Imitation and the Emergence of Nash Equilibrium Play in Games with Many Players

Edward Cartwright

No 684

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

Imitation and the emergence of Nash equilibrium play in games with many players*

Edward Cartwright
Department of Economics,
EUREQua,
Les Maisons des Sciences Economiques,
106-112 boulevard de l'Hopital,
75647 Paris Cedex 13,
France.
e.cartwright@tiscali.co.uk.
http://myweb.tiscali.co.uk/ejcartwright.

October 2003

Keywords: imitation, best reply, convergence, Nash equilibrium

Abstract

We model a learning dynamic in which players imitate and innovate. Of interest is to question whether Nash equilibrium play emerges, and if so, the role that imitation plays in this emergence. Our main result provides a general class of coordination game for which approximate Nash equilibrium play does emerge. Important conditions include that players imitate 'similar' individuals. The role of imitation in learning is discussed in the context of two examples where it is shown that imitation can lead to Pareto superior outcomes.

*This paper is an outgrowth of Chapter 3 of the PhD Thesis 'Learning, bounded rationality and conformity in games with many players' completed at the University of Warwick.

1 Introduction

It is widely accepted that individuals have a tendency to imitate or conform to the actions of others.[1] This behavior may result from 'social influences', such as desires for popularity and acceptance (Jones 1984 and Bernheim 1994). Alternatively, in the presence of imperfect information and bounded rationality, imitation may be a way of drawing on the experiences of others (Kandori, Mailath and Rob 1993, Schlag 1998). Does it follow, however, that imitation is consistent with individual rationality? The current paper addresses this issue by modelling a learning dynamic in which players imitate and innovate and by questioning whether Nash equilibrium play emerges in the long run. Our main results provide a family of games for which the learning dynamic will indeed converge to an approximate Nash equilibrium. These results suggest that imitation can be consistent with individual rationality.

A key aspect of our model will be that players use interchangeably *two different* behavioral heuristics. In particular players are assumed to use interchangeably an imitation heuristic and an innovation heuristic. When imitating a player refers to a subset of the population – his reference group – and imitates the strategy of the most successful player in that group. When innovating a player selects a strategy that will, ceteris paribus, increase his payoff. Both of these heuristics have close parallels in the prior literature: In particular, the imitation heuristic is similar to the imitative behavior modelled by Selten and Ostmann (2000) (see Section 2.2) and nests as special cases the behavior modelled by Vega-Redondo (1997) and Alos-Ferrer, Ania and Schenk-Hoppe (2000). The innovation heuristic gives rise to, what is commonly called, 'a better reply dynamic' as modelled, for example, by Ritzberger and Weibull (1995) (see Section 2.3). The approach of the current paper, however, is distinguished from this prior literature by allowing players to use *both* heuristics.

Different learning heuristics will naturally have different advantages and disadvantages for the players who use them. That players will mix and match different heuristics is thus to be expected (Gigerenzer and Todd 1999). As noted above the prior literature, in treating imitation, has frequently assumed that players solely imitate, with some experimentation or error (e.g.

[1] Experimental evidence of social influence and imitation in the economic literature is provided by, amongst others, Selten and Apesteguia (2002) and Offerman, Potters and Sonnemans (forthcoming). The importance of conformity and imitation has long been recognized in psychology and sociology (see, for example Asch 1952, Deutsch and Gerard 1955 and for a more modern discussion Gross 1996).

Vega-Redondo 1997, Alos-Ferrer, Ania and Schenk-Hoppe 2000, Selten and Ostmann 2000). This research has led to results demonstrating that imitation leads to the emergence of non-Nash equilibrium states (e.g. Vega-Redondo 1997 and Selten and Ostmann 2002). These results suggest that imitation need not be consistent with individual rationality. Experimental studies provide, however, evidence to question this conclusion. Selten and Apesteguia (2002) and Offerman, Potters and Sonnemans (2002), for example, while finding evidence that individuals imitate, find that imitation is not all they do, and importantly, as a consequence, play may fail to converge on the outcomes predicted by imitative dynamics. Reflecting on this we feel it is crucial, as we do in the current paper, to model learning dynamics where imitation is only one of the heuristics that players use.[2]

We shall focus for the most part on 'games with many players'; that is, games where no one individual or small group of individuals can significantly alter the payoff of anyone but themselves. Our motivation for focussing on this type of game stems from a belief that imitation is most likely to occur in these games.[3] To formally model 'games with many players' we make use of a pregame framework introduced by Wooders, Cartwright and Selten (2003). By imposing a 'large game property' we are able to model a family of games satisfying the desired properties. The pregame framework also has the advantage that each player is characterized by an attribute; the attributes of players provide a natural metric with which to gauge the similarity of players. One interesting consequence is that we can model imitation in populations where everyone is different. By contrast, the prior literature typically restricts attention to symmetric games (e.g. Vega-Redondo 1997 and Gale and Rosenthal 1999). We are not restricted to symmetric games but our main results do require, as an assumption, that each player imitates similar individuals. Thus, instead of imposing symmetry we find that (approximate) symmetry emerges as a necessary condition.

Our first result treats an imitation dynamic and provides conditions for convergence to an imitation equilibrium; an imitation equilibrium need not be a Nash equilibrium.[4] Our main result, Theorem 2, provides sufficient

[2] Other authors who take this approach are Levine and Pesendorfer (2000, 2001) and Gale and Rosenthal (1999); we discuss these papers in more detail below.

[3] These are potentially complex games: predicting the actions of others may be difficult if not impossible and there appears little to be gained from trying to do so. There are, however, many players whom one can observe and learn from.

[4] If players value equality or 'fairness' then an imitation equilibrium may be an intuitively appealing concept of equilibrium. There is evidence that equality and 'fairness' are important to individuals (e.g. Clark and Oswald 1996 or Chapter 4 of Kagel and Roth 1995).

conditions for convergence of the imitation with innovation dynamic to an approximate Nash, imitation equilibrium. One condition is that the game be a coordination game with bound L and that each player refer to at least L players. A coordination game of bound L has the property that when L or more players change strategy and each of these L players achieves a payoff increase then the average payoff of the population also increases. A complementary Theorem 3 is presented where, unlike in Theorem 2, we require no upper bound on the size of reference groups.

Our results, given that approximate Nash equilibrium play emerges, suggest that imitation (in combination with innovation) can be consistent with individual rationality. Similar results were obtained by Gale and Rosenthal (1999) in the context of interaction in a Cournot like model.[5] An appealing aspect of our research is the generality of game modelled. For instance, the previous literature on learning has typically focussed on games where the existence of a Nash equilibrium is trivial (e.g. Vega-Redondo 1997, Levine and Pesendorfer 2000, 2001 and Gale and Rosenthal 1999) but this is not the case in the game we model.[6] Complementary results are also due to Wooders, Cartwright and Selten (2003) who demonstrate that, in large games, there exists approximate Nash equilibria in which 'similar players play similar strategies'. Note, however, that the question of whether players learn to play one of these equilibria is not addressed by Wooders et. al.; for a less general class of game, our Theorems 2 and 3 demonstrate that such an equilibrium will indeed emerge.

There is a large literature, not mentioned above, on the convergence of learning dynamics to Nash equilibrium play (Fudenberg and Levine 1998). Our results clearly add to this literature. One interesting issue is whether imitation can enable convergence where 'more rational' learning algorithms do not. It has long been recognized that adaptive learning dynamics (such as an innovation dynamic) need not converge to Nash equilibrium. Hart and Mas-Colell (2003) go further by demonstrating that there are no *uncoupled* dynamics that are guaranteed to converge to Nash equilibrium. Uncoupled dynamics, however, have the property that a player is not influenced by the payoffs of others. Imitation does not, therefore, give rise to an uncoupled dynamic. The question of the extent to which imitation can enable convergence to Nash equilibrium play remains therefore an interesting and open question. In a final section we provide two examples to discuss this issue and

[5] Other results in which variants of imitative learning lead to 'optimal actions' are due to Schlag (1998) and Ellison and Fudenberg (1993, 1995).

