6. Structural Robustness of Evolutionary Models in Game Theory[♣]

6.1. Introduction

Naturally, the method that scientists have traditionally followed to advance our formal understanding of evolutionary social interactions has been to design and study models that were tractable with the tools of analysis available at the time. Until not long ago, such tools have derived almost exclusively from the realm of mathematics, and they have given rise to mainstream Evolutionary Game Theory (EGT). Mainstream EGT has proven to be tremendously useful (Weibull, 1995), but it is founded on many assumptions made to ensure that the resulting models could be mathematically analysed (e.g. infinite and homogeneous populations, random encounters, infinitely repeated interactions…). The aim of this chapter is to assess the extent to which some of these assumptions are affecting the conclusions obtained in mainstream EGT.

The assumptions made in EGT for the sake of mathematical tractability have had important implications both in terms of the *classes of systems* that have been investigated, and in terms of the *kind of conclusions* that have been drawn concerning such systems.

In terms of *classes of systems*, in order to achieve mathematical tractability, EGT has traditionally analysed *idealised systems*, i.e. systems that *cannot* exist in the real world (e.g. a system where the population is assumed to be infinite). Typically, mainstream EGT has also imposed various other assumptions that simplify the analysis, but which do not necessarily make the system ideal in our terminology (i.e. unable to exist in the real world). Some examples of common

[♣] Some parts of the material presented in this chapter have been published in Izquierdo, L. R., Izquierdo, S. S., & Polhill, J. G. (2006), "EVO-2x2: a modelling framework to study the evolution of strategies in 2x2 symmetric games under various competing assumptions", in *Proceedings of the First World Congress on Social Simulation*, Kyoto, Japan, Vol. 2, pp. 273-280, and in Izquierdo, S.S. and Izquierdo, L.R. (2006). On the Structural Robustness of Evolutionary Models of Cooperation. *Lecture Notes in Computer Science* 4224, pp. 172-182.

assumptions in EGT are: populations are *well-mixed* (each individual is equally likely to interact with any other individual), interactions are *infinitely repeated*, strategies are *deterministic* and there is a *finite* set of them, individuals are selected with probabilities *proportional* to their fitness, and invasions are *homogenous* and *arbitrarily small*. Applying mainstream EGT to non-idealised systems can be very problematic because the validity for non-idealised systems of conclusions drawn from extremely similar idealised systems is not as straightforward as one may think. As an example, Beggs (2002) demonstrates that when analysing some types of evolutionary idealised systems, results can be widely different depending on the order in which certain limits are taken: if one takes the limit as population size becomes (infinitely) large and then considers the limit as the force of selection becomes strong, then one obtains different results from those attained if the order of the limits is inverted. Thus, Beggs (2002) warns that "care is therefore needed in the application of these approximations".

The need to achieve mathematical tractability has also influenced the *kind of conclusions* obtained in mainstream EGT. Thus, mainstream EGT has focused on analysing the stability of incumbent strategies to arbitrarily small mutant invasions, but has not paid much attention to the overall dynamics of the system in terms of e.g. the size of the basins of attraction of different evolutionary stable strategies, or the average fraction of time that the system spends in each of them.

Nowadays it has just become possible to start addressing the limitations of mainstream EGT outlined above. The current availability of vast amounts of computing power through the use of computer grids is enabling us to conduct formal and rigorous analyses of the dynamics of non-idealised systems through an adequate exploration of their sensitivity both to basic parameters and to their structural assumptions. These analyses can complement previous studies by characterising dynamic aspects of (idealised and non-idealised) systems beyond the limits of mathematical tractability. It is this approach that we follow in this chapter.

The structure of this chapter is as follows: section 6.2 outlines the general research question that EGT is mainly concerned with, and explains how our approach can complement the work conducted in mainstream EGT. Section 6.3 describes EVO-2x2, a computer simulation modelling framework designed to formally assess the impact of various assumptions commonly made in mainstream EGT. The subsequent two sections illustrate the use and the usefulness of EVO-2x2 with a particular example. The specific application selected here is a study of the structural robustness of evolutionary models of cooperation. To put our work into context, section 6.4 provides a brief and critical review of some of the most relevant work conducted on the evolutionary emergence of cooperation within the realms of game theory. Section 6.5 summarises some of the most interesting results we have obtained and the method we followed to analyse and summarise them. Finally, section 6.6 presents the conclusions of this investigation.

6.2. Overall research question and approach

In very broad terms, the question that EGT tries to answer is usually of the form: "In a population of individuals who repeatedly interact with each other, what sort of behavioural traits are likely to emerge and be sustained under evolutionary pressures?". Naturally, the answer to such a question may depend on a number of assumptions regarding population size, population structure (i.e. how individuals meet to interact), the specific nature of each interaction, the mechanisms through which natural selection occurs, and how mutations take place. In this chapter we present a formal modelling framework (EVO-2x2) designed to address this general question from different angles, i.e. using various different assumptions. EVO-2x2 provides a single coherent framework within which results obtained from different models can be contrasted and compared with analytical approaches. Thus, EVO-2x2 can be used to investigate the impact of various assumptions which may all be valid when trying to answer the general question posed above.

EVO-2x2 implements a wide range of competing plausible assumptions, all of which are fully consistent with the most basic principles of the theory of evolution. Logically, the assumptions embedded in EVO-2x2 limit its applicability. The most stringent assumption in EVO-2x2 is arguably the fact that interactions are modelled as 2-player 2-strategy (2x2) symmetric games. We will see in the next section, however, that individuals in EVO-2x2 are explicitly and individually represented, so any simulation conducted in EVO-2x2 is a non-idealised system (i.e. a system that could potentially exist in the real world). This move towards greater realism implies some loss of mathematical tractability, e.g. closed-form analytical solutions for the systems modelled in EVO-2x2 are not readily available. Nevertheless, EVO-2x2 is simple enough so many insights can be gained by using the theory of stochastic processes to analyse the results obtained by performing many simulation runs with it, as will be shown later. The following section explains all the assumptions embedded in EVO-2x2 in detail. Subsequently we illustrate the use of EVO-2x2 by studying the structural robustness of evolutionary models of cooperation.

6.3. Description of EVO-2x2

EVO-2x2 is a computer simulation modelling framework designed to formally investigate the evolution of strategies in 2x2 symmetric games under various competing assumptions. EVO-2x2 enables the user to set up and run many computer simulations (effectively many different models) aimed at investigating the same question using alternative assumptions. The specific question to be addressed is: "In a population of individuals who interact with each other by repeatedly playing a certain 2x2 symmetric game, what strategies are likely to emerge and be sustained under evolutionary pressures?".

