Modern Game Theory:
Deduction vs. Induction

Amy Greenwald

January 20, 1997

Abstract

The aim of this paper is twofold: firstly, to present a survey of the theory of games, and secondly, to contrast deductive and inductive reasoning in game theory. This report begins with an overview of the classical theory of strategic form games of complete information. This theory is based on the traditional economic assumption of rationality, common knowledge of which yields Nash equilibrium as a deductive solution to games in this class. In the second half of this paper, modern game-theoretic ideas are introduced. In particular, learning and repeated games are analyzed using an inductive model, in the absence of common knowledge. In general, inductive reasoning does not gives rise to the Nash equilbrium when learning is deterministic, unless initial beliefs are somehow fortuitously chosen. However, computer simulations show that in the presence of a small random component, repeated play does indeed converge to Nash equilibrium. This research is of interest to computer scientists because modern game theory is a natural framework in which to formally study multi-agent systems and distributed computing.

1 Introduction

If all the players in a game are rational[1], and all the players know that all the players are rational, and all the players know that all the players know that all the players are rational, and so on, then the players deduce the optimal moves of the other players and their own optimal responses; this is precisely a Nash equilibrium. In this way, common knowledge of rationality gives rise to deductive solutions of strategic form games of complete information. Without the assumption of common knowledge of rationality, games are not conducive to deductive solutions. Modern game theory is concerned with inductive reasoning in repeated games, which is described clearly in Arthur [1]:

> Each agent keeps track of the performance of a private collection of belief-models. When it comes time to make choices, he acts upon his currently most credible (or possibly most profitable) one. The others he keeps at the back of his mind, so to speak. Alternatively, he may act upon a combination of several...Once actions are taken, agents update the track record of all their hypotheses.

While this model does have intuitive appeal, it does not in fact compute a Nash equilibrium in repeated games of complete information, unless the set of initial beliefs of all players is already a Nash equilibrium. However, in repeated games of incomplete information, inductive reasoning leads to convergence to Nash equilibrium from a more general set of initial beliefs. Moreover, computer simulations demonstrate that if agents' beliefs contain a small random component, then this model gives rise to an equilibrium.

This report is organized as follows. The first half is an overview of the classical theory of strategic form games of complete information, including examples as well as the formal definition of Nash equilibrium. In spite of some well-known paradoxes, common knowledge of rationality is assumed. In the second half of this paper, this assumption is dropped, and an inductive model of reasoning in repeated game playing is considered. In particular, theoretical results on convergence and non-convergence to Nash equilibrium are discussed, concluding with an interesting example due to Arthur known as the Santa Fe Bar Problem.

---

[1] Throughout this exposition, a rational player is one who optimizes by choosing actions which maximize payoffs.

2 Strategic Form Games

As a subfield of economics, game theory provides a framework in which to model situations where intelligent agents, or *players*, with different and often conflicting interests, make decisions from a set of possible *strategies*, while aiming to maximize individual performance, or *payoffs*. However, in contrast to the extreme cases of either perfect competition, where it is assumed that the actions of individual agents have negligible effects on the market, or monopoly, where a single entity usurps all market power, *the payoffs of a game are jointly determined by the strategies of all players.*

Presently, we develop the classical theory of non-cooperative, non-zero-sum games of complete information, based on an assumption of common knowledge of rationality. In non-cooperative games, there is a possibility of negotiation prior to play, but there are no coalitions or binding agreements. In non-zero-sum games, there is in general a choice between contentious and cooperative behavior – this is in contrast to zero-sum games, where the interests of players are diametrically opposed. In games of complete information, all relevant information is known *a priori* to all players.

2.1 Examples

The most well-known game-theoretic scenario is the paradoxical situation known as the *Prisoner's Dilemma*, which was popularized by Axelrod [3] in his popular science book. The following is one variant of the story from which this game derives it name.[2]

> A crime has been committed for which two suspects are being held incommunicado. The district attorney questions the two suspects. If both suspects confess, they will be punished, but not terribly severely, as the D.A. will reward them for their honesty. If only one suspect confesses, the confessor will be severely punished for carrying out the crime singlehandedly, while the other suspect will be let off scot free. Lastly, if neither suspect confesses, the D.A. will have no choice but to convict both suspects.

---

[2] For the usual interpretation, see Rapoport [7], the two-time winner of the Prisoner's Dilemma computer tournament organized by Axelrod.

