2. Main assumptions in game theory

This chapter is a critical dissection of the main assumptions embedded in each of the most advanced branches of deductive game theory at this time. We distinguish between game theory as a framework (which makes no assumptions about individuals' behaviour or beliefs), classical game theory, evolutionary game theory, and learning game theory. Given the breadth and depth of game theory work, this thesis cannot present an exhaustive list of all the assumptions considered in the field. We focus on the most prevalent and relevant ones. The critical review of deductive game theory in this chapter is meant to serve as a framework where the assumptions whose impact is investigated in the subsequent chapters of this thesis can be precisely identified. It will also serve to identify what assumptions are retained in the models developed in this thesis. The last section of this chapter briefly describes some of the branches of game theory that are not purely deductive.

2.1. Game theory as a framework

Game theory as a framework is a methodology used to build models of real-world social interactions. The result of the modelling exercise is a game, i.e. a formal abstraction of the social interaction which is typically defined by[2]:

- the set of individuals who interact (called *players*),
- the different choices available to each of the individuals (called *strategies*),
- and a *payoff* function that assigns a (usually numerical) value to each individual for each possible combination of choices made by every individual.

Importantly, the abstract model developed within this framework does not make any assumptions about the players' behaviour, neither in a normative nor in a positive sense.

---

[2] We use here the representation of a game in strategic form for the sake of clarity. If the sequential structure of the game is complex (in terms of order of movement, players' asymmetries and information flow), the representation of the game in extensive form (which *explicitly* details the order of events, the order of moves, and the information sets) may be more adequate (see chapter 1 in Vega-Redondo (2003) for details).

Game theory as a framework is particularly useful to describe and analyse decision-making in social interactions where the outcome potentially depends on the decisions made by several individuals (i.e. interdependent decision-making processes). According to the Stanford Encyclopaedia of Philosophy, "game theory is the most important and useful tool in the analyst's kit whenever she confronts situations in which what counts as one agent's best action (for her) depends on expectations about what one or more other agents will do, and what counts as their best actions (for them) similarly depend on expectations about her" (Ross, 2006).

As with any formal model, some of the complexity of the real-world situation represented will be lost in the process of abstraction. The rationale to undertake such a process of abstraction, which implies loss of descriptive accuracy to some extent, is that it will yield insights beyond those that could be achieved without the model. Furthermore, the knowledge acquired from the analysis of the abstract formal model can still be valid in other real-world situations whose important features are captured by the same formal model even though the model was not initially developed with such situations in mind. To the extent that the formal model captures the essence of the situation under study, enables us to establish inference processes that we could not undertake otherwise, and yields insights that can be transferred to other domains, we consider that the formal model is useful (Colman, 1995, pg. 6).

Game theory as a framework makes two important assumptions. The first one is ontological and it refers to how social interactions are modelled in game theory. The framework used in game theory makes a clear distinction between structure (i.e. rules of the game) and action. The rules of the game fully constrain the set of possible actions that can be taken, i.e. there is no room for action to change structure. Obviously this is not the only ontological view that one can take when trying to distil the essence of social interactions. This clear cut between structure and action will prove useful in many circumstances, but it may not always be adequate; therefore it is important to be aware that there are many other ways of modelling social interactions (Hargreaves Heap and Varoufakis, 1995, chapter 1).

Assuming that the essence of the social interaction to be modelled is captured by the formal abstraction to a satisfactory extent (in terms of context, interplay between action and structure, history effects…) the most important assumption made when using game theory as a framework relates to the definition of the *payoff* function. In most branches of game theory, payoffs are meant to represent individuals' preferences for each possible outcome of the social interaction. The notable exception is evolutionary game theory, where payoffs most often (but not always) represent Darwinian fitness. The following two sections explain this in detail.

2.1.1. Payoffs interpreted as preferences

The payoff function for each player is effectively a preference ordering over the set of possible outcomes. Behind the concept of "payoff function" is the implicit assumption that preferences will guide action (otherwise there would not be much point in defining a payoff function). While seemingly innocuous, this underlying assumption does have certain philosophical implications which, though fascinating, fall out of the scope of this thesis (Hargreaves Heap and Varoufakis, 1995, pg. 12).

A common misconception about game theory relates to the roots of players' preferences. There is no assumption in game theory (not even as a framework) that players' preferences are formed in complete disregard of each other's interests. On the contrary, preferences in game theory are assumed to account for everything, i.e. they may include altruistic motivations, moral principles, and social constraints, for example (Colman, 1995, pg. 301; Vega-Redondo, 2003, pg. 7).

Game theory as a framework assumes that players' preference order is well defined, i.e. it satisfies the conditions of reflexivity, completeness, and transitivity (Hargreaves Heap and Varoufakis, 1995, pg. 6); and that their preference order does not change. If no further assumption is made on individuals' preferences, these are said to be *ordinal*. Ordinal preferences provide no information about the strength of preferences, so arithmetic operations on ordinal payoffs are not meaningful. An admittedly obvious point, but one which may be worth noting, is that direct comparisons of ordinal preferences between different players (e.g. “player A likes outcome X more than player B does”) are meaningless.

In almost all game theoretical models, however, preferences are assumed to be cardinal, i.e. payoffs take numerical values on an interval scale. With this assumption, payoffs give a measure of the strength of the preferences, and therefore payoff differences are indeed meaningful. If nothing more than cardinality is assumed, comparisons of preferences between different players are still meaningless.

Most game theoretical models go beyond the assumption of cardinal preferences: they interpret payoffs as von Neumann-Morgenstern utilities (Colman, 1995, section 2.1; Hargreaves Heap and Varoufakis, 1995, section 1.2; Vega-Redondo, 2003, pg. 7). The benefit of making such a strong assumption is that it allows game theorists to use expected utility theory to evaluate probability distributions over possible outcomes of the game. (Note that payoffs relate to outcomes that are certain). It is important to remember that these models are –implicit or explicitly– assuming considerably more about players’ preferences than just cardinality: cardinality by itself is not enough to formally justify models where individuals maximise expected payoffs. Expected payoff maximisation requires preferences to be well defined (see above) and three extra assumptions: continuity, preference increasing with probability, and independence (Hargreaves Heap and Varoufakis, 1995, pg. 10). When all these assumptions hold, payoffs embody players’ attitudes to risk, and then it is true that an individual who acts on her preference ordering acts *as if* she is maximising her expected payoff (see chapter 2 in Colman (1995) for details).

