the expected payoff of playing against the mixed strategy corresponding to the frequencies of the pure strategies in the population. At the end of each period, each individual "learns", in the sense that he changes strategy to the best response to the current frequencies vector (which is dealt with as a mixed strategy of the opponent in each iteration). It is not necessary that all individuals change their strategies to the best responses to the frequencies, so long as there are some of them that do so. After learning, "mutations" take place: Each individual with some (small) probability λ decides to drop its current strategy and choose a new pure strategy uniformly at random.

The KMR model is actually a Markov process whose state space is $\{0, 1, \ldots, N\}^{N-1}$, ie, all the possible assignments of individuals to the available pure strategies. It is easily seen that the transition matrix of this process is strictly positive. It suffices to observe that there is a positive probability that at the end of a period, after learning, all individuals choose to mutate and it happens that they choose those strategies that produce a given configuration. Thus, the KMR model has a unique stationary distribution, that is independent of the initial conditions and that will attach positive probabilities to all states in the state space. Yet, it is shown that when the mutation probability λ is sufficiently small, all the probability mass in the stationary distribution is accumulated only on a few states. By letting $\lambda \to 0$ Kandori Mailath and Rob direct our attention to a *limiting stationary distribution*, which is often easily computable and produces powerful results.

These alternative models for studying dynamical systems are in many cases complementary to other dynamics models and help us in equilibrium refinements in cases where the latter models are actually indecisive.

1.3.4 Computational aspects of game-theoretic evolution

There has been an explosion in recent years of computational simulation results in evolutionary games. In such models complex and interesting phenomena is often evidenced but (as discussed later in section 2.3) proof and full understanding (i.e. analytical tractability) is often difficult to establish.

In the following section a flavour of these models is presented through a small set of recent examples relevant to DELIS.

We outline in section 3 some of our on-going work to extend game theoretic analysis in order to support these more complex scenarios. This involves the ambitious task of advancing the state-of-the-art in game theoretical analysis of evolving systems.

Chapter 2

The Simulation approach

As stated in chapter 1, in many target domains the behaviour of systems are not understandable using the traditional tools of classical game theory. By this we mean that game theoretical analysis does not give any predictive understanding or prescriptive guidance for the behaviour of actors (or agents) within the system, or indeed of system level outcomes.

This is because the strong classical assumptions concerning the rationality, computational capacity and limited communication abilities of actors do not always hold. Also the focus on equilibria may be of little value when systems are always out of equilibrium.

In general the classical assumptions breakdown when systems contain:

- significant levels of noise
- adaptive agents using heuristic algorithms
- agents with limited computational capacity
- agents with limited knowledge and information concerning the system as a whole
- delayed and / or faulty feedback concerning utility of actions
- agents with the ability to communicate directly with each other
- heterogeneous agents with different learning algorithms and initial behavioural predispositions.

Although it may be the case that specific forms of the above can be incorporated into some classical treatments under some given scenarios, in general these kinds of properties make such treatments very difficult. Essentially, without strong assumptions it is not possible to know how other agents will behave in given situations. Hence, it is not possible for agents to know, *a priori*, what the best way to act is.

In the context of DELIS-like systems, where dynamic and evolving systems are considered, at least some of the behaviour will be beyond the current state-of-the-art in analytical game theoretical understanding. One approach to address these relaxed assumptions is to use evolutionary metaphors, drawing on ideas from biology. These can go a long way to addressing the issues. The application of ESS analysis and the Replicator Dynamics (see previous chapter) provide powerful mechanisms for dealing with bounded rationality, noise and trajectories between equilibiria - however, this is an on-going area of research (see section 3).

Recently there has been a trend towards the explicit behavioural modelling (in computer simulations) of multiple interacting entities. Each entity (or agent) is specified as a computer program sharing a common environment when the simulation is executed the modeller observes the resultant emergent behaviour. In these approaches models can be constructed of arbitrary complexity (perhaps employing A.I. techniques) and size (up to millions of agents). Currently these approaches are still within the research domain and as such no accepted or mature methodology is established. In order to indicate the kind of approach and techniques used we will (in the following sections) give some specific examples the examples have been selected because they have relevance to the DELIS objectives and also because they evidence a particular level, technique and / or approach.

When researchers move from analytical tractability in mathematical models to computational constructability in algorithmic (agent-based) models, new kinds of questions need to be addressed namely, how do agents actually behave on an individual level (micro-level behaviours) in day-to-day interactions? Additionally, what kinds of micro-behaviours are sufficient to produce the higher-level norms, institutions, practices and dynamics that we observe in human economic worlds?

Interestingly, there are bodies of recent research that has attempted to answer these questions, both empirically and theoretically, called Evolutionary Economics and "Organisational Science" we consider these areas later in section 2.2.

2.1 Individual-Based Modelling

It is well know that the ESS and Replicator Dynamics approaches give very good (analytically provable) results for systems where the general assumptions on which they rest hold (generally, *a priori* pay-off matrices, mean field mixing, large populations, minimal effects of mutation etc). However, where populations are small and susceptible to noise or large amounts of mutation (change) or the payoffs are not known in advance or where local interaction is significant (i.e. some kind of topology is imposed on partner selection) then results can significantly differ from what the previous results suggest even when pure strategies are used and the evolutionary assumption holds (i.e. proportions of strategies in the next generation proportional to fitness in current generation).

In order to explore the results of these kinds of systems researchers have made use of individual based modelling in which each individual agent move (and the resulting payoff) is simulated. Generally after some period of time these "action sequences" are followed by a process of evolution using techniques from Genetic Algorithms [52] or Genetic Programming [56].

In the now famous work of Nowak and May [71] [63], they showed that by placing agents on nodes in a two-dimensional lattice where they interact and evolve locally only (i.e. only playing with and copying neighbours based on utility) that under many parameter values cooperation persisted in the population (often dynamically and cyclically). The only way to determine these conditions was via the individual modelling approach. Although post hoc theoretical arguments can then be made generally such systems produce very complex dynamics which are hard to analytically capture and often tend towards the chaotic.

Recently the use of tags [53], both abstractly [74] [37] [75] [29] [40] and applied to P2P type random overlay networks [42] [41] have been studied with simulation. From these studies it has been found that simple mechanisms of local interaction combined with reproduction of strategy based on fitness can produce high levels of cooperation even in the single-round Prisoner's Dilemma (the single-round Prisoner's Dilemma (PD) captures a minimal form of a commons dilemma - see below). These results are particularly relevant to P2P random overlay networks since interactions are by definition local to the neighbours of each peer and the interactions between peers can be viewed as a commons dilemma of sorts since each node in a P2P is likely to be in the system for what they can get out of it rather than for altruistic reasons although in reality some altruism exists. Initial empirical work on the behaviour of nodes in actual file sharing P2P systems has been perform [1].

2.1.1 Example 1: Prisoner's dilemma on the network

The single-round two player PD captures a form of commons dilemma where cooperation would benefit both participants but there is always an incentive for for each to not cooperative (to free-ride) and get a higher individual score. Figure 2.1 shows the pay-off matrix for the PD and explains the dilemma.