1 \ 2	C	D
C	4,4	0,5
D	5,0	1,1

Figure 1: The Prisoner's Dilemma

The Prisoner's Dilemma is a two player, strategic (or normal) form game. Such games are easily described by payoff matrices, where the strategies of player 1 and player 2 serve as column and row labels, respectively, and the corresponding payoffs are listed as pairs in matrix cells such that the first (second) number is the payoff to player 1 (2). The payoff matrix which describes the Prisoner's Dilemma is depicted in Figure 1, with C denoting "confess", or "cooperate", and D denoting "don't confess", or "defect".

This game is known as the Prisoner's Dilemma because the outcome of the game is (D, D), which yields suboptimal payoffs of $(1, 1)$. This equilibrium solution arises as a direct consequence of rationality. The reasoning is as follows. If player 1 plays C, then player 2 is better off playing D, since D yields a payoff of 5, whereas C yields only 4; but if player 1 plays D, then player 2 is again better off playing D, since D yields a payoff of 1, whereas C yields 0. Hence, regardless of the strategy of player 1, a rational player 2 plays D. By a symmetric argument, a rational player 1 also plays D. Thus, the outcome of the game is (D, D).

A second well-known example of a two-player game is a game called *Matching Pennies*. In this game, each of the two players flips a coin, and the payoffs are determined as follows (see Figure 2). Let player 1 be the *matcher*, and let player 2 be the *mismatcher*. If the coins come up matching (*i.e.*, both heads or both tails), then player 2 pays player 1 the sum of $1. Otherwise, if the coins do not match (*i.e.*, one head and one tail), then player 1 pays player 2 the sum of $1. This is an example of a zero-sum game where the interests of the players are diametrically opposed; this class of games is so-called because the payoffs in the matrix do indeed sum to zero.

1 \ 2	*H*	*T*
H	1,-1	-1,1
T	-1,1	1,-1

Figure 2: Matching Pennies

1 \ 2	*B*	*F*
B	2,1	0,0
F	0,0	1,2

Figure 3: Battle of the Sexes

Another popular two-player game is called the *Battle of the Sexes*. A man and a woman would like to spend an evening out together; however, the man prefers to go to a football game (strategy F), while the woman prefers to go to the ballet (strategy B). Both the man and the woman prefer to be together, even at the event that is not to their liking, rather than go out alone. The payoffs of this coordination game are shown in Figure 3; the woman is player 1 and the man is player 2.

The final example of a game is an ecological game which was studied by Maynard Smith [8] in his analysis of the theory of evolution in terms of games. The game is played between animals of the same size who live in the wilderness and encounter one another in their search for prey. During an encounter between two animals, each animal has a choice between behaving as a hawk: *i.e.*, fighting for the prey; or as a dove: *i.e.*, running away peacefully. If both animals decide to play like hawks, then each animal has an equal chance of winning the value v of the prey or of losing the fight at cost c, where $0 < v < c$; thus, the expected payoff to both players is $(v - c)/2$. Alternatively, if both animals act as doves, then the prey is shared with equal payoffs $v/2$. Finally, if one animal behaves like a hawk and the other behaves like a dove, then the hawk gets a payoff worth the full value of the prey and the other gets nothing. In this game, the animals prefer to choose opposing strategies: if one animal plays hawk, then it is in the best interest of the other to play dove; and inversely, if one animal plays dove, then it is in the best interest of the other to play hawk.

1 \ 2	*H*	*D*
H	(v-c)/2, (v-c)/2	v,0
D	0,v	v/2,v/2

Figure 4: Hawks and Doves

This section included several popular examples of strategic form games. The next section presents the general theory of strategic form games and the Nash equilbrium solution concept, which follows directly from rationality. Although rationality is generally accepted in the deductive analysis of this class of games, it should be noted that this assumption leads to peculiar solutions in the examples that were presented. For example, in the Prisoner's Dilemma, (individually) rational behavior yields a solution which is sub-optimal. Moreover, in the game of Matching Pennies, rationality gives rise to a probabilistic solution. Finally, in the Battle of the Sexes and the game of Hawks and Doves, rationality does not generate a unique solution. In view of these quirky outcomes, one of the goals of the inductive approach to modern game theory is to give an alternative justification for Nash equilibrium.