Finally, the strongest assumption on preferences relates to social comparisons. There are (relatively few) models where payoffs interpreted as preferences are compared across players. This is a very strong assumption which finds its roots in the social philosophy of utilitarianism, and is not commonly observed in game theoretical models that interpret payoffs as preferences; however, it can certainly be found in the literature (see e.g. Bendor et al. (2004)). In stark contrast, it will be shown in the next section that most models in evolutionary game theory interpret payoffs as fitness, and they actually *require* comparing the payoffs obtained by different players (and often performing arithmetic operations with them).

2.1.2. Payoffs in evolutionary models

In evolutionary game theory models, the emphasis is not so much on the players, but on the strategies. In fact, it is most often understood that each player is pre-programmed to play a certain (pure or mixed) strategy, thus establishing equivalence between players and strategies. The interest then lies in studying the evolution of large *populations* of players who repeatedly interact to play a game. The aim is identifying which strategies (i.e. type of players) are most likely to thrive in this "ecosystem" and which will be wiped out by selection forces. In this context, payoffs are not interpreted as preferences, but as a value that measures the success of a strategy in relation to the others. Selection forces then act to favour strategies with higher payoffs. Thus, in models of biological (as opposed to cultural) evolution, payoffs are most often interpreted as Darwinian fitness. The crucial point here is that payoffs obtained by different players will be compared and used to determine the relative frequency of different types of players (i.e. strategies) in succeeding generations. This may not be a major assumption when modelling biological evolution, but it is one that cannot be ignored if evolution is interpreted in cultural terms.

2.2. Classical game theory

Classical game theory is devoted to the study of how instrumentally rational players should behave in order to obtain the maximum possible payoff in a formal game. Thus, as a deductive and normative branch of game theory, one could argue that classical game theory itself is incapable of being empirically tested and falsified (Colman, 1995, pg. 6). What we can clearly infer from the combination of empirical research and game theory is that, if empirical observations clash with game theoretical solutions, then (a) the observed real-world situation does not correspond to the abstracted game, or (b) at least one assumption made by game theory does not hold (or both (a) and (b)). Hence the importance of clearly identifying the assumptions made in classical game theory. The following sections analyse the two most relevant ones: complete availability of information and instrumental rationality.

2.2.1. Availability of information

A major assumption embedded in classical game theory (CGT) relates to information availability. This is a key issue, since information availability crucially affects what course of action may be regarded as rational. As an example, if players did not know anything about the game (not even its strategic nature) beyond the payoff they obtain after playing certain actions, many very simple learning models could be regarded as rational. CGT is mostly concerned with games of complete information. In these games, it is assumed that players not only know the rules of the game and their own payoffs, but also their counterparts' payoff functions. Furthermore, complete availability of information is assumed to be common knowledge. Common knowledge (CK) in game theory often comes with a certain order: zero-order CK of X is just the assumption that X prevails for every player (e.g. zero-order common knowledge of complete information (CKCI) means that every player has complete information); first-order CK is the assumption that every player knows that X prevails for every player (e.g. first-order CKCI means that every player knows that every player has complete information); in general, (n)th-order CK is the assumption that (n-1)th-order CK is known by every player. If no order is specified, it is assumed that the order is infinite (this produces an infinite recursion of shared assumptions). For different accounts of the meaning of common knowledge see Vanderschraaf and Sillari (2007).

CGT also considers games of incomplete information. As a matter of fact, if one is happy to accept certain (strong) conditions on what may count as a "rational belief", the distinction between complete and incomplete information is not essential, since games of incomplete information can be easily transformed into games of complete information (Harsanyi, 1967a, b, 1968). The basic idea behind this transformation consists in assuming that there are different "types of players", each of them with a different payoff function. Then, one must see each player's uncertainty about her counterparts' payoff functions as deriving from the player's uncertainty about which types of players her counterparts are. Finally, the transformation requires applying Harsanyi and Aumann's argument about the impossibility of players with mutual knowledge of rationality "agreeing to disagree" (Aumann, 1976). This last step ensures that rational players hold common beliefs about the probabilities that their counterparts will turn out to be of one type or another. Once this assumption is made, the analysis of the game with incomplete information is essentially the same as one of complete information.

2.2.2. Instrumental rationality

The concept of instrumental rationality in classical game theory finds its clearest roots in Hume's *Treatise on Human Nature*. In CGT rationality is understood as the capacity of identifying the actions that best satisfy the person's predefined objectives (Hargreaves Heap and Varoufakis, 1995, pg. 7), i.e. rationality plays no role in setting objectives. This basically means that instrumentally rational players have unlimited computational capacity devoted to maximise their individual payoff function, which is defined in advance. The assumption of rationality in CGT has been widely challenged. One of the alternatives that has received great attention is Simon's (1957) original concept of procedural rationality, later recast as bounded rationality (Simon, 1982) mainly for modelling purposes. Simon (1982) emphasises that people have limited knowledge of their situations, limited ability to process information, and limited time to make choices.

In any case, the main challenge within CGT comes from the fact that in most games there is no maximising strategy for any given player regardless of her counterparts' actions, i.e. rationality remains undefined in the absence of beliefs about what the other players will do. Naturally, this belief-dependency of rationality has led to different concepts of rationality based on different assumptions about what beliefs about other players' behaviour are allowed. The following sections explain the three most important approaches, namely:

1. Dominance reasoning.
2. Rationalisable strategies.
3. Consistently aligned beliefs: Nash equilibrium.

It is worth mentioning at this point that –most often– *the three* approaches outlined above make use of two extra assumptions, namely: common knowledge of complete information (CKCI; explained in the previous section), and common knowledge of rationality (CKR). Following the definition of common knowledge outlined in the previous section, first-order CKR is the assumption that every player knows that every player is rational (rationality is understood following one of the 3 interpretations mentioned above); (n)th-order CKR is the assumption that (n-1)th-order CKR is known by every player. If no order is specified, it is assumed that the order of CKR is infinite (see Aumann (1976) for a formal definition). CKCI and CKR are embedded in the definitions of approaches (2) and (3) mentioned above. Without assuming CKCI and CKR, most games are not solvable regardless of the approach taken. For the sake of clarity the following subsections will discuss the role of CKR assuming that CKCI comes with it.

Dominance reasoning

Rationality can be minimally identified with "not playing (strictly) dominated strategies"[3] (Vega-Redondo, 2003, pg. 32). This view of rationality does not require any assumption about the behaviour of other players: there is no belief that a player could hold about the other players' strategies such that it would be optimal to select a dominated strategy. In general, one has the option to reject only those strategies that are dominated by other pure strategies or, alternatively, choose to reject the (potentially larger) set of strategies that are dominated by some mixed strategy.