We found that we could produce high levels of cooperation within networks (modelled as unstructured, undirected graphs) with each node either a cooperator or defector, where nodes play single rounds of PD against randomly chosen neighbours if we mutate not only strategies on reproduction but neighbour lists also. We do not go into details here, however, figure 2.2

	Cooperate	Defect
Cooperate	C, C	S, T
Defect	T, S	P, P

Figure 2.1: A payoff matrix for the two-player single round Prisoner's Dilemma (PD) game. Given $T > C > P > S \wedge 2C > T + S$ the Nash equilibrium is for both players to select Defect but both selecting Cooperate would produce higher social and individual returns. However, if either player selects Cooperate they are exposed to Defection by their opponent - hence the dilemma

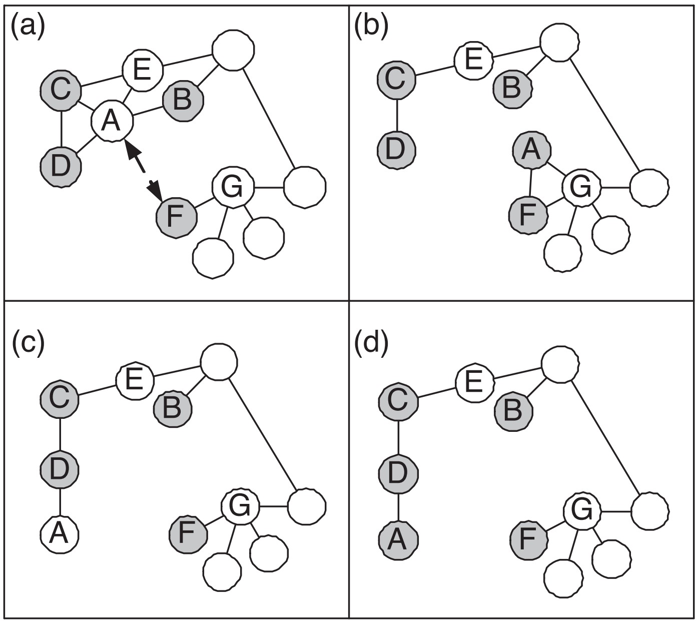

Figure 2.2: An illustration of replication and mutation as applied to the network. Dark nodes play D and light nodes play C. In (a) the arrowed link represents a comparison of utility between A and F. Assuming F has higher utility then (b) shows the state of the network after A copies Fs links and strategy and links to F. A possible result of applying mutation to As links is shown in (c) and the strategy is mutated in (d).

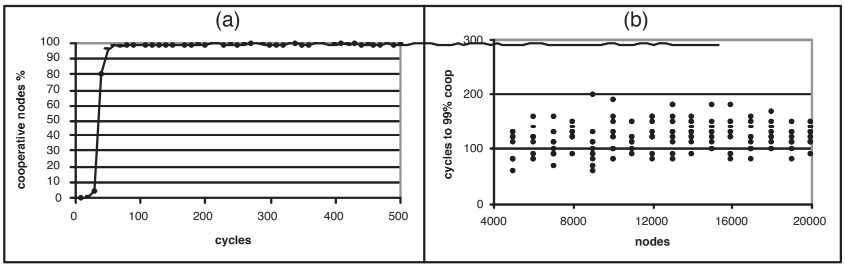

Figure 2.3: Charts show results from network model. Chart (a) shows a typical time series (N = 10,000 nodes) giving the percentage of cooperative nodes. Chart (b) shows the number of cycles before 99 percent of nodes are cooperative for various values of N. Each dot is an individual run.

gives an overview of the way that reproduction and mutation is applied to a network structure. Figure 2.3 shows some results from our model. Very recently we applied the same process to a more realistic query-answering scenario as advanced by [80]. We found that the mechanism suppressed the tendency of individual nodes to flood the system with queries [41].

These initial results, although encouraging still require refinement before they can be used to produce deployable systems. Where techniques exist as individual based models then the behavioural rules (i.e. the algorithms that nodes use) are by definition already specified. This means that for the purposes of adapting such results for application to computer networks we have an initial starting point from which to work. Initial work proposes [43] moving from an original abstract model to more realistic scenarios via a succession of models. Using this general technique, algorithms can be imported and adapted from biological or economic disciplines / scenarios into engineering problem scenarios.

2.1.2 Example 2: Reputation and image

Lai et al [57] build on both the reciprocity work of Axelrod [7] and work by theoretical biologists Nowak and Sigmund [72] in which the importance of image and reputation was investigated in the context of repeated interactions in the Prisoner's Dilemma. However, Lai et al modify the interaction payoffs of the Prisoners Dilemma game to represent a file-sharing scenario in a P2P (see figure 2.4). They then investigate the effectiveness of both private histories and shared histories in supporting high levels of cooperation.

Peers using private histories store their own histories of interaction with all they encounter whereas those using shared histories have access to a shared store of all past interactions of all peers. They show that as the size of the population increases private histories become less functional because peers are more likely to meet strangers. They use their results to argue

	Allow download	Ignore request
Request file	7, -1	0, 0
Don't request	0, 0	0, 0

Figure 2.4: A payoff matrix capturing a possible client / server interaction pattern in a simulated file-sharing scenario (taken from [57]).

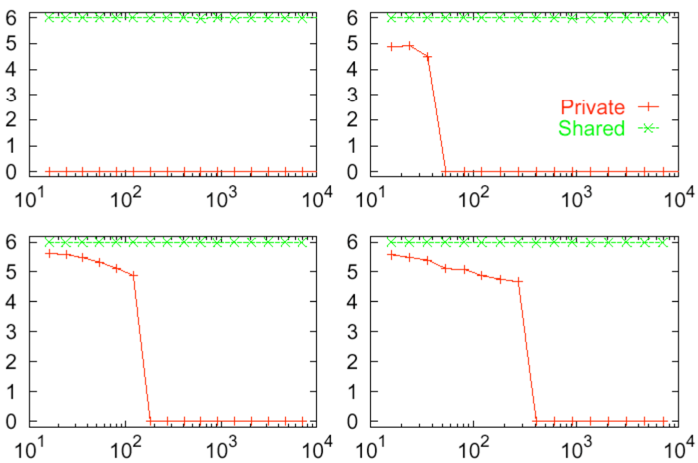

Figure 2.5: From Lai et al [57]. Each chart shows a comparison of the shared and private history methods over a number of different parameters (including different proportions of initial cheaters and numbers of games per round). The x-axis represents size of population, the y-axis level of cooperation. Each point represents an independent run. As can be seen in each presented case, shared histories outperform private histories.

for the development of robust mechanisms that can support shared histories (though they acknowledge such mechanisms are costly and difficult to implement). Figure 2.5 gives an example some of their results (taken from [57]).

What is interesting in this work is that Lai et al draw on the methods, results and techniques of both a (political scientist) [7] and (biologists) Nowak and Sigmund [72] and attempt to adapt them to a P2P scenario. By applying simulation techniques and algorithms rather than proofs and formalisms, they manage to import and integrate knowledge from diverse disciplines and apply it (using simulation) to the problem in hand. The work compares different possible reputational mechanism for effectiveness. In this sense the simulation can be seen as a kind of "socio-biologically" informed pre-prototyping exercise.

2.1.3 Agent-based simulation (ABS)

In both the previous examples a simple payoff matrix in which agents received payoffs based on their actions modelled interaction. The matrix was specified *a priori* and exogenously and stands-in for some kind of interaction process (e.g. file downloading). However, as demonstrated in the previous example, how can one decide on a payoff matrix for some target phenomena? There are many possible plausible values but what about more realistic situations, where payoffs are not explicit or are dynamic depending on other properties of the system? In some circumstances it is not realistically possible to derive an *a priori* payoff matrix to capture some target phenomena. How can such target domains be analysed?