2.2 Nash Equilibrium

This section develops the formal theory of finite games in strategic form. Let $\mathcal{I} = \{1, \ldots, I\}$ be a set of *players*, where $I \in \mathbb{N}$ is the number of players. The (finite) set of *pure strategies* available to player $i \in \mathcal{I}$ is denoted by S_i, and the set of pure strategy profiles is the cartesian product $S = \prod_i S_i$. By convention, write $s_i \in S_i$ and $s = (s_1, \ldots, s_I) \in S$. In addition, let $S_{-i} = \prod_{j \neq i} S_j$ with element $s_{-i} \in S_{-i}$, and write $s = (s_i, s_{-i}) \in S$. The payoff function $r_i : S \to \mathbb{R}$ for the i^{th} player is a real-valued function on S; in this way, the payoffs to player i depend on the strategic choices of all players. This description is summarized in the following definition.

Definition 2.1 *A* strategic form game *is a tuple* $(\mathcal{I}, \{S_1, \ldots, S_I\}, \{r_1, \ldots, r_I\})$, *where* $\mathcal{I} = \{1, \ldots, I\}$ *is a set of players, and for all* $i \in \mathcal{I}$,

- S_i *is the strategy set of the* i^{th} *player*
- $r_i : S \to \mathbb{R}$ *is the payoff function for the* i^{th} *player*

Example 2.1 *Formally, the Prisoner's Dilemma consists of a set of players* $\mathcal{I} = \{1, 2\}$, *with strategy sets* $S_1 = S_2 = \{C, D\}$, *and payoffs as follows:* $r_1(C, C) = r_2(C, C) = 4$, $r_1(C, D) = r_2(D, C) = 0$, $r_1(D, C) = r_2(C, D) = 5$, *and* $r_1(D, D) = r_2(D, D) = 1$.

An implication of the assumption of rationality is that a rational player always plays a *best response* to the strategies of the other players, where a best response is an optimizing strategy. A Nash equilibrium is a strategy profile in which all players choose strategies that are best responses to the choice of strategies by the other players. As such, at a Nash equilibrium, no player has any incentive to deviate from the prescribed strategy, given that no other players deviate. The formal definition of Nash equilibrium is given in terms of the best response set.

Definition 2.2 *The set of best responses for player* i *to strategy profile* s *is:*

$$\text{BR}_i(s) = \{s_i^* \in S_i \mid \pi_i(s_i^*, s_{-i}) \geq \pi_i(s_i, s_{-i}), \textit{ for all } s_i \in S_i\}$$

Let $\text{BR}(s) = \prod_i \text{BR}_i(s)$.

A Nash equilibrium is a strategy profile in which the players all play best responses to all the other players' strategies.

Definition 2.3 *A Nash equilibrium is a strategy profile* s^* s.t. $s^* \in \text{BR}(s^*)$.

It is apparent from this definition that a Nash equilibrium is a fixed point of the best response relation. In particular, in the Prisoner's Dilemma, (D, D) is a Nash equilibrium because given that player 1 plays D, the best response of player 2 is to play D; and, given that player 2 plays D, the best response of player 1 is to play D. The Battle of the Sexes has two pure strategy

Nash equilibria, namely (B, B), and (F, F), by the following reasoning. If the woman chooses B, then the best response of the man is B; and if the man chooses B, then the best response of the woman is B. Analogously, If the woman chooses F, then the best response of the man is F; and if the man chooses F, then the best response of the woman is F. Note that the game of Matching Pennies does not have a pure strategy Nash equilibrium. If player 1 plays H, then the best response of player 2 is T; but if player 2 plays T, the best response of player 1 is not H, but T. Moreover, if player 1 plays T, then the best response of player 2 is H; but if player 2 plays H, then the best response of player 1 is not T, but H. However, this game does have a mixed strategy Nash equilibrium. A mixed strategy is a probability distribution over a set of pure strategies.

Definition 2.4 *A mixed strategy set for player i (notation Q_i) is the set of probability distributions over the pure strategy set S_i:* i.e.,

$$Q_i = \{q_i : S_i \to [0,1] \mid \sum_{s_i \in S_i} q_i(s_i) = 1\}$$

The notational conventions extend to mixed strategies, *e.g.*, $q = (q_i, q_{-i}) \in Q$. In the context of mixed strategies, the payoff function $r_i : Q \to \mathbb{R}$ for the i^{th} player is given by:

$$r_i(q) = \sum_{s_i \in S_i} r_i(s_i, q_{-i}) q_i(s_i)$$

As usual, the payoffs to player i depend on the mixed strategies of all players.

