8. Conclusions

This thesis was initiated with the overall aim of advancing game theory by formally studying the implications of dropping some of its most stringent assumptions, which have been made for the sake of tractability and are not generally supported by empirical evidence.

Naturally, the first part of this research consisted in clearly identifying the most relevant and prevalent assumptions made in the different branches of game theory. This investigation led to the critical dissection of deductive game theory presented in chapter 2, which served as a guiding framework to structure the rest of the research conducted in this thesis. In particular, this critical review enabled a precise identification of those assumptions of game theory that are abandoned and those that are retained in the models developed in this thesis. Specifically, all the research conducted here abandons the strong assumptions made in classical game theory regarding player's rationality, players' beliefs about their counterparts' behaviour, and the alignment of such beliefs across players. The research conducted in this thesis also abandons the assumption of one single *infinite* population, which is commonly made in evolutionary game theory, and which was shown in chapter 2 to have wider implications than may be initially suspected.

The abandonment of several assumptions that are made in game theory to allow for mathematical tractability has meant that new methodologies were needed to formally analyse the models developed in this thesis. In particular, computer simulation has proven to be particularly useful to enhance and complement mathematical derivations. The combined use of analytical work and computer simulation has enabled me to draw some methodological conclusions that are also included in this chapter.

The structure of this final chapter is particularly simple. Section 8.1 summarises the main contributions of this thesis to the advancement of game theory. These are presented at two different levels of abstraction for the sake of clarity: subsections 8.1.1 and 8.1.2 present the specific contributions of this thesis to the advancement of learning and evolutionary game theory respectively (and the implications of these for the study of social dilemmas), whereas subsection 8.1.3 discusses in more general terms the wider implications of the research conducted here for game theory as a whole. The methodological conclusions derived from the symbiotic use of computer simulation and mathematical analysis are then summarised in section 8.2. Finally, the last section of this chapter (8.3) identifies areas for future research.

8.1. Contributions to the advancement of game theory

8.1.1. Specific contributions to learning game theory

Chapter 4 of this thesis provided an in-depth analysis of the transient and asymptotic dynamics of the Bush-Mosteller reinforcement learning algorithm, whereas chapter 5 explored cased-based reasoning as decision-making process in strategic contexts. The specific insights obtained for each of these learning algorithms were summarised in sections 4.10 and 5.8 respectively. The following presents the main conclusions that can be drawn from this investigation in more general terms:

- The transient dynamics of models in learning game theory can be substantially different from their asymptotic behaviour. Moreover, some systems may take an extraordinarily long time to reach their asymptotic dynamics (see e.g. Figure 4-8 and Figure 5-8). This is especially important because most theoretical research focuses on the characterisation of asymptotic equilibria exclusively, whereas studies using computer simulation tend to explore only the short-term dynamics of models.
- The transient dynamics of models in learning game theory tend to be very complex and highly path-dependent (see e.g. section 5.4). Players learn from each other's actions in a very dynamic fashion, and their individual responses affect every player's payoff (and –consequently– their subsequent behaviour). This means that one single decision made by one player may change the evolution of the whole system substantially and have a permanent effect on its overall dynamics (especially in models without "trembling hands noise").

- It has been long known that the inclusion of “trembling hands noise” can affect the dynamics of models in learning game theory. This thesis has illustrated that this type of noise can *completely* change the dynamics of a model by showing that some outcomes that are observed with arbitrarily high probability in unperturbed models can effectively lose all their attractiveness if players make occasional mistakes in selecting their actions (see e.g. sections 4.8 and 5.7).
- In general, occasional mistakes made by players can destabilise outcomes in two different ways: by giving the deviator a higher payoff, or by giving any of the non-deviators a lower payoff. Thus, outcomes where unilateral deviations hurt the deviator (strict Nash) but not the non-deviators (protected) tend to be the most stable (see sections 4.8 and 5.7.3).