6.3.1. The conceptual model

In this section we explain the conceptual model that EVO-2x2 implements. The information provided here should suffice to re-implement the same conceptual model on any platform. Figure 6-1 provides a snapshot of EVO-2x2 interface, which is included here to clarify the explanation of the model. The reader may also want to consider following the explanation of the model using it at the same time; EVO-2x2 is included in the Supporting Material of this thesis. We use bold red italicised arial font to denote ***parameter*** names.

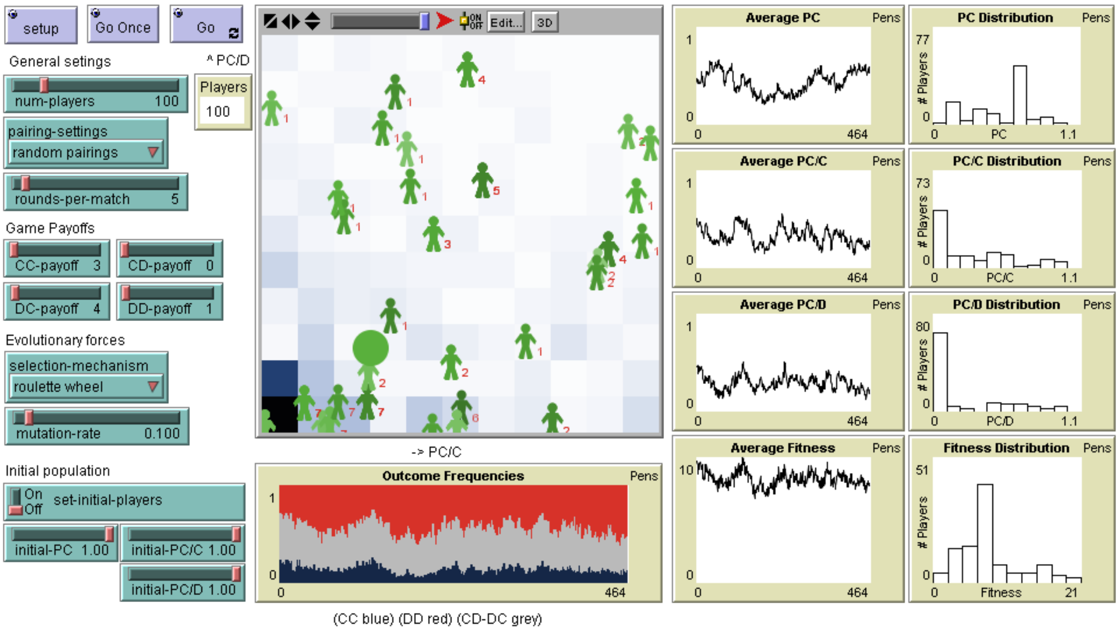

Figure 6-1. Snapshot of the interface in EVO-2x2.

Overview of EVO-2x2

In EVO-2x2, there is a population of ***num-players*** players. Events occur in discrete time-steps, which can be interpreted as successive generations. At the beginning of every generation every player's payoff (which denotes the player's fitness) is set to zero. Then, every player is paired with another player, according to some customisable procedure (***pairing-settings***), to play a 2-player match.

Each match consists of a number of sequential rounds (***rounds-per-match***). In each round, the two members of the pair play a symmetric 2x2 game once, where each of them can undertake one of two possible actions. These two possible actions are called cooperate (C) and defect (D). The action selected by each of the players determines the magnitude of the payoff that each of them receives in that round (***CC-payoff***, ***CD-payoff***, ***DC-payoff***, ***DD-payoff***). The total payoff that a player obtains in a match is the sum of the payoffs obtained in each of the rounds.

Players differ in the way they play the match, i.e. they generally have different strategies. The strategy of a player is determined by three numbers in the interval [0 , 1]:

- *PC*: Probability to cooperate in the first round.

- *PC/C*: Probability to cooperate in round n ($n > 1$) given that the other player has cooperated in round ($n - 1$).
- *PC/D*: Probability to cooperate in round n ($n > 1$) given that the other player has defected in round ($n - 1$).

Once every player has played one –and only one– match (except when the pairing mechanism is *round robin*, as explained below), two evolutionary processes (i.e. natural selection (***selection-mechanism***) and mutation (***mutation-rate***)) come into play to replace the old generation with a brand new one. Successful players (those with higher payoffs) tend to have more offspring than unsuccessful ones. This marks the end of a generation and the beginning of a new one, and thus the cycle is completed.

Parameters

The value of every parameter in EVO-2x2 can be modified at run-time, with immediate effect on the model. This enables the user to interact closely with the model by observing the impact of changing various assumptions during the course of one single run.

Population parameters

num-players: Number of players in the population. This number is necessarily even for pairing purposes.

set-initial-players: This is a binary variable that is either *on* or *off*. If *on*, every player in the initial population will have the same strategy, which is determined using the following parameters: ***initial-PC***, ***initial-PC/C***, and ***initial-PC/D***. If *off*, the initial population of strategies will be created at random using a uniform distribution.

Rounds and Payoffs

rounds-per-match: Number of rounds in a match.

CC-payoff: Payoff obtained by a player who cooperates when the other player cooperates too.

CD-payoff: Payoff obtained by a player who cooperates when the other player defects.

DC-payoff: Payoff obtained by a player who defects when the other player cooperates.

DD-payoff: Payoff obtained by a player who defects when the other player also defects.

Pairing settings

This parameter (**pairing-settings**) determines the algorithm that should be used to form pairs of players. There are three options:

- *random pairings*: Pairs are made at random, without any bias. Every player plays one and only one match in a generation.
- *round robin*: Every player is paired with every other player once, so every player plays exactly (**num-players** – 1) matches per generation.
- *children together*: Players are paired preferentially with their siblings (and at random among siblings). Once all the possible pairs between siblings have been made, the rest of the players are paired at random. Every player plays one and only one match in a generation. This procedure was implemented because it seems plausible in many biological contexts that individuals belonging to the same family tend to interact more often among them than with individuals from other families. The algorithm is formally equivalent to simple applications of tags (Holland, 1993) in evolutionary models (see Hales, 2000).

Evolutionary forces

selection-mechanism: This parameter determines the algorithm used to create the new generation. There are four options:

- *roulette wheel*: This procedure involves conducting **num-players** replications, which form the new generation. In each replication, players from the old generation are given a probability of being chosen to be replicated that is proportional to their total payoff (which denotes their fitness).
- *Moran process*: In each time-step (i.e. generation), one player is chosen for replication with a probability proportional to its fitness. The offspring replaces a randomly chosen player (possibly its parent). Payoff totals are set to zero at the beginning of every time-step.