Theorem 2.1 (Nash, 1951) *Every finite, strategic form game has a mixed strategy Nash Equilibrium.*

In the classical theory of strategic form games, common knowledge of rationality, together with common knowledge of the underlying structure of a game, implies Nash equilibrium (whenever the Nash equilibrium is unique). If all the players are rational, and if they all know that they are all rational, then they all know that the others all play best responses, and as a result, they all play best responses to those best responses; this is precisely the definition of Nash equilibrium. However, common knowledge of rationality is a very strong assumption. In the second half of this paper, we investigate whether or not Nash equilibrium can arise in the absence of common knowledge of rationality.

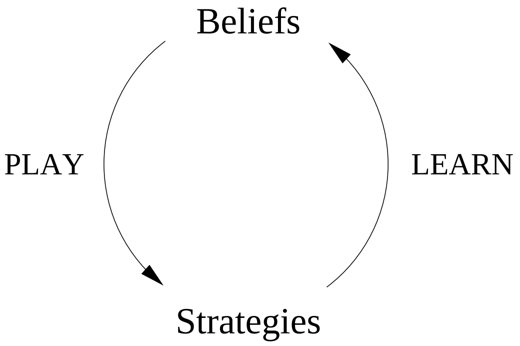

Figure 5: Model of Inductive Reasoning

3 Learning in Repeated Games

In strategic form games of complete information, deductive reasoning gives rise to Nash equilibrium, assuming common knowledge of rationality. One of the aims of modern game theory is to give an alternative justification for the concept of Nash equilibrium by investigating the following question: Is Nash equilibrium a necessary consequence of repeated play of strategic form games, in the absence of common knowledge? In other words, in addition to being understood as the result of a deductive reasoning process, is Nash equilibrium also the outcome of an inductive reasoning process? This question is studied using the above model of inductive reasoning (see Figure 5). Players have an initial set of beliefs. Given these beliefs, players *play* the game: *i.e.*, players choose a strategy from their respective strategy sets. Then, based on the observed strategies (or payoffs), players *learn*: *i.e.*, update their beliefs.

3.1 On Convergence to Nash Equilibrium

Recall that a strategic form game is specified by a set $\mathcal{I}$ of players, together with, for all $i \in \mathcal{I}$, a pure strategy set S_i and a payoff function r_i. A mixed strategy $q_i \in Q_i$ for the ith player is a probability distribution over the set of pure strategies. On the other hand, a belief is a probability distribution over the set of opponents' pure strategies (notation $p_{-i} \in Q_{-i}$). Intuitively, a strategy implicitly encodes how players behave as they learn from their opponents' past actions. Likewise, a belief records how players think other players will behave as the other players learn.

Definition 3.1 *A* history *of length t is a sequence of play drawn from S: $h^t = (s^1, \ldots, s^t)$. Let H^t denote the set of all t-histories, and let $H = \bigcup_t H^t$.*

The play of the game is prescribed by behavioral strategies, and how players learn is given by a learning rule. Formally, a behavioral strategy for the ith player is a function $g_i : H \to Q_i$ from the set of all possible histories to the ith mixed strategy set. Likewise, a learning rule is for the ith player is a function $f_i : H \to Q_{-i}$ from the set of all possible histories to the set of opponents' mixed strategy sets. For example, if the ith player learns according to Bayes' rule, then f_i computes player i's posterior belief about other players' actions in period t after observing h^{t-1}.

Definition 3.2 *A* learning process *is a family of functions $\{(f_i, g_i)\}_{1 \le i \le I}$, where f_i is a learning rule and g_i is a behavioral strategy for the ith player.*

A learning process induces a learning path.

Definition 3.3 *A* learning path *is a sequence $\{(h^t, p^t_{-i}, q^t_i) \mid t = 1, 2, \ldots\}$, where*

- *h^t is a t-history*
- *$p^t_{-i} = f_i(h^{t-1})$ is player i's belief at time t*
- *$q^t_i = g_i(h^{t-1})$ is player i's strategy at time t*

A learning path converges to Nash equilibrium if two conditions hold, namely *optimization* and *prediction*. Optimization means players maximize with respect to their beliefs. Prediction says that beliefs converge to other players' actual strategies: *i.e.*, player i's beliefs about player j's strategy become closer and closer to player j's actual strategy over time. Together, optimization and prediction imply that all players eventually play their best responses to all other players' actual strategies, which is the definition of Nash equilibrium.

