LEARNING TO PLAY APPROXIMATE NASH EQUILIBRIA

IN GAMES WITH MANY PLAYERS

Edward Cartwright

No 671

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF WARWICK

Learning to play approximate Nash equilibria in games with many players

Edward Cartwright
Department of Economics,
University of Warwick,
Coventry.
CV4 7AL

E.J.Cartwright@warwick.ac.uk.
www.warwick.ac.uk/~ecrhs.

March 2003

Keywords: imitation, best reply, convergence, Nash equilibrium

Abstract

We illustrate one way in which a population of boundedly rational individuals can learn to play an approximate Nash equilibrium. Players are assumed to make strategy choices using a combination of imitation and innovation. We begin by looking at an imitation dynamic and provide conditions under which play evolves to an imitation equilibrium; convergence is conditional on the network of social interaction. We then illustrate, through example, how imitation and innovation can complement each other; in particular, we demonstrate how imitation can 'help' a population to learn to play a Nash equilibrium where more rational methods do not. This leads to our main result in which we provide a general class of large game for which the imitation with innovation dynamic almost surely converges to an approximate Nash, imitation equilibrium.

1 Introduction

Dynamic models of learning in games can provide insights on when and how a population of boundedly rational players can learn to play a Nash equilibrium. The limits to individual rationality and the importance of the Nash equilibrium concept in economics and game theory make an understanding of such issues fundamental. In this paper we study learning in games with many players. The complexity of these games, as suggested by the large number of players, makes anything approaching rational behavior seem unlikely. We provide, however, sufficient conditions on behavior to ensure that play will converge to an approximate Nash equilibrium for a general class of large game.

We model a learning dynamic in which players are assumed to imitate and innovate. More precisely, each player uses interchangeably two decision making heuristics - an imitation heuristic and an innovation heuristic. Before detailing these heuristics and our results, we briefly outline our principle motivations for assuming such behavior. These are twofold; first, a belief that these two heuristics capture key aspects of individual behavior in large games, and second, a belief that learning through the combination of imitation with innovation is likely to lead to the emergence of Nash equilibrium play. We expand on these motivations in turn.

It is widely accepted that individual behavior is partly motivated by 'social influences', such as desires for popularity or acceptance, and that such behavior can lead to imitation (see, for example, Jones 1984 and Bernheim 1994). An individual may also be motivated to learn through imitation when he has imperfect information about his payoff function or his strategy set. When faced with such incomplete information, imitation is a means through which a player can draw on and learn from the collective experience of others (Young 2001b). Note that a player's lack of information may or may not reflect bounds on his rationality. 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).

An obvious limitation of imitation is that it leaves little room for novelty or originality. This suggests that imitation is not and cannot be the sole constituent of learning. Novelty could be seen to arise from experimentation or mistakes but individual behavior appears more purposeful than this (even in complex games). For example, Selten and Apesteguia (2002) and Offerman et. al. (forthcoming), in running experiments with 'Cournot interaction type games', find evidence of both imitation and of attempts to initiate cooperation or collusion; purposeful attempts by some subjects to increase individual payoffs appeared to be apparent (even though subjects were not aware of the payoff structure of the game). The innovation heuristic is motivated to capture such unilateral behavior whereby a player attempts to increase his payoff.

Not only do we believe an imitation with innovation dynamic can capture key aspects of individual behavior, we also feel it is likely to lead to the emergence of Nash equilibrium play. Intuitively, different heuristics can be associated with different advantages and disadvantages. Imitation, for example, appears to be a dynamic in which the actions of individuals will become coordinated in the sense that one strategy profile emerges as a convention or focal point. The lack of innovation, however, implies that such a strategy profile need not be individually rational. Vega-Redondo (1997) and Selten and Ostmann (2000), for example, model variants on an imitation dynamic and demonstrate that play may converge to a strategy profile that is not a Nash equilibrium. By contrast, any stable state of an innovation dynamic should be individually rational. Given, however, that each individual acts in isolation there is less opportunity for the actions of individuals to become congruent. In particular, a player may neither directly or indirectly predict the behavior that can be expected of others. Illustrations of how adaptive play (similar to an innovation dynamic) need not converge to a Nash equilibrium are provided by, amongst others, Young (1993, 2001a). Suppose that a player uses more than one heuristic. The advantages of one heuristic could potentially compensate for the disadvantages of another. Imitation and innovation appear to be two types of behavior that are particularly suited to complement each other. Gale and Rosenthal (2001) provide some evidence for this in demonstrating how Nash equilibrium outcomes can arise from imitation and innovation. The results of this paper provide further evidence in a relatively more general context.

As stated above, the dynamic we model assumes that players use interchangeably and imitation and an innovation heuristic. The imitation heuristic is based in part on a model of imitation used by Selten and Ostmann (2000). A player imitates by referring to a subset of the population - his reference group - and by copying the action of the most successful player. The sophistication in this behavior comes from referring to a specific subset of the population (which may have been carefully selected) and in only imitating the most successful players referred to. These two properties of the imitation heuristic distinguish our approach from much of the previous literature on imitation. For example, many authors (e.g. Kandori, Mailath and Rob 1993, Ellison and Fudenberg 1993, Vega-Redondo 1997 and Alos-Ferrer, Ania and Schenk-Hoppe 2000) model a dynamic in which each player can be seen to refer to the total player set. An alternative (e.g. Kirman 1993 and Ellison and Fudenberg 1995) is to assume players refer to a random sample of the population; under such an assumption a player will only refer to a subset of the population in any one period but, over time, may refer to everyone within the population. We also note that many authors (e.g. Kirman 1993, Levine and Pesendorfer 2000, 2001 and Gale and Rosenthal 2001) model a dynamic in which players do not necessarily 'imitate the successful' but instead, 'conform' to the actions of others in the sense that a player chooses the strategy he observes being played most often. The literature on imitation is considered in more detail in section 2.2.

In using the innovation heuristic a player chooses an action that will, ceteris paribus, increase his payoff. This suggests that a player acts on the basis that other players will not subsequently change strategy. In games with many players attempts to 'second guess' the behavior of opponents may be difficult if not impossible. Also, much experimental evidence supports the notion that individuals act on the basis of recent past experience (Selten 1998). Thus, it seems reasonable that a player should act on the assumption that the actions of other players will not change. We highlight that innovation is similar to but not the same as a best response or myopia dynamic, as commonly defined and much studied (see Fudenberg and Levine 1998). A player behaving myopically chooses a strategy that, ceteris paribus, maximizes their payoff. Thus, in behaving myopically a player chooses the 'best' strategy; this may differ from someone innovating who merely has to choose a 'better' strategy. Innovation requires less rationality on the part of players than myopia.

Our analysis of the imitation with innovation learning dynamic begins be assuming that players only imitate. This leads to the definition of an imitation equilibrium - a state stable under the imitation dynamic. An imitation equilibrium has the property that players who refer to each other typically play the same strategy. It need not be a Nash equilibrium. Note, however, that if players have a desire for equality, or what they may perceive as fairness, an imitation equilibrium may be an intuitively appealing concept of equilibrium. Individuals do appear to be influenced by 'fairness' considerations. For example, wages may be judged in relation to the wages of others (Clark and Oswald 1996). Also, fairness appears to influence bargaining in experimental studies (see Chapter 4 of Kagel and Roth 1995).