The elimination of dominated strategies by each player rarely leads to one single profile of strategies (the one-shot Prisoner's Dilemma is an exception for this), so CKR is usually brought into play. CKR allows the process of successive elimination of dominated strategies: with this interpretation of rationality, first-order CKR means that players assume that no player will select a dominated strategy. The elimination of certain strategies when assuming (n)th-order CKR may open the door to eliminate more strategies by assuming (n+1)th-order CKR.

---

[3] For a player A, strategy S_A is (strictly) dominated by strategy S^*_A if for each combination of the other players' strategies, A's payoff from playing S_A is (strictly) less than A's payoff from playing S^*_A (Gibbons, 1992, p. 5).

This iterative process goes on until no strategies can be further eliminated. When this process leads to one single strategy for every player (i.e. one single outcome) then the game is said to be dominance solvable.

Rationalisable strategies

A stronger interpretation of rationality dictates that rational players maximise their expected payoff on the basis of *some* expectations about what the others will do (clearly this interpretation prevents players from playing dominated strategies). Using this concept of rationality and assuming CKR leads to the definition of rationalisable strategies: rationalisable strategies are those that remain after making such assumptions (Bernheim, 1984; Pearce, 1984). The term rationalisable derives from the fact that every player can defend choosing such a strategy (i.e. rationalise it) on the basis of beliefs that are *consistent* with the assumption of CKR. However, given that each player may have many different rationalisable strategies (by holding different beliefs about her counterparts' beliefs), it could well be the case that once the game is played (i.e. once every player has selected a specific rationalisable strategy), some of these beliefs are proven wrong. To be clear, a set S of rationalisable strategies (one for each player) may derive from beliefs where one of the players is assuming that one of her counterparts will select a (rationalisable) strategy different from the one assigned to this counterpart in the set S itself. Informally, this would occur if one of the players presumes that one of her counterparts will "make a mistake" by expecting something that the player does not intend to do (even though this "mistaken belief" is perfectly consistent with CKR). In other words, the beliefs underlying rationalisable strategies must be consistent with the assumption of CKR for each individual player independently, but they may be inconsistent across players. Hargreaves Heap and Varoufakis (1995, pp. 51-52) give a 2-player example where both players select a rationalisable strategy on the basis of beliefs that are inconsistent across players. The following section explains that imposing consistency of beliefs across players leads to the (stronger) concept of Nash equilibrium.

Let us conclude this section by relating the concept of rationality explained here and that assumed when conducting dominance reasoning (see previous section).

As mentioned above, rationalisable strategies are necessarily undominated; a natural question is then: are iteratively undominated strategies always rationalisable? The answer to this question for 2-player games is yes (Pearce, 1984). In other words, for two player games these two concepts are equivalent. This is not true, however, for games involving more than two players. In such games, there can be iteratively undominated strategies that are not best response to any strategy profile. The subtle difference between these two concepts of rationality is brilliantly explained by Vega-Redondo (2003, pp. 66-68).

Consistently aligned beliefs: Nash equilibrium

The previous section showed that if players select rationalisable strategies, the outcome of the game may be such that the beliefs of some players are proven wrong by the choices actually made by other players. The concept of Nash equilibrium derives from imposing the additional constraint that beliefs must be consistently aligned across players. Thus, a Nash equilibrium is a set of rationalisable strategies (one for each player) whose implementation confirms the expectations of each player about the other players' choices (Hargreaves Heap and Varoufakis, 1995, pg. 53). A *corollary* of this definition is that Nash equilibria are formed by sets of strategies that are best replies to each other. This simple shortcut through the cumbersome web of players' beliefs over their counterparts' beliefs is probably one of the main factors that explain the success of the Nash equilibrium (NE) in the social sciences. Another reason is that NEs can be interpreted in a number of meaningful and useful ways (Holt and Roth, 2004). The concept of NE, however, is not free from problems. There are many games without any NE in pure strategies, and many others with more than one. In these cases, the assumption of consistently aligned beliefs is particularly problematic. How can players coordinate their beliefs in the absence of communication? The problem of multiple NE is particularly acute in repeated games, as illustrated by the extensive variety of "folk theorems" available in the literature. In broad terms, "folk theorems" demonstrate that repeated interactions typically allow for a wide range of equilibrium behaviour. Vega-Redondo (2003, chapter 8) reviews several "folk theorems", differing in their time horizon (finite or infinite), information conditions (complete or incomplete information, and perfect or imperfect observability), and equilibrium concept (Nash or subgame perfect).

Let us conclude this section by stating that the concept of NE is significantly stronger than that of rationalisable strategies. In particular, Bernheim (1984) showed by example that one can find rationalisable strategies that are not part of any NE (i.e. there is no NE that assigns a positive weight to them). In other words, there are outcomes where all players are selecting a rationalisable strategy, and which cannot be interpreted as the result of a mis-coordination among players that were hoping to arrive at a NE. This clearly indicates that the notion of rationalisability embodies something broader than equilibrium mis-coordination (Vega-Redondo, 2003, pg. 65).

Refinements of Nash equilibrium

The problem of multiple Nash equilibria outlined in the previous section has led to the proposal of countless refinements aimed at eliminating those NEs that are not considered plausible or desirable for not fulfilling some additional condition (see van Damme (1987) for a comprehensive study). Unfortunately, so many refinements have been developed by now that "in many games which have multiple Nash equilibria, each equilibrium could be justified by some refinement present in the literature" (Alexander, 2003). In this section we briefly present only one, namely "trembling hand perfection" in its strategic-form version (see Vega-Redondo, 2003, chapter 4), since the idea underlying this refinement will be used extensively in this thesis.

The "trembling hand perfect" refinement, which was proposed by Selten (1975), eliminates those Nash equilibria that are not robust to small mistakes. The refinement process assumes that players' hands may tremble, i.e. players may select an unintended action (i.e. deviate from the equilibrium) with small probability. An alternative view of the same phenomenon is that players may experiment with small probability. Some NEs may resist the possibility of these trembles and some may not: those NEs that do not survive arbitrarily small trembles are eliminated. Slightly more formally, the set of trembling hand perfect equilibria in a game is the limit of the sequence of Nash equilibria in perturbed versions of the game (i.e. versions of the game played with trembles) as the probability of trembles goes to zero. In 2-player strategic-form games, an equilibrium is perfect if and only if it is a Nash equilibrium that involves no weakly dominated strategies by either player (Van Damme, 1987, Theorem 3.2.2). The reasoning behind this refinement will prove to be very useful to reduce the set of possible outcomes of the game in the models developed in chapters 4 and 5 of this thesis.