The application to social dilemmas of the models developed in this thesis (and the review of similar models in the literature) has enabled me to draw the following general conclusions in this regard:

- Cooperation in social dilemmas is not only a common outcome in models where players learn from each other’s behaviour, but also the unique asymptotic outcome in many cases (see sections 4.1, 4.5 and 5.4).
- Cooperative outcomes are most commonly observed in models where players satisfice to some extent: they have an aspiration threshold that divides the set of outcomes into two classes: satisfactory and unsatisfactory outcomes. Naturally, aspiration thresholds that make the cooperative outcome satisfactory and the non-cooperative outcome unsatisfactory tend to promote the highest rates of cooperation (see sections 4.7 and 5.4).
- Cooperative outcomes tend to be particularly susceptible to be destabilised by small trembles. This is so because deviations have two undesirable effects: they favour the deviator *and* they hurt the non-deviators. Therefore trembles in cooperative outcomes encourage all cooperating players to change their behaviour. On the other hand, non-cooperative outcomes are particularly robust to trembles because deviations from them hurt the

deviator *and* benefit the non-deviators, thus encouraging everyone to keep defecting (see sections 4.8 and 5.7.3).

8.1.2. Specific contributions to evolutionary game theory.

Chapter 6 described EVO-2x2, the modelling framework developed in this thesis to assess the impact of various assumptions made in mainstream evolutionary game theory for the sake of mathematical tractability. The following summarises the main conclusions that can be drawn from this investigation in general terms (for more specific conclusions see section 6.6):

- The study of the evolution of finite populations is significantly different from that of infinite populations (both in terms of the methods that are adequate for their analysis and on the results obtained with them). This fact has serious implications, since most of our intuitions about evolutionary dynamics come from analyses of models where populations are infinite.
- Stochastic effects (e.g. the potential occurrence of two or more mutations at the same time) play an important role in the analysis of finite evolutionary systems (see sections 2.3.4 and 6.5).
- The type of strategies that are likely to emerge and be sustained in finite evolutionary contexts is strongly dependent on assumptions that traditionally have been thought to be unimportant or secondary (e.g. number of players, continuity of the strategy space, mutation rate, and population structure). See results presented in section 6.5.2.
- There seems to be great value in developing general frameworks that facilitate rigorous and transparent comparisons between different stochastic finite models and the results obtained with them.

The use of EVO-2x2 was illustrated by conducting an investigation on the structural robustness of evolutionary models of cooperation. The results obtained in that research (and other papers in the literature – see e.g. Imhof et al., 2005) showed that stochastic evolution of *finite* populations need not select the strict Nash equilibrium (as is the case when making the assumptions of mainstream evolutionary game theory) and can therefore favour cooperation over defection. Stochastic finite systems exhibit dynamics over the strategy space with time averages that –for some parameterisations– are concentrated around cooperative strategies (e.g. TFT; see section 6.5.2).

8.1.3. General contributions to game theory

The dissection of game theory made in chapter 2 of this thesis (and some of the issues discussed in section 7.5) showed that classical game theory is founded on rather problematic assumptions that may have deeper philosophical implications than commonly assumed. Fortunately, this has been increasingly acknowledged in the last few years, and several models that abandon the demanding assumptions of classical game theory on players' rationality and beliefs have been put forward and analysed in depth. This reasonably new programme of research, to which the present thesis contributes, is starting to provide fruitful insights.