Our first main result provides sufficient conditions under which an imitation dynamic almost surely converges to an imitation equilibrium. We recall that players may imitate those in their reference group. A reference network details the reference group of every player. Theorem 1 states that if the reference network has a clustering coefficient of one then play will evolve, almost surely, to an imitation equilibrium. A reference network has a clustering coefficient of one if whenever a player i refers to players j and k, both players j and k refer to each other. Many social and economic networks have clustering coefficients near one (Granovetter 1973 and D. Watts 1999). Note that Theorem 1 requires no assumption on the game being played.

Having looked at an imitation dynamic in some detail we turn our attention to the imitation with innovation dynamic. A stable state of such a dynamic is an approximate Nash, imitation equilibrium. We begin with three examples that demonstrate how learning through imitation and learning through innovation may, or may not, complement each other. Example 5, for instance, provides a game and reference network where (1) an imitation dynamic need not converge to an imitation equilibrium, (2) an innovation dynamic need not converge to a Nash equilibrium, yet (3) an imitation with innovation dynamic will converge, almost surely, to a Nash, imitation equilibrium.

For our main result we use the concept of a pregame satisfying the large game property as introduced by Wooders, Cartwright and Selten (2001). A principle component of a pregame is a set of player attributes. In games induced from a pregame satisfying a large game property the payoff of a player is essentially a function of the proportions of players with each attribute playing each strategy (and his own strategy). Our Theorem 2 states that, subject to relatively mild assumptions, in any sufficiently large game induced from a pregame satisfying the large game property the imitation with innovation dynamic converges, almost surely, on an approximate Nash, imitation equilibrium. We note how players learn not only to play an approximate Nash equilibrium but also an imitation equilibrium. Indeed, players use pure strategies throughout and so play converges to an approximate Nash, imitation equilibrium in pure strategies.

Our main result demonstrates how approximate Nash equilibrium play can emerge in large games if players learn through imitation and innovation. Similar results were obtained by Gale and Rosenthal (1999) in the context of interaction in a Cournot like model. An appealing aspect of our results are the generality of game modelled. 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). This is not the case in the game we model.

This is highlighted by the fact that through a corollary of Theorem 2 we are able to contribute to the literature on the existence of pure strategy Nash equilibrium in large games (e.g. Schmeidler 1973, Mas-Colell 1984 and Wooders et. al. 2001). In particular, the fact that play converges to an approximate Nash, imitation equilibrium demonstrates that one must exist; this complements existence results due to Wooders et. al. (2001).

A second aspect of our main result is the suggestion that imitation can be consistent with individually rational play in games with many players. This complements results due to Wooders et. al. (2001) who demonstrate that, in large games, there exists an approximate Nash equilibrium in which 'similar players play similar strategies'. Note, that the question of whether players learn to play this equilibrium is not addressed by Wooders et al.; for a slightly less general class of game, Theorem 2 demonstrates that this equilibrium will indeed emerge. Related results on the individual rationality of imitation are due to Schlag (1998, 1999) and Ellison and Fudenberg (1993, 1995). In varying contexts these authors show how imitative learning can lead to the adoption of 'optimal actions'.

There are many further relationships between this paper and the literature on learning in games. We highlight two. First, there is a large literature, not mentioned above, on the convergence of learning dynamics to Nash equilibrium play. Much of this literature considers learning dynamics very different from ours such as fictitious play or the replicator dynamic (see Fudenberg and Levine 1998). Often the differing choice of dynamic reflects the type of game to be studied (see, for example Kalai and Lehrer 1993). The literature that has used learning dynamics more comparable to ours has principally addressed the issue of equilibrium selection (e.g. Young 1993, Robson and Vega-Redondo 1996 and Levine and Pesendorfer 2000, 2001). More precisely, learning has been modelled in games where the convergence of play to a Nash equilibrium appears trivial, the question of interest has been which type of equilibrium is more likely to emerge. We have relatively little to say on the issue of equilibrium selection other than suggesting that an imitation equilibrium may be more likely to emerge.

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 concludes. Two appendices present generalizations of our main results.

2 The model

Let $N = \{1, ..., n\}$ denote a finite *player set* and let $S = \{s^1, ..., s^K\}$ denote a finite *strategy set*. A *strategy vector* is given by $\sigma = (\sigma_1, ..., \sigma_n) \in S^n$ where σ_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 $\{\sigma(t)\}_{t\geq 0}$ on state space Σ. The transition matrix of the Markov chain will be denoted by P. The value $P_{\sigma\sigma'}$ is interpreted as the probability of state σ' immediately following state σ.

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. We introduce in turn the imitation and innovation heuristics before formally defining the imitation with innovation dynamic. First, however, we define a reference network; the imitation heuristic makes use of such a network.

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. We will also refer to a reference matrix R 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.[1]

[1] Given the reference matrix R the reference group R_i of player i could be thought of as the ith row of R.

We will assume that the reference network remains constant throughout the evolution of play. It will become clear, as we proceed, that the reference network can be crucial in determining how play evolves. This suggests that a player may wish to change his reference group as he learns more about the game and his fellow players. In an Appendix we model this possibility by assuming that players use a *good advice heuristic* to choose a reference group (as well as a strategy) in each period. We are able to show that the main conclusions of the paper are unaffected by this freedom in reference group choice.

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* $p_i : \Sigma \rightarrow \Delta(S)$ where the value $p_i(s^k|\sigma)$ is interpreted as the probability that a player i, using the imitation heuristic, would select the strategy s^k if strategy vector σ was played in the previous period. When using the imitation heuristic a player can be seen to progress through three stages. These are outlined below for a player i choosing a strategy conditional on strategy vector σ. A reference network R is assumed.

1. *Identify costrategists:* the set of *costrategists* of player i, denoted $C_i(\sigma)$, are those players $l \in R_i$ such that $\sigma_l = \sigma_i$.

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

3. *choose strategy:* player i chooses strategy $s^k \in S$ with probability $p_i(s^k|\sigma)$ where (a) if there is a success example j of player i such that $\sigma_j = s^k$ then $p_i(s^k|\sigma) > 0$, and (b) if every success example of player i is a costrategist of player i then $p_i(\sigma_i|\sigma) = 1$.

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 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.

We highlight that the imitation heuristic is fairly vague about a player's behavior. In particular, if player i has the option of changing strategy (because she has a success example who is not a costrategist) then the possibility is left open for her to potentially choose any strategy. This means she may, for example, experiment, make mistakes or choose the same strategy as in a previous period. Many authors (e.g. Young 1993 and Vega-Redondo 1997) assume that players either choose strategies sequentially, i.e. one person per period, or have some positive probability of not changing strategy. Our results apply to these types of dynamic. We note, however, that a player using the imitation heuristic may *always* imitate success examples. Thus the possibility of mistakes or experimentation etc. is not required for our results.

The 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. Consider an *imitation heuristic with inertia*. 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 $\sigma_j \neq \sigma_i$ if and only if

$$u_j(\sigma) = \max_{l \in R_i} u_l(\sigma) > \max_{k \in C_i(\sigma)} u_k(\sigma).$$

In this case player i may only change strategy through imitation if there is a success example earning a *strictly higher* payoff than any of her own costrategists. This creates inertia in that a player is less likely to change strategy. In the main body of the paper we assume throughout that the imitation heuristic is used by players (as opposed to the imitation heuristic with inertia). This has the advantage of simplifying the analysis. In an appendix we consider in more detail possible differences if players use the imitation heuristic with inertia. We demonstrate, through example, that the type of heuristic used can significantly alter the evolution of play. Despite this, however, we show how analogs to our two main theorems can still be derived.