2.3. Evolutionary game theory

Many biological and socio-economic systems are governed, at least to some extent, by evolutionary pressures. Such evolutionary systems may be composed of entities of very different nature, such as molecules, cells, genes, animals, organisations, ideas, behaviours… but they all share the three common features that characterise any evolutionary system: diversity, selection, and replication.

Diversity: entities in the system are not all the same; they show dissimilarities that affect their so-called individual fitness. Fitness is just a measurable indicator that determines how a population of entities evolves: entities with higher fitness will tend to spread relatively more than those with lower fitness. The precise mechanism that links current fitness with future population composition is the selection mechanism, which is explained in the next point. Note that in general this selection mechanism reduces the diversity of the system, since it favours some existing entities over others. There may be, however, mechanisms that tend to preserve the heterogeneous nature of the system: most evolutionary systems are subject to processes that create and maintain diversity. This diversity-generating mechanism acts in the opposite direction to the selection force, and it is the only mechanism that may preclude the system from locking-in. In biological systems, diversity generally stems from genetic mutations whereas in many socio-economic systems, it is innovations, asymmetries in the flow of information, or even simple mistakes, which are often responsible for the incessant appearance of different forms of behaviour. The process by which new entities appear in an evolutionary system is usually called mutation in biological contexts and experimentation or innovation in socio-economic contexts.

Selection: The mechanism of selection is a discriminating force that favours some specific entities rather than others. By selecting only certain entities from the population, this selection force diminishes the heterogeneity of the system. As mentioned above, the criterion by which some entities are selected among the population rather than others is usually called fitness. In evolutionary game theory strategies (which may be seen as behavioural phenotypes) are selected on the basis of the payoff they obtain, i.e. the relative frequency of strategies which obtained higher payoffs in the population will increase at the expense of those which obtained relatively lower payoffs.

<u>Replication</u> / <u>Inheritance</u> / <u>Preservation</u>: The properties of the entities in the system (or the entities themselves) are preserved, replicated or inherited from one generation to the next at least to some extent. Replication mechanisms can be carried out through a range of processes, from genetic transmission in biological systems to social learning processes such as imitation in some socio-economic contexts.

The main assumption underlying evolutionary thinking is that the entities which are *more successful*[4] at a particular time will have the best chance of being present in the future. In biological and economic contexts, this assumption often derives from competition among entities for scarce resources or market shares. In social contexts, evolution is often understood as *cultural* evolution, where this refers to changes in behaviour, beliefs, or social norms over time (Alexander, 2003), and may be justified by "the tendency of human behaviour to adjust in response to persistent differentials in material incentives" (Sethi and Somanathan, 1996, pg. 783).

Evolutionary game theory (EGT) is devoted to the study of the evolution of strategies. In biological systems, players are most often assumed to be pre-programmed to play one given strategy, so studying the evolution of a population of strategies becomes formally equivalent to studying the evolution of a population of players. By contrast, in socio-economic models, players are usually assumed to live forever, and switch their strategy following evolutionary pressures. The role of players relative to the role of strategies is irrelevant for the formal analysis of the system, where –in both cases– it is strategies that are

[4] Note that this is a measure of *relative* performance.

actually subjected to evolutionary pressures. Thus, without loss of generality and for the sake of clarity, we take here the biological stand and assume that players select always the same strategy.

Thus, EGT is devoted to the study of large *populations* of players who repeatedly interact to play a game. Strategies are subjected to selection pressures in the sense that the relative frequency of strategies which obtain higher payoffs in the population will increase at the expense of those which obtain relatively lower payoffs. The aim is to identify which strategies (i.e. type of players) are most likely to thrive in this "evolving ecosystem of strategies" and which will be wiped out by selective forces. As mentioned before, payoffs in evolutionary contexts are not interpreted as preferences; instead they provide the value that is used to measure the relative success of one strategy in relation to the others.

2.3.1. Evolutionary stability: evolutionary stable strategies

The study of dynamic systems often begins with the identification of their stable states. This is often called static analysis, as it does not consider the dynamics of the system explicitly, but only its rest points. The most important concept in the static analysis of EGT is the concept of Evolutionary Stable Strategy (ESS), proposed by Maynard Smith and Price (1973). Very informally, a population playing an ESS is *uninvadable* by any other strategy (Weibull, 2002). To be more precise, consider a very large population of players who are repeatedly drawn at random to play a 2-player symmetric game. Initially all players are selecting the same (incumbent) strategy. That strategy is an ESS if there exists a positive invasion barrier such that for any given mutation that may occur and assuming that the population share of individuals playing the mutant strategy falls below this barrier, the incumbent strategy earns a higher payoff than the mutant strategy (Weibull, 1995, pg. 33). The original concept of ESS has proven to be tremendously useful, but it is important to be aware of the assumptions underpinning its theoretical framework: the ESS is derived for a system composed of a *single infinite* population of individuals who are repeatedly *randomly* drawn to play a *2-player symmetric* game; furthermore, it only considers *monomorphic* populations (all individuals are playing the same strategy) which can be invaded by only *one type of mutant strategy at a time*.

In particular, the assumption of one single *infinite* population has a number of important implications. For a start, this assumption is in effect a mean-field approximation used to equate the average payoff actually obtained by a population with the expected value of a probability distribution of payoffs (which would be obtained by explicitly modelling players' interactions). It is also the assumption that justifies treating as equivalent a mixed strategy and a population profile where pure strategies are played in the population with the frequency induced by the corresponding probability in the mixed strategy (see Vega-Redondo, 2003, pp. 356-7). Finally, it effectively eliminates the impact of arbitrarily small invasions on the incumbent population. This last point is best explained with a simple example. Consider a 2-player population where player i can impose a punishment of magnitude P on player j at a cost of $C < P$. Clearly, punishing j would give a relative advantage to i over j, so this behaviour would be evolutionary favoured. Now consider a large population of potentially punishable players j, and think of the effect of the same single punishment conducted by one mutant i on one of the players in the incumbent population. Player i will incur the cost C, but the average payoff of the incumbent population will only decrease in P divided by the size of the population n. If n is infinite, then the effect of i's punishment on the incumbent population is just zero. This reasoning is important because it is behind the (correct) argument that the concept of ESS is a refinement of (symmetric 2-player games) Nash equilibrium. Without the assumption of infinite populations, the argument does not necessarily hold (see Galán and Izquierdo (2005) for an illustration). To avoid this issue without having to impose infinite populations, an alternative is to make sure that the smallest invasion barrier expressed as a population share exceeds $1/n$ (Weibull, 1995, pp. 33-34).