One way to tackle these issues is to model the underlying interaction process (at some level of abstraction) and measure performance. So, for a routing problem a simulation model would be produced in which individual nodes are represented within a network executing their routing algorithms on simulated messages. In the context of a file-sharing scenario, nodes (peers) are modelled sharing data across a network. In such an approach no *a priori* payoffs are required, rather, a specification of a target problem scenario is required and then a simulation model capturing the salient aspects of it.

When interact is modelled at this level of detail it is often termed agent-based modelling. In such models each individual agent is represented in the system as a fully functioning and interacting rule based system (a semi-autonomous sub-system with the system). The ways the agents may interact are therefore not completely preset and new and unexpected (emergent) collective behaviours can often be observed.

2.1.4 Example 3: Incentives on the network

Qixiang Sun and Garcia-Molina [80] use a simulation model of the underlying process of passing and processing queries in a P2P file-sharing scenario. They perform experiments over a number of parameters in order to test the effectiveness of their Selfish Link-based InCentive (SLIC) algorithm.

In this algorithm, peers maintain weights against their links to other peer nodes in the network. If they receive good service from a link they increase its weight, if they receive bad service the decrease the weight. Peers then share their resources in proportion to the weights. Hence producing a form of reciprocity (similar to tit-for-tat see [7]). Hence the node level algorithm designs have been influenced by work from previous pay-off based models.

To test SLIC a number of experiments are performed with different unstructured overlay network topologies (randomly generated). A system of flood-fill query passing is simulated, where queries for files are passed to all neighbours recursively until either some time-out depth is reached or a peer finds a hit for the query. The process of file downloading is not modelled and notes do not store files (in the simulation) - their file storage is represented by a probability of finding a match for any query (a single real number). So, neither actual files nor the downloading process is simulated directly, but the other node (peer) level mechanisms are simulated in full over a simulated network.

The behaviour of the nodes is described textually and with pseudo-code algorithms (figure 2.6 gives a flavour of the kind and level of description that is typical of agent-based simulation approaches) and results are given over a number of different parameters and runs. Their method of analysis is to set all to cooperative and then to vary a single node (a probe node) reducing its cooperative level and analysing the dynamic result over time in the network. They demonstrate that their SLIC algorithm quickly reduces the service given the free-riders from the network since the free-riders neighbour peers quickly reduce their link weights to the free-rider and hence do not deliver service.

During round i, node u performs the follow actions:

1: generate $\rho_u \cdot C$ new queries
2: $W_{total} = \sum_{(u,v) \in E_u} W_i(u,v)$
3: **for** each edge $(u,v) \in E_u$ **do**
4:     Process and forward $(\mathbf{1} - \rho_u)C \cdot \frac{W_i(u,v)}{W_{total}}$ queries with the highest TTL from the link (u,v)
5: **end for**
6: if there are still spare capacity, repeat steps 2 through 5 to divide the spare capacity.
7: Tally number of hits for newly expired queries generated by u, i.e., $Q_i(u)$.
8: **for** each edge $(u,v) \in E_u$ **do**
9:     $W_{i+1}(u,v) = compute\_weight(W_i(u,v), Q_i(u))$
10: **end for**

Figure 2.6: Pseudo-code for node (peer) operation (from [80]).

Although Sun and Garcia-Molina don‘t use the term “agent-based” to described their model it is consistent with the methodology. The model is not in fact an evolutionary one in the sense that neither agents (nodes) nor strategies evolve in any sense. However, the links between nodes are changed – so in that sense there is change over time.

2.2 Evolutionary Economics and Organisational Science

Evolutionary thinking aims at explaining behaviour in systems that, though not endowed with global rationality, exhibit teleological development. For example, such systems as species tend to progress but we cannot assign to species per se the will to evolve (we may only assign the will to survive to individuals belonging to species).

Evolutionary processes are articulated in three sub-processes: variation in the characteristics of the elements belonging to a system, selection of those elements which best fit the environment where the system live, and retention of the characteristics of best fitted elements at the level of the system.

The emerging result of the three interacting sub-processes is a learning process through which systems evolve. Variation sub-processes generate a repertoire of possible alternative solution to a survival problem. Selection sub-processes choose the best alternative and retention sub-processes ensure that the selected alternative will be adopted at system level.

Evolutionary processes display a number of characteristics that were deemed germane to metaphorical description of mechanisms of evolution of economic systems. First, search of alternatives is not driven by a maximising behaviour; on the contrary is the result of blind variation and chance. Thus, selection of best available alternatives does not necessarily lead to a global maximum. Second, within evolutionary processes, time is an important explanatory variable.

The key role of time in evolutionary thinking emerges in three ways. First, evolutionary processes are path dependent. Favoured species diffuse in a self-reinforcing mechanisms in which rate of reproduction increases as the base of individuals carrying the winning phenotypes increases. In this light, as time goes by, competition among species is increasingly unbalanced and potentially dominant variations have different chance of survival depending on the point in time in which they appeared.

Second, evolutionary processes are irreversible. If an individual is selected out, in the future, the evolving system will not be able to change its mind and retrieve the discarded alternative.

Third, connected to irreversibility, evolutionary processes are characterised by hysteresis. That is, once the cause that generated an effect is removed, the effect remains an the system will not go back to the previous state. For example, if a particular species is favoured within a population, given to transitory environmental conditions, once removed the conditions, the initially favoured species will continue to diffuse. The self-sustained diffusion takes place because a larger base of individuals carrying the favoured phenotypes has a larger probability to reproduce, until environmental conditions will favour different phenotypes.

The before mentioned reasons suggest that the evolutionary metaphor may be useful to capture the core of economic dynamics because these latter both unfold as boundedly rational, trial-and-error learning processes and show path-dependent behaviours. In other words, evolutionary theorising becomes useful when we relax hypotheses of full rationality in decision-making and historical efficiency, that is, the assumption that observed behaviours at any point in time reflect unique outcomes of underlying systematic processes, independently of historical details.

Evolutionary concepts have fertilised recent theorising on economic change [69] and, as Hodgson suggests [50], is deeply entrenched in economic thought. In particular, evolutionary thinking supported analysis of industry evolution, organisational dynamics and technological change taking different stances in respect to the definition of the unit upon which selection processes operate. The choice of different units of selection implies different descriptions of variation, selection and retention processes. We suggest that evolutionary theorising in economic and organisational theory may be grouped in, at least, three main threads.

2.2.1 Population ecology

Despite the pray by practicing managers for continual adaptation and survival, the history demonstrates that it is quite difficult for the established firm to remain successful in the face of environmental changes brought by, for example, new technologies and deregulation. It is often the case that the firms which establish the leadership in the new environment are often armed with a set of new resources and competencies and with a set of new management practices. This simple observation basically supports the emergence of an ecological perspective of organizational adaptation during the 1980s [45] [46] [47] [48] [49] [44] [19].