The imitation heuristic can be compared to similar behavioral rules in the literature. Imitation heuristics can differ primarily in two aspects - first, who a player refers to, and second, how a player interprets the information he receives. We discuss each of these aspects in turn. Before doing so we highlight that the heuristic used by Selten and Ostmann (2000) is equivalent to the imitation heuristic with inertia, while 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 $i \in N$.[2] We note that these authors assume that players use varying forms of experimentation in supplement to imitation. This contrasts with the approach of this paper where players use innovation.

Most of the literature assumes that players refer to the entire player set, that is $R_i = N$ for all $i \in N$ (for example Kandori et. al. 1993, Vega-Redondo 1997, Gale and Rosenthal 1999, Levine and Pesendorfer 2000, 2001 and Alos-Ferrer et. al. 2000). Ellison and Fudenberg (1993) consider a model in which players refer to those 'close to them' in terms of some spatial distribution; we will use a similar notion in Section 5. Another alternative, as used by Kirman (1993), Ellison and Fudenberg (1995) and Schlag (1997, 1999) is that a player refers to a random sample of the population. In this way a player only refers to a subset of the population in any one period but can potentially refer to the entire player set. This random sampling is not permitted according to the imitation heuristic. In Section 7, however, we allow players to change their reference group thus permitting random sampling.

There are various ways that a player can interpret the information he receives. As with the imitation heuristic modelled in this paper, Vega-Redondo (1997) and Alos-Ferrer et. al. (2000), amongst others, model a heuristic in which a player can be said to imitate the most successful *player* that he observes. Ellison and Fudenberg (1995) consider a heuristic in which a player could be said to imitate the most successful *strategy* that he observes in the sense that a player chooses the strategy that he observed as giving the highest average payoff.[3] By contrast, the imitation heuristics modelled by Kirman (1993), Gale and Rosenthal (1999) and Levine and Pesendorfer (2000, 2001) assume that players conform to the 'average strategy of the population'; thus, players does not imitate strategies according to their success but according to their popularity. Ellison and Fudenberg (1993) consider a heuristic in which players imitate strategies only if they are both successful and popular. Other possibilities and a discussion of this issue is

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

[3] Suppose player i refers to three players - himself and players k and j. Further, suppose players i and k play strategy A and get payoffs of 0 and 100 respectively while player j plays strategy B and gets payoff 90. If player i imitates the most successful player he will imitate player k. If he imitates the most successful strategy he will play strategy B.

provided by Schlag (1997, 1999).

We make one final comment. Ellison and Fudenberg (1995), Robson and Vega-Redondo (1996) and Schlag (1997, 1999) model games of imperfect information. Players are assumed to imitate on the basis of *observed* or *realized* payoffs. Our framework permits games of imperfect information. We implicitly assume, however, that players imitate on the basis of *expected* payoffs and not *realized* payoffs (see Robson and Vega-Redondo 1996 for a discussion of this issue).

2.3 Innovation heuristic

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

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

$$u_i(s^k, \sigma_{-i}) > u_i(\sigma) + \varepsilon.$$

2. *choose strategies:* player i chooses strategy $s^k \in S$ with probability $m_i(s^k|\sigma)$ where (a) if there are no innovation opportunities for player i then $m_i(\sigma_i|\sigma) = 1$, and, (b) if there is an innovation opportunity for player i then $m_i(s^k|\sigma) > 0$ for some strategy s^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. It is important to note that $m_i(s^k|\sigma)$ can be zero even if s^k is an innovation opportunity. For example, a player need not, necessarily, choose the innovation opportunity that would have maximized her payoff in the previous period. This contrasts with the imitation heuristic where it is assumed that every success example is imitated with some positive probability. We note that the possibility for mistakes, experimentation and inertia exist in the innovation heuristic to the same extent as they did in the imitation heuristic.

The innovation heuristic is similar to best response or myopic behavior as modelled by many authors (see Fudenberg and Levine 1998). There are, however, important differences. First, ε is commonly assumed to be zero. Second, when using myopia a player always chooses a strategy that would have maximized her payoff in the previous period. As we have noted, when using an innovation heuristic the probability that she play such a strategy may be zero. This would suggest that the innovation heuristic requires less computation to perform. This suggests, in turn, that a 'less rational' player is capable of innovating.

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 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 $\lambda_i = 1$ player i always uses the innovation heuristic to select a strategy while 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 similarly $\lambda = 1$ if $\lambda_i = 1$ for all $i \in N$. We say that $\lambda \neq 0, 1$ if $\lambda_i \in (0, 1)$ for all $i \in N$.[4]

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. The resulting stochastic process is referred to as the *imitation with innovation dynamic* which we indicate as $\mathcal{I}(p; m; \lambda)$. 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.[5]

We highlight that the imitation with innovation dynamic does not have persistent randomness. That is, there are stable states of the dynamic (as

[4] 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.

[5] 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.

will be demonstrated in Sections 3 and 4). The approach we use contrasts with much of the existing literature. Typically, there is assumed to be some positive probability that a player experiments by randomly selecting an arbitrary strategy. This persistent randomness implies the system can never be absorbed into a stable state. Dynamics for which there is not persistent randomness are studied by Ellison and Fudenberg (1993, 1995) and Blume (1993, 1995). Blume (1995) discusses this issue in more detail.

3 The dynamics of imitation

We begin our analysis of the imitation with innovation dynamic by assuming that $\lambda = 0$. That is, by assuming that players only ever use the imitation heuristic to select a strategy. We define a static equilibrium concept.[6]

Imitation Equilbirum: The strategy vector σ is an *imitation equilibrium* of stage game Γ relative to reference network R if

$$\max_{l \in R_i / C_i(\sigma)} u_l(\sigma) < \max_{l \in C_i(\sigma)} u_l(\sigma)$$

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

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, as such, no player will wish to change strategy. This immediately suggests Lemma 1, which we state without proof. 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.

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

This result demonstrates that the Markov process described by the imitation with innovation dynamic when $\lambda = 0$ is not irreducible. That is, there are many absorbing states. This follows from the observation that any strategy vector σ in which every player $i \in N$ plays the same strategy is an

[6] 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.

imitation equilibrium. If all communication classes of the dynamic are singletons then Lemma 1 implies that the imitation with innovation dynamic will converge, almost surely, to an imitation equilibrium. In general, however, there may exist non-singleton communication classes. That is, there may exist a communication class Ψ where $|\Psi| > 1$ and where $\sum_{q\in\Psi} p_{\sigma q} = 1$ for all $\sigma \in \Psi$. An 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 $R_3 = \{2, 3\}$. Thus, player 2, for example, refers to players $1, 2$ and 3. Two strategy vectors are of interest.

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

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 players earning a payoff of 4.⧫

The cycle of play that we observe in Example 1 appears to reflect the reference network. One important characteristic of a network is its clustering coefficient. This is a measure of the cliquishness of the network.[7]

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 $j \in R_k$.[8]
2. $|R_i| \geq 3$ for every player $i \in N$.[9]

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

[7] See D. Watts (1999) and references there in for a definition and discussion.