2.3.2. Evolutionary dynamics: the replicator dynamics

Naturally, to study the dynamics of an evolutionary system explicitly (i.e. beyond the analysis of its rest points), it becomes necessary to specify the particular process that governs such dynamics. The most extensively studied dynamic process in EGT is the replicator dynamics, proposed by Taylor and Jonker (1978). In the replicator dynamics (RD), payoffs are interpreted as the number of viable offspring that inherit the same behavioural phenotype (i.e. strategy) as their (single) parent. The theoretical model underpinning the basic RD also assumes a *single infinite* population of individuals who are repeatedly *randomly* drawn to play a *2-player symmetric* game. Furthermore, individuals can only play *one out of a finite set of pure strategies*, and *mutations (and random drift) are not allowed*[5]. This set of assumptions is enough to fully determine a *deterministic* dynamic process in which the rate of change in the frequency of any given strategy is equal to the relative difference between its average payoff and the average payoff obtained across all strategies in the population. Most often, time is treated as a continuous variable, and this allows the formalisation of the dynamic process as a system of ordinary differential equations.

With these assumptions in place, game theorists have been able to derive a chain of useful mathematical results that link the concept of ESS, the dynamics of the basic RD and the concept of NE. The logical chain is as follows: the population profile induced by an ESS is asymptotically stable in terms of the RD (Hofbauer et al., 1979); the mixed strategy corresponding to an asymptotically stable equilibrium of the RD is in (symmetric) perfect Nash equilibrium with itself (see proof in e.g. Weibull, 1995, section 3.4); and finally, a mixed strategy played at a symmetric Nash equilibrium (in a 2-player symmetric game with a finite set of pure strategies) induces a stationary population state of the RD (see proof in e.g. Vega-Redondo, 2003, pg. 367).

2.3.3. Further developments

While undoubtedly extremely useful, the assumptions embedded in the original concept of ESS and in the basic RD limit the applicability of the analytical results obtained with them, particularly in social (rather than biological) contexts (see e.g. Probst, 1999; Gotts et al., 2003b; Vega-Redondo, 2003, pg. 372). These concerns led to the development of more general frameworks which would encompass as particular cases not only the RD but also a wider range of dynamic processes, and could be applied not only to 2-player symmetric games, but also to general games. Of special interest are the *multi-population* models with *regular* and *payoff monotonic* dynamics.

[5] Mutations can be superimposed as a separate component of the dynamic process (see e.g. Imhof et al. 2005).

- *Multi-population* models study *n*-player games, where each player is randomly drawn from a distinct (infinite) population. This setting allows modelling any finite game in normal form where players in different positions are subjected to independent evolutionary pressures.
- *Regularity* ensures that the proportional rates of change of strategies are well defined and are continuously differentiable.
- Finally, *payoff monotonicity* is a mild condition which imposes that for any given pair of strategies in any particular population, their proportional rates of change are ordered in the same way as their respective average payoffs (Vega-Redondo, 2003, pg. 377).

It turns out that most of the analytical results linking the concepts of ESS, NE, and the dynamics of the basic RD can be carried over to this general framework (once the appropriate generalisations for these concepts have been defined; see e.g. Weibull (1995, chapter 5) and Vega-Redondo (2003, chapter 10)). This type of general framework[6] represents a remarkable step forward in generality and, consequently, the applicability of the analytical results obtained with them is greatly increased. However, these general models still make two assumptions that somewhat limit their applicability to social contexts (Probst, 1999): regularity and infinite populations.

As pointed out by Probst (1999), the assumption of regularity rules out many adaptation mechanisms that are considered of much interest in modelling social systems (e.g. best-response dynamics). This assumption, which is rarely made in learning game theory (LGT), is one of the main differences between EGT models and LGT models, in terms of the mathematical properties of the induced formal systems.

The assumption of infinite populations effectively averages out the stochasticity of the system, so the obtained deterministic dynamics can be formalised as a system of differential equations. This assumption has greater implications than one may initially suspect. As Traulsen et al. (2006) point out, “the finiteness of

[6] There are various similar versions (see Weibull, 1995).

populations may indeed lead to fundamental changes in the conventional picture emerging from deterministic replicator dynamics in infinite populations". To be more precise, any model with finite populations can be formalised as a Markov process, and the system of differential equations is the approximation of the Markov process in the limit as the population tends to infinity. Also, one is often interested in studying the behaviour of the system in the long run, which involves calculating the limit of the dynamics as time goes to infinity. The problem in doing this is that results can be dramatically different depending on the order in which one takes these two limits. This will be clearly illustrated in a somewhat different context in chapter 4. Fortunately, our theoretical knowledge of these issues has progressed immensely in the last few years. In particular, the seminal paper by Benaim and Weibull (2003) is a breakthrough in the field of stochastic approximation in EGT. In any case, it is clear that "care is therefore needed in the application of these approximations" (Beggs, 2002).

In summary, 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); thus, it is not surprising that the analysis of finite evolutionary systems is nowadays a field of great scientific dynamism (see e.g. Nowak et al., 2004; Taylor et al., 2004; Imhof et al., 2005; Santos et al., 2006; Traulsen et al., 2006).

2.3.4. Stochastic finite systems

Once it has been acknowledged that stochasticity plays an important role in the analysis of finite evolutionary systems, the main challenge for current EGT seems to lie in understanding the impact of the various other assumptions made in traditional EGT on these finite stochastic systems.