[6] This is assuming that players use pure strategies.

suggest why imitation can potentially 'help' play to converge to 'optimal' outcomes.

We proceed as follows; in Section 2 we outline the model and introduce the imitation and innovation heuristics. In Section 3 we analyze a dynamic in which players only use imitation. In Section 4 we add innovation before looking at learning in large games in Section 5. Section 6 provides a discussion on imitation and Section 7 concludes.

2 The model

Let $N = \{1, ..., n\}$ denote a finite *player set* and let $S = \{1, ..., K\}$ denote a finite *strategy set*. A *strategy vector* is given by $s = (s_1, ..., s_n) \in S^N$ where s_i is interpreted as the strategy of player i. Throughout it will be assumed that players do not play mixed strategies. Let Σ denote the set of strategy vectors. A *stage game* is given by a tuple $(N, S, \{u_i\}_{i=1}^n)$ consisting of a finite player set N, finite strategy set S and a *payoff function* $u_i : \Sigma \to \mathbb{R}$ for each player $i \in N$.

Given a stage game Γ, play is assumed to evolve over discrete time periods, indexed, $t = 0, 1, 2, ...$ In each period t the stage game Γ is played. Every player $i \in N$ is assumed to choose a strategy for period t conditional on the strategy vector of the previous period $t - 1$. The evolution of play is therefore modelled as a discrete time homogenous Markov chain $\{s(t)\}_{t\geq 0}$ on state space Σ. The transition matrix of the Markov chain will be denoted by P. The value P_{sm} is interpreted as the probability of state m immediately following state s.

We model the behavior of players using an imitation with innovation dynamic. This dynamic postulates that players use a combination of imitation and innovation in choosing a strategy to play. If a player decides to imitate then he uses an *imitation heuristic* while if he decides to innovate he uses an *innovation heuristic*. A player's *probability of innovation* details the likelihood that he will innovate. In using the imitation heuristic players make use of a *reference network*. These concepts are formalized below.

2.1 Reference network

Given a player set N a *reference matrix* R is an $N \times N$ Boolean matrix $R = [r_{ij}]$. If element $r_{ij} = 1$ we say that player i *refers to* player j while if $r_{ij} = 0$ we say that player i does not refer to player j. We set $r_{ii} = 1$ for all $i \in N$. That is, a player is assumed to refer to themselves. We do not assume that R is symmetric. The matrix R will also be referred to as a *reference network*. Given a reference network R, for each player $i \in N$, let $R(i)$ be the subset of N such that $j \in R(i)$ if and only if $r_{ij} = 1$. We refer to $R(i)$ as the *reference group* of player i. Thus, player j belongs to the reference group of player i if and only if player i refers to player j.[7]

Note that in the current paper we take a reference network as given and assume that this network does not change as play evolves. An alternative approach, whereby a player may change his reference group over time, is considered by Cartwright (2003).

2.2 Imitation heuristic

The imitation heuristic represents a procedure that a player i can use to choose a strategy for current period t conditioning on the strategy vector of the previous period $t-1$. This heuristic closely resembles an imitation dynamic introduced by Selten and Ostmann (2000). The heuristic can be summarized under an *imitation probability function* $p_i : \Sigma \to \Delta(S)$ where the value $p_i(k|s)$ is interpreted as the probability that a player i would select strategy k if strategy vector s was played in the previous period. When using the imitation heuristic a player can be seen to progress through three stages.

1. *Identify costrategists:* the set of *costrategists* of player i, denoted $C_i(s)$, are those players $l \in R(i)$ such that $s_l = s_i$.

2. *Identify success examples:* a *success example* of player i is a player $j \in R(i)$ such that
$$u_j(s) = \max_{l \in R(i)} u_l(s)$$

3. *Choose strategy:* player i chooses strategy $k \in S$ with probability $p_i(k|s)$ where (a) if there is a success example j of player i where $s_j = k$ then $p_i(k|s) > 0$, and (b) if every success example of player i is a costrategist of player i then $p_i(s_i|s) = 1$.

In identifying a set of costrategists player i identifies those players to whom she refers and who play the same strategy as herself. Note that player i must belong to the set of costrategists of player i. A success example of player i is any player j who earns the highest payoff of any player referred to by i. Note that player i may be a success example for player i. In choosing a strategy player i may choose the same strategy as a success example. That

[7] Given the reference matrix R the reference group $R(i)$ of player i could be thought of as the i'th row of R.

is, she may *imitate a success example*. If every success example of player i is also a costrategist then player i will play the same strategy as in the previous period.

The imitation heuristic permits a good deal of flexibility in agent behaviour. In particular, it may be that $p_i(k|s) > 0$ even if there is no success example of player i playing strategy k. Thus, player i need not imitate a success example but may experiment, make a mistake or simply keep to the strategy of a previous period. This latter possibility means that our results apply to dynamics, common in the literature (e.g. Young 1993 and Vega-Redondo 1997), where players are assumed to change strategies sequentially, i.e. one person per period, or have some positive probability of not changing strategy. Importantly, however, the imitation heuristic may be such that $p_i(k|s) = 0$ for any strategy k that is *not* played by a success example; that is a player may *always* imitate success examples. Thus, while experimentation and inertia are permitted in our model they are not necessary to derive our results.

The heuristics used by Kandori, Mailath and Rob (1993), Vega-Redondo (1997) and Alos-Ferrer, Ania and Schenk-Hoppe (2000) can be seen as a special case of the imitation heuristic for which $R(i) = N$ for all $i \in N$.[8] In comparing our imitation heuristic with that of Selten and Ostmann (2000) we note that our imitation heuristic allows the possibility that a player i may imitate a non-costrategist who is earning the same payoff as one of her costrategists. This implies, in particular, that she may imitate a non-costrategist who is earning the same payoff as herself. The imitation dynamic of Selten and Ostmann (2000) differs in that a player j can only be a success example of i if she is earning *strictly more* than the costrategists of i.[9] Cartwright (2003) also considers this form of imitation and demonstrates that analogues of the main theorems of the current paper hold.[10]

[8] All these dynamics assume a player has the option to choose the same strategy as in the previous period.

[9] Formally, this heuristic is identical to that of the imitation heuristic with one modification: a player j can be a success example of player i when $s_j \neq s_i$ if and only if

$$u_j(s) = \max_{l \in R(i)} u_l(s) > \max_{k \in C_i(s)} u_k(s).$$

[10] Less closely related models of imitation are due to, amongst others, Kirman (1993), Ellison and Fudenberg (1993, 1995), Levine and Pesendorfer (2000, 2001) and Gale and Rosenthal (2001). The principal difference between our approach and those taken in these papers is who people choose to imitate. Kirman (1993), Levine and Pesendorfer (2000, 2001) and Gale and Rosenthal (2001), for example, assume that an imitator is more likely to choose the strategy being played by the most people - irrespective of payoffs.

2.3 Innovation heuristic

In a similar way to the imitation heuristic, the innovation heuristic can be summarized by an innovation probability function $m_i : \Sigma \rightarrow \Delta(S)$. The value $m_i(k|s)$ is interpreted as the probability that a player i would select strategy k if strategy vector s was played in the previous period. Let $\varepsilon \geq 0$ be a real number referred to as an *inertia parameter*.