- *winners take all*: This method selects the player(s) with the highest total payoff (i.e. the “winners”). Then, for **num-players** times, a random player within this “winners set” is chosen to be replicated. The **num-players** replications constitute the new generation. Note that this mechanism (which is sometimes called “cultural imitation”, e.g. see Traulsen et al., 2006) violates the proportional fitness rule.
- *tournament*: This method involves selecting two agents from the population at random and replicating the one with the higher payoff for the next generation. In case of tie, one of them is selected at random. This process is repeated **num-players** times. The **num-players** replications form the new generation.

mutation-rate: This is the probability that any newly created player is a mutant. A mutant is a player whose strategy (the 3-tuple formed by *PC*, *PC/C*, and *PC/D*) has been determined at random.

6.3.2. Displays

EVO-2x2 provides various displays which are shown in Figure 6-1. Some of these displays are time-series plots showing the historical evolution of the value of a particular variable throughout generations (e.g. frequency of outcomes and population average values of *fitness*, *PC*, *PC/C*, and *PC/D*), whereas others refer only to the last generation (e.g. population distributions of *fitness*, *PC*, *PC/C*, and *PC/D*).

The large square in the middle of the interface is the representation in the strategy space of every individual player in a generation. This representation is 2-dimensional in EVO-2x2 due to constraints in the modelling platform (NetLogo 3.0.2), but we also provide in the Supporting Material a 3D version of EVO-2x2, called EVO-2x2-3D (implemented in NetLogo 3-D Preview 1), where the three dimensions of the strategy space (*PC*, *PC/C*, and *PC/D*) are explicitly represented. This is the only difference between EVO-2x2-3D and EVO-2x2: EVO-2x2-3D represents players in the *PC–PC/C–PC/D* 3-dimensional strategy space, while EVO-2x2 displays the projection of such a space on the *PC/C–PC/D* plane (Figure 6-2). In Figure 6-2, the sphere (in the left-hand image) and its circular projection (in the right-hand image) indicate population averages.

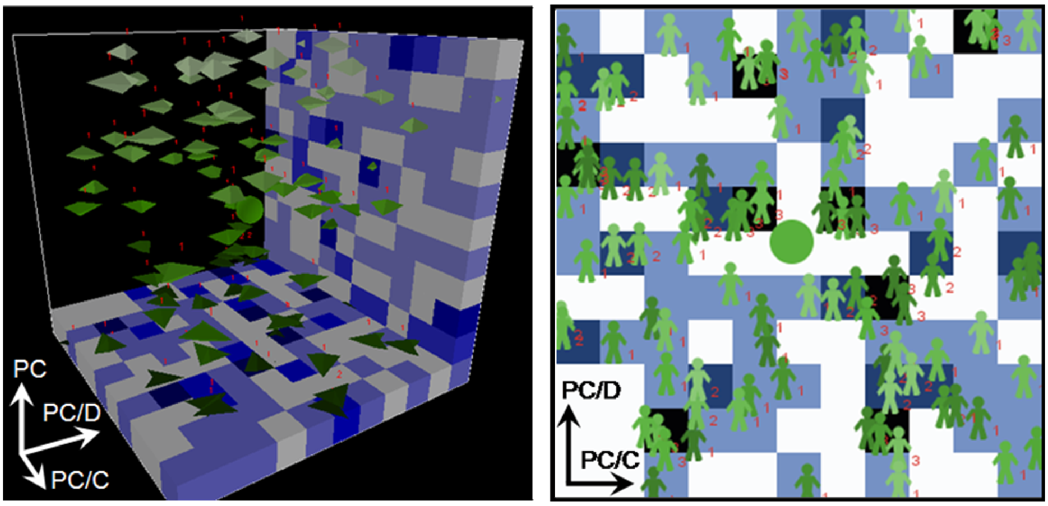

Figure 6-2. Representation of players in the strategy space using EVO-2x2-3D (left) and EVO-2x2 (right). The image on the right shows the top-down projection of the representation on the left.

The cells in the background of the 2-dimensional projections of the strategy space are coloured in shades of blue according to the number of players that have spent some time on them. Each player that has visited a certain part of the strategy space leaves a mark that is used to create the density plots shown in Figure 6-2. The more players who have stayed for longer in a certain area, the darker its shade of blue.

6.3.3. Exploration of the parameter space

The rationale behind EVO-2x2 was to conduct a systematic exploration of the impact of various competing assumptions. An exploration of the parameter space is something that can be easily conducted within NetLogo using a tool called BehaviorSpace. This tool allows the user to set up and run experiments. Running an experiment consists in running a model many times, systematically varying the model's settings and recording the results of each model run.

The problem when undertaking experiments that involve large parameter sweeps is to organise, analyse, and summarise the vast amount of information obtained from them so the results can be meaningfully interpreted. To do that, we have created a set of supporting scripts (written in Perl and Mathematica, and available in the supporting Material) that are able to read in the definition of the experiment setup and all its results in the format used by NetLogo. The output of these scripts is:

- an automatically generated directory structure that reflects all the combinations of parameter values explored in the experiment (e.g. /100/random-pairings/roulette-wheel/0.001/…/), and
- a customisable summary of the results of each model run, which is placed in the appropriate folder.

An example of a useful summary of the results produced in a simulation run is the accumulated frequency of different types of strategies throughout the course of a simulation run. This is something that can be plotted in a 3D contour plot, and in complementary 2D density plots, as shown in Figure 6-3. The relationship between the 3D contour plot and the accompanying 2D density plots is sketched in Figure 6-4.