According to the perspective, the established firm exhibits strong structural inertia. During the process of earlier growth, a firm accumulates learning about the technologies, customers and the market of its core business so as to establish a presence in the market and achieve operational efficiency. Yet, it is this learning and resultant routines and capabilities that often hinders the established firm from experimenting new business ideas and searching for new operational and management solutions. Core capabilities of the established firm become core rigidities [58]. In the presence of inertia, Darwinian evolutionary dynamics apply to the competition among firms. Variation processes generate the emergence of new organisational forms with new structural characteristics. In the course of environmental changes, firms with structural forms that do not fit environment are selected out and are replaced by firms with the new structural form which fits in better with new environment. Retention processes take place at the level of populations of organizations by the means of foundation of new organizations showing the winning structural form. Organizational change takes place not in the form of voluntary adaptation by the established firm but in the form of change in the population of firms through selection of unfitted populations.

2.2.2 Dynamics of organisational routines

Nelson and Winter [70] proposed an evolutionary theory of organisational capabilities adopting an economic analogue of natural selection that operates at market level as inefficient firms are winnowed out. In Nelson and Winters theory, a firm genotype is constituted by the set of operating organisational routines.

Variation processes work within organisations when new routines are generated. Nelson and Winter describe variation processes as moulding both boundedly rational problem-solving activities and random events. More precisely, organisational routines change because they are revised in order to solve emerging anomalies or to imitate competitors better routines. In reshaping routines, decision-makers adopt heuristics that recombine parts of existing routines. In this light, variation impinges upon typical behavioural assumptions such as local search in the neighbourhood of past solutions [23] and mimics natural processes of recombination of gene endowments through sexual reproduction. The stochastic element of variation emerges as errors intervene in reshaping routines and results may be partially unintended.

Selection processes occur at market level when endowments of organisational routines determine firms competitive ability and markets select out unprofitable firms. Units of selection, thus, are firms, which may be regarded as the phenotype of an endowment of organisational routines [69].

Finally, in the Nelson and Winters framework, processes of retention of successful organisational routines encapsulate elements from Lamarckian evolutionary framework since successful organisational routines diffuse not only because are crystallised within successful firms but also because are imitated by other firms. Of course, in the process of imitation, errors may intervene that trigger further variation processes.

Connected to the work of Nelson and Winter is the body of work economic change interpreted as the consequence of random activity and evolution of technology [26] [27] [28]. These studies focus on random activity as the kernel of search behaviour through which firms hone their organisational routines and improve their adaptation chances. In this light, innovations fuel technical change and economic evolution.

This area of contributions has also close connections with works on self-reinforcing mechanisms in economics. Indeed, evolutionary processes, in their assigning a key role of time, provide theoretical lenses for the analysis of positive feedback and path-dependence in economics. In this light, Arthur

[3] [4] [5] explained how events emerging at the beginning of the history of an industry may generate lock-in and dominance of inferior technological standards.

2.2.3 Intra-organisational ecology

Although rooted in the tradition of strategic management research, the theoretical contributions of the IOE theory need to be understood in the context of the adaptation versus selection debate on organisational change, which has been, and still perhaps is, one of the central debates in the fields of organizational studies [2] [6] [11].

The adaptation versus selection debate on organizational change concerns the sharp contrast population between the population ecology perspective, presented in the foregoing, and the traditional strategic perspective taken by business policy and management scholars. These latter believe in the discretion of management of a firm, particularly a large, resource-rich firm, in controlling the fate of their organization. They argue that managers are able to, and indeed do, achieve the re-alignment between their organization and environment over time by developing new purposes, and policies, by designing new organizational and administrative systems, and even by changing performance standards [82] [20].

Although these two perspectives, deterministic environmental selection and voluntary organizational adaptation, present contrasting views on organizational change, the IOE theory, pioneered by Robert Burgelman [16] [14] [15] [17] [66] [18] and shared by some organizational theorists [65], aims to integrate ecological and strategic perspectives by showing how a large established firm simultaneously deals with internal adaptive forces and external selection pressures when adapting to changes in its environment.

The intra-organizational ecology theory of organizational change, presents a unique synthesis to adaptation versus selection debate on organizational change. The perspective views a firm as the ecology of strategic initiatives, which fall both within and outside the scope of the firms corporate strategy, and argues that organizational adaptation could be realized through a process of internal competition among these strategic initiatives, which are units upon which selection processes operate.

New initiatives, which fall outside the scope of the firms current corporate strategy, represent potential variations in the strategic behaviour of a firm. Variation, thus, is a deterministic process according in which front-line managers configure business opportunities that are disconnected from a firm official strategy and core competencies.

Some features of the firm such as administrative systems, organizational culture, and most importantly, the concept of corporate strategy exhibit inertia and often resist against apparent needs for changes. These inertial features function as an internal selection environment, guide resources allocation to competitive strategic initiatives and may pave the way for strategic renewal.

Strategic initiatives may gradually acquire corporate resources and eventually challenge the validity of a firms corporate strategy. Retention process functions by adjusting a firms corporate strategy so to include those initiatives which, though originally excluded, proved to be successful.

Throughout the process of organizational change, the firm is both adaptive and inertial.

2.3 Challenges for modelling approaches

Individual and agent based models, although tackling some of the shortcomings of the analytical approach (i.e. classical game theory, ESS or replicator dynamics approaches), introduce their own problems. The main problem is the status of results obtained from such models and the interpretation of key assumptions with respect to the real world or target phenomena. In the context of engineering, the question is: How can we be sure that mechanisms discovered and tested in simulation wont fall over at an future time? Since we have no proof, just simulation results, our results are always contingent. One method of tackling this is to simulate many possible scenarios to a high degree of realism - but this is expensive and non-trivial and still only delivers contingent confidence. A better approach is to formulate a theory of how the mechanism works and produce simpler models that become analytically tractable. However there is no general methodology for doing this currently - we will attempt to make some progress on this (see section 3.3).

2.3.1 Where's the proof?

A major limitation of the evolutionary approach is the difficulty in producing anything close to the kinds of proofs that analytical game theory can deliver. Essentially, out of equilibrium systems, composed of complex interconnected, evolving non-linear units (under noise) offer many problems for attempts at proof. Given this, the next best option is the rigorous analysis of results from computer simulations. The difficulty with such models, however, is in determining what properties are significant in producing given behaviours i.e. identifying necessary and sufficient conditions for observed behaviour of interest rather than merely advancing the model as an existence proof and therefore only giving contingent results.

Ultimately, general results derived from simulation results can always be overturned if new simulation results contract them. Essentially, it is not possible to identify general properties with the certainty of a proof. It is of course possible to identify specific results and properties with certainty (assuming the computer program implementing the simulation is correct in relation to the specification). One way to increase certainty in this latter case is to compare models implemented by different authors on different platforms (the so called "docking" or "aligning" of models [8]).

Ideally, simulation models should reinforce, inform and test the predictions of analytical proofs (possibly of highly simplified / abstracted systems) in order to give those proofs more credence this kind of approach [9] [10], [72] although not always possible, provides an ideal way to combine approaches in a mutually beneficial way. It would appear that the best kinds of work in the context of engineering application would supply both proofs and analytical arguments and simulation models demonstrating how practical applications might pan out.

2.3.2 What does fitness mean?

In the context of biological interpretations of evolutionary models the concept of fitness (i.e. replicative power) has a direct interpretation number of offspring produced. However, when evolutionary models are applied to human social processes, such as economic systems, it is unclear what exactly fitness might mean. It would seem that each different model application domain would require a careful justification of the nature of fitness the model for example, in a model of competing firms profit might be considered as fitness. However, on the whole such interpretations are problematic.