1. *Identify innovation opportunities:* an innovation opportunity for player i is a strategy $k \in S$ such that

$$u_i(k, s_{-i}) > u_i(s_{-i}, s_i) + \varepsilon.$$

2. *Choose strategies:* player i chooses strategy $k \in S$ with probability $m_i(k|s)$ where (a) if there are no innovation opportunities for player i then $m_i(s_i|s) = 1$, and, (b) if there is an innovation opportunity for player i then $m_i(k|s) > 0$ for some strategy k that is an innovation opportunity.

If a player could have improved upon her payoff by more than ε in the previous period then she has an innovation opportunity. If she has no innovation opportunities then she uses the same strategy as in the previous period. If, however, a player does have an innovation opportunity then there must be a positive probability that she plays at least one of her innovation opportunities. The possibility for mistakes, experimentation and inertia exist in the innovation heuristic to the same extent as it does in the imitation heuristic. The innovation heuristic clearly suggests a 'best response' or 'myopia' dynamic as much studied in the literature (e.g. Young 1993, Blume 1993, 1995). Note, however, that $m_i(k|s)$ can be zero even if k is an innovation opportunity. Thus, a player need not, necessarily, choose the innovation opportunity that would have maximized her payoff in the previous period.[11] The innovation heuristic is thus a 'better response' dynamic (see, for example, Ritzberger and Weibull 1995).

2.4 The imitation with innovation dynamic

It remains to combine the imitation and innovation heuristics to form the imitation with innovation dynamic. The final element we introduce is the

[11] This contrasts with the imitation heuristic where it is assumed that every success example is imitated with some positive probability.

vector of *innovation probabilities* $\lambda \in \mathbb{R}^N$ where $\lambda_i \in [0,1]$ is referred to as the *innovation probability of player i*. The value λ_i is the probability with which player i uses the innovation heuristic with the imitation heuristic used otherwise. Thus, if $\lambda_i = 0$ player i always uses the imitation heuristic. We say that $\lambda = 0$ if $\lambda_i = 0$ for all $i \in N$ and say that $\lambda \neq 0, 1$ if $\lambda_i \in (0,1)$ for all $i \in N$.[12]

Given a set of imitation probability functions $\{p_i\}_{i=1}^n$, a set of innovation probability functions $\{m_i\}_{i=1}^n$ and vector of innovation probabilities λ we can derive the transition matrix P of the Markov chain. The resulting stochastic process is referred to as the *imitation with innovation dynamic* which we indicate as $\mathcal{I}(p; m; \lambda)$. If $\lambda = 0$ then we refer to an *imitation dynamic*. It proves more convenient to characterize the imitation with innovation dynamic according to the inertia parameter ε, innovation probabilities λ and reference matrix R. We thus denote by $\mathcal{I}(\varepsilon; \lambda; R)$ any imitation with innovation dynamic that is consistent with the three characteristics indicated.[13]

As with any Markov process the state space Σ can be partitioned into a set of transient states T and a set of recurrent states Ψ. The probability of observing a state $s \in T$ converges to zero over time. Set Ψ can be further partitioned into *communication classes* $\Psi_1, ..., \Psi_C$. A communication class Ψ_c has the property that $\sum_{m \in \Psi_c} p_{sm} = 1$ for all $s \in \Psi_c$. Thus, if some strategy vector $s \in \Psi_c$ is played in period t then in every subsequent period a strategy vector belonging to the set Ψ_c will be played. If $\Psi_c = \{s\}$ then we refer to strategy vector s as an *absorbing state*. If $\Psi_c = \Sigma$ (i.e. the set of strategy vectors) then we say the dynamic is *irreducible*. As will become clear, without being more specific about the game and reference network, we cannot know the nature of the innovation with imitation dynamic; for some games and reference networks it many, for example, be irreducible while for others it may have only singleton communication classes. We note that in the current paper we will look for and provide conditions under which there exist only singleton communication classes.[14]

[12] The value of λ_i could be made conditional on the strategy vector and our results still apply. That is, the probability a player innovates could depend on the strategy vector of the previous period.

[13] The value of ε and a reference network R are insufficient to identify the set of functions p and m. Note, however, that the set of funtions p and m may be consistent with a unique value for ε and a unique reference matrix R.

[14] A motivation for this type of approach is suggested in the introduction of Fudenberg and Levine (1998).

3 The dynamics of imitation

We begin our analysis of the imitation with innovation dynamic by assuming that $\lambda = 0$. That is, by assuming an imitation dynamic. We define a static equilibrium concept.[15]

Imitation equilibrium: The strategy vector m is an *imitation equilibrium* relative to reference network R if:

$$\max_{l \in R(i)/C_i(m)} u_l(m) < \max_{l \in C_i(m)} u_l(m)$$

for all $i \in N$, where we recall that $C_i(m)$ denotes the set of costrategists of player i for strategy vector m.

If the state of the system is an imitation equilibrium then no player $i \in N$ has a success example who is not a costrategist and thus no player will wish to change strategy. This immediately suggests Lemma 1, which we state without proof.

Lemma 1: A state m is an absorbing state of the imitation dynamic $\mathcal{I}(\varepsilon; \lambda = 0; R)$ if and only if it is an imitation equilibrium relative to R.

Any strategy vector m in which every player $i \in N$ plays the same strategy is an imitation equilibrium.[16] Thus, there are many absorbing states of the innovation dynamic. In general there may also exist non-singleton communication classes as a very simple example illustrates.

Example 1: There are 3 players and 2 strategies, labelled A and B. The reference network is such that $R(1) = \{1, 2\}, R(2) = \{1, 2, 3\}$ and $R(3) = \{2, 3\}$. Two strategy vectors are of interest.

strategy vector	payoff vector
A, B, B	$4, 0, 2$
A, A, B	$2, 0, 4$

[15] An imitation equilibrium as defined in this paper is essentially equivalent to a *destination* as defined by Selten and Ostmann (2000). Selten and Ostmann (2000) require that an imitation equilibrium also be robust to possible deviations by success leaders.

[16] We note that an imitation equilibrium need not be such that every player plays the same strategy. Indeed a player need not play the same strategy as those he refers to.

There exists a communication class in which we see constant repetition of the strategy vectors (A, B, B) and (A, A, B). Basically, players 1 and 3 do not change strategy while player 2, by contrast, switches between strategies B and A, motivated by observing others earning a payoff of 4.⧫

We suggest that the cycle of play observed in Example 1 results for two reasons. First, player 2 is referring to, and imitating, players that are not 'similar' to himself. Second, the reference network is not sufficiently clustered. We shall have more to say on the first point in subsequent sections; here we pursue the second point. One important characteristic of a network is its clustering coefficient - a measure of the cliquishness of the network.[17]

Clustering coefficient: We say that a reference network R has a *clustering coefficient of one* when:

1. for any three distinct players $i, j, k \in N$ if $j, k \in R(i)$ then $k \in R(j)$ and $j \in R(k)$.[18]

2. $|R_i| \geq 3$ for every player $i \in N$.[19]

Thus, if a player $i \in N$ refers to both players j and k and the network R has a clustering coefficient of one then player j must refer to player k and player k refer to player j. We note that the reference network in Example 1 does not have a clustering coefficient of one; player 2 refers to players 1 and 3 but player 3 does not refer to player 1, nor player 1 refer to player 3. This lack of clustering allows the cycle of Example 1 to emerge, as demonstrated by our first result.

Theorem 1: For any stage game Γ and any reference network R that has a clustering coefficient of one the imitation dynamic $\mathcal{I}(\varepsilon; \lambda = 0; R)$ almost surely converges on an imitation equilibrium.[20]

[17] The clustering coefficient ranges between a value of 0 and 1. If the clustering coefficient is zero there is no clustering and if the clustering coefficient is one then there is much clustering. See D. Watts (1999) and references there in for a full definition and further discussion.