Lemma 3.1 *A learning path $\{(h^t, p^t_{-i}, q^t_i) \mid t = 1, 2, \ldots\}$* converges to Nash equilibrium *iff the following two mathematical conditions are satisfied:*

- OPTIMIZATION*:* $\forall t,\ q_i^t \in \mathrm{BR}_i(p_{-i}^t)$
- PREDICTION*:* $\forall i \neq j,\ \lim_{t\to\infty}\ |(p_{-i}^t)_j - q_j^t\ | = 0$

In this inductive model of play, just as the notion of Nash equilibrium separates into two complementary components, the definition of rationality breaks down into two components, which are referred to herein as *rational play* and *rational learning*. Rational play is an extension of the economic understanding of rationality: players choose strategies which maximize their immediate payoffs, given their beliefs. Rational learning is generally accepted to mean Bayesian updating, namely the updating of prior beliefs in light of new information to obtain posterior beliefs.

Kalai and Lehrer [5] prove that repeated play of strategic form games among two rational players (in the two-fold sense described above) converges to approximate Nash equilibrium (specifically, ϵ-Nash equilibrium), provided that players' initial belief sets contain a *grain of truth*: *i.e.*, players assign positive probability to the strategies their opponents actually play.

Theorem 3.1 (Kalai and Lehrer, 1990) *Consider a game of incomplete information:* i.e.*, opponents' payoffs are unknown. Given a learning process* $\{(f_i, g_i)\}_{1\leq i\leq I}$*, there exists a time* T s.t. *for all* $t \geq T$*, there exists an* ϵ*-Nash equilibrium* s.t. *almost all learning paths* $\{(h^t, p_{-i}^t, q_i^t) \mid t = T, T+1, \ldots\}$ *converge to the* ϵ*-Nash equilibrium.*

Proof 3.1 (Sketch) The theorem states that if players are rational, and if players beliefs satisfy the grain of truth condition, then play converges to a Nash equilibrium: *i.e.*, optimization and prediction both eventually hold. Rational play is optimizing by definition. Moreover, Bayesians eventually learn to predict, if initial beliefs contain a grain of truth, by the following. The sequence of posterior probabilities that result from Bayesian updating satisfies a property like consistent estimation, which says that this sequence tends to get closer to the actual distributions of play, as t increases, provided that actual strategies are in the support of players' prior belief distibutions.

This section revealed positive results on convergence to Nash equilibrium in an inductive setting. This topic was further investigated by Nachbar, whose research produced negative results on the same question. His work is the subject of the next section.

3.2 On Non-Convergence to Nash Equilibrium

Nachbar [6] shows that repeated play of strategic form games among rational players does not in general converge to a Nash equilibrium, unless players' initial beliefs coincide with a Nash equilibrium. Let us introduce his approach by considering prediction. The prediction property states that via a learning mechanism, players' beliefs eventually converge to the other players' actual strategies. However, there does not exist an initial belief set from which a player can learn to predict the continuation path of play, for every possible opposing strategy. Intuitively, countably many finite histories serve as data from which players learn; however, there are uncountably many possible continuation strategies. The proof proceeds by a diagonalization argument.

Claim 3.2 *There is no learning rule which is aymptotically accurate against all strategies.*

Proof 3.2 Suppose not. Let f be such a learning rule. Construct strategy TROUBLE as follows (for simplicity of exposition, assume the pure strategy set consists of L and R). TROUBLE: Let f predict an arbitrary strategy g. If g plays L with probability $\leq 1/2^n$, TROUBLE plays L with probability 1. If g plays L with probability $> 1/2^n$, TROUBLE plays R with probability 1.

Since g was arbitrary, let f predict TROUBLE. Now, if TROUBLE plays L with probability $\leq 1/2^n$, then TROUBLE plays L with probability 1. *This is a contradiction.* Moreover, if TROUBLE plays L with probability $> 1/2^n$, then TROUBLE plays R with probability 1. *This is also a contradiction.* Therefore, learning rule f does not exist. Since f was arbitrary, there is no learning rule which is aymptotically accurate against all strategies.

It follows from the above observation that it is necessary to restrict either the set of initial beliefs or the set of possible strategies in order to achieve prediction. Kalai and Lehrer restricted the set of initial beliefs to those which contained a grain of truth. Nachbar introduces the notion of *conventional strategy sets*: the idea is that players can learn to predict the play of the game when strategies are conventional. Conventional strategies contain a *grain of reflection*: *i.e.*, players assign positive probability to their opponents' best responses to the strategies they actually play. The concepts of optimization and prediction are restated in conventional terms.