[8] 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.

[9] 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 $i \in R_i$.

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. We state our first main result.

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

Proof: Given an arbitrary state σ we demonstrate that there exists states, indexed, $\sigma(2), ..., \sigma(T)$ where $P_{\sigma\sigma(2)} > 0$, $P_{\sigma(t)\sigma(t+1)} > 0$ for all $T-1 \geq t \geq 2$ and where $\sigma(T)$ 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 $k \in N$ such that $k \in R_i$, given that the reference network R has a clustering coefficient of one, it must be the case that $k \in R_j$ and $j \in R_k$. This, in turn, implies that $i \in R_k$. Similarly, if there exists a player $l \in R_j$ then $l \in R_i$ and $i, j \in R_l$. 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 $\sigma(2)$ there must exist some $s^k \in S$ such that $\sigma_j = s^k$ 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 $\sigma(2), \sigma(3), ...$ who is not a costrategist. Thus, no player i belonging to a clique can change strategy between states $\sigma(2), \sigma(3), ...$

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 $k \in N$ such that $i \in R_k$. Given that the network R has a clustering coefficient of one there must exist a player $j \neq i$ such that $j \in R_k$. Further, if $i, j \in R_k$ 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, k who must refer to each other. Given that player i refers to a subset of a clique in states $\sigma(2), \sigma(3), ...$ 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 $\sigma(t_i)$ 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 $\sigma(t)$, $t \geq t_i$, player i does not have a success example who is not a costrategist. This completes the proof.■

Given that a reference network which has a clustering coefficient of one is sufficient to guarantee convergence on an imitation equilibrium we may ask whether or not it is necessary. 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. We cannot go any further this, however, as the following example demonstrates.

Example 2: There are 3 players and the reference network is such that $R_1 = \{1, 2, 3\}$, $R_2 = \{2, 3\}$ and $R_3 = \{3\}$. The network R does not have a clustering coefficient of one. For any game Γ, however, the imitation with innovation dynamic $\mathcal{I}(\varepsilon; \lambda = 0; R)$ almost surely converges to an imitation equilibrium. To demonstrate, we proceed by contradiction. We note that player 3 cannot change strategy, so suppose player 3 is playing some strategy A. Player 2 can either be playing strategy A or not. If at any point player 2 imitates then he will play strategy A for all subsequent periods. Finally, we consider player 1. If play is not to converge on an imitation equilibrium then player 1 must repeatedly change strategy through imitation. Note, however, that if player 1 imitates player 3 then so can player 2. Thus, both players 2 and 3 will almost surely end up playing strategy A. This leads to a contradiction.

This example could be objected to on the grounds that player 3 only refers to himself. The example can, however, easily be amended, with the same conclusions, to one in which every player refers to at least three other players.♦

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

In the previous section we looked in some detail at the long run convergence properties of the imitation with innovation dynamic on the assumption that players solely use imitation. We have provided conditions for which the dynamic converges on an imitation equilibrium. 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. This provides ample motivation for a player to use an innovation heuristic. We now turn to consider what happens when players use such a heuristic. Let us begin with two definitions,

Nash ε-Equilibrium: The strategy vector σ is a *Nash ε-equilibrium* of stage game Γ if

$$u_i(s^k, \sigma_{-i}) \leq u_i(\sigma) + \varepsilon$$

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

Nash, Imitation ε-Equilibrium: The strategy vector σ a *Nash, Imitation ε-Equilibrium* of stage game Γ relative to reference network R if σ is both a Nash ε-equilibrium and an imitation equilibrium relative to R.

We refer to a Nash, imitation 0-equilibrium as a Nash, imitation equilibrium and a Nash 0-equilibrium as a Nash equilibrium. These definitions should need no explanation and lead to the following result which we state without proof,

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

There is an extensive literature on the convergence, and non-convergence, of best response dynamics (See Fudenberg and Levine 1998 and references therein). Thus, given the similarities between best response and innovation, we do not look specifically at the at the case where $\lambda = 1$. It is, however, interesting to look at the interaction between innovation and imitation. We illustrate with three examples. In each example we evaluate whether or not the imitation with innovation dynamic converges on an absorbing state for the three possibilities of $\lambda = 1$ (innovation), $\lambda = 0$ (imitation) and $\lambda \neq 0, 1$ (imitation with innovation). The results of these examples can be summarized by the following table,[10]

Example	innovation	Imitation	innovation and imitation
3	converges	converges	need not converge
4	need not converge	converges	converges
5	need not converge	need not converge	converges

Before discussing any conclusions let us set out the examples where we assume throughout that $\varepsilon = 0$.

Example 3: There are two players and three strategies A, B and C. Both players refer to each other. The payoff matrix is as follows where player 1 chooses a row and player 2 a column,[11]

	A	B	C
A	$3,1$	$3,2$	$0,0$
B	$0,0$	$0,0$	$0,0$
C	$0,0$	$0,0$	$10,10$

[10] Examples can easily be derived to illustrate the other five possible combinations of convergence in the three dynamics.

[11] The first entry in the payoff matrix is that of the row player and the second that of the column player.

Strategy vector (C, C) is the unique Nash, imitation equilibrium. Suppose, however, that the current state is (A, A). This is not a Nash equilibrium as player 2 may innovate and choose strategy B. Strategy vector (A, B) is not an imitation equilibrium as player 2 may imitate and choose strategy A. Thus, if $\lambda \neq 0, 1$ the imitation with innovation dynamic need not converge on an absorbing state. It is easily checked, however, that if $\lambda = 0$ or if $\lambda = 1$ the imitation with innovation dynamic does converge on an absorbing state.♦

Example 4: There are two players and four strategies A, B, C and D. Both players refer to each other. The game can be represented by the following payoff matrix.

	A	B	C	D
A	$1,0$	$0,0$	$0,1$	$0,0$
B	$0,1$	$0,0$	$1,0$	$0,0$
C	$0,0$	$10,10$	$0,0$	$0,0$
D	$0,0$	$20,0$	$0,0$	$100,100$

There exists a unique Nash, imitation equilibrium (D, D). If $\lambda = 1$ the imitation with innovation dynamic need not converge on a Nash equilibria. Suppose for example the current state is (A, A). Play may evolve through the cycle of states $(A, C) \rightarrow (B, C) \rightarrow (B, A) \rightarrow (A, A)$. By contrast, the imitation with innovation dynamic will clearly converge on an imitation equilibrium if $\lambda = 0$. Similarly, the imitation with innovation dynamic converges almost surely to a Nash, imitation equilibria if $\lambda \neq 0, 1$. This is apparent after considering what may happen if the current state is (A, C); player 1 may imitate player 2, implying play evolves to state (C, C); at this point, player 1 may imitate and player 2 may use innovation in which case play evolves to state (C, B) and ultimately (D, D).♦

Example 5: There are four players and four strategies A, B, C and D. The reference network is such that $R_1 = \{1, 2\}$, $R_2 = \{1, 2, 3, 4\}$, $R_3 = \{1, 2, 3, 4\}$ and $R_4 = \{3, 4\}$. Play revolves around the following matrix game,

	A	B	C	D
A	$4,0$	$0,0$	$3,4$	$0,0$
B	$0,1$	$0,0$	$4,0$	$0,0$
C	$0,0$	$0,0$	$0,0$	$0,0$
D	$0,0$	$20,0$	$0,0$	$100,100$

Players 1 and 2 choose a row and players 3 and 4 choose a column. There are then four plays of the above matrix game as player 1 plays the matrix game against both players 3 and 4 and player 2 plays the matrix game against both players 3 and 4. Thus, if the strategy vector is (A, A, A, C) the payoff vector is $(7, 7, 0, 8)$ while if the strategy vector is (A, C, C, C) the payoff vector is $(6, 0, 4, 4)$.