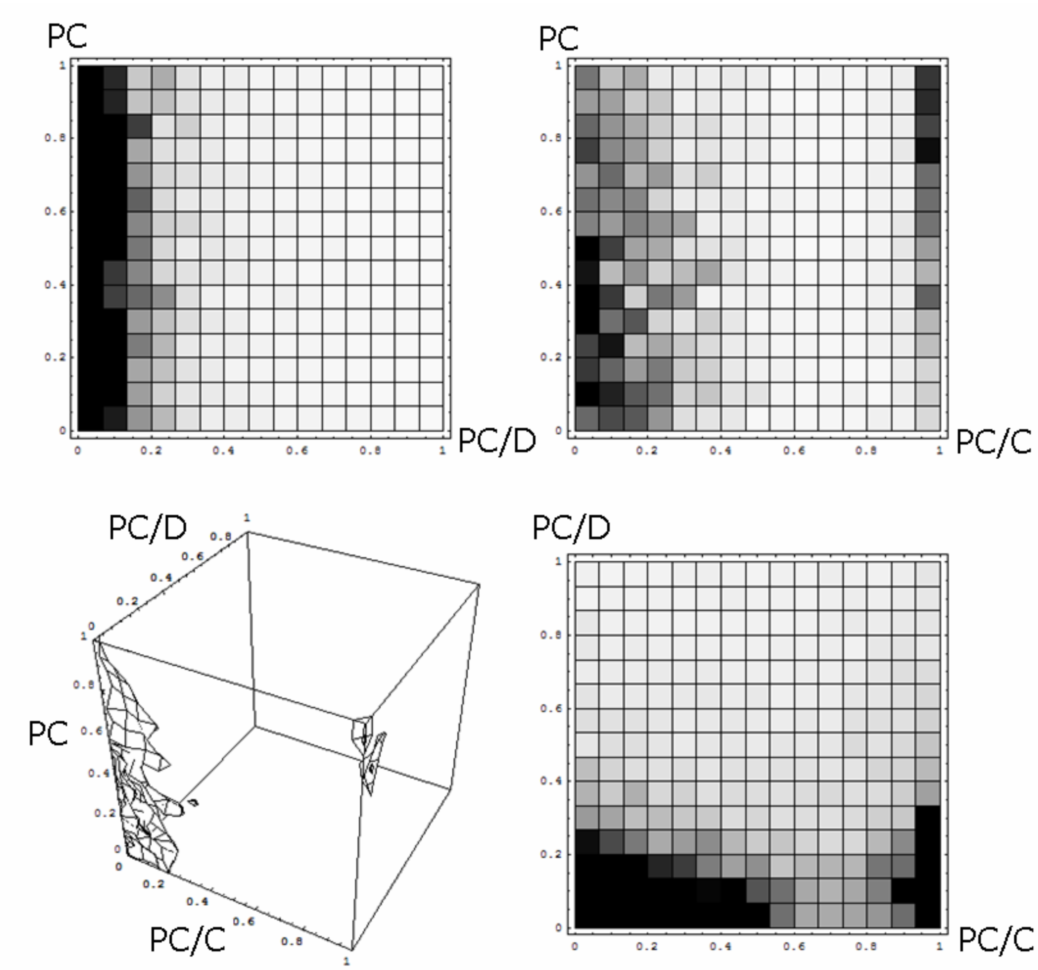

Figure 6-3. Example of a graphical summary of the results obtained with EVO-2x2. This figure is automatically created and placed in the appropriate folder by the supporting scripts.

6.3.4. Implementation details

EVO-2x2 has been implemented in NetLogo 3.0.2 (Wilensky, 1999). We also provide a 3-D version of EVO-2x2, called EVO-2x2-3D, which has been implemented in NetLogo 3-D Preview 1 (Wilensky, 1999). The two programs are available in the Supporting Material together with a user guide under the GNU General Public Licence.

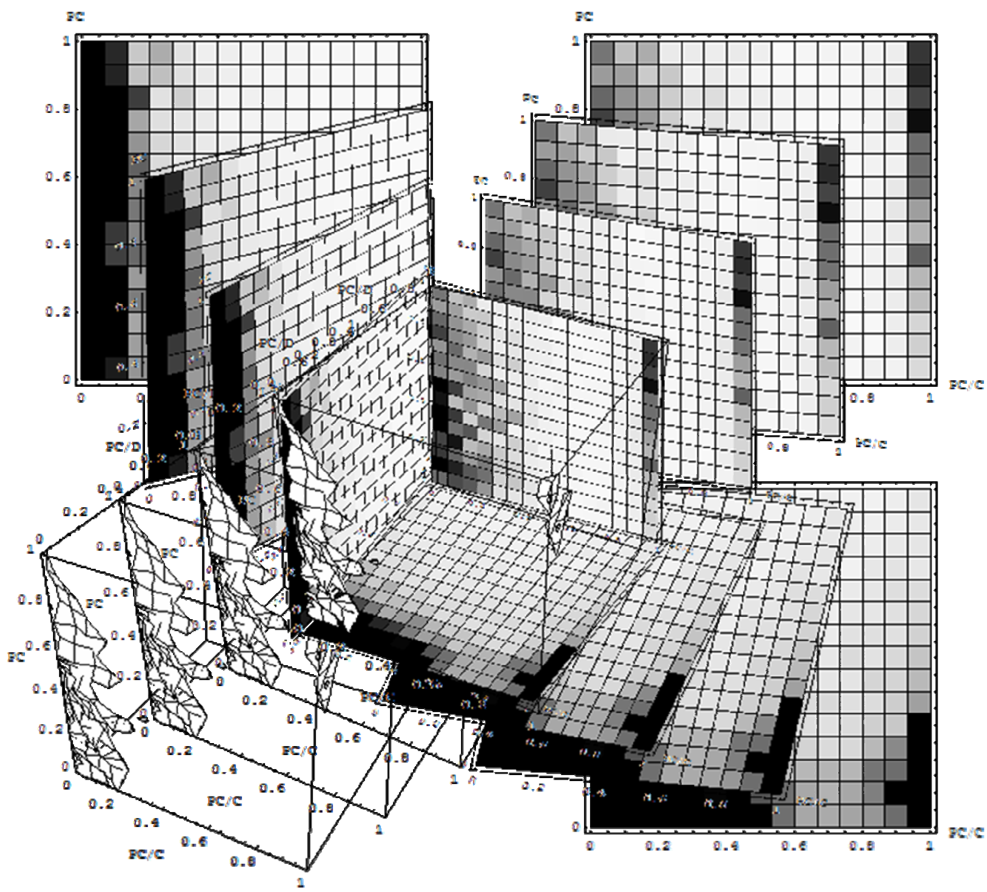

Figure 6-4. Sketch showing the relationship between the 3D contour plot and the accompanying 2D density plots created by the supporting scripts.

6.4. Evolutionary emergence of cooperation

The fundamental challenge of understanding the evolutionary emergence and stability of cooperation can be illuminated, at the most elementary level, by identifying the conditions under which a finite number of units that interact by playing the Prisoner's Dilemma (PD) may cooperate. These units might be able to adapt their individual behaviour (i.e. learn), or the population of units as a whole may adapt through an evolutionary process (or both). While formalizing the problem of cooperation in this way significantly decreases its complexity (and generality), the question still remains largely unspecified: how many units form the population? How do they interact? What strategies can they use? What is the value of each of the payoffs in the game? and, crucially, what are the processes governing the dynamics of the system?