[18] Given that $i \in R_i$ it may appear that this condition implies symmetry of the network R whereby if $j \in R_i$ it must be the case that $i \in R_j$. The fact, however, that players i, j, k must be distinct means that the network need not be symmetric.

[19] The requirement that $|R_i| \geq 3$ is a minor assumption to rule out problems in defining the clustering coefficient if $|R_i| < 3$. We recall that $i \in R_i$.

[20] Given that a reference network has a clustering coefficient of one is sufficient to guarantee convergence on an imitation equilibrium we may ask whether or not it is necessary.

Proof: Given an arbitrary state s we demonstrate that there exists states, indexed, $s(2), ..., s(T)$ where $P_{ss(2)} > 0$, $P_{s(t)s(t+1)} > 0$ for all $T-1 \geq t \geq 2$ and where $s(T)$ is an imitation equilibrium. Assume that every player $i \in N$ in every period always chooses the same strategy as a success example. Furthermore, assume that there is an ordering to strategies (the same for all players) whereby if a player i has more than one success example he selects the strategy of the success example playing the 'smallest' strategy. This behavior is consistent with a deterministic process that occurs with positive probability under the imitation with innovation dynamic.

Consider an arbitrary player $i \in N$ for whom there exists a player $j \in R(i)$, $j \neq i$ such that $i \in R(j)$. For any player $l \in N$ such that $l \in R(i)$, given that the reference network R has a clustering coefficient of one, it must be the case that $l \in R(j)$ and $j \in R(l)$. This, in turn, implies that $i \in R(l)$. Similarly, if there exists a player $h \in R(j)$ then $h \in R(i)$ and $i, j \in R(h)$. Thus, $R(j) = R(i)$ for all $j \in R(i)$. We refer to the set $R(i)$ as a *clique*; every player within a clique refers to, and only to, all other players in the clique. Given the behavior assumed of players, in state $s(2)$ there must exist some $k \in S$ such that $s_j = k$ for all $j \in R(i)$. That is, all players in the clique play the same strategy. This implies that no player $j \in R(i)$ can have a success example in states $s(2), s(3), ...$ who is not a costrategist. Thus, no player i belonging to a clique can change strategy between states $s(2), s(3), ...$

Consider an arbitrary player $i \in N$ for whom there does not exist a player $j \in R(i)$, $j \neq i$ such that $i \in R(j)$. Suppose that there exists a player $l \in N$ such that $i \in R(l)$. Given that the network R has a clustering coefficient of one there must exist a player $j \neq i$ such that $j \in R(l)$. Further, if $i, j \in R(l)$ this implies that $i \in R(j)$ and $j \in R(i)$. This is a contradiction. Thus, $i \notin R(k)$ for all $k \in N \backslash \{i\}$. We say that player i does not belong to a clique. Player i does, however, refer to a subset of a clique. This is immediate from the analysis of the previous paragraph and the fact that i refers to at least two distinct players j, l who must refer to each other. Given that player i refers to a subset of a clique in states $s(2), s(3), ...$ every player referred to by player i (with the possible exception of themselves) must be playing the same strategy. Thus, if there is a success example of player i who is not a costrategist in some state $s(t_i)$ there cannot be a success example of player i in any subsequent state unless they are costrategists of i. Given that the player set is finite there must exist some t_i such that for every state $s(t)$,

Example 1 demonstrates that for any reference network R in which there are three players i, j, k where $j \in R(i)$ and $k \in R(i)$ but $k \notin R(j)$ or $j \notin R(k)$, a game Γ can be constructed for which the imitation with innovation dynamic has a non-singleton communication class. As shown by Cartwright (2003) we cannot go any further than this.

$t \geq t_i$, player i does not have a success example who is not a costrategist. This completes the proof.■

We conclude this section with a discussion of the likelihood that economic and social networks have a clustering coefficient of one. An illustration of a familiar economic network may be useful - consider firms competing in a market. Many markets, such as food retail, are composed of a small number of large, 'dominate' firms and a large number of small, 'fringe' firms. Firms can be expected to refer to the actions of competitors in order to gauge variables such as prices and marketing strategy. The following type of reference network seems plausible - (a) the large firms refer to each other, ignoring the small firms, while (b) the small firms refer solely to a subset of the large firms. This network would have a clustering coefficient of one.

Speaking more generally, it is unlikely that a network should have a clustering coefficient of one. It is, however, not unlikely that economic and social networks should have clustering coefficients that are 'near to one' (D. Watts 1999 and references therein) or have 'a tendency to converge to one' (Granovetter 1973). While definitive results seem unlikely, Theorem 1 is suggestive that play will converge to an imitation equilibrium when the reference network has a clustering coefficient that is close to one. Future work hopes to address this issue.

4 Adding Innovation

It should be apparent that an imitation equilibrium need not be a Nash equilibrium. Indeed a player may be able to significantly improve her payoff by selecting a different strategy than that consistent with an imitation equilibrium. Given that our principal interest is in imitation we move immediately to modelling the imitation with innovation dynamic. We begin with a definition.

Nash ε-equilibrium: Strategy vector m is a *Nash ε-equilibrium* if:

$$u_i(m_i, m_{-i}) \geq u_i(k, m_{-i}) - \varepsilon$$

for all $i \in N$ and for all $k \in S$.

We refer to a Nash 0-equilibrium as a Nash equilibrium and informally a Nash ε-equilibrium (for $\varepsilon > 0$) as an approximate Nash equilibrium. We refer to a state m that is both an Nash ε-equilibrium and an imitation equilibrium as a Nash, imitation ε-equilibrium or informally an approximate Nash, imitation equilibrium. The following result is stated without proof:

Lemma 2: A state m is an absorbing state of the imitation with innovation dynamic $\mathcal{I}(\varepsilon; \lambda \neq 0, 1; R)$ if and only if it is a Nash, imitation ε-equilibrium.

The principal implication of Lemma 2 is that *if* an imitation with innovation dynamic converges to an absorbing state then that state must be an approximate Nash equilibrium. Generally, however, the imitation with innovation dynamic will have *non-singleton* communication classes. There are two basic reasons for this. First, the dynamic may have no absorbing states; this follows from the principle that games need not have an approximate Nash equilibrium (in pure strategies); this issue is treated in detail by Wooders, Cartwright and Selten (2003) and Cartwright and Wooders (2003). Second, the dynamic may have absorbing states but also non-singleton communication classes giving rise to 'cyclical dynamics'; this possibility should be familiar from Example 1. Simple examples of both of these possibilities can easily be derived (Cartwright 2003). Clearly, therefore, to guarantee convergence to an approximate Nash equilibrium conditions will be required on the game and reference network. We provide such conditions in the following section.

5 Games with Many Players

We draw on a pregame framework used by Wooders, Cartwright and Selten (2003). This framework enables us to model a family of games that can all be seen to share a common structure of strategic interaction. A further advantage of the framework is that it permits us to compare the similarity of players - which proves important in looking at imitation.

Take as given a compact metric space of player *attributes* Ω and a finite set of strategies S. An attribute is interpreted as a complete description of a players characteristics and payoff function. Let W denote the set of functions mapping $\Omega \times S$ into $\mathbb{Z}_+$. A member of w is called a *weight function.* A *universal payoff function* h maps $\Omega \times S \times W$ into $\mathbb{R}_+$. In interpretation $h(\omega, k, w)$ is the payoff of a player of attribute ω from playing strategy k when the strategies of the complementary player set are summarized by w. A *pregame* is given by the triple (Ω, S, h).

Let N be a finite set and let α be a mapping from N to Ω, called an *attribute function*. The pair (N, α) is a *population.* As we shall formal explain, a population induces through the pregame a stage game $\Gamma(N, \alpha)$.