If $\lambda = 1$ the imitation with innovation dynamic need not converge on a Nash equilibria; as in Example 4, if neither player 1 or 2 is playing strategy C or D and neither player 3 or 4 is playing strategy B or D then play cannot evolve to the unique Nash equilibrium (D, D, D, D). Similarly, if $\lambda = 0$ the imitation with innovation dynamic need also not converge on an absorbing state; there exists a cycle of states $(A, A, A, C) \to (A, C, C, C) \to (A, A, A, C)$.

If $\lambda \neq 0, 1$ then the imitation with innovation dynamic does converge to a Nash, imitation equilibrium. To appreciate this assume an initial state (B, B, C, C). All players may use the imitation heuristic in the subsequent two periods leading to state (B, B, B, C) and then (B, B, B, B). If players 1 and 2 use the innovation heuristic and players 3 and 4 use the imitation heuristic then play may evolve to (D, D, B, B) and ultimately the unique Nash, imitation equilibrium (D, D, D, D).♦

In discussion perhaps the most interesting point to note is how the combination of imitation with innovation can imply convergence on a Nash equilibrium when the use of imitation or innovation in isolation do not imply such convergence. In particular, in both examples 4 and 5 there are Nash equilibria that seem to be appropriate long run outcomes but to which the imitation with innovation dynamic need not converge if players solely use the innovation heuristic.[12] These examples illustrate how imitation may 'help' players to learn to play a Nash equilibria. We discuss this possibility in more detail in the next section and in the conclusion. Another interesting point illustrated, in particular by example 3, is how, even if play converges, when players use innovation, it may not converge to a state that is stable under an imitation dynamic. This is significant if players do have desires for 'fair' outcomes in which they are treated 'equally' with those players to whom they refer.

[12] We note that examples 4 and 5 are fairly robust to changes in the innovation and imitation heuristics. For example, the conclusions are unaltered if there is a positive probability that a player will play the same strategy as in the previous period.

5 Large games and convergence

In this section we look to provide sufficient conditions for the imitation with innovation dynamic to converge on an approximate Nash, imitation equilibrium. In doing so we impose conditions on both the stage game being played and on the reference network. The notion of a pregame satisfying a large game property, as introduced and defined by Wooders, Cartwright and Selten (2001), will be used.

5.1 Pregames

A *pregame* is given by a triple (Ω, S, h) consisting of a compact metric space of player attributes Ω, a finite strategy set S and a function $h : \Omega \times S \times W \to \mathbb{R}$ where W is a set of weight functions. A function w from $\Omega \times S$ into $\mathbb{R}$ is said to be a *weight function* if it satisfies $\sum_{s^k \in S} w(\omega, s^k) \in \mathbb{Z}$ for all $\omega \in \Omega$.

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*. We say that a weight function w_α *corresponds* to population (N, α) when it satisfies

$$\sum_{s^k \in S} w_\alpha(\omega, s^k) = \left|\alpha^{-1}(\omega)\right|$$

for all $\omega \in \Omega$. We let W_α denote *the set of weight functions corresponding to the population* (N, α). Given a population (N, α) and a strategy vector σ we say that weight function $w_{\alpha,\sigma}$ is *relative to strategy vector* σ *and attribute function* α if,

$$w_{\alpha,\sigma}(\omega, s^k) = \sum_{i \in N:\ \alpha(i)=\omega \text{ and } \sigma_i = s^k} 1$$

for all $s^k \in S$ and all $\omega \in \Omega$. Thus, $w_{\alpha,\sigma}(\omega, s^k)$ denotes the number of players of attribute ω (as determined by α) who are playing strategy s^k (as determined by σ).

Given population (N, α) and player $i \in N$, define α_{-i} as the restriction of α to $N \backslash \{i\}$. Let $w_{\alpha_{-i},\sigma}$ be a weight function defined by its components as follows

$$w_{\alpha_{-i},\sigma}(\omega, s^k) = \begin{cases} w_{\alpha,\sigma}(\omega, s^k) - 1 \text{ if } \alpha(i) = \omega \text{ and } \sigma_i = s^k \\ w_{\alpha,\sigma}(\omega, s^k) \text{ otherwise.} \end{cases}$$

for all $\omega \in \Omega$ and for all $s^k \in S$. We will use $W_{\alpha-\omega}$ to denote the set of weight functions corresponding to population $(N \backslash \{i\}, \alpha_{-i})$ where $\omega = \alpha(i)$.

Given a population (N, α), a *game*

$$\Gamma(N, \alpha) = ((N, \alpha), S, \{h_\omega : S \times W_{\alpha-\omega} \longrightarrow \mathbb{R} | \omega \in \alpha(N)\})$$

is induced from the pregame (Ω, S, h) by defining, for each $\omega \in \alpha(N)$,

$$h_\omega(t, w) = h(\omega, t, w)$$

for all $t \in S$ and all $w \in W_{\alpha-\omega}$. In interpretation, $h_{\alpha(i)}(t, w)$ is the payoff received by a player $i \in N$ of attribute $\alpha(i)$ from playing the strategy t when the strategies of other players are summarized by w. Note that players of the same attribute have the same payoff function, inherited from the pregame. A players payoff function is thus indexed by their attribute type - a departure from the notation used in the first half of the paper.

We should perhaps highlight how in this section we have changed from considering one game in isolation to considering a set or family of games. This family of games is determined by the pregame. We focus on pregames that satisfy a large game property.

5.2 Large games

A pregame satisfies the *large game property* if it satisfies both continuity of payoff functions in attributes and global interaction.

Continuity of payoff functions: The pregame $\mathcal{G} = (\Omega, S, h)$ satisfies *continuity of payoff functions in attributes* if for any $\varepsilon > 0$ there exists real numbers $\eta_c(\varepsilon)$ and $\delta_c(\varepsilon) > 0$ such that for any two games $\Gamma(N, \alpha)$ and $\Gamma(N, \overline{\alpha})$ where $|N| > \eta_c(\varepsilon)$, if, for all $i \in N$,

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

then, for any $i \in N$ and for any strategy vector σ,

$$\left| h_{\alpha(i)}(s^k, w_{\alpha_{-i},\sigma}) - h_{\overline{\alpha}(i)}(s^k, w_{\overline{\alpha}_{-i},\sigma}) \right| < \varepsilon$$

for all $s^k \in S$, where $w_{\alpha,\sigma}$ and $w_{\overline{\alpha},\sigma}$ are the weight functions relative to strategy vector σ and, respectively, attribute functions α and $\overline{\alpha}$.