This thesis in particular has thoroughly analysed the dynamics of two models of learning that have received notable empirical support (see chapters 4 and 5). In this way, the work reported here enhances game theorists' toolkit of models that can be usefully employed to study real-world systems. One of the main challenges that game theory faces nowadays derives from the need of managing and synthesising the various insights obtained with a number of disparate models that abandon the stringent assumptions of game theory through different avenues. This diversity of new assumptions and results calls for the creation of frameworks aimed at facilitating a clear and transparent comparison between models and the results obtained with them. This thesis has tried to meet this challenge by placing its contributions in an overall framework that can encompass, in admittedly very broad terms, most of the research conducted in game theory until now (see chapter 2). In the particular context of evolutionary game theory, the modelling framework developed in chapter 6, i.e. EVO-2x2, represents a step forward in this direction too. Using EVO-2x2, it has been demonstrated here that some of the assumptions made in mainstream evolutionary game theory for the sake of mathematical tractability can have a greater effect than has been traditionally thought. Specifically, 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 (see section 6.5).

Thus, in general terms, this thesis has contributed to game theory (a) by examining the formal implications of replacing some of the unsupported assumptions in mainstream game theory with assumptions that stem from empirical research, and (b) by creating frameworks aimed at making differences between models explicit and at facilitating the comparison of results obtained with different models.

8.2. Methodological contributions

Before the development of computational modelling, the formal analysis of game theoretical models could be conducted using mathematical analyses only, and this may have distorted our understanding of such models to some extent. This thesis has shown that computer modelling can greatly enhance and complement mathematical derivations. These two techniques to analyse formal systems are both extremely useful, and they are complementary in the sense that they can provide fundamentally different insights on the same issue. Chapter 4 is a clear illustration of the fact that the level of understanding gained by using these two techniques together could not have been obtained using either of them on their own. Thus, the use of only one of these techniques may lead to an incomplete picture of the dynamics of a model. Chapter 4 also illustrates how each technique can produce both problems and hints for solutions for the other.

This thesis has also shown that most models in learning and evolutionary game theory can be usefully formalised as Markov processes. In the absence of noise, these tend to have many different recurrent classes (i.e. areas of the state space that cannot be escaped once entered). In such cases, one single (stochastic) decision made by one player may lead the system to one or another recurrent class (and completely change the properties of the resulting dynamics), making the formal analysis of these models very challenging (see e.g. section 5.4). The inclusion of some kind of noise (e.g. mutations or trembling hands) tends to simplify the analysis to a great extent, since it often means that all the states of the system communicate (and this most often implies that the stochastic process is ergodic). On a slightly more negative note, this fact also demonstrates that very small changes in the assumptions of a model may have quite an important effect on its dynamics. In any case, this thesis has illustrated that the theory of Markov processes can be particularly useful to analyse formal models of social interactions, and it has also provided various indications on which specific mathematical results may be most valuable depending on the properties of the system to be analysed (see e.g. sections 3.2.2 and 4.1).

8.3. Areas for future work

8.3.1. Assessment of the philosophical foundations of game theory

As noted by some authors (see e.g. Hargreaves Heap and Varoufakis, 1995, pp. 14-18), game theory is rooted in philosophical foundations that are not free from controversy. One of the most contentious issues in this regard concerns the concept of instrumental rationality used in classical game theory (see section 2.2.2). Critically studying the philosophical foundations of game theory seems to be a matter of great importance for at least two reasons: because most economists and many game theorists seem to be almost unaware that the foundations of game theory are at the very least debatable, and because a richer notion of rationality may provide game theory with the intuitive appeal and logical coherence that some of its analyses lack (Hargreaves Heap and Varoufakis, 1995, p. 14). This thesis in particular (see section 7.5) has outlined the basis of a potential line of future research based on a new form of reasoning, i.e. reasoning by outcomes. This proposed area of research could potentially lead to more plausible solution concepts that could capture more of the intuitional knowledge (i.e. heuristics) that people seem to implicitly use in their social interactions.

8.3.2. Learning algorithms vs. Rationality

As explained in section 2.4.1, a current limitation of learning game theory is that most models assume that every player in the game follows the same decision-making algorithm. Thus, in many of these models the observed dynamics may be very dependent on the fact that the game is played among “cognitive clones”, and the extent of this effect is not often evaluated. Confronting the investigated learning algorithms with alternative decision-making algorithms seems to be a promising way forward in learning game theory. In particular, confronting learning algorithms with highly rational players seems to have the potential to be very illuminating.