We recall that a stage game is given by a tuple $(N, S, \{u_i^\alpha\}_{i=1}^n)$. The player set N and strategy set S are clearly given; it remains, therefore, to define the payoff functions $\{u_i^\alpha\}_{i=1}^n$ where u_i^α maps from the set of strategy vectors S^N into $\mathbb{R}_+$. Given a population (N, α) and a strategy vector m we say that weight function $w_{\alpha,m}$ is *relative to strategy vector* m if and only if:

$$w_{\alpha,m}(\omega, k) = |\{i \in N : \alpha(i) = \omega \text{ and } m_i = k\}|$$

for all $\omega \in \Omega$ and $k \in S$. Thus, $w_{\alpha,m}(\omega, k)$ denotes the number of players of attribute ω who are playing strategy k. For each $i \in N$ payoff functions are defined to satisfy:

$$u_i^\alpha(s_i, s_{-i}) \equiv h(\omega, k, w_{\alpha,s})$$

for all $s \in S^N$.

Different populations may clearly induce different games. Thus, the pregame framework allows us to model a family of games. Very little structure is, however, imposed on the games that may be induced - we merely require that players of the same attribute be 'identical' both in terms of their payoff functions and how they influence others.

5.1 Large games

Following Wooders, Cartwright and Selten (2003) we make two assumptions on pregames. The first is a relatively mild assumption requiring players with 'similar' attributes to be 'similar'. Formally, it can be stated:

Continuity in attributes: The pregame $\mathcal{G} = (\Omega, S, h)$ satisfies *continuity in attributes* when: for any $\varepsilon > 0$ and any two populations (N, α) and $(N, \overline{\alpha})$, if,

$$dist(\alpha(i), \overline{\alpha}(i)) < \varepsilon$$

for any $i \in N$ then for any strategy vector m it holds that,

$$\left|u_i^\alpha(m_i, m_{-i}) - u_i^{\overline{\alpha}}(m_i, m_{-i})\right| < \varepsilon$$

for any $i \in N$.

An important point to note is that the assumption of continuity in attributes treats a change in the attributes of players while the strategies they play are held constant. For this reason it appears mild.

Our second assumption is reflective of the type of game we wish to consider - namely games with many players where no one player has a significant influence on the payoffs of others.

Large game property: The pregame $\mathcal{G} = (\Omega, S, h)$ satisfies *the large game property* when: for any $\varepsilon > 0$, any game $\Gamma(N, \alpha)$ and any two strategy vectors s and m if:

$$\frac{1}{|N|}\sum_{k\in S}\sum_{\omega\in\alpha(N)}|w_{\alpha,s}(\omega,k) - w_{\alpha,m}(\omega,k)| < \varepsilon$$

then:

$$|u_i^\alpha(k, m_{-i}) - u_i^\alpha(k, s_{-i})| < \varepsilon \tag{1}$$

for all $i \in N$ and all $k \in S$.

The large game property dictates that a player achieves approximately the same payoff given any two strategy vectors where the *proportions* of players of each attribute playing each strategy are approximately the same. This means that no one player or small group of players can have a large influence on the payoff of anybody but themselves. As discussed in the introduction we feel that imitation with innovation is most likely to be observed in these types of games.

An example in Cartwright (2003) illustrates that play may fail to converge on an absorbing state in games induced from a pregame satisfying continuity in attributes and the large game property.

5.2 Coordination Games

We provide a definition of a coordination game. For any two strategy profiles m and s let $X(m, s) \subset N$ be those players $j \in N$ such that $m_j \neq s_j$.

Coordination game: Game $\Gamma(N, \alpha)$ is a *coordination game with bound* L when: for any two strategy profiles m and s where $|X(m, s)| > L$ if:

$$u_i^\alpha(m_i, m_{-i}) > u_i^\alpha(s_i, s_{-i})$$

for all $i \in X(m, s)$ then,

$$\sum_{i\in N} u_i^\alpha(m_i, m_{-i}) > \sum_{i\in N} u_i^\alpha(s_i, s_{-i}). \tag{2}$$

Let $\mathcal{CG}(L)$ denote the set of coordination games with bound L that can be induced from pregame $\mathcal{G}$. A coordination game with bound L has the property that when more than L players change strategy and each player who changes strategy gets a payoff increase then the 'total payoff of the population' increases. We note that any game $\Gamma(N, \alpha)$ belongs to set $\mathcal{CG}(|N|)$.

5.3 Large Game Reference Networks

We briefly turn our attention to reference networks. Given population (N, α) and player $i \in N$ we denote by $B_i(\delta)_\alpha$ the subset of player set N such that player $j \in B_i(\delta)_\alpha$ if and only if $dist(\alpha(i), \alpha(j)) \leq \delta$. That is, if we draw a ball in attribute space around $\alpha(i)$ of diameter δ then $B_i(\delta)_\alpha$ is those players within the ball.

Large game reference networks: We say that reference network R is a *large game reference network with bounds L, U and δ* if:

1. R is symmetric[21] and has a clustering coefficient of one,
2. $R(i) \subset B_i(\delta)_\alpha$ for all $i \in N$, and,
3. $L \leq |R(i)| \leq U$ for all $i \in N$.

We denote by $\mathcal{LR}(L, U, \delta)$ the set of large game reference networks with bounds L, U and δ.

Behind the concept of a large game reference network are basically two refinements on reference networks studied in Section 3. First, R is assumed to be symmetric. As discussed by Cartwright (2003) this is not necessary for our results - it does, however, significantly simplify the analysis. Symmetry is also a common simplifying assumption in modelling social networks (e.g. D. Watts 1999). The second, and most important refinement, is that players are assumed to only refer to those with 'similar' attributes as themselves. This proves crucial to deriving our main results.[22]

5.4 Main result

We have now introduced all the necessary concepts to state our second result. This result provides conditions under which the imitation with innovation dynamic converges to an absorbing state.

Theorem 2: Let $\mathcal{G}$ be any pregame satisfying continuity in attributes and the large game property. Given $\varepsilon > 0$ and positive integer U there exists real number $\eta_2(\varepsilon, U)$ such that for any population (N, α) where $|N| > \eta_2(\varepsilon, U)$ if

[21] That is, if $i \in R_j$ then $j \in R_i$ for all $i, j \in N$.

[22] Note that the bound L on the minimum size of reference group is not independent of the value δ on similarity of players in the same reference group. The smaller is δ then the smaller may have to be L.

$\Gamma(N,\alpha) \in \mathcal{CG}(L)$ and $R \in \mathcal{LR}\left(L, U, \frac{\varepsilon}{3}\right)$, for some L, then the imitation with innovation dynamic $\mathcal{I}(\varepsilon; \lambda \neq 0, 1; R)$ almost surely converges on a Nash, imitation ε-equilibrium.[23]

Proof: Suppose that the statement of the Theorem is false. Then there exists some $\varepsilon > 0$ and some U such that, for each integer ν there is a population (N^ν, α^ν) where $|N^\nu| > \nu$, where $\Gamma(N^\nu, \alpha^\nu) \in \mathcal{CG}(L^\nu)$ and $R^\nu \in \mathcal{LR}\left(L^\nu, U, \frac{\varepsilon}{3}\right)$ for some L^ν, and for which there exists a non-singleton communication class of the imitation with innovation dynamic $(\lambda \neq 0, 1)$.

From the proof of Theorem 1 it is immediate that the population (N^ν, α^ν), for any ν, can be partitioned into a set of cliques. That is, the player set N^ν can be partitioned into subsets $c_1^\nu, ..., c_{Q^\nu}^\nu$ with the property, for all $i \in N^\nu$, that if $i \in c_q^\nu$ then $R^\nu(i) = c_q^\nu$.

For any game $\Gamma(N^\nu, \alpha^\nu)$ and any initial state s^ν suppose that play evolves according to the following process,

1. all players $i \in N^\nu$ use the imitation heuristic, and imitate any success example, until the process evolves to an imitation equilibrium.