In the context of DELIS, again, the concept of fitness requires analysis within each application domain. By taking an engineering stance and prescribing fitness as some desirable property at processing client/server nodes (such as download speed, query answering rates, jobs done etc) we can sidestep some problems of interpretation but there is still the problem assuming that nodes with higher fitness (or more precisely the behaviour of nodes with higher fitness) will tend to spread through a population nodes.

The traditional interpretation is as given in [57] is simply that it is assumed that the behaviour of nodes with higher fitness spreads because users somehow copy client software that gives better perceived results (results being equated with fitness). In a file-sharing scenario, for example, this might be percentage of queries producing a downloadable hit. However, the problem with this approach is that it assumes that somehow (in an un-modelled way) users update their client software based on self-interest utilising a homogenous utility metric. Such models rarely address where such homogenous metrics might come from or how and when users swap client software. It would appear that such behaviour would be more complex than current models posit.

2.3.3 Are the agents too dumb?

It is often claimed that the classical game theory assumption that all actors (or agents) have perfect information and unlimited computation (and know that all other actors have this) is unrealistic giving actors god like powers. But in many evolutionary models, actors have, what could be argued as, a overly bounded rationality. Generally only a few degrees of freedom are aloud in the strategy space and behaviour change is through some kind of random mutation and myopic copying without any underlying theory or model of how behaviour is affecting outcomes. Evolutionary actors often appear to be blind and dumb as opposed to the god-like vision of the classical actors. An explanation for this probably originates in two factors, firstly, that many of the evolutionary models hark back to those produced by biologists (in which evolution is blind and dumb) and secondly as a counter to the overly cleaver agents in the classical approach.

However, it would seem that even (perhaps especially) within DELIS-like networks specifically when we want to deal with potentially cleaver cheaters as well as simple bounded optimisers, we need a range of behaviour strategies not just all cleaver or all dumb. Also, it would appear that in addition to the heterogeneity of behaviour at any given time there would be heterogeneity of the adaptive mechanisms employed - different actors (or agents) may adapt their behaviour using different mechanisms (or maybe not at all).

Chapter 3

Combining Approaches

Our aim is to develop the application of classical, evolutionary and modelling approaches to evolving networks. Additionally we intend to combine and link these approaches something that is currently very rarely done. We aim to draw on the diversity of expertise within the group to focus on common problems and combining techniques. We see this aim as ambitious and timely. Our aim is to make substantive contributions to methodologies for combined approaches.

It is widely recognised that classical approaches offer solid proofs but that more complex computational models sometimes capture new and interesting phenomena all agree that a combination of the two, proof and computational implementation in the form of simulations in more realistic task domains, is the gold standard to aim for. However, there is currently little work, theory or methodology that guides the combination of the two approaches.

In this section we propose a number of on-going projects and ideas that the partners within WP4.3 will work on over the next 12 months. We outline each project briefly in sections 3.1 to 3.8. Some of the projects will be collaborative and some will independently support each other. For example, the ideas outlined in section 3.8 are directly relevant to the larger goals of section 3.1. All the projects relate to the methodological goals given in section 3.3.

We do not envisage solving *all* the problems and issues raised in these project overviews. We will focus on those avenues of investigation that become productive. We therefor do not present this section as a complete specification or plan of work but rather as the current on-going work which will almost certainly be pruned over-time.

We now briefly discuss issues of WP co-ordination and dissemination. In the context of co-ordination the WP partners have already had a brief, yet highly productive and encouraging, meeting (discussing the content of this deliverable along with wider long-term goals for the WP has outlined here). The meeting took place in Bertinoro, Italy on 1/6/04 and minuets of that meeting are available on the DELIS partners website. We aim to continue this and arrange further meetings in order to facilitate on-going co-ordination and cooperation. Relevant partners, in addition to regular e-mail discussions and electronic collaboration, will meet physically, to discuss progress and on-going work in two informal workshops between month 6 and month 18 of the project. We expect the workshops to be at least one full day and will comprise of presentations and detailed technical discussions.

In the context of WP dissemination when new techniques are identified we aim to disseminate results to relevant academic and industrial communities as quickly as possible by presenting and publishing results in recognised conferences, workshops and well-cited relevant journals. If considered necessary we will organise workshop events with relevant industrial partners in addition to academic communities to aid dissemination.

3.1 Cooperation through network dynamicity

As discussed in chapter 2 we have already produced models using individual-based [37] [42] and agent-based [41] models demonstrating how high cooperation can be maintained under assumptions of boundly rational (selfish) greedy behaviour within both the single-round Prisoners Dilemma and the file sharing scenario given in section 2.1.4. Over the next 12 months we will pursue the following goals and tasks. At this stage we feel it appropriate to be ambitious (but not unrealistic).

Goals:

- To develop both practical and theoretical understanding and reusable engineering techniques related to this unique cooperation forming process for application to the problems of self-organisation of cooperative interactions within decentralised and dynamic networks
- Where possible we wish to demonstrate working systems and formal proofs we consider this to be a gold standard which, while non-trival, is worth pursuing

Tasks:

- Identify application areas and produce simulations that address those areas with a higher degree of realism
- Prototype at least one application to the level of proof-of-concept, running on typical Internet infrastructure
- Develop analytical tools from Game Theory to help understand the process, identify necessary conditions for cooperation

- Link work from the social simulation literature on social networks to results and models where appropriate

3.2 The emergence of organisations

Of considerable current interest in the distributed systems engineering world is the idea of reusable, run-time decentralised component based programming (or even agent based mobile code). In such visions code is highly encapsulated and abstracted into a “component” or an “agent”. Such units combine at run-time dynamically over networks to achieve system level (user specified) goals. In agent-based systems the agents (often forms of mobile code) have their own goals and proactively seek other agents to form cooperative groups (or organisations) to achieve their mutual goals.

Goals:

- Identify, refine and apply behavioural ideas from organisational science (including empirical work on human systems) in order to understand how individual agents come to form organisations and to explore what kinds of organisation are stable, equitable and productive.
- Align this with both realistic human social phenomena and potential application areas with agent-based software systems over computer networks such that results can be of use to bother social scientists and engineers.
- Gain a deeper theoretical understanding of such models and methods by linking them to analytically tractable (abstracted) forms of the same process.

Tasks:

- Construction of (a set of) agent-based simulation models that capture the dynamics of organisational (firm) formation based on boundly rational behaviours of agents. The design will be based on existing Evolutionary Economics literature including empirical work. The primary focus of the model(s) will be on the interaction of skill-sets (i.e. knowledge and intellectual skills) in the formation of productive organisations.
- Identify an application domain within agent-based engineering Identify an application domain within social scientific analysis
- Where possible link to existing game theoretic notions Nash Equilibrium and evolutionary game theoretic tools such as ESS and the replicator dynamics.

3.3 Linking and deriving model sequences

It has been evident in the nature of this report that we have a problem akin to the excluded middle in our modelling approaches. Analytical tractability often calls for simplifying assumptions that tend to exclude some of the more interesting phenomena produced by (some might say overly complex) agent-based or individual-based simulation models.

However, as we have noted, simulation results dont count for proof. Specifically when a model is complex then the number of significant parameters and mechanisms hidden within it are too numerous to experimentally search. Does this mean that the two approaches are stuck on disparate islands? We think not. Indeed we look to work where a link is made between analytical proof and complex simulation results [9] [10] [7] [72]. Although often tentative and of a limited nature we consider such work as important in demonstrating best practice.