Global interaction: The pregame $\mathcal{G} = (\Omega, S, h)$ satisfies *global interaction* if for any $\varepsilon > 0$ there exists real numbers $\eta_g(\varepsilon)$ and $\delta_g(\varepsilon) > 0$ such that for any game $\Gamma(N,\alpha)$ where $|N| > \eta_g(\varepsilon)$ and for any two weight functions w_α and g_α, both relative to attribute function α, if,

$$\frac{1}{|N|}\sum_{s^k\in S}\sum_{\omega\in\alpha(N)}\left|w_\alpha(\omega,s^k)-g_\alpha(\omega,s^k)\right| < \delta_g(\varepsilon)$$

then,

$$\left|h_{\alpha(i)}(s^k,w_{\alpha_{-i}})-h_{\alpha(i)}(s^k,g_{\alpha_{-i}})\right| < \varepsilon \tag{1}$$

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

We denote by $\mathcal{G}(\eta_c,\delta_c,\eta_g,\delta_g)$ a pregame that satisfies continuity of payoff functions as demonstrated by functions η_c and δ_c and satisfies global interaction as demonstrated by functions η_g and δ_g where η_c,δ_c,η_g and δ_g map $\mathbb{R}_+$ into $\mathbb{R}_+$. A pregame $\mathcal{G}(\eta_c,\delta_c,\eta_g,\delta_g)$ satisfies the large game property.

The notion of a pregame satisfying the large game property is discussed in some detail by Wooders, Cartwright and Selten (2001). Here we provide a brief summary. The definition of continuity of payoff functions in attributes compares two populations in which the attributes of players are slightly perturbed. As such, two different games $\Gamma(N,\alpha)$ and $\Gamma(N,\overline{\alpha})$ are compared. Continuity of payoff functions in attributes requires that a player's payoff function should be approximately the same in both games. A global interaction assumption suggests that a player's payoff is a function primarily of the number of people of each attribute playing each strategy, relative to the total population. As such, a player's payoff is largely dependent on the *proportions* of players of each attribute type playing each strategy (and, of course, on their own strategy choice).

Our interest in large game property is motivated by two considerations. First, an existing result from Wooders et al. (2001) states that if the large game property holds, plus certain other mild assumptions, then for sufficiently large populations there exists an approximate Nash equilibrium σ that partitions the population into a relatively small number of societies; players belonging to the same society play the same strategy and have similar attributes. To see the importance of this result it must first be appreciated that in general a game will not have an approximate Nash, imitation equilibrium. Indeed the existence of an approximate Nash equilibrium is, generally speaking, unlikely.[13] The result due to Wooders et al (2001) suggests that a Nash, imitation equilibrium may exist for large games. A second motivation

[13] Remember that players choose pure strategies and so we are questioning the existence of an Nash ε-equilibrium in pure strategies.

for introducing the large game property is how it appears to capture the type of games for which the modelled behavior appears most appropriate. In particular, in large games both imitation and innovation appear sensible decision making heuristics. A large player set, for instance, makes imitation seem appropriate given the greater potential to learn from the experience of others. Also, a large player set suggests that predicting the actions of others may be difficult and thus innovation (based on a ceteris paribus assumption) appears appropriate.

The imitation with innovation dynamic need not converge to an approximate Nash, imitation equilibrium in large games. We illustrate with the following example.

Example 6: The attribute space is given by $\Omega = \{R, C\}$. There are two strategies A and B. Payoffs are calculated according to the following matrix game M,

	A	B
A	$1,0$	$0,1$
B	$0,1$	$1,0$

In interpretation, a player with attribute R chooses a row in game M and a player with attribute C chooses a column. For any population (N, α) the game $\Gamma(N, \alpha)$ is such that every player of attribute R is matched to play game M against every player of attribute C; a player must play the same strategy (of game M) against all opponents. The payoff of a player equals his total accumulated payoff from playing game M divided by $|N|$, the size of the population. Depending on the level of ε there exists a set of Nash, imitation equilibria in which approximately half of the players of attribute C choose strategy A and in which half of the players with attribute R choose strategy A. This pregame satisfies the large game property.

If players only refer to players of the same attribute then the imitation with innovation dynamic need not converge on an absorbing state for games induced from this pregame (for small ε). Two remarks help illustrate this. First, if $\lambda = 1$ (i.e. just innovation) the imitation with innovation dynamic will not converge on a Nash ε equilibrium unless play commences at one.[14] This is a familiar result. Second, stated informally, in this game the imitation heuristic and innovation heuristic are essentially equivalent. In particular,

[14] Except for a few trivial games that could be induced from this pregame - every player having attribute C, for example.

if the imitation with innovation dynamic does not converge on an absorbing state when $\lambda = 1$ then it will not if $\lambda \neq 0, 1$.♦

5.3 Coordination games and large game reference networks

We provide sufficient conditions on both the game and reference network to guarantee the convergence of the imitation with innovation dynamic on an absorbing state. We begin by defining the concept of a coordination game and below define large game reference networks.

For any two strategy profiles $\sigma, \overline{\sigma}$ let $X(\sigma, \overline{\sigma}) \subset N$ be those players $j \in N$ such that $\sigma_j \neq \overline{\sigma}_j$.

Coordination game: Given a pregame $\mathcal{G}$, the game $\Gamma(N, \alpha)$ is a *coordination game with bound* L when for any two strategy profiles $\sigma, \overline{\sigma}$ if, $|X(\sigma, \overline{\sigma})| \geq L$ and,

$$h_{\alpha(i)}(\overline{\sigma}_i, w_{\alpha_{-i},\overline{\sigma}}) > h_{\alpha(i)}(\sigma_i, w_{\alpha_{-i},\sigma})$$

for all $i \in X(\sigma, \overline{\sigma})$ then,

$$\sum_{i \in N} h_{\alpha(i)}(\overline{\sigma}_i, w_{\alpha_{-i},\overline{\sigma}}) > \sum_{i \in N} h_{\alpha(i)}(\sigma_i, w_{\alpha_{-i},\sigma}).$$

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 $\mathcal{CG}(|N|)$.

It appears relatively mild to assume that a game induced from a pregame satisfying the large game property should be a coordination game. In particular the nature of a large game is that a player's actions will typically influence their own payoff much more than the payoffs of others. Thus, if a player changes strategy to his own benefit it appears relatively mild to assume that the total payoff of the population increases. We note, however, that in a game with many players small individual losses can accumulate to big population wide losses. Reflecting this, a game may not be a coordination game with bound L for small L; examples include n-firm Cournot quantity setting competition and n-player Prisoners Dilemma. The larger is L, however, the more likely it should be that a game is a coordination game with bound L.[15] We note that games induced from the pregame of Example

[15] Note if $\Gamma(N, \alpha) \in \mathcal{CG}(L)$ then $\Gamma(N, \alpha) \in \mathcal{CG}(L^*)$ for any $L^* > L$.

6 are not coordination games with bound L for any $L < N$; in these games the total payoff of the population is fixed independently of the strategies of the players; thus, one players gain is another players loss.

We turn our attention to reference networks. In games induced from a pregame it seems intuitive that a player's reference group should be determined by his attribute and by the attribute function. Given a pregame $\mathcal{G}$ a *reference network function* RN is a function mapping attribute functions to reference networks. In interpretation, $RN(\alpha)$ is the reference network of population (N, α). We define a particular form of reference network after introducing some notation. Given the 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: Given a pregame $\mathcal{G}$ and reference network function RN the reference network $RN(\alpha) \equiv R$ is a *large game reference network with bounds* L, U and δ if

1. R is symmetric[16] 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 three refinements on reference networks studied in Section 3. First, there is an upper and lower bound on the size of a player's reference group as given by U and L. Second, players only refer to those players with 'similar' attributes to themselves where δ measures the similarity. Third, the reference network is symmetric. These three refinements seem relatively mild but the implications are worth exploring a little further.