2. in the following period a unique player $i \in N^\nu$ uses the innovation heuristic and chooses an innovation opportunity. All other players use the imitation heuristic.

3. the process returns to stage 1 and repeats.

Fix a value for ν and consider the evolution of play. By Theorem 1 play will, almost surely, converge to an imitation equilibrium s during the first stage of the process. For each clique c_q there must exist some strategy $k_q \in S$ such that $s_j = k_q$ for all $j \in c_q$. That is, any two players in the same clique play the same strategy.

If a contradiction is to be avoided there must exist some player $i \in N^\nu$ who has an innovation opportunity given strategy vector s. Suppose, that in stage 2 of the process player i chooses an innovation opportunity. This implies that strategy vector $\overline{s}$ is observed in the next period (say period t) where $\overline{s}_j = s_j$ for all $j \in N^\nu \backslash \{i^\nu\}$ and

$$u_i^{\alpha^\nu}(\overline{s}_i, \overline{s}_{-i}) > u_i^{\alpha^\nu}(s_i, s_{-i}) + \varepsilon. \tag{3}$$

Let $\delta \equiv \frac{\varepsilon}{3}$. In period $t+1$, all players use the imitation heuristic. We note that if $i \in c_{\overline{q}}$ then no player $l \in c_q$ where $c_q \neq c_{\overline{q}}$ can have a success

[23] It is apparent from the proof that the statement of the Theorem can be relaxed to $\mathcal{LR}(L, U, *)$ where $* > \frac{\varepsilon}{2}$ is arbitrarily close to $\frac{\varepsilon}{2}$.

example who is not a costrategist. Thus, if strategy vector $\overline{\overline{s}}$ is observed in period $t+1$ then $\overline{\overline{s}}_l = \overline{s}_l$ for all $l \in N^\nu \backslash c_{\overline{q}}$. Given continuity in attributes for any $j \in c_{\overline{q}}$ we have:

$$\left| u_i^{\alpha^\nu}(s_i, s_{-i}) - u_j^{\alpha^\nu}(s_j, s_{-j}) \right| < \delta. \tag{4}$$

By the large game property, if ν is chosen sufficiently large, for any player $j \neq i$ we have:

$$\left| u_j^{\alpha^\nu}(\overline{s}_j, \overline{s}_{-j}) - u_j^{\alpha^\nu}(s_j, s_{-j}) \right| < \delta. \tag{5}$$

By (3), (4) and (5) it holds that:

$$u_i^{\alpha^\nu}(\overline{s}_i, \overline{s}_{-i}) > u_j^{\alpha^\nu}(\overline{s}_j, \overline{s}_{-j}).$$

for all $j \in c_{\overline{q}} \backslash \{i\}$. This implies that player i is the unique success example for those players $j \in c_{\overline{q}} \backslash \{i\}$. Note that player i will be their own and only success example. Thus, $\overline{\overline{s}}_j = \overline{s}_i$ for all $j \in c_{\overline{q}}$.

Given the large game property and the fact that U is independent of ν, for ν sufficiently large:

$$\left| u_i^{\alpha^\nu}(\overline{s}_i, \overline{s}_{-i}) - u_i^{\alpha^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) \right| < \delta. \tag{6}$$

This implies that by (3) that:

$$u_i^{\alpha^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) > u_i^{\alpha^\nu}(s_i, s_{-i}) + 2\delta. \tag{7}$$

Continuity in attributes implies:

$$\left| u_i^{\alpha^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) - u_j^{\alpha^\nu}(\overline{\overline{s}}_j, \overline{\overline{s}}_{-j}) \right| < \delta \tag{8}$$

for all $j \in c_{\overline{q}}$. Thus, by (4), (7) and (8) we have:

$$u_j^{\alpha^\nu}(\overline{\overline{s}}_j, \overline{\overline{s}}_{-j}) > u_j^{\alpha^\nu}(s_j, s_{-j}) \tag{9}$$

for all $j \in c_{\overline{q}}$.

Compare strategy vectors s and $\overline{\overline{s}}$. We note that $X(s, \overline{\overline{s}}) = c_{\overline{q}}$. It is immediate from (7) and (9), given that $\Gamma(N^\nu, \alpha^\nu) \in \mathcal{CG}(L^\nu)$ and $R \in \mathcal{LR}\left(L^\nu, U, \frac{\varepsilon}{3}\right)$, that, for sufficiently large ν:

$$\sum_{j \in N^\nu} u_j^{\alpha^\nu}(\overline{\overline{s}}_j, \overline{\overline{s}}_{-j}) > \sum_{j \in N^\nu} u_j^{\alpha^\nu}(s_j, s_{-j}).$$

Also note that $\overline{\overline{s}}$ is an imitation equilibrium. Thus, as play evolves repeatedly as above the total payoff of the population increases and never decreases. Given that the state space is finite this gives the desired contradiction.■

A corollary of Theorem 2 (and of Theorem 3 to follow) is that, given the appropriate conditions, there must exist an approximate Nash equilibrium and imitation equilibrium.[24] This complements results due to Wooders, Cartwright and Selten (2003) and Cartwright and Wooders (2003) in providing sufficient conditions for the existence of an approximate Nash equilibrium 'consistent with conformity'.

5.5 Large reference groups

We offer a complementary result to that of Theorem 2 in which we place no upper bound on the maximum reference group size. In doing so we place a third assumption on the pregame.[25]

Coordination property: The pregame $\mathcal{G} = (\Omega, S, h)$ satisfies the *coordination property* when: for any induced game $\Gamma(N, \alpha)$, any two strategy vectors m and s and any player i such that $\alpha(i) = \overline{\omega}$ and $m_i = s_i = \overline{k}$ if:

$$\begin{aligned} w_{\alpha,m}(\overline{\omega}, \overline{k}) &> w_{\alpha,s}(\overline{\omega}, \overline{k}) \text{ and} \\ w_{\alpha,m}(\omega, k) &= w_{\alpha,s}(\omega, k) \text{ for all } \omega \neq \overline{\omega} \text{ and } k \in S \end{aligned}$$

then:

$$u_i^\alpha(m_i, m_{-i}) \geq u_i^\alpha(s_i, s_{-i}).$$

If a pregame satisfies the coordination property then, ceteris paribus, a player's payoff cannot decrease if the number of players with the same attribute as himself playing the same strategy as himself increases.

We state our third result.

Theorem 3: Consider pregame $\mathcal{G}$ that satisfies the large game property, continuity in attributes and the coordination property. Given $\varepsilon > 0$ there exists real number $\eta_3(\varepsilon)$ such that for any population (N, α) where $|N| > \eta_3(\varepsilon)$ if $\Gamma(N, \alpha) \in \mathcal{CG}(L)$ and $RN(\alpha) \in \mathcal{LR}\left(L, |N|, \frac{\varepsilon}{4}\right)$, for some L, then the

[24] Note that a Nash equilibrium need not exist in coordination games even for large populations. Consider, for example a population of players matched to play a 'two strategy, off diagonal coordination game'. The unique Nash equilibrium is 'half the population play one strategy and the other half play the other strategy'. There can only exist a Nash equilibrium when there are an even number of players.

[25] It is clear that further assumptions are required than used in Theorem 2. Consider for example the case where every player has the same attribute and refers to all players in the population. Also suppose that any approximate Nash equilibrium has the property that half the population play one strategy and the other play another strategy.

imitation with innovation dynamic $\mathcal{I}(\varepsilon; \lambda \neq 0, 1; R)$ almost surely converges to a Nash, imitation ε-equilibrium..