A feature of the system that has been long known to play a crucial role is the mechanism by which individuals pair to play the game. The pairing algorithm does not necessarily have to be imposed by a fixed population structure, but may be actively conducted by the players themselves (Eshel and Cavalli-Sforza, 1982). Naturally, the impact of the standard assumption (random encounters) is investigated by considering other pairing mechanisms. One of the first studies to show the relevance of different population structures in finite systems was conducted by Nowak and May (1992; 1993), who used a spatial model (where local interactions occurred between individuals occupying neighbouring nodes on a square lattice) to show that stable population states for the prisoner's dilemma depend upon the specific form of the payoff matrix. For a review of several studies in the context of social dilemmas that consider populations where some pairs of agents are more likely to interact than others see Gotts et al. (2003b). Of particular interest is the field of study on tags (Holland, 1993). Tags are arbitrary social marks that, in principle, are not linked to any particular form of behaviour, but they do 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). In chapter 6 we investigate various pairing mechanisms and, in particular, we analyse one which is formally equivalent to the use of tags. For a recent illustration of the latest developments in the field of structured populations in finite systems, see Santos et al. (2006), who study social dilemma games played in (fixed) networks with various degrees of heterogeneity in the degree distributions. The most recent literature in this field is mainly focused on studying the emergence of cooperation in *spatially* structured populations (see e.g. Hauert and Doebeli, 2004; Doebeli and Hauert, 2005; Németh and Takács, 2007). For a recent illustration of the fact that allowing players to selectively choose their partners can have dramatic effects on the emergence of cooperation in finite systems see e.g. Joyce et al. (2006).

In chapter 6 we also investigate various selection mechanisms (i.e. algorithms that determine how the population composition varies as a function of the payoffs obtained by each individual). This is another area of research where a substantial amount of work has been conducted in the last few years. In a recent paper, Traulsen et al. (2006) develop a framework within which one can explore various intensities of selection, i.e. different ways in which payoffs relate to fitness (where fitness is the function that determines the potential to reproduce). This selection framework makes use of the Fermi distribution function from statistical mechanics to control the balance between selection and random drift in finite populations. Using this function, Traulsen et al. (2006) explore different intensities of selection –ranging from neutral, random drift, up to the extreme limit of cultural imitation dynamics– in the three 2-player 2-strategy social dilemma games (these are explained in detail in section 3.1). Traulsen et al. (2006) are able to calculate the fixation probabilities of different strategies, and they also use stochastic approximation theory to relate their results on finite populations to those obtained with infinite populations.

An assumption that –to our knowledge– has not been investigated in depth in evolutionary stochastic finite systems is the one relating to the properties of the set of strategies that players are allowed to select. In chapter 6 of this thesis we show that this assumption may have wider implications than one may initially suspect.

There are many other ways in which several authors have addressed some of the limitations of EGT outlined above. Unfortunately (but probably inevitably), the study of the implications of various assumptions made in mainstream EGT is being undertaken in a somewhat disorganised fashion. This inconvenience is probably a consequence of the dynamism of this field, and it will hopefully be corrected in time through the creation of general frameworks that facilitate rigorous and transparent comparisons between different models and the results obtained with them. Chapter 6 of this thesis is meant to be a step in this direction, by providing a single coherent framework within which results obtained from different stochastic finite models can be contrasted and compared.

2.4. Learning game theory

Like evolutionary game theory, learning game theory (LGT) abandons the demanding assumptions of classical game theory on players' rationality and beliefs. However, unlike evolutionary game theory –where players are often assumed to be pre-programmed to play a fixed strategy–, LGT assumes that players are able to learn over time about the game and the behaviour of others (through e.g. reinforcement, imitation or belief updating), and this learning process is *explicitly* modelled (Vega-Redondo, 2003, pg. 398). This distinction means that the level at which dynamic processes are defined in EGT and LGT is fundamentally different (Fudenberg and Levine, 1998). Models in EGT are aggregate in the sense that they describe the aggregate behaviour of a population of players through various generations; the population is subject to evolutionary pressures (and therefore the *population* adapts), but the individual components of the population have a predefined fixed behaviour. On the other hand, models in LGT comprise players who *individually* adapt through learning, and it is this learning process that is formally described. Models in LGT explicitly represent the learning processes that each individual player carries out, and the dynamics that are generated at the aggregate level (which are most often stochastic and non-regular) emerge out of the strategic interactions among the players.

Another fundamental difference between LGT and EGT relates to the relationship between the number of players in the game and the number of players in the population. Models in LGT tend to focus on one very small population of n players (most often $n = 2$), who play an n-player game (all individuals in the population play the game at once). This is in stark contrast with EGT models, where individuals within a large (usually infinite) population are drawn to play a 2-player game. As explained in section 2.3.1, this distinction can have very important implications.

Despite these differences, theoretical work linking results from EGT and LGT seems to indicate that we may be close to a point where the integration of the two approaches is within reach (Weibull, 1998). This is a question that is further discussed in section 7.4.

Interestingly, there seem to be two fundamentally different motivations to study learning models in the LGT literature. One is mainly concerned with identifying learning algorithms that will lead to NE or, ideally, to refinements of NE. The following quote by Vega-Redondo nicely summarises this motivation: "In particular, our concern is to identify different classes of games in which the corresponding learning processes bring about long-run convergence to some Nash equilibrium. As we shall see, many of the proposed models *fare reasonably well* for certain games but induce quite *unsatisfactory performance* for some others." [our emphasis] (Vega-Redondo, 2003, pg. 398).

This thesis follows another motivation: we are mainly concerned with identifying the strategic implications of decision-making algorithms that have received support from cognitive science research. Work following this second rationale has sometimes been labelled “cognitive game theory” (CogGT) in the literature (e.g. Flache and Macy, 2002). Nowadays, an increasing number of researches use CogGT to investigate animal –often human– behaviour in strategic contexts using models that seem more plausible than those deriving from classical game theory. Thus, CogGT models are often used to identify learning mechanisms that will lead to patterns of behaviour observed in real-world interactions (and these patterns often do not correspond to NE). The following summarises some features that characterise the way players are modelled in CogGT (Flache and Macy, 2002; Macy and Flache, 2002), in contrast with classical game theory:

- Players base their decisions on experience of past events as opposed to logical deductions about the future. This inductive approach requires fewer assumptions about other players and may be more adequate to model animal (including human) behaviour. Since inferences about other players' strategies –or about future payoffs– is made in the light of the history of the game, they can only lead to probable −rather than necessarily true− conclusions (even if the evidence used is accurate).
- Players have feedback on their actions; otherwise learning cannot occur. Learning takes many forms, depending on the available feedback, the available knowledge, and the way these are used to modify behaviour.
- The fact that players learn from experience means that they often cannot undertake an optimal behaviour (since inferences about other players' behaviour cannot be guaranteed to be true). An optimal approach requires knowledge that sometimes has to be inferred from experience. In the process of acquiring the necessary knowledge, suboptimal behaviour can occur as a result of exploring different actions or having drawn imperfect conclusions from experience. When modelling players who learn from experience, it often seems reasonable to assume that they satisfice rather than optimise. The concept of 'satisficing' was introduced by Simon (1957) to indicate that agents often seek for a solution to a problem until they have found one which is 'good enough', rather than persisting in the hope of finding an optimal solution (which could be nonexistent,

incalculable, or unidentifiable). The 'good enough' solution is usually defined by setting a certain aspiration threshold.