- CONVENTIONAL STRATEGIES: a subset of repeated game strategies that contain a grain of reflection
- CONVENTIONAL OPTIMIZATION: there always exists a conventional strategy that is a best response
- CONVENTIONAL PREDICTION: players learn to predict any profile of conventional strategies

In addition to considering conventional strategy sets: *i.e.*, those which are predictable, Nachbar is also interested in strategy sets which are neutral in the sense that they are not suggestive of a Nash equilibrium. In particular, a conventional set that consists only of the unique Nash equilibrium of a game is not neutral. This assumes away the problem of learning the Nash equilibrium. Standard examples of neutral strategies are those which satisfy some standard bound on complexity. For example, a strategy that can be described by a finite program is a neutral strategy. Another example of a neutral strategy is a memoryless strategy such as play L with probability p and play R with probability $1 - p$, regardless of history. Intuitively, play that consists of only neutral conventional strategies suggests that players have a preference for simple strategies.

Theorem 3.3 (Nachbar, 1995) *Consider a repeated strategic form game of complete information. For any neutral set of conventional strategies, there does not exist an initial belief set such that conventional optimization and conventional prediction hold simulataneously.*

Proof 3.3 (Sketch) Consider a neutral set of conventional strategies. By definition, conventional prediction holds if the strategy set is conventional. As a result, players learn to predict the others' strategies. However, their possible responses are limited to those which are neutral and conventional. These simple strategies are not best responses. For players to take advantage of their predictive abilities requires flexible, complex strategies. If we expand the set of conventional strategies to include more complicated strategies, this simply generates more and more complicated best responses. In fact, the only way to obtain conventional optimization is to include enough strategies such that conventional prediction fails. In particular, if the set of conventional strategies includes all strategies, then conventional optimization holds, but conventional prediction fails.

Example 3.4 Consider the game of Matching Pennies (see Figure 2). Let the set of neutral conventional strategies be $\{H^*, T^*, \text{ RANDOMIZE } (50:50)\}$, where H^* is HEADS ALWAYS and T^* is TAILS ALWAYS. Assume grain of truth; in particular, let both players assign positive probability to all conventional strategies. By results in Kalai and Lehrer [5], conventional prediction holds. In particular, if player 2 plays H^*, then player 1 eventually learns that player 2 is playing H^*. However, conventional optimization fails since none of the conventional strategies have the flexibility to adjust to learned information. We might try adding more sophisticated strategies such as: randomize in the first period; thereafter, matcher plays match to mismatcher's first play, and mismatcher plays mismatch to matcher's first play. Now, after a single period of play the possible continuations are $(H^*, H^*), (H^*, T^*), (T^*, H^*)$, and (T^*, T^*), none of which approaches the Nash equilibrium path of play. In this example, the only set of initial beliefs which leads to convergence to the Nash equilibrium is the Nash equilibrium strategy itself, since otherwise both players strictly prefer H^* or T^*, depending on their initial beliefs.

Nachbar's result does not contradict that of Kalai and Lehrer. Nachbar simply says that in the absence of common knowledge, this inductive model of repeated games does not give rise to Nash equilibrium play, unless initial beliefs are fortuitously aligned. Foster and Young [4] build on the work of Nachbar by isolating the culprit that drives this negative result. They point out that learning in this model of inductive reasoning is deterministic: the sequence of beliefs that it generates is fully determined by the initial conditions and the observed sequence of strategies. Intuitively, if there is no randomization in beliefs, then beliefs cannot converge to a randomized strategy Nash equilibrium. In particular, repeated play cannot converge to a Nash equilibrium in a game such as Matching Pennies which has a unique mixed strategy Nash equilibrium.

Theorem 3.4 (Foster and Young, 1996) *Consider a game of incomplete information:* i.e.*, opponents' payoffs are unknown. Assume common priors:* i.e.*, common knowledge of prior beliefs. Given learning process* $\{(f_i, g_i)\}_{1 \le i \le I}$ *that is deterministic in beliefs, every learning path either converges to a pure-strategy Nash equilibrium or does not converge in either beliefs or strategies.*

3.3 Santa Fe Bar Game

This section describes an interesting example of a repeated game, namely the Santa Fe bar problem. This game is one of complete information that affords a deductive solution, assuming common knowledge of rationality. In addition, computer simulations of inductive reasoning give rise to convergent behavior, in the absence of common knowledge, but in the presence of a small random component.