Proof: A proof proceeds in an almost identical fashion to that of Theorem 2. It is only with respect to (6) that we observe any significant difference. Thus, using the notation of the proof of Theorem 2, we will merely demonstrate an analogue of (6) in the new context of Theorem 3. Set $\delta \equiv \frac{\varepsilon}{4}$. Informally, our objective is to demonstrate that the payoff of player i falls by at most δ when the players in the clique $c_{\overline{q}}$ imitate player i and play strategy $\overline{s}_i$.

Given the population (N^ν, α^ν) consider the population $(N^\nu, \overline{\alpha}^\nu)$ where $\overline{\alpha}^\nu(j) = \alpha^\nu(i)$ for all $j \in c_{\overline{q}}$ and where $\overline{\alpha}^\nu(j) = \alpha^\nu(j)$ for any other $j \in N^\nu$. By continuity in attributes:

$$\left| u_i^{\alpha^\nu}(\overline{s}_i, \overline{s}_{-i}) - u_i^{\overline{\alpha}^\nu}(\overline{s}_i, \overline{s}_{-i}) \right| < \delta$$

and

$$\left| u_i^{\alpha^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) - u_i^{\overline{\alpha}^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) \right| < \delta.$$

By the continuity property:

$$u_i^{\overline{\alpha}^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) \geq u_i^{\overline{\alpha}^\nu}(\overline{s}_i, \overline{s}_{-i}).$$

Thus, we obtain:

$$\left| u_i^{\alpha^\nu}(\overline{s}_i, \overline{s}_{-i}) - u_i^{\alpha^\nu}(\overline{\overline{s}}_i, \overline{\overline{s}}_{-i}) \right| < 2\delta. \tag{10}$$

Substituting (10) for (6) and recognizing the change in the value of δ a proof of Theorem 3 should now be apparent from that of Theorem 2.■

Theorem 3 complements Theorem 2 in providing a convergence result when no player is bounded in the number of players he refers to. Indeed, it may be the reference network has the property that every player refers to every other player in the population; if all 'players are sufficiently similar' then Theorem 3 could be applied to show the existence of an approximate Nash equilibrium where every player in the population plays the same strategy. We also note that the large game property is much less important in the proof of Theorem 3 than in the proof of Theorem 2: in the proof of Theorem 3 it is merely required that the actions of *one* player have only a limited effect on the payoffs of others.

6 Imitation and Rationality

A principal motivation for the current paper was to address whether imitation can be consistent with 'rational play'. Theorems 2 and 3 suggest that, within our framework, imitation can be consistent with individual rationality; this follows from the fact that players imitate and yet play converges to an approximate Nash equilibrium.[26]

A further possibility is that imitation is not only consistent with individually rational play but actually enables or 'helps' boundedly rational players to behave optimally. This may happen, for example, because convergence to a Nash equilibrium is only observed when there is imitation. Alternatively, it may be that in the presence of imitation play converges on Pareto superior outcomes. Both of these possibilities can indeed occur within the types of games considered in this paper. To illustrate we pursue through two examples the second possibility given above.[27]

Before considering the two examples we briefly question why it may be that imitation 'enables learning'. We offer three reasons: (1) Imitation may 'speed up' the learning process in that imitation is easy and quicker than innovation (Levine and Pesendorfer 2000, 2001). This is surely the case but not of issue in the current paper given our focus on the long run properties of play and not rates of convergence. (2) By observing the actions of others a player may become aware of a strategy that he would not have realized was possible otherwise. This occurs to some extent in Example A. (3) Imitation has the potential to create a sense of 'collective action' in a way that unilateral behaviour, such as innovation, does not. We see this occurring in both examples but particularly in Example B.

We now turn to our two examples. Both examples could be thought of in terms of technological or scientific evolution where strategy B is a superior technology or technique to A etc.

Example A: There exists a unique attribute. There are five strategies

[26] Under the dynamics used in the proof of Theorem 2 a player who imitates always increases his payoff. This clearly helps explain why imitation is consistent with individual rationality. It also suggests that we could equate imitation with innovation. This is, however, not the case because a player through imitation can realize individual gains of less than ε - which is not possible through innovation. This, in turn, suggests that we could just set $\varepsilon = 0$. If, however, $\varepsilon = 0$ there need not exist a Nash ε-equilibrium.

[27] The first possibility is clearly also of interest and considered by Cartwright (2003). Examples necessarily, however, have to be more complex than the examples used here with a consequent loss of intuition. A greater understanding of the convergence properties of the innovation dynamic would also be useful to fully treat this issue.

$S = \{A, B, C, D, E\}$. Denote by $w_s(k)$ the *proportion* of the population playing strategy k given a strategy vector s. For population (N, α) and any $i \in N$ the payoff function is given by:[28]

$$\begin{aligned} u_i^\alpha(A, s_{-i}) &= 1; \\ u_i^\alpha(B, s_{-i}) &= 1 + w_s(B) + w_s(C) \\ u_i^\alpha(C, s_{-i}) &= 1 + w_s(B) + w_s(C); \\ u_i^\alpha(D, s_{-i}) &= 3(w_s(B) + w_s(D)); \\ u_i^\alpha(E, s_{-i}) &= 3(w_s(C) + w_s(E)) \end{aligned}$$

It is simple to see that this pregame satisfies the large game property and coordination property.[29] We note that the payoff from strategies B and C are identical. We also note that 'everybody play D' and 'everybody play E' are the Pareto optima giving a payoff of 3. Suppose that $R(i) = N$ for all $i \in N$ and let $\varepsilon = 0$. Consider a game $\Gamma(N, \alpha)$ and let the initial state be 'everybody play A'. Strategies B and C represent innovation opportunities. Suppose that play evolves to a state s whereby half of the population are playing B and half C. State s is a Nash equilibrium and thus an absorbing state of the innovation dynamic. Players would receive a payoff of 2. State s is not, however, an imitation equilibrium. Given an imitation with innovation dynamic play would converge to a Pareto optima whereby each play earned a payoff of 3.♦

In Example A we see that innovation may fail to converge to the Pareto optima because some players choose strategy B and some strategy C. Initially this divergence appears to come at no loss to payoffs. Ultimately, however, to realize the gains from playing strategy D or E it is necessary that a significant majority of the population is playing either B or C. With imitation such problems do not arise because the divergence in strategy choice between B and C would not persist.

Example B: The strategy space is given by $S = \{A, B\}$ and the attribute space by $\Omega = [0, 2]$. Given a population (N, α) and strategy vector s let $w_s(k)$ be the proportion of the population playing strategy k. Given population (N, α) the payoff function of player $i \in N$ is defined:

$$u_i^\alpha(A, s_{-i}) = w_s(A); \quad u_i^\alpha(B, s_{-i}) = \omega w_s(A) + 2w_s(B)$$

[28] If both Example A and Example B we use the following notational convention: when we write, for example, $u_i^\alpha(B, s_{-i}) = 1 + w_s(B) + w_s(C)$ the strategy vector s is taken to be the strategy vector (B, s_{-i}) or in other words we set $s_i = B$.

[29] Given a suitable rescaling of payoffs.

for all $k \in S$ where $\alpha(i) = \omega$. Thus, if all players play A they each receive a payoff of 1 and if they all play B they each receive a payoff of 2. If a player plays strategy B when all others are playing strategy A then his payoff will depend on ω. It is easily checked that this pregame satisfies the large game property, continuity in attributes and the coordination property.[30] Consider population (N, α) where $\alpha(1) = 2$ and $\alpha(i) = 0$ for all other $i \in N$. Suppose that the reference network has the property that $R(i) = N$ for all $i \in N$ and let $\varepsilon = 0$. Given an initial state 'everybody play A', player 1 is the unique player with an innovation opportunity. Thus, play will evolve to strategy vector s where player 1 plays B and everybody else A; player 1 receives a payoff of 2 and all others a payoff of approximately 1. Strategy vector s is a Nash equilibrium and thus an absorbing state of the innovation dynamic. It is clearly, however, not an imitation equilibrium. The imitation with innovation dynamic would converge to the Pareto optimal state 'everybody play B'.♦

In Example B the 'group action' that is created by imitation enables gains that would not be achievable by unilateral action. In particular given the Nash equilibrium strategy vector s it is not in any players interests to change to strategy B. It is, however, in everyone's interest that all players should change to strategy B. Imitation enables this 'group shift' from strategy A to strategy B.