Symmetry is a common simplifying assumption in modelling social networks (e.g. Jackson and Wolinsky 1996 and D. Watts 1999). It can, however, be a strong assumption; in markets, for example, small firms may refer to big firms but big firms not refer to small firms. The assumption of symmetry can be weakened and the conclusions of Theorem 2 still hold but this comes at the cost of significantly complicating the analysis; a requirement that reference networks be 'predominantly' symmetric is still required.

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

The assumption that a player refers to those with similar attributes to herself is intuitively appealing. If, however, a player has an attribute that is relatively scarce then this implies she must refer to relatively few people. For this to be reasonable would seem to require that a player has a specific preference for referring to players with similar attributes to herself; that is, to be willing to trade referring to relatively few players in order to refer only to those players who are similar to herself. We note how the above remarks demonstrate that the possible values of δ and L are not independent.

5.4 Main result

We have now introduced all the necessary concepts to state our second result. A sketch proof and discussion is provided in Section 5.6.

Theorem 2: Let $\mathcal{G}(\eta_c, \delta_c, \eta_g, \delta_g)$ be any pregame satisfying the large game property and RN any reference network function. Given any $\varepsilon > 0$ and any real number U there exists real numbers $\eta_2(\varepsilon, U)$ and $\delta_2(\varepsilon, U) \geq \delta_c\left(\frac{\varepsilon}{3}\right)$ such that for any population (N, α) where $|N| > \eta_2(\varepsilon, U)$ if $\Gamma(N, \alpha) \in \mathcal{CG}(L)$ and $RN(\alpha) \in \mathcal{LR}(L, U, \delta_2(\varepsilon, U))$, for some L, then the imitation with innovation dynamic $\mathcal{I}(\varepsilon; \lambda \neq 0, 1; R)$ almost surely converges to a Nash, imitation ε-equilibrium.[17]

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 $RN(\alpha^\nu) \in \mathcal{LR}\left(L^\nu, U, \delta_c\left(\frac{\varepsilon}{3}\right)\right)$ for some L^ν, and for which there exists a non-singleton communication class of the imitation with innovation dynamic $(\lambda \neq 0, 1)$. Let $\delta = \delta_c\left(\frac{\varepsilon}{3}\right)$ and let $R^\nu = RN(\alpha^\nu)$ for all ν.

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 $R_i^\nu = c_q^\nu$.

For any game $\Gamma(N^\nu, \alpha^\nu)$ and any initial state σ 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.

[17] It is apparent from the proof that the statement $\delta_2(\varepsilon, U) \geq \delta_c\left(\frac{\varepsilon}{3}\right)$ can be relaxed to $\delta_2(\varepsilon, U) \geq \delta_c(*)$ where $* > \frac{\varepsilon}{2}$ is arbitrarily close to $\frac{\varepsilon}{2}$.

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 σ during the first stage of the process. For each clique c_q^ν there must exist some strategy $s_{\nu q} \in S$ such that $\sigma_i = s_{\nu q}$ for all $i \in c_q^\nu$. 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 $i^\nu \in N^\nu$ who has an innovation opportunity given strategy vector σ. Suppose, that in stage 2 of the process player i^ν chooses strategy s^k. This implies that strategy vector $\overline{\sigma}$ is observed in the next period (say period t) where $\overline{\sigma}_j = \sigma_j$ for all $j \in N^\nu \backslash \{i^\nu\}$ and $h_{\alpha^\nu(i^\nu)}(\overline{\sigma}_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\overline{\sigma}}) > h_{\alpha^\nu(i^\nu)}(\sigma_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\sigma}) + \varepsilon$.

In period $t+1$, all players use the imitation heuristic. We note that if $i^\nu \in c_{\overline{q}}^\nu$ then no player $l \in c_q^\nu$ where $c_q^\nu \neq c_{\overline{q}}^\nu$ can have a success example who is not a costrategist. Thus, if strategy vector $\overline{\overline{\sigma}}$ is observed $\overline{\overline{\sigma}}_l = \overline{\sigma}_l$ for all $l \in N^\nu \backslash c_{\overline{q}}^\nu$. Given the value of δ and continuity of payoff functions, for sufficiently large ν and for any $j \in c_{\overline{q}}^\nu$,

$$\left| h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) - h_{\alpha^\nu(i^\nu)}(\sigma_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\sigma}) \right| < \frac{\varepsilon}{3}.$$

By the assumption of global interaction, for sufficiently large ν and for any player $j \neq i^\nu$,

$$\left| h_{\alpha^\nu(j)}(\overline{\sigma}_j, w_{\alpha^\nu_{-j},\overline{\sigma}}) - h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) \right| < \frac{\varepsilon}{3}.$$

Thus,

$$\begin{aligned} h_{\alpha^\nu(i^\nu)}(\overline{\sigma}_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\overline{\sigma}}) &> h_{\alpha^\nu(j)}(\overline{\sigma}_j, w_{\alpha^\nu_{-j},\overline{\sigma}}) + \varepsilon \\ &\quad - \left| h_{\alpha^\nu(j)}(\overline{\sigma}_j, w_{\alpha^\nu_{-j},\overline{\sigma}}) - h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) \right| \\ &\quad - \left| h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) - h_{\alpha^\nu(i^\nu)}(\sigma_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\sigma}) \right| \\ &> h_{\alpha^\nu(j)}(\overline{\sigma}_j, w_{\alpha^\nu_{-j},\overline{\sigma}}) + \frac{\varepsilon}{3} > h_{\alpha^\nu(j)}(\overline{\sigma}_j, w_{\alpha^\nu_{-j},\overline{\sigma}}). \end{aligned}$$

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

Given the assumption of global interaction and the fact that U is independent of ν, for sufficiently large ν

$$\left|h_{\alpha^\nu(i^\nu)}(\overline{\overline{\sigma}}_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\overline{\overline{\sigma}}}) - h_{\alpha^\nu(i^\nu)}(\overline{\sigma}_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\overline{\sigma}})\right| < \frac{\varepsilon}{3}. \tag{2}$$

This implies that

$$h_{\alpha^\nu(i^\nu)}(\overline{\overline{\sigma}}_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\overline{\overline{\sigma}}}) > h_{\alpha^\nu(i^\nu)}(\sigma_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\sigma}) + \frac{2}{3}\varepsilon.$$

The choice of δ and continuity of payoff functions implies that for sufficiently large ν

$$\left|h_{\alpha^\nu(j)}(\overline{\overline{\sigma}}_j, w_{\alpha_{-j},\overline{\overline{\sigma}}}) - h_{\alpha^\nu(i^\nu)}(\overline{\overline{\sigma}}_{i^\nu}, w_{\alpha_{-i^\nu},\overline{\overline{\sigma}}})\right| < \frac{\varepsilon}{3}$$

for all $j \in c^\nu_q$. Thus,

$$\begin{aligned} h_{\alpha^\nu(j)}(\overline{\overline{\sigma}}_j, w_{\alpha^\nu_{-j},\overline{\overline{\sigma}}}) &> h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) + \frac{2}{3}\varepsilon \\ &\quad - \left|h_{\alpha^\nu(j)}(\overline{\overline{\sigma}}_j, w_{\alpha^\nu_{-j},\overline{\overline{\sigma}}}) - h_{\alpha^\nu(i^\nu)}(\overline{\overline{\sigma}}_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\overline{\overline{\sigma}}})\right| \\ &\quad - \left|h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) - h_{\alpha^\nu(i^\nu)}(\sigma_{i^\nu}, w_{\alpha^\nu_{-i^\nu},\sigma})\right| \\ &> h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}) \end{aligned} \tag{3}$$

for all $j \in c^\nu_q$.