We want to generalise the method by which a complex intractable model may be simplified (via a series of steps or transitions) towards a more tractable version while retaining the phenomena of interested. Conversely we also see the reverse process is valuable i.e. how to take a proven technique and, via a series of simulation models, come to a technique deployable in a real computing environment (with all the noise and lack of control therein). Some previous work [38] [39] applied techniques from machine learning to automate certain simplifications method with limited success and researchers in social simulation have proposed model alignment or docking [8] [29] as away of checking different models for equivalence but as far as we are aware nobody has proposed a convincing general methodology for performing these operations.

Goals:

- Develop, through application to our outstanding problems (as previously highlighted), a general methodology or technique for generating new models from existing ones, either making them more relevant to an application (more complex) or more amenable to tractability (less complex) while retaining the phenomena of interest.

- Ideal result would be a generalised method for moving any given model of an emergent phenomena both towards application and towards analytic proof we do not envisage such a result is realistic hence our realistic goal is to start a process of investigation with this goal in mind.

Tasks:

- Generate a series of models (both simulation and analytical) for both 3.1 and 3.2 (above) capturing phenomena of interest at various levels of abstraction

- Develop sufficient methods for aligning these models and/or generating new models towards alignment if necessary

- Examine the use of machine learning techniques to automate some aspects of the process

- Produce results in the form of examples and generalised methods and their evaluation

3.4 Equilibrium Selection.

The most popular notion of stability in game theory is the Nash equilibrium (NE) [68]. Although the notion of NE is quite natural, if we restrict ourselves to pure strategies, we may not be able to reach a pure Nash equilibrium (PNE) at all. For example, in the Matching Pennies game shown in figure 3.1 there is no PNE at all. More importantly, even if we have PNE in a game, we may face a situation where multiple PNE exist, of different payoffs for the players and quite different aggregate performances. The problem then is for a rational player, how to decide which of the several NE is the "right" one to settle upon. To this direction, numerous refinements of the space of (P)NE have been proposed in the literature. In fact, there are so many refinements, that practically every NE may be shown to be the "right choice" of a proper refinement! There is a strong hope that evolutionary game theory will assist this kind of choices. The reason is the systematic way by which the evolution is modeled, as a kind of a reasonable game between an individual and some other individuals (or even against all other individuals). The point is to be able to construct computationally efficient algorithms for quantifying the values of all the NE, or even solving the corresponding optimization problem in the space of NE wrt a given objective function. Recall that for an evolutionary game the reachable NE are a subset of the rest points for the dynamics of this game, which in turn are the endpoints of (continuous in the limit) trajectories starting from some strategy initially adopted by (ie, prevailing among) the users. This should not be confused with the (computationally hard) task of determining the worst/best NE in a traditional game.

Additionally it would be extremely interesting to devise an algorithm for detecting the existence of an ESS in an evolutionary game, that only takes as input the $n \times n$ matrix U determining the payoffs of the game and responds with YES/NO, or even better, constructs an ESS in case of existence. Recall that an ESS is a NE of the corresponding bimatrix game with an additional

	Heads	Tails
Heads	0,1	1, 0
Tails	1,0	0, 1

Figure 3.1: The Matching Pennies game: The row player bets on different outcomes showing up, while the column player bets on the same outcomes showing up. For this game there is no PNE.

stability property, which does not necessarily hold in a NE. Indeed, one can construct simple examples where a small payoff matrix admits no ESS at all (although the game has at least one NE). A similar computational problem is the description of a computationally efficient evolutionary process (ie, the proper rule for evolution) that will lead as fast as possible to a rest point which is also a NE for the corresponding traditional game, or even an ESS for the evolutionary dynamics.

3.5 Bounded Transitions between stable points.

Apart from the problem of constructing Nash Equilibria or Evolutionary Stable Strategies, or any other stable points of a dynamical system, it is in many scenarios crucial to be able to move from one point to another within a bounded movement cost. For example, one might want, starting from a NE of given price of anarchy, to jump in polynomial time to a strictly better NE with specific characteristics, with, eg, better-reply movements, without ever reaching a profile with price of anarchy more than a times worse than that of the starting point. The nashification techniques are of similar flavour, but in this case we are in a position to start from an initial NE and we would like to explore the space of equilibria for an equilibrium with specific characteristics, avoiding to pay too much on the trajectory towards the destination. A similar issue is of interest for the ESS in evolutionary games: Starting from a given ESS and a given bound on the deterioration of the price of anarchy that we may cause temporarily, we would like to determine trajectories that will lead to new ESS with some special characteristics (eg, excluding undesirable actions).

3.6 Bounded vs. Unbounded Rationality.

A central assumption of the traditional game theory is that each player adopts a kind of *selfish behaviour*, exploiting the complete knowledge of other players' decisions. For example, in order to be able to assign a cardinal utility function to individual players, one typically assumes that each player has a well-defined, consistent set of preferences over the set of "lotteries" over the outcomes which may result from individual choice. Since the number of different lotteries over outcomes is uncountably many, this requires each player to have a well-defined, consistent set of uncountably many preferences, which is typically considered to be infeasible.

In many cases though this is not a feasible situation, either because it is not computationally efficient to handle such an amount of information, or because each player only knows the actual strategies of those players in its own neighborhood. Despite the fact that traditional game theory does not deal with such situations, there is a strong hope that evolutionary game theory will manage eventually to successfully describe and predict the behaviour of such players since it is better equipped to handle these *weaker rationality assumptions* and yet causing the same effects on players' behaviours in the long run as complete-knowledge traditional games. A typical example of such a scenario is when we wish to cut off the iteratively dominated strategies of the players. Depending on the model of evolution, we are able in some cases to assure that all these strategies, that a rational individual should never play in a traditional game, will eventually vanish in the evolutionary version of the game.

The computational point of view demands again to design algorithms dealing with bounded rationality and achieving in polynomial time (possibly good approximations of) the same outcome as the one expected by the corresponding traditional game.

3.7 Trajectory Prediction.

Due to its nature, evolutionary game theory explicitly models the dynamics present in interactions among individuals in a population. One might try to capture the dynamics of the decision-making process in traditional game theory by modeling the game in its extensive (rather than its strategic) form. However, for most games of reasonable complexity, the extensive form of the game quickly becomes unmanageable. Moreover, in the extensive form of a game, traditional game theory represents an individual's strategy as a specification of what choice that individual would make at each information set in the game. A selection of strategy then corresponds to a selection, prior to game play, of what that individual will do at any possible stage of the game. This representation of strategy selection clearly presupposes hyperrational players and fails to represent the process by which a player observes its opponents' behaviours, learns from these observations and makes a best/better response, replication, imitation, etc choice to what it has learned so far.

We would like to be able to decide in polynomial time whether two phenomenically different evolutionary dynamics actually produce the same trajectories, up to some kind of homeomorphism. This would enable a classification scheme of the evolution schemes into broader categories according to their orbital characteristics. Such kind of classification for replicator dynamics exist only when there are 2 or 3 distinct types of individuals in the population.

3.8 Imposing structural properties of a game into the evolutionary dynamics.

Our last question deals with some new ways of interaction in the evolutionary dynamics of a game, that will also depict the special structure of the corresponding traditional game. For example, when a virus spreads in a network, the architecture of the network itself and the starting points of the virus in it should affect somehow the success of the virus. The proposed game theoretic models of evolution proposed so far in the literature, mainly focus on the case where the individuals in a population collide with each other in a random fashion. Ie, the underlying "interaction" infrastructure is represented by a clique. What if this is not the case, and we have instead some special graph representing the interactions? We need new evolutionary models to capture such cases, that will somehow encode the structure of this graph in the dynamics, via elementary properties (eg, the connectivity or the expansion of the graph).