It has been well known since the early years of the study of the evolution of cooperation that, in general, the question of how –if at all– cooperation emerges in a particular system significantly depends on all of the above defining characteristics of the system (see e.g. Axelrod, 1984; Bendor and Swistak, 1995, 1997, 1998; Gotts et al., 2003b). Here we report previous work that has shed light on the robustness of evolutionary models of cooperation. We find it useful to place these models in a fuzzy spectrum that goes from mathematically tractable models with strict assumptions that limit their applicability (e.g. work on idealised systems), to models with the opposite characteristics. The rationale behind the construction and use of such a spectrum is that when creating a formal model to investigate a certain question (e.g. the evolution of cooperation), there is often a trade-off between the applicability of the model (determined by how constraining the assumptions embedded in the model are) and the mathematical tractability of its analysis (i.e. how deeply the functioning of the model can be understood given a certain set of available tools of analysis).

The former end is mostly populated by models *designed to* ensure mathematical tractability. Near this end we find papers that study the impact of some structural assumptions, whilst still keeping others which ensure the model remains tractable and which, unfortunately, also tend to make the model retain its idealised nature. Gotts et al. (2003b) review many such papers in sections 2 and 4. Some of these investigations have considered finite vs. infinite populations (Nowak et al., 2004; Taylor et al., 2004; Imhof et al., 2005), different pairing settings or population structures (see section 6 in Gotts et al. (2003b) for a review, and Santos et al. (2006) for the most recent advances in this field), deterministic vs. stochastic strategies (Nowak, 1990; Nowak and Sigmund, 1990; Nowak and Sigmund, 1992), finite vs. infinitely repeated games (Nowak and Sigmund, 1995), and arbitrary intensities of selection (Traulsen et al., 2006). While illuminating, the applicability of most of these studies is somewhat limited since, as mentioned before, the models investigated there tend to retain their idealised nature.

Near the opposite end, we find models that tend to be slightly more applicable (e.g. they consider non-idealised systems), but they are often mathematically intractable. It is from this end that we start in our investigation. To our knowledge, the first relevant study with these characteristics was conducted by Axelrod (1987). As explained in section 3.1, Axelrod had previously organized two open tournaments in which the participant strategies played an iterated PD in a round robin fashion (Axelrod, 1984). Tit for Tat (TFT) was the winner in both tournaments, and also in an *ecological analysis* that Axelrod (1984) conducted after the tournaments. Encouraged by these results, Axelrod (1987) investigated the generality of TFT's success by studying the evolution of a randomly generated population of strategies (as opposed to the arguably arbitrary set of strategies submitted to the tournament) using a particular genetic algorithm. The set of possible strategies in this study consisted of all deterministic strategies able to consider the 3 preceding actions by both players. From this study, Axelrod (1987) concluded that in the long-term, "reciprocators […] spread in the population, resulting in more and more cooperation and greater and greater effectiveness". However, the generality of Axelrod's study (1987) is doubtful for two reasons: (1) he used a very specific set of assumptions, the impact of which was not tested, and (2) even if we constrain the scope of his conclusions to his particular model, the results should not be trusted since Axelrod only conducted 10 runs of 50 generations each. As a matter of fact, Binmore (1994, p. 202; 1998) cites unpublished work by Probst (1996) that contradicts Axelrod's results.

In a more comprehensive fashion, Linster (1992) studied the evolution of strategies that can be implemented by two-state Moore machines in the infinitely repeated PD. He found a strategy called GRIM remarkably successful. In particular, GRIM was significantly more successful than TFT. GRIM always cooperates until the opponent defects, in which case it switches to defection forever. Linster (1992) attributed the success of GRIM over TFT to the fact that GRIM is able to exploit poor strategies while TFT is not. Linster's investigation was truly remarkable at its time, but technology has advanced considerably since then, and we are now in a position to expand his work significantly by conducting parameter explorations beyond what was possible before. As an example, note that Linster (1992) could only consider deterministic strategies and one specific value for the mutation rate; furthermore, in the cases he studied where the dynamics were not deterministic, there is no guarantee that his simulations had reached their asymptotic behaviour.

Another important part of the literature on the study of the evolutionary emergence of cooperation using computer simulation comes from the use of tags. Tags are socially recognisable marks or signals that, in principle, are not necessarily linked to any particular form of behaviour (Holland, 1993). Tags do, however, influence the way individuals interact: individuals with similar tags have a preference to interact with each other (see e.g. Riolo (1997), Hales (2000), Riolo et al. (2001), Edmonds and Hales (2003)). Tags, like strategies, are also assumed to be passed from parents to their kin. Thus, tags and strategies follow a very similar evolutionary process. The resulting correlation between tags and strategies leads to a tendency for individuals with similar strategies to interact with each other. In the context of social dilemmas this correlation clearly favours cooperative behaviours, as it effectively diminishes the chances of exploitation.

Riolo (1997) developed the first tag model in the study of the evolutionary emergence of cooperation in the PD. He showed that real-valued tags can promote high levels of cooperation in the iterated PD. Hales (2000) developed Riolo's work and studied discrete tags, with preferential pairings occurring only if tags matched exactly. With this exact tag matching constraint, cooperation can emerge even when players interact for only one round. Hales' pairing mechanism is formally equivalent to "children-together" in EVO-2x2 (see section 6.3.1). Tags as a useful mechanism to promote cooperation were further explored by Riolo et al. (2001). This piece of work, however, turned out to be flawed, as it relied upon the fact that individuals were forced to donate to others with an identical tag (see Roberts and Sherratt (2002) and Edmonds and Hales (2003) for a much more in-depth investigation). Since then research using tags has worked towards making this cooperation mechanism more robust, so it can be usefully applied in real-world contexts (see e.g. Hales and Edmonds (2005), and Edmonds (2006)).

In the following section we use EVO-2x2 to conduct a consistent and systematic exploration of the impact of competing assumptions in non-idealised evolutionary models of cooperation.

6.5. Robustness of evolutionary models of cooperation

In this section we illustrate the usefulness of EVO-2x2 by applying it to advance our formal understanding of the structural robustness of evolutionary models of cooperation. To do this, we analyse simple non-idealised models of cooperation and we study their sensitivity to small structural changes (e.g. slight modifications in the way players are paired to interact, or in how a generation is created from the preceding one). Specifically, we aim to determine what behavioural traits are likely to emerge and be sustained under evolutionary pressures in the Prisoner's Dilemma (PD). To do this rigorously, we have run many computer simulations (effectively many different models) aimed at addressing the same question: "In a population of individuals who interact with each other by repeatedly playing the PD, what strategies are likely to emerge and be sustained under evolutionary pressures?". Given the amount of computing power required to conduct this research, all the simulations have been run on computer grids.