The *Sante Fe bar problem* was invented by Brian Arthur [1], an economist at the Sante Fe Institute, in order to illustrate how one might model inductive reasoning. Here is the scenario, in a nutshell:

> There is one bar in Sante Fe. Every night, the city dwellers make a decision as to whether or not they will go to the bar. They would all like to be at the bar if and only if it is not too crowded. More specifically, if there is an available seat, they all prefer to go to the bar, but if not, they all prefer to stay at home.

The Sante Fe bar problem is a non-cooperative game. As such, it can be expressed formally as a repeated strategic form game. The players in this game are the inhabitants of Santa Fe. The strategy set of the players consist of two strategies, namely go to the bar or stay home. The payoffs of the game are determined by the number of players that choose to go to the bar.

Definition 3.4 *During round t, for a bar of capacity c, the Santa Fe bar game is defined as follows:*

- $\mathcal{I} = \{1, \ldots, I\}$
- $S_i = \{\text{G, H}\}$, *where* G *means go to the bar and* H *means stay home*
- $r_i(s)$ *is given by the following matrix:*

i \ $-i$	$<c$	$>c$
G	+	-
H	-	+

The Nash equilibrium solution of this game is a mixed strategy profile. If there are 100 inhabitants of Santa Fe ($I = 100$), and if the capacity of the bar is 60 ($c = 60$), then the Nash equilibrium strategy for all players is to go to the bar 60% of the time and to stay home 40% of the time. Like the game of Matching Pennies, the Santa Fe bar game is a game of complete information with a mixed strategy Nash equilibrium. If all players deduce this Nash equilibrium *a priori*, and if all players solemnly play this Nash equilibrium strategy at all times, then indeed play converges to the Nash equilibrium, and attendance at the bar converges to capacity. This unlikely outcome occurs only under the assumption of common knowledge of rationality.

In an attempt to find a solution to this game in the absence of common knowledge, Arthur noted the following. On a given night, the inhabitants of Sante Fe must decide whether or not to go to the bar. Intuitively, all the players act in such a way as to maximize their degree of happiness (*i.e.*, everyone is always out to have a good time). At one extreme, if everyone predicts that everyone else will go to the bar, then no one goes to the bar; but then, no one is happy with their decision. On the contrary, if everyone predicts that no one is going to go to the bar, then everyone will indeed go to the bar; once again, no one is happy with their decision. In either case, no one is happy; thus, neither of the above situations will persist. In particular, the conclusion is that there is no uniform *rational* behavior which maximizes overall happiness. (This outcome is reminiscent of the ecological game of Hawks and Doves in which it is also preferable to behave in contrast with the behavior of others.)

Based on these observations, Arthur simulated the behavior of *boundedly rational* agents who use inductive reasoning to build expectational models. These models are formed from a pool of simple functions which aim to predict attendance at the bar. Initially, the agents randomly select a fixed number of predictor functions from the pool: $p_1, \ldots, p_K$. All agents maintain a vector of non-negative weights associated with their predictor functions. In particular, $w_i^t(p_k)$ is the weight associated with the kth predictor function by agent i at time t. Initially, all weights are uniform. Throughout the simulation, the agents monitor the accuracy of their predictor functions and update the corresponding weights accordingly. The play of the game is prescribed by the most accurate predictor functions, which is indicated by their relative weights. More specifically, agent i utilizes the following algorithm, during round t of the game:

1. PLAY: $s_i^t = p^*(\{s^{t-1}\})$, where $w_i^{t-1}(p^*) \geq w_i^{t-1}(p_k)$, for all p_k
2. LEARN: for all predictors p_k, $w_i^t(p_k) = f(w_i^{t-1}(p_k))$, where f is defined *s.t.*:
   - $w_i^t(p_k) > w_i^{t-1}(p_k)$, if $p_k(\{s^{t-1}\})$ is accurate,
   - $w_i^t(p_k) < w_i^{t-1}(p_k)$, otherwise.

In this model of inductive reasoning, agents exhibit rational play and boundedly rational learning. Their play is rational in the sense that it is optimal with respect to their beliefs, where beliefs are described by predictor functions. Their capacity for learning is bounded by their small pool of simple predictor functions. Arthur simulated this repeated game and obtained an efficient solution in which the overall attendance at the bar stabilized near the capacity of the bar.

The outcome of this simulation is consistent with the theoretical results reported in the last section. Recall that in a purely determinstic model of inductive reasoning, a mixed strategy Nash equilibrium must be known *a priori* and cannot be learned via repeated play. However, Arthur built a simulation with small random component that is seen in the initial choice of predictor functions: *i.e.*, beliefs. Randomized beliefs appears to be the primary factor which leads to convergence, but this remains to be proven.