A smooth way of inserting the network affection in the evolution of a game, might be to use the power distribution on the graph distances of nodes, instead of the uniform distribution, in order to determine the contending nodes in each round. More specifically, we could let each user to choose its own opponent according to the probability distribution $\forall u \in V \setminus \{v\}, \mathbb{P}\{u \text{ is the opponent of } v\} = \mu \cdot [d(u,v)]^a$, where $d(u,v)$ is the length of the shortest path from u to v in the network, $a \geq 0$ is a tuning parameter and $\mu \equiv \sum_{w \in V \setminus \{v\}} [d(w,v)]^a$ is a normalizing constant. Observe that if we set $a = 0$ then we get the uniform distribution, while for $a \to \infty$ we force v to contend only with its $1-$neighbours in the graph. Such an approach may lead to phase transition phenomena in the appearance of equilibria or stable strategies with specific properties. This approach actually allows the smooth injection of the network structure into the evolutionary game.

Chapter 4

Summary / Conclusion

We present this report (D4.3.1) as an outline of our different approaches and traditions along with what we see as some common problems we can tackle with high relevance to DELIS objectives. We have also outlined our specific on-going projects with their desired goals and associated tasks.

We are particularly keen to attempt to bring together our approaches since this is rarely done elsewhere (we have made reference to work that we feel does this) but we have a realistic understanding of the difficult problems we are endeavouring to solve.

Our basic approach then, given this is the start of the DELIS project, is to be optimistic and ambitious at this stage and to confront what we view as our main obstacles head-on rather than attempt to explain them away or ignore them.

However, we still believe we are being realistic and hope that we have defined our work such that even negative results will be of value to the wider community and the on-going project.

We will report on our progress at the 18-month stage in deliverable D4.3.2.

Bibliography

[1] Adar E., Huberman B. A., *Free riding on gnutella*, First Monday **5** (2000), no. 10.

[2] Aldrich H.E., *Organizations and environments*, Prentice-Hall, 1979.

[3] Arthur W.B., *Self-reinforcing mechanisms in economics*, The Economy as an Evolving Complex System, Addison-Wesley, 1988.

[4] ______, *Competing technologies, increasing returns, and lock-in by historical events*, Economics Journal **99** (1989), no. 394, 116–131.

[5] ______, *Increasing returns and path-dependence in the economy*, University of Michigan Press, 1994.

[6] Astley G.W., Van De Ven A.H., *Central perspectives and debates in organization theory*, Administrative Science Quarterly **28** (1983), 245–273.

[7] Axelrod R., *The evolution of cooperation*, Basic Books, 1984.

[8] Axelrod R., Epstein J.M., Cohen M.D., *Aligning simulation models: A case study and results*, Computational and Mathematical Organization Theory **1** (1996).

[9] Binmore K., *Game theory and the social contract, i: Playing fair*, The MIT Press, 1984.

[10] ______, *Game theory and the social contract, ii: Just playing*, The MIT Press, 1994.

[11] Bourgeios, L.J.III, *Strategic management and determinism*, Academic Management Review **9** (1984), 586–596.

[12] Boylan R., *Laws of large numbers for dynamical systems with randomly matched individuals*, Journal of Economic Theory **57** (1992), 473–504.

[13] ______, *Continuous approximations of dynamical systems with randomly matched individuals*, Journal of Economic Theory **66** (1995), 615–625.

[14] Burgelman R.A., *Corporate entrepreneurship and strategic management: insights from a process study*, Management Science **29** (1983), 1349–1364.

[15] ______, *A model of interaction of strategic behaviour, corporate context, and the concept of strategy*, Academy of Management Review **8** (1983), no. 1, 61–70.

[16] ______, *A process model of internal corporate venturing in the diversified major firms*, Administrative Science Quarterly **28** (1983), 223–244.

[17] ______, *Intraorganizational ecology of strategy making and organizational adaptation: theory and field research*, Organization Science **2** (1991), 239–262.

[18] ______, *Fading memories: A process theory of strategic business exit in dynamic environments*, Administrative Science Quarterly **39** (1994), 24–56.

[19] Carroll G.R. Hannan M.T, *Demography of corporations and industries*, Princeton University Press, 2000.

[20] Child J., *Organizational structure, environment and performance: The role of strategic choice*, Sociology **6** (1972), 2–22.

[21] Conitzer V., Sandholm T., *Complexity results about nash equilibria*, IJCAI, 2003, pp. 765–771.

[22] Cressman R., *Evolutionary dynamics and extensive form games*, MIT Press, 2003.

[23] Cyert R.M., March J.G., *A behavioral theory of the firm*, Prentice-Hall, 1963.

[24] Czumaj A., Krysta P., Voecking B., *Selfish traffic allocation for server farms*, Proc. of the 34th ACM Symp. on Theory of Computing (STOC '02), 2002, pp. 287–296.

[25] Czumaj A., Voecking B., *Tight bounds for worst-case equilibria*, Proc. of the 13th ACM-SIAM Symposium on Discrete Algorithms (SODA '02), 2002, pp. 413–420.

[26] David P., *Technical choice, innovation, and economic growth*, Cambridge University Press, 1975.

[27] Dosi G., *Technological paradigms and technological trajectories: A suggested interpretation of the determinants and direction of technical change*, Research Policy **11** (1982), 147–62.

[28] Dosi G., Freeman C., Nelson R., Silverberg G., Soete L., *Technical change and economic theory*, Pinter Publishers, 1988.

[29] Edmonds B., Hales D., *Replication, replication and replication - some hard lessons from model alignment*, Journal of Artificial Societies and Social Simulation **6** (2003), no. 4.

[30] Fabrikant A., Papadimitriou C., Talwar K., *The complexity of pure nash equilibria*, Proc. of the 36th ACM Symp. on Theory of Computing (STOC '04), 2004.

[31] Feldmann R., Gairing M., Luecking T., Monien B., Rode M., *Nashification and the coordination ratio for a selfish routing game*, Proc. of the 30th International Colloquium on Automata, Languages and Programming (ICALP '03), Springer-Verlag, 2003, pp. 514–526.

[32] Fotakis D., Kontogiannis S., Koutsoupias E., Mavronicolas M., Spirakis P., *The structure and complexity of nash equilibria for a selfish routing game*, Proc. of the 29th International Colloquium on Automata, Languages and Programming (ICALP '02), Springer-Verlag, 2002, pp. 123–134.

[33] Fotakis D., Kontogiannis S., Spirakis P., *Selfish unsplittable flows*, Proc. of the 31st International Colloquium on Automata, Languages and Programming (ICALP '04), Springer-Verlag, 2004, p. to appear.

[34] Gairing M., Luecking T., Mavronicolas M., *The price of anarchy for polynomial social cost*, 27th International Symposium on Mathematical Foundations of Computer Science (MFCS'04), Springer-Verlag, 2004, p. to appear.

[35] Gairing M., Luecking T., Mavronicolas M., Monien B., *Computing nash equilibria for scheduling on restricted parallel links*, Proc. of the 36th ACM Symp. on Theory of Computing (STOC '04), 2004.