6.5.1. Method followed to analyse the simulation results

Defining a state of the system as a certain particularisation of every player's strategy, it can be shown that all simulations in EVO-2x2 with positive mutation rates can be formulated as irreducible positive recurrent and aperiodic discrete-time finite Markov chains. Thus, ergodicity is guaranteed. This observation enables us to say that there is a unique long-run distribution over the possible states of the system, *i.e.* initial conditions are immaterial in the long-run (Theorem 3.15 in Kulkarni (1995)). Although calculating such (dynamic) distributions analytically is infeasible, we can estimate them using the computer simulations. The problem is to make sure that a certain simulation has run for long enough, so the limiting distribution has been satisfactorily approximated. To make sure that this is the case, for each possible combination of parameters considered, we ran 8 different simulations starting from widely different initial conditions. These are the 8 possible initial populations where every individual has the same pure strategy (the 8 corners of the strategy space). Then, every simulation run is conducted for 1,000,000 generations. Thus, in those cases where the 8 distributions are similar, we have great confidence that they are showing a distribution close to the limiting distribution[37]. As an example, consider Figure 6-5, where distributions starting from the 8 different initial conditions are compared.

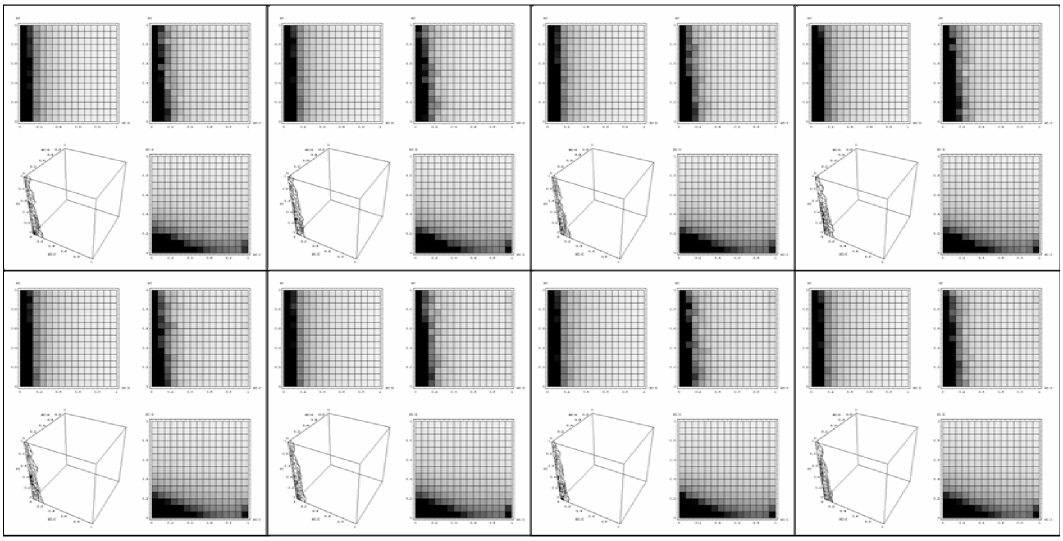

Figure 6-5. Accumulated frequency of different types of strategies in 8 simulation runs starting from different initial conditions. Axes are as in Figure 6-3.

6.5.2. Results and discussion

In this section we report several cases where it can be clearly seen that some of the assumptions in EGT that are sometimes thought to have little significance (e.g. mutation-rate, number of players, or population structure) can have a major

[37] The appropriateness of the inductive method used here (which is not formal proof) to infer the asymptotic distribution of the system can be qualitatively checked by thinking what would happen if this method were to be applied to study the system characterised in chapter 4. In that case, the method would consist in running 4 simulations starting from the corners of the strategy space. Clearly, simulations starting in an SRE would stay there forever. Thus, only in those cases where there is really a unique asymptotic distribution, would the 4 simulations eventually look similar, and only when very close to the limiting distribution. In other words, the method used here would work perfectly well for the system characterised in chapter 4: the 4 cumulative distributions would look similar if and only if they were close to the limiting distribution.

impact on the type of strategies that emerge and are sustained throughout generations. The following are parameter values that are common to all the simulations reported here[38]:

CC-payoff = 3; **CD-payoff** = 0; **DC-payoff** = 5; **DD-payoff** = 1;

selection-mechanism = *roulette wheel*;

Consider first the two distributions in Figure 6-6, which only differ in the value of the mutation rate used (0.01 on the left, and 0.05 on the right). The distribution on the left shows the evolutionary emergence and (dynamic) permanence of strategies similar to TFT (*PC* ≈ 1, *PC/C* ≈ 1, and *PC/D* ≈ 0; average time ≈ 3.3%). Such strategies are observed one order of magnitude less frequently for slightly higher mutation rates (distribution on the right; average time ≈ 0.3%). The other parameter values used were **num-players** = 100; **pairing-settings** = *random pairings*; **rounds-per-match** = 50.

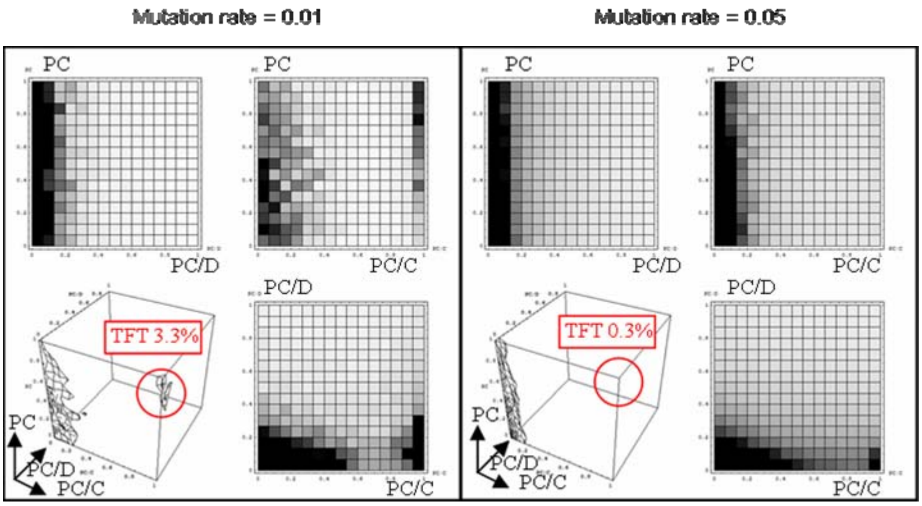

Figure 6-6. Influence of the mutation rate on the dynamics of the system. TFT measures the average time that strategies with *PC* ≥ (13/15), *PC/C* ≥ (13/15) and *PC/D* ≤ (2/15) were observed.