4 Conclusion

In the first half of this paper it is established that in strategic form games of complete information, assuming common knowledge of rationality, as well as common knowledge of the strategy space and payoff functions, deductive reasoning give rises to the Nash equilibrium solution concept, whenever the Nash equilibrium is unique. Deduction is a reasoning process by means of which conclusions are drawn from principles (rules of logic) and information (set of premises), all of which is known *a priori*. Given valid rules of logic and a true set of premises, the conclusions of deductive inference must be true: *i.e.*, they cannot be rejected without contradicting the premises. In this sense, deductive conclusions are implicit in the premises. In particular, in classical game theory, the notion of Nash equilibrium is implicit in the assumption of common knowledge of rationality, since if all the players are rational, and if they all know that they are all rational, and so on, then they all know that the others all play best responses, and as a result, they all play best responses to those best responses, which is precisely a Nash equilibrium.

Common knowledge of rationality, together with common knowledge of strategies and payoffs, are generally accepted to be the sufficient conditions for Nash equilibrium. While these strategies do ensure the choice of the Nash equilibrium strategy profile (whenever the Nash equilibrium is unique), these conditions are far stronger than the minimal set of epistemic conditions which are sufficient for Nash equilibrium. Aumann and Brandenburger [2] observed the following: if all players are rational, and if all the players know their own (private) payoff functions and the strategy choices of all the other players, then the strategies constitute a Nash equilibrium. While common knowledge of rationality indicates *why* the players all play Nash equilibrium strategies, it suffices to simply *know* that the players all play Nash equilibrium stratgies without knowing why. Far from a need for common knowledge of rationality, these conditions simply require mutual knowledge of strategies. Furthermore, Parikh has pointed out in conversation that it is in fact not even necessary that players *know* the strategies of the other players, in the philosophical sense of justified true belief, but instead the weaker condition of *true belief* is all that is required. This final point further motivates the inductive approach to game theory.

The second half of this paper concentrated on the modern game-theoretic approach to Nash equilibrium, which asks the question: is Nash equilibrium a necessary consequence of inductive reasoning, in repeated play of strategic form games of complete information, in the absence of common knowledge of rationality, or in the presence of private information about the underlying payoff structure of a game? In contrast to deductive reasoning, inductive reasoning yields conclusions that are not contained in the premises, and are more general than the premises. A common example of inductive inference is mathematical induction on the structure of the natural numbers. Note that this type of inductive reasoning in fact contains both deductive and inductive counterparts; in particular, deduction is used to show that $P(n)$ implies $P(n+1)$, for all $n \in \mathcal{N}$. Similarly, the model of inductive reasoning that was considered in this paper is also part deductive and part inductive: rational players *deduce* their optimal play, but learning is an *inductive* process.

Inductive reasoning in repeated games is appealing because it gives rise to robust solutions, in the sense that players are capable of learning any type of behavior. For example, if an opponent who seemed rational in the past suddenly turns "crazy", then the other players modify their beliefs to incorporate the observation that the said opponent is now crazy. This being so, inductive conclusions are by definition merely *likely* to be true, given an initial set of beliefs and a set of observations, for if they were true beyond a shadow of a doubt, then the aforementioned robustness property would cease to hold. This implies that the inductive process can at best give rise to Nash equilibrium, with some positive probability. As beliefs converge to strategies at Nash equilibrium, this interpretation suggests that randomization over beliefs is a necessary component in any model of inductive reasoning that guarantees convergence to Nash equilibrium.

References

[1] Brian Arthur. Inductive reasoning and bounded rationality. *Complexity in Economic Theory*, 84(2):406–411, May 1994.

[2] Robert Aumann and Adam Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 63:1161–1180, 1995.

[3] Robert Axelrod. *The Evolution of Cooperation.* Basic Books, New York, 1984.

[4] Dean Foster and Peyton Young. Learning with hazy beliefs. *Preprint*, 1996.

[5] Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica*, 61:1019–1045, 1993.

[6] John Nachbar. Prediction, optimization, and learning in games. *Preprint*, 1996.

[7] Anatol Rapoport. *Two-Person Game Theory.* University of Michigan Press, Ann Arbor, 1966.

[8] John Maynard Smith. *Evolution and the Theory of Games.* Cambridge University Press, New York, 1982.