7 Conclusion

This paper has consider a dynamic model of agent learning through imitation with innovation. Sufficient conditions were provided for play to converge to an approximate Nash, imitation equilibrium. The principal focus was on games with many players reflecting our belief that imitation is most likely to be observed in such games. Our results suggest that imitation can be consistent with individually rational behavior. Through example we demonstrate that imitation may even enable learning of 'more efficient' strategy vectors (than innovation alone).

Two potential applications of our results appear to be in modelling technological or scientific evolution or in modelling market interaction. In terms of technological and scientific evolution the notion of learning through imitation and innovation is a natural one (see, for example Kuhn 1996 and Ziman 2000). Models of learning in 'Cournot like' market interaction games have

[30] Given a suitable metric on Ω and rescaling on payoffs.

been the subject of a number of related papers (e.g. Vega-Redondo 1997, Alos-Ferrer, Ania and Schenk-Hoppe 2000, Selten and Ostmann 2000, Selten and Apesteguia 2002). To apply the imitation with innovation dynamic in studying such learning processes remains a goal for future research.

Other avenues for future research include the possibility that imitation can increase the rate of convergence to an absorbing state - something this paper did not address given our focus on long run dynamics. Also of interest, given its importance to the learning dynamic, would be to study the possible evolution of the reference network. A related literature concerns network formation (Dutta, B. and M. Jackson 2003). This literature treats the network as the game in the sense that a players payoff is directly dependent upon the links that he has in the network. In the model of this paper the network is merely a medium through which the game is played and so the effect of the network on a player's payoffs is indirect. It may be interesting to apply the ideas from the network formation literature in modelling the evolution of an endogenised interaction network.

As a final remark we note that any interpretation of our results must take into account the realism of our model of learning. One way to test this is through experimental work. There has been some experimental work on imitation and the importance of social learning (e.g. Offerman, Potters and Sonnemans 2002 and Selten and Apesteguia 2002). There has also been experimental work on learning in 'large games' (e.g. Van Huyck 1997, Rapoport, Seale and Winter 2001). A particular interesting question that arises from this paper is whether players do indeed imitate and only imitate players that are similar to themselves.

References

[1] Alos-Ferrer, C., A.B. Ania and K.R. Schenk-Hoppe, 2000, An evolutionary model of Bertrand oligopoly, *Games and Economic Behaviour*, 33: 1-19.

[2] Asch, S., 1952 *Social Psychology*, Prentice Hall, New York.

[3] Bernheim, B.D., 1994, A Theory of conformity, *Journal of Political Economy*, 102: 841-877.

[4] Blume, L., 1995, The Statistical Mechanics of Best-Response Strategy Revision, *Games and Economic Behavior*, 11: 111-145.

[5] Cartwright, E., 2003, Learning to play approximate Nash equilibria in games with many players, University of Warwick Working Paper no. 671.

[6] Cartwright, E. and M. H. Wooders, 2003, Social conformity and bounded rationality in games with many players, University of Warwick Working Paper no. 672.

[7] Clark, A. and A. Oswald, 1996, Satisfaction and Comparison Income, Journal of Public Economics 61: 359-381.

[8] Deutsch, M. and H.B. Gerard, 1955, A study of normative and informational social influences upon individual judgement, *Journal of Abnormal and Social Psychology* 51: 629-636.

[9] Dutta, B. and M. Jackson (Eds), 2003, *Networks and Groups: Models of Strategic Formation*, Springer-Verlag, Heidelberg.

[10] Ellison, G. and D. Fudenberg, 1993, Rules of thumb for social learning, *Journal of Political Economy* 101: 612-643.

[11] Ellison, G. and D. Fudenberg, 1995, Word-of-mouth communication and social learning, *Quarterly Journal of Economics* 110: 93-125.

[12] Fudenberg, D. and D.K. Levine, 1998, *The Theory of Learning in Games*, MIT Press.

[13] Gale, D. and R. Rosenthal, 1999, Experimentation, imitation and strategic stability, *Journal of Economic Theory*, 84: 1-40.

[14] Gigerenzer, G., P.M. Todd and the ABC Research Group, 1999, *Simple Heuristics That Make Us Smart*, Oxford University Press.

[15] Granovetter, M., 1973, The strength of weak ties, *American Journal of Sociology*, 83: 1420-1443.

[16] Gross, R., 1996, *Psychology: The Science of Mind and Behaviour*, Hoddon and Stoughton.

[17] Hart, S., 2003, 'Uncoupled dynamics do not lead to Nash equilibrium, working paper.

[18] Jones, S., 1984, *The Economics of Conformism*, Oxford: Blackwell.

[19] Kandori, M., G.J. Mailath and R. Rob, 1993, Learning, mutation and long run equilibria in games, *Econometrica* 61: 29-56.

[20] Kirman, A., 1993, Ants, rationality, and recruitment, *Quarterly Journal of Economics*, 137-156.

[21] Kuhn, T., 1996, *The Structure of Scientific Revolutions, Third Edition*, The University of Chicago Press.

[22] Levine, D. and W.Pesendorfer, 2000, Evolution through imitation in a single population, Working Paper.

[23] Levine, D. and W.Pesendorfer, 2001, The evolution of cooperation through imitation, Working Paper.

[24] Morris, S., 2000, Contagion, *Review of Economic Studies*, 67: 57-78.

[25] Offerman, T., J. Potters and J. Sonnemans, 2002, Imitation and belief learning in an oligopoly experiment, *Review of Economic Studies*, 69: 973-997.

[26] Rapoport, A, D. Seale and E. Winter, 2001, Coordination and learning behavior in large groups with asymmetric players, *Games and Economic Behavior*, 1-26.

[27] Ritzberger, K. and J.W. Weibull, 1995, Evolutionary selection in normal form games, *Econometrica*, 63: 1371-1399.

[28] Schlag, K.H., 1998, Why imitate, and if so, how? A bounded rational approach to multi-armed bandits. *Journal of Economic Theory*, 78: 191-209.

[29] Schlag, K.H., 1999, Which one should I imitate? *Journal of Mathematical Economics*, 31: 493-522.

[30] Selten, R., 1998, Features of experimentally observed bounded rationality, *European Economic Review*, 42: 413-436.

[31] Selten, R. and J. Apesteguia, 2002, Experimentally observed imitation and cooperation in price competition on the circle, University of Bonn Economic Discussion Paper 35-2002.

[32] Selten, R. and A. Ostmann, 2000, Imitation Equilibrium, University of Bonn Economic Discussion Paper 16-2000.

[33] Vega-Redondo, F., 1997, The evolution of Walrasian behavior, *Econometrica*, 65: 375-384.

[34] Watts, D., 1999, *Small Worlds*, Princeton University Press.

[35] Wooders, M., E.J. Cartwright and R. Selten, 2003, Social Conformity in Games with Many Players, University of Warwick working paper no. 589.

[36] Young, H.P., 1993, The Evolution of Conventions, *Econometrica*, 61: 57-84.

[37] Young, H.P., 2001, *Individual Strategy and Social Structure*, Princeton University Press.

[38] Ziman, J., 2000, *Technological Innovation as an Evolutionary Process*, Cambridge University Press.