The two distributions in Figure 6-7 only differ in the number of players in the population (100 on the left, and 10 on the right). The distribution on the left shows

[38] The payoffs used in this chapter are those employed by Axelrod (1984), and consequently those used in most simulation papers on the evolution of cooperation. They are used here too to facilitate comparisons with previous research.

the evolutionary emergence and (dynamic) permanence of strategies similar to TFT (average time ≈ 3.3%), whereas –again– such strategies are observed one order of magnitude less frequently in smaller populations (average time ≈ 0.4%). The other parameter values are: **pairing-settings** = *random pairings*; **rounds-per-match** = 50; **mutation-rate** = 0.01.

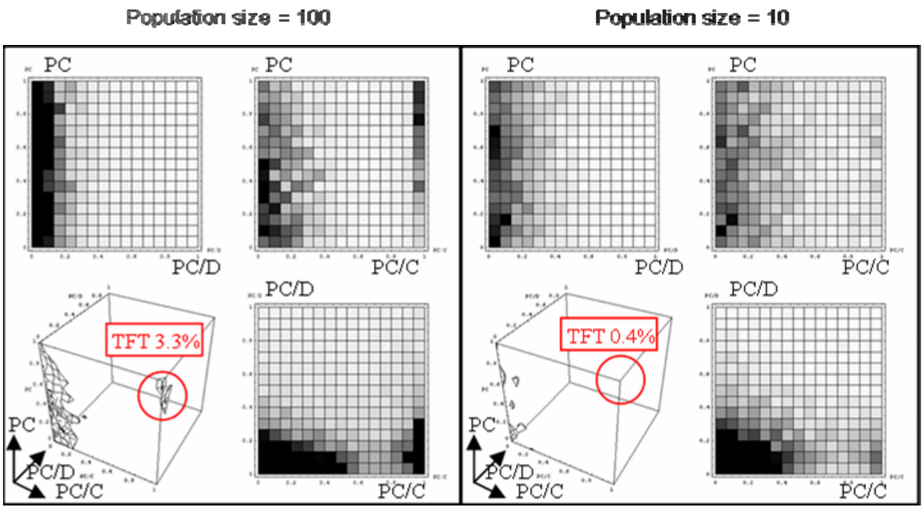

Figure 6-7. Influence of the number of players in the population. TFT measures the average time that strategies with $PC \geq (13/15)$, $PC/C \geq (13/15)$ and $PC/D \leq (2/15)$ were observed.

The two distributions in Figure 6-8 only differ in the algorithm used to form the pairs of players (*random pairings* on the left, and *children together* on the right). On the left, strategies tend to be very similar to ALLD ($PC \approx 0$, $PC/C \approx 0$, and $PC/D \approx 0$), i.e. strongly uncooperative (average time ALLD ≈ 72%). In stark contrast, the distribution on the right is concentrated around strategies similar to TFT (average time TFT ≈ 23%; average time ALLD ≈ 1%). The other parameter values used were: **num-players** = 100; **rounds-per-match** = 5; **mutation-rate** = 0.05. The underlying reason behind the dramatic increase in cooperation when using the pairing algorithm "children together" (which is formally equivalent to simple applications of tags, see e.g. Hales, 2000) is that this mechanism promotes mimicry. Children, who have inherited the same strategy from their parents, tend to be paired together. This confers a great evolutionary advantage to cooperation, since it effectively rules out the possibility of exploitation: cooperators (and defectors) play only with each other.

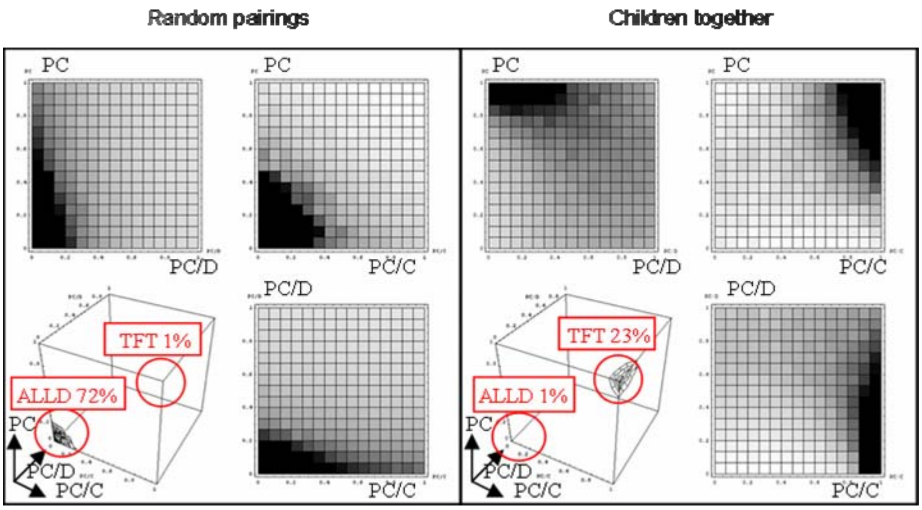

Figure 6-8. Influence of different pairing mechanisms. TFT measures the average time that strategies with $PC \geq (10/15)$, $PC/C \geq (10/15)$ and $PC/D \leq (5/15)$ were observed; ALLD measures the average time that strategies with $PC \leq (5/15)$, $PC/C \leq (5/15)$ and $PC/D \leq (5/15)$ were observed.

Figure 6-9 shows a very interesting result. The two distributions in Figure 6-9 only differ in the set of possible values that *PC*, *PC/C* or *PC/D* can take. For the distribution on the left the set of possible values is any (floating-point) number between 0 and 1, and the strategies are mainly uncooperative, similar to ALLD (average time ALLD ≈ 60%). For the distribution on the right, the set of possible values is only {0, 1}, and the distribution is concentrated in TFT (average time TFT ≈ 58%). The other parameter values used were: **num-players** = 100; **mutation-rate** = 0.05; **rounds-per-match** = 10;
pairing-settings = *random pairings*.