The distinction between the two different motivations outlined above becomes clear when one considers social dilemmas. In most single-stage social dilemma games, the cooperative strategy is dominated (i.e. it cannot lead to NE); however empirical studies have generally found that, while it is not easy to establish cooperation, levels of cooperation tend to be higher than would be expected if the assumptions made in CGT held true. Thus, when studying social dilemmas, researchers in LGT following the "NE motivation" would presumably consider models leading to cooperative solutions generally unsatisfactory. In stark contrast, in the context of social dilemmas, CogGT has been mainly concerned with identifying a set of model-independent learning principles that are necessary and sufficient to generate cooperative solutions (Flache and Macy, 2002). Interestingly –if unsurprisingly–, it seems that researchers more inclined towards CogGT tend to use computer simulation (instead of mathematical analysis) relatively more than those researchers following the "NE motivation".

2.4.1. Different learning algorithms

As mentioned above, the process of learning can take many different forms, depending on the available knowledge, the available feedback, and the way these are used to modify behaviour. The assumptions made in these regards give rise to different models of learning. In most models of LGT, players use the history of the game to decide what action to take. In the simplest models (e.g. reinforcement learning) this link between acquired information and action is direct (e.g. in a stimulus-response fashion); in more sophisticated models players use the history of the game to form expectations about the other players' behaviour, and they then react optimally to these inferred expectations. Following Vega-Redondo (2003, chapter 11) we briefly present here some of the most studied learning models in ascending order of sophistication, according to the amount of information that players use and their computational capabilities.

Reinforcement learning

Reinforcement learning models will be discussed at length in section 4.1. Let us say for now that they are arguably the simplest family of learning algorithms investigated in LGT. Reinforcement learning is also one of the most widespread adaptation mechanisms in nature. Reinforcement learners use their experience to choose or avoid certain actions based on their immediate consequences. Actions that led to satisfactory outcomes (i.e. outcomes that met or exceeded aspirations) in the past tend to be repeated in the future, whereas choices that led to unsatisfactory experiences are avoided. In general, reinforcement learners do not use more information than the immediately received payoff, which is used to adjust the probability of the conducted action accordingly. The specific details of how this general principle is implemented in different models can lead to substantially different dynamics, as explained in section 4.1.

Static perceptions; better and best (myopic) response

In this more sophisticated family of learning models, each player is assumed to know not only the payoff she receives in each possible outcome of the game, but also the actions that every player selected at a certain time t. When making her decision for time $(t + 1)$ every player assumes that every other player will keep her strategy unchanged (i.e. static perception of the environment); then, each individual player, working under such assumption and knowing the payoff structure of the game in what pertains to her own payoff, can identify the set of strategies that will lead to an improvement in her current payoff (if possible). In better-response models, one of these payoff-improving strategies is selected at random; in best-response models, only those strategies that give the highest payoff given the prevailing assumptions are considered for selection. In these models players assume that their environment is static and deterministic, and respond to it in a myopic fashion, i.e. ignoring the implications of current choices on future choices and payoffs. Vega-Redondo (Vega-Redondo, 2003, pp. 415-420) summarises several results for this type of learning algorithm.

Fictitious play

Fictitious play models were first proposed by Brown (1951). Fudenberg and Levine (1998) provide a recent and comprehensive account of this family of models. As in best (myopic) response models, players in fictitious play (FP) models are assumed to have a certain model of the situation and decide optimally on the basis of it. The higher level of sophistication introduced in FP models concerns the (still stationary) model of the environment that players hold. FP players assume that the mixed strategy played by every other player at a certain time is equal to the frequency with which they have selected each of their available actions up until that moment. Thus, instead of considering the actions taken by every other player only in the immediately preceding time-step (as in the models explained in the previous section), they implicitly take into account the full history of the game. After forming her beliefs about every other player's strategy, a FP player (myopically) responds optimally to them.

In 2-player games, the belief sequence induced by FP is known to converge to a profile that defines a Nash equilibrium. This result, however, may be somewhat misleading, as it does not imply that players will play the strategy profile induced by such a sequence of beliefs in an *uncorrelated* fashion (Fudenberg and Kreps, 1993), randomising their decisions *independently* from each other as the definition of a Nash equilibrium requires. As an example, imagine that the belief sequence in a 2x2 game converges to a strategy profile (i.e. an assignment of frequencies to all the strategies available to a player) where fictitious player 1 selects action A_1 with frequency 1/3 (and action B_1 with frequency 2/3) and fictitious player 2 selects action A_2 with frequency 1/3 (and action B_2 with frequency 2/3). The mathematical result mentioned above guarantees that there is a Nash equilibrium with the strategy profile FP converges on. This would seem to suggest that the pattern of play in fictitious play will be the same as the pattern of play induced by a Nash equilibrium, but this is not necessarily the case. Thus, in our example, the Nash equilibrium in mixed strategies would imply that any outcome has a positive probability of occurring (e.g. outcome [A_1, B_2] would occur with probability 2/9). On the contrary, by setting players' initial beliefs appropriately (which are determined by numerical weights, one for each of the other player's pure strategies) one can construct examples where player 1 selects action A_1 if and only if player 2 selects action A_2 (Fudenberg and Kreps, 1993). This, in particular, would imply that outcome [A_1, B_2] would never occur. Thus, the payoff obtained by each player in this latter case can be completely different from the expected payoff obtained if players selected action A_i or B_i in an uncorrelated fashion. Therefore, each component of the belief sequence in FP must be understood as a *marginal* distribution for each player separately; the *joint* distribution may be very different from that resulting from Nash equilibrium play.