8.3.3. Evolution of learning algorithms

As explained in section 2.4, one of the main differences between evolutionary and learning game theory is the level at which adaptation takes place[41]. Adaptation processes in evolutionary models occur at the population level: populations are subject to evolutionary pressures (and therefore the population adapts), but the individual components of populations may not adapt at all (i.e. they may have a predefined fixed behaviour). On the other hand, adaptation processes in learning models take place at the individual level through learning, and it is this learning process that is formally described[42]. Most current efforts to integrate these two branches of game theory aim at drawing similarities between the (mean-field) dynamics of certain learning algorithms and an appropriate version of the replicator dynamics (see e.g. Börgers and Sarin (1997), Laslier et al. (2001), Hopkins (2002), Laslier and Walliser (2005), Hopkins and Posch (2005), Beggs (2005)). A complementary (and less pursued) way in which these two branches can be integrated to some extent consists in analysing models that incorporate adaptation processes both at the individual and at the population level, i.e. studying the evolution of different learning algorithms (Kirchkamp, 1999, 2000). Playing with the relative strength of these two levels at which adaptation may take place is likely to offer new insights on the conditions that may favour the evolutionary emergence of certain reasoning processes over others.

8.3.4. Stochastic approximation theory

This thesis and a significant number of papers in the literature (see the brief review presented in section 4.1) have benefited immensely from recent developments in the theory of stochastic approximation. This theory is devoted, in particular, to identifying the conditions under which the actual dynamics of a stochastic system can be approximated by an appropriately constructed deterministic model. Further developments in the theory of stochastic

[41] Another important difference relates to the interpretation of payoffs in each of these branches of game theory (see section 2.1).

[42] Another difference between these two branches of game theory relates to the *nature* of the adaptation process that is modelled. Adaptation in evolutionary models takes place through processes of selection and mutation (see section 2.3), while this is not necessarily the case in learning models (see section 2.4).

approximation theory will undoubtedly enable game theorists to better understand their models, and also to analyse the dynamics of models that were previously intractable. Furthermore, developing our understanding of the relations between stochastic and deterministic models is likely to provide new insights on the relation between learning and evolutionary game theory (Weibull, 2002).

8.3.5. Development of frameworks

This thesis has extensively argued for the value of frameworks at several points (see e.g. sections 2.4.1, 6.1 and 7.4). The wide variety of models developed in the last few years in game theory calls for the creation of frameworks aimed at facilitating the process of model comparison, both in terms of their assumptions and in terms of the results obtained with them. As argued in section 7.4, the development of frameworks is useful not only to assess the impact of various assumptions in theoretical terms, but also to inform experimental research. Thus, the use of frameworks may facilitate the interaction between game theorists and other social scientists, an area for future work that is outlined below.

8.3.6. Greater interaction with other social sciences

There is clearly a lot to gain from the interaction of game theory and other social sciences. Traditionally, game theory has developed almost entirely from introspection and theoretical concerns. Whilst the work developed in game theory up until now has proven to be tremendously useful, it seems clear that game theory will not fulfil all its potential as a useful practical tool to analyse real-world social interactions unless a greater effort is made to interact with other social sciences. In particular, a closer interaction with more empirically-driven social scientists is likely to increase the applicability and relevance of game theory for the study of real-world social interactions. Ideally, this interaction should not be postponed until the stage in the research where a theoretical model is to be validated; on the contrary, empirical research (both experimental and field work) can suggest exciting and relevant avenues where theoretical research may be most needed. In this way, empirical and theoretical work can usefully drive, shape, and benefit from each other. As Weibull (2002) says, “perhaps this is the beginning of a new phase in economic research where economists get together with psychologists, sociologist, and social anthropologists”. Let us make it happen.