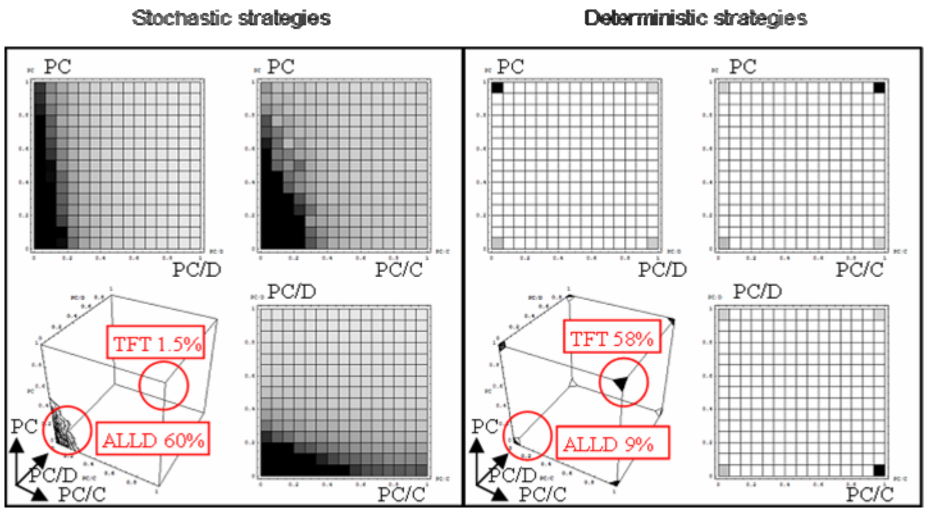

Figure 6-9. Stochastic (mixed) strategies vs. deterministic (pure) strategies: influence in the system dynamics. TFT measures the average time that strategies with $PC \geq (10/15)$, $PC/C \geq (10/15)$ and $PC/D \leq (5/15)$ were observed; ALLD measures the average time that strategies with $PC \leq (5/15)$, $PC/C \leq (5/15)$ and $PC/D \leq (5/15)$ were observed.

Given the clarity and importance of the results presented in Figure 6-9 we investigated this issue further. In Figure 6-10 and Figure 6-11 we show the effect of gradually increasing the set of possible values for *PC*, *PC/C* and *PC/D* (i.e. ***num-strategies***). Figure 6-10 shows the (average) number of each possible outcome of the game (CC, CD/DC or DD) in observed series of 10^6 matches (this number of matches is selected so the effect of changing the initial state is negligible, i.e. results are close to the stationary limiting distribution).

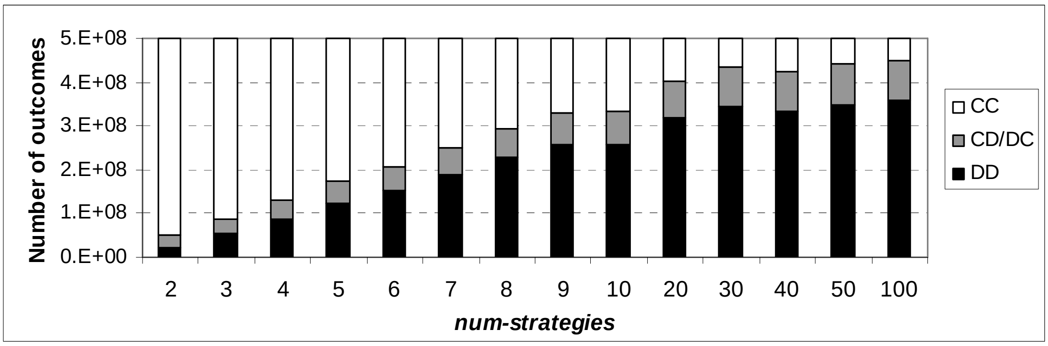

Figure 6-10. Influence in the distribution of outcomes (CC, CD/DC or DD) of augmenting the set of possible values for *PC*, *PC/C* and *PC/D*.

Figure 6-11 shows the average values of *PC*, *PC/C* and *PC/D* observed in the same series. Augmenting the set of possible values for *PC*, *PC/C* and *PC/D* undermines cooperation and favors the emergence of ALLD-like strategies. The other parameter values used were: **num-players** = 100; **mutation-rate** = 0.01; **rounds-per-match** = 10; **pairing-settings** = *random pairings*.

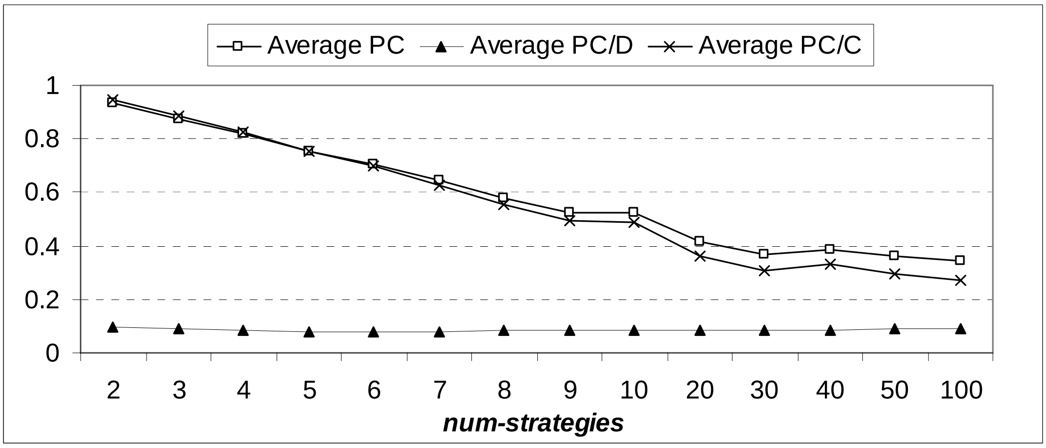

Figure 6-11. Influence of augmenting the set of possible values for *PC*, *PC/C* and *PC/D* in the average values of these variables in the population.

Thus, it is clear that the number of possible strategies has a tremendous effect on the evolutionary stability of cooperation. This is mainly due to the fact that the emergence of TFT-like behaviour crucially relies on perfect reciprocation. A single defection in a contest between two TFT-like strategies with high –but lower than 1– values of *PC/C* will result in a chain of uncoordinated outcomes CD-DC, thus losing much of their evolutionary advantage over ALLD.

6.6. Conclusions of this chapter

In this chapter we have shown by example that some of the assumptions made in mainstream evolutionary game theory for the sake of mathematical tractability can have a greater effect than what has been traditionally thought. In particular, the granularity of the strategy space and the assumption of well-mixed populations have proved to be critical in determining the type of strategies that are likely to emerge and be sustained in evolutionary contexts.

More specifically, this chapter has studied the structural robustness of evolutionary models of cooperation, i.e. their sensitivity to small structural changes. To do this, we have focused on the Prisoner’s Dilemma game and on the set of stochastic strategies that are conditioned on the last action of the player's counterpart. Strategies such as Tit-For-Tat (TFT) and Always-Defect (ALLD) are particular and classical cases within this framework; here we have studied their potential appearance and their evolutionary robustness, as well as the impact of small changes in the model parameters on their evolutionary dynamics. Our results show that strategies similar to ALLD tend to be the most successful in most environments, whereas strategies similar to TFT tend to spread best in large populations, where individuals with similar strategies tend to interact more frequently, when only deterministic strategies are allowed, with low mutation rates, and when interactions consist of many rounds.