Smooth fictitious play

The perverse correlation effects outlined in the previous section motivated a stochastic version of the original fictitious play named *smooth* fictitious play (SFP, Fudenberg and Kreps, 1993). As in the original fictitious play, players in SFP assume that the mixed strategy played by every other player at a certain time is equal to the frequency with which they have selected each of their available actions up until that moment. In SFP models, however, players are no longer assumed to respond to their beliefs about the other players' strategies in the knife-edge fashion implied by the best-response correspondence; instead they respond in a continuous, differentiable way. The step-like determinism of the best-response correspondence used in FP is replaced by a smooth-looking function that returns a probabilistic response to the other players' inferred strategies in SFP. In SFP (as in FP), the rate of adjustment of behaviour slows down at a rate that permits the use of stochastic approximation theory, and this has facilitated the derivation of several theoretical results. In particular, SFP players' strategies are guaranteed to converge to Nash equilibrium in 2x2 games (Fudenberg and Levine, 1998).

Rational learning

The most sophisticated model of learning in LGT was proposed by Kalai and Lehrer (1993a; 1993b). Players in this model are assumed to be fully aware of the strategic context they are embedded in. They are also assumed to have a set of subjective beliefs over the behavioural strategies of the other players. Informally, as put by Vega-Redondo (2003, pg. 434), the only assumption made about such beliefs is that players cannot be "utterly surprised" by the course of the play, i.e. players must assign a strictly positive probability to any belief that is coherent with the history of the game. Finally, players are assumed to respond optimally to their beliefs with the objective of maximising the flow of future payoffs discounted at a certain rate. A detailed explanation of the (very powerful) results obtained with this model seems to fall out of the scope of this brief account of learning models. We refer the interested reader to Vega-Redondo (2003, pp. 433-441), who provides a brilliant account of this part of the literature, and concludes that "some of the assumptions underlying the rational-learning literature […] should be interpreted with great care".

Let us conclude this section by pointing out a common weakness of most current models in LGT (including those developed in this thesis): they almost invariably assume that every player in the game follows the same decision-making algorithm. This seems to be the natural first step in exploring the implications of a decision-making algorithm; however, it is clear that in many of these models the observed dynamics are very dependent on the fact that the game is played among "cognitive clones". Confronting the investigated learning algorithm with other decision-making algorithms seems to be a promising second step in LGT studies.

2.4.2. Assumptions in the learning models developed in this thesis

Reinforcement learning

Chapter 4 is an in-depth analysis of the transient and asymptotic dynamics of the Bush-Mosteller reinforcement learning algorithm for 2-player 2-strategy games. The following summarises the main assumptions made in this model in terms of the nature of the payoffs, the information players require and the computational capabilities that they have.

- Payoffs: In this model, payoffs and aspiration thresholds are not interpreted as von Neumann-Morgenstern utilities (for which the distinction between positive and negative values is irrelevant), but as a set of variables measured on an interval scale that is used to calculate stimuli (this is explained in detail in section 4.2).
- Information: Each player is assumed to know the range of possible actions available to her, and the maximum absolute difference between any payoff she might receive and her aspiration threshold. Players do not use any information regarding the other players.

- Memory and computational capabilities: Players are assumed to know their own (potentially) mixed strategy at any given time. They need to be able to conduct arithmetic operations.

Case-based reasoning

Chapter 5 is an exploration of cased-based reasoning as a decision-making algorithm in strategic contexts. The following summarises the main assumptions made in this model in terms of the nature of the payoffs, the information players require and the computational capabilities that they have.

- Payoffs: In this model, payoffs can be interpreted as preferences measured on an ordinal scale.
- Information: Each player is assumed to know the range of possible actions available to her, and her own aspiration threshold. Players do not use any information regarding the other players.
- Memory and computational capabilities: For each possible state of the world they may perceive, players are assumed to store in memory the last payoff they received for each of the possible actions available to them. They need to be able to rank their preferences.

2.5. Non-strictly-deductive branches of game theory

This thesis aims to be an advancement in the field of *deductive* game theory. It is important to note that there are other branches of game theory which are not purely deductive; these non-strictly-deductive branches tend to use game theory as a framework to fit observed empirical data and understand the underlying mechanisms that may be producing the observed results. There is clearly a lot to gain from the interaction of deductive and non-deductive game theory. Traditionally, deductive game theory has developed almost entirely from introspection and theoretical concerns. Unless this is corrected in the coming years, deductive game theory may suffer the danger of becoming practically irrelevant or, in less dramatic terms, not fulfilling all its potential as a useful tool to analyse real-world social interactions. On the other hand, if the objective is to find a model that fits empirical data to a satisfactory extent, it is crucial to understand the behaviour of different models in detail; if one is not content with fitting only, but some level of understanding is also pursued, then it becomes fundamental to know the implications of various cognitive mechanisms (i.e. assumptions) for the development of the game. Thus, it seems very clear that empirical studies have also a lot to gain from theoretical analyses. These issues will be discussed in chapter 7, but let us say for now that the work reported in this thesis has tried to be relevant by (a) studying the strategic implications of decision-making algorithms that have received empirical support from the cognitive sciences and (b) building frameworks to clearly identify the factors (i.e. types of assumption) that may have the greatest impact in the outcome of a social interaction (i.e. a game).

There are a number of learning models that have been proposed to explain experimental data (see chapter 6 in Camerer, 2003), and many of them have been investigated in purely theoretical terms. The transition from theoretical learning models to non-strictly deductive branches of game theory is very smooth. Here we mention two: psychological game theory and behavioural game theory. Psychological game theory is a term coined by Colman (2003).

> "Psychological game theory […] overlaps behavioral game theory but focuses specifically on non-standard reasoning processes rather than other revisions of orthodox game theory such as payoff transformations. Psychological game theory seeks to modify the orthodox theory by introducing formal principles of reasoning that may help to explain empirical observations and widely shared intuitions that are left unexplained by the othodox theory" (Colman, 2003).

Overlapping psychological game theory, behavioural game theory is completely driven by empirical (especially experimental) data, and models are assessed according to how well they are fitted to data. While models in cognitive game theory are designed to help us *reflect on* a certain process, behavioural game theory builds on models which are usually designed to *represent* the actual process.

> "Behavioral game theory is about what players *actually* do. It expands analytical theory by adding emotion, mistakes, limited foresight, doubts

> about how smart others are, and learning to analytical game theory. Behavioral game theory is one branch of behavioral economics, an approach to economics which uses psychological regularity to suggest ways to weaken rationality assumptions and extend theory." (Camerer, 2003, p.3)

Let us finish the chapter by stating that learning models have been reported to outperform classical game-theoretic predictions on experimental data (see Macy, 1995; Roth and Erev, 1995; Erev and Roth, 1998; Camerer, 2003, chapter 6). The empirical support of learning models in game theory will be expanded for reinforcement learning and case-based reasoning in the following chapters.