[36] ———, *Nash equilibria in discrete routing games with convex latency functions*, Proc. of the 31st International Colloquium on Automata, Languages and Programming (ICALP '04), Springer-Verlag, 2004, p. to appear.

[37] Hales D., *Cooperation without space or memory: Tags, groups and the prisoner's dilemma*, Proc. of the 2ndInternational Conference on Multi-Agent-Based Simulation (MABS2000), Springer-Verlag, 2000.

[38] ———, *Tag based cooperation in artificial societies*, PhD Thesis, University of Essex, Dept. of Computer Science, 2001.

[39] ———, *Understanding our creations: Using machine learning to understand machine models*, Centre for Policy Modelling Working Paper 02-96, Centre for Policy Modelling, Manchester Metropolitan University, UK, 2002.

[40] ———, *Change your tags fast - a necessary condition for cooperation?*, Joint Workshop on Multi-Agent and Multi-Agent-Based Simulation (MAMABS), Springer-Verlag, 2004.

[41] ———, *From selfish nodes to cooperative networks - emergent link-based incentives in peer-to-peer networks*, To appear in Proc. of the 4th IEEE International Conference on Peer-to-Peer Computing (P2P2004), IEEE Comput. Soc. Press, 2004.

[42] ———, *Self-organising, open and cooperative p2p societies - from tags to networks*, To appear in Proc. of the 2nd Workshop on Engineering Self-Organsing Applications (ESOA2004), Springer-Verlag, 2004.

[43] Hales D., Edmonds B., *Fashioning social simulations into engineering tools - the case of cooperation on p2p networks*, Centre for Policy Modelling Working Paper 04-134, Centre for Policy Modelling, Manchester Metropolitan University, UK, 2004.

[44] Hannan M.T, Carroll G.R., *Dynamics of organizational populations: Density, legitimation, and competition*, Oxford University Press, 1992.

[45] Hannan M.T, Freeman J., *The population ecology of organization*, American Journal of Sociology **82** (1977), 929–64.

[46] ———, *Structural inertia and organizational change*, American Sociological Review **49** (1984), 149–64.

[47] ———, *The ecology of organizational mortality: American labor unions, 1836-1985*, American Journal of Sociology **92** (1987), 910–43.

[48] ———, *The ecology of organizational founding: American labor unions, 1836-1985*, American Journal of Sociology **94** (1988), 25–52.

[49] ———, *Organizational ecology*, Harvard University Press, 1989.

[50] Hodgson G., *Economics and evolution; bringing life back into economics*, Polity Press, 1993.

[51] Hofbauer J., Sigmund K., *Evolutionary game dynamics*, Bulletin of the American Mathematical Society **40** (2003), no. 4, 479–519.

[52] Holland J., *Adaptation in natural and artificial systems*, The MIT Press, 1992.

[53] ______, *The effect of labels (tags) on social interactions*, SFI Working Paper 93-10-064, Santa Fe Institute, 1993.

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

[55] Koutsoupias E., Papadimitriou C., *Worst-case equilibria*, Proc. of the 16thAnnual Symposium on Theoretical Aspects of Computer Science (STACS '99), Springer-Verlag, 1999, pp. 404–413.

[56] Koza J., *Genetic programming: On the programming of computers by means of natural selection*, The MIT Press, 1992.

[57] Lai K., Feldman M., Stoica I., Chuang J., *Incentives for cooperation in peer-to-peer networks*, Proc. of the 1st Workshop on Economics of Peer-to-Peer Systems (June 5-6, 2003, Berkeley, CA), http://www.sims.berkeley.edu/research/conferences/p2pecon/, 2003.

[58] Leonard-Barton D., *Core capabilities and core rigidities: a paradox in managing new product development*, Strategic Management Journal **13** (1992), 111–125.

[59] Luce R., Raiffa H., *Games and decisions*, John Wiley and Sons, 1957.

[60] Luecking T., Mavronicolas M., Monien B., Rode M., *A new model for selfish routing*, Proc. of the 21stAnnual Symposium on Theoretical Aspects of Computer Science (STACS '04), Springer-Verlag, 2004, pp. 547–558.

[61] Matsui A., Gilboa I., *Social stability and equilibrium*, Econometrica **59** (1991), 859–867.

[62] Mavronicolas M., Spirakis P., *The price of selfish routing*, Proc. of the 33rd ACM Symp. on Theory of Computing (STOC '01), 2001, pp. 510–519.

[63] May R. et al, *Spatial games and evolution of cooperation*, Advances in Artificial Life, Proc. of the 3rd European Conference on Artificial Life, Springer-Verlag, 1995.

[64] Milchtaich I., *Congestion games with player-specific payoff functions*, Games and Economic Behavior **13** (1996), 111–124.

[65] Miner A.S., *Seeking adaptive advantage: evolutionary theory and managerial action*, Evolutionary Dynamics of Organizations, Oxford University Press, 1994.

[66] Miner A.S., Mittman B.S., *An intraorganizational ecological perspective on managerial risk behaviour, performance, and survival: individual, organizational, and environmental effects*, Evolutionary Dynamics of Organizations, Oxford University Press, 1994.

[67] Monderer D., Shapley L., *Potential games*, Games and Economic Behavior **14** (1996), 124–143.

[68] Nash J. F., *Noncooperative games*, Annals of Mathematics **54** (1951), 289–295.

[69] Nelson R., *Recent evolutionary theorizing about economic change*, J. of Economic Literature **33** (1995), 48–90.

[70] Nelson R., Winter S., *An evolutionary theory of economic change*, Harvard University Press, 1982.

[71] Nowak M., May R., *Evolutionary games and spatial chaos*, Nature **359** (1992), 826–929.

[72] Nowak M., Sigmund K., *Evolution of indirect reciprocity by image scoring*, Nature **359** (1998), 573–577.

[73] Papadimitriou C., *Algorithms, games and the internet*, Proc. of the 33rd ACM Symp. on Theory of Computing (STOC '01), 2001, pp. 749–753.

[74] Riolo R., *The effect of tag-mediated selection of partners in evolving populations playing the iterated prisoner's dilemma*, SFI Working paper 97-02-016, Santa Fe Institute, 1997.

[75] Riolo R., Cohen M., Axelrod R., *Evolution of cooperation without reciprocity*, Nature **414** (2001), 441–443.

[76] Rosenthal R.W., *A class of games possessing pure-strategy nash equilibria*, International Journal of Game Theory **2** (1973), 65–67.

[77] Roughdarden T., Tardos E., *How bad is selfish routing?*, J. Assoc. Comput. Mach. **49** (2002), no. 2, 236–259.

[78] Roughgarden T., *The price of anarchy is independent of the network topology*, Proc. of the 34th ACM Symp. on Theory of Computing (STOC '02), 2002, pp. 428–437.

[79] Samuelson L., *Evolutionary games and equilibrium selection*, Series on Economic Learning and Social Evolution, MIT Press, 1997.

[80] Sun Q., Garcia-Molina H., *SLIC: A selfish link-based incentive mechanism for unstructured peer-to-peer networks*, Proc. of the 24th IEEE International Conference on Distributed Systems (ICDCS2004), IEEE Comput. Soc. Press, 2004.

[81] Taylor P. D., Jonker L., *Evolutionary stable strategies and game dynamics*, Mathematical Biosciences **40** (1978), 145–156.

[82] Thompson J.D., *Organization in action*, McGraw-Hills, 1967.