Compare strategy vectors σ^ν and $\overline{\overline{\sigma}}^\nu$. We note that $X(\sigma^\nu, \overline{\overline{\sigma}}^\nu) = c^\nu_q$. It is immediate from (3), given that $\Gamma(N^\nu, \alpha^\nu) \in \mathcal{CG}(L^\nu)$ and $RN(\alpha^\nu) \in \mathcal{LR}\left(L^\nu, U, \delta_c\left(\frac{\varepsilon}{3}\right)\right)$, that, for sufficiently large ν

$$\sum_{j \in N^\nu} h_{\alpha^\nu(j)}(\overline{\overline{\sigma}}_j, w_{\alpha^\nu_{-j},\overline{\overline{\sigma}}}) > \sum_{j \in N^\nu} h_{\alpha^\nu(j)}(\sigma_j, w_{\alpha^\nu_{-j},\sigma}).$$

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.■

Theorem 2 demonstrates that for a broad class of games with many players the imitation with innovation dynamic almost surely converges on an approximate Nash, imitation equilibrium. A corollary of Theorem 2 (and of Theorem 3 to follow) is that there must exist an approximate Nash, imitation equilibrium in sufficiently large coordination games induced from a pregame satisfying the large game property. This complements a result due to Wooders, Cartwright and Selten (2001). They demonstrate that all sufficiently large games induced from a pregame satisfying the large game property have an approximate Nash, imitation equilibrium provided there is a bound, independent of population size, on the number of players of each attribute. We require no such restriction on the dispersal of players in attribute space.[18] Before discussing Theorem 2 in more detail we provide a complementary result.

5.5 Bounding the number of societies

Define a society as a group of players who (1) refer to, and only to, all other members of the society and (2) play the same strategy. The bound on reference group size in Theorem 2, as given by U, implies that the number of societies in any approximate Nash, imitation equilibrium, will grow arbitrarily large as the size of the population increases. A principle motivation of Wooders et. al. (2001) was to demonstrate the existence of a Nash equilibrium that partitioned the player set into a bounded number of societies where the bound is independent of population size. Thus, as the population size increases societies become arbitrarily large.

We offer a complementary result to that of Theorem 2 in which the number of societies can be bounded independently of the population size. Before doing so we refine the notion of a coordination game.

Coordination game: Given a pregame $\mathcal{G}$, the game $\Gamma(N, \alpha)$ is a *coordination game with bounds L and δ* if $\Gamma(N, \alpha)$ is a coordination game with bound L and if for any player $i \in N$, any strategy $s^k \in S$ and any two weight functions w_α and g_α, if

$$\sum_{\omega \in B_i(\delta)_\alpha} w_\alpha(\omega, s^k) > \sum_{\omega \in B_i(\delta)_\alpha} g_\alpha(\omega, s^k)$$

and $w_\alpha(\omega, s^k) = g_\alpha(\omega, s^k)$ for all $\omega \notin B_i(\delta)_\alpha$ then

$$h_{\alpha(i)}(s^k, w_{\alpha_{-i}}) \geq h_{\alpha(i)}(s^k, g_{\alpha_{-i}}).$$

We denote by $\mathcal{CG}(L, \delta)$ the set of coordination games with bound L and δ.

A coordination game with bounds L and δ has the additional property (over a coordination game with bound L) that a player gets a higher payoff

[18] 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.

when there are more players with 'similar' attributes to himself who are playing the same strategy as himself. This seems an intuitively plausible characteristic of a coordination game.

We state our third main result.

Theorem 3: Let $\mathcal{G}(\eta_c, \delta_c, \eta_g, \delta_g)$ be any pregame satisfying the large game property and let RN be any large game reference network function. Given any $\varepsilon > 0$ there exists real numbers $\eta_3(\varepsilon)$ and $\delta_3(\varepsilon) \geq \delta_c\left(\frac{\varepsilon}{3}\right)$ such that for any population (N, α) where $|N| > \eta_3(\varepsilon)$ if $\Gamma(N, \alpha) \in \mathcal{CG}(L, \delta_3(\varepsilon))$ and $RN(\alpha) \in \mathcal{LR}(L, |N|, \delta_3(\varepsilon))$, for some L, then the 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 (2) that we observe any significant difference. This changes to

$$h_{\alpha^\nu(i^\nu)}(\overline{\overline{\sigma}}_{i^\nu}, w_{\alpha^\nu_{-i^\nu}, \overline{\overline{\sigma}}}) \geq h_{\alpha^\nu(i^\nu)}(\overline{\sigma}_{i^\nu}, w_{\alpha^\nu_{-i^\nu}, \overline{\sigma}})$$

as implied by the fact $\Gamma(N, \alpha) \in \mathcal{CG}\left(L, \delta_c\left(\frac{\varepsilon}{3}\right)\right)$.■

The convergence result of Theorem 3 is not dependent upon each player referring to a bounded number of players. Thus, the number of societies need not grow large as the size of the population grows large. Indeed suppose there exists a reference network where every player refers to every other player in the population and where 'every player in the population is approximately similar'; Theorem 3 could be applied to show the existence of an approximate Nash, imitation equilibrium where every player in the population plays the same strategy.

5.6 Discussion

A sketch of the proof of Theorem 2 (and Theorem 3) provides some intuition and allows us to highlight some additional issues. A simple example is also provided.

Take as given a population (N, α). Define a *clique* C as a subset of player set N with the property that every player $i \in C$ refers to, and only to, the clique, thus $R_i = C$ for all $i \in C$. Suppose that $RN(\alpha) \in \mathcal{LR}(L, U, \delta)$ for some L. This implies that the reference network $RN(\alpha)$ will have the property that the player set can be partitioned into a set of *cliques* $C_1, ..., C_Q$, where each clique is of size L or greater. In the long run, given that players imitate, it is to be expected that players in the same clique will play the same strategy. Thus, assume, for the moment, that players in the same clique always play the same strategy. Further suppose that a 'clique only changes strategy' if doing so would, ceteris paribus, increase the payoff of each member of the clique. Finally, assume that only 'one clique at a time changes strategy'. If the game $\Gamma(N, \alpha) \in \mathcal{CG}(L)$ then it is clear, if play evolves as above, that the per-capita payoff will increase and never decrease. Play must therefore evolve to an absorbing state and thus an approximate Nash, imitation equilibrium. We highlight that cliques are clearly related to societies as defined in Section 5.5.[19]

In sketching the proof of Theorems 2 and 3 it remains for us to argue why cliques could be seen as behaving in the way outlined in the previous paragraph. Consider a clique C and a period t where every player $i \in C$ is playing some strategy A. We note that all members of the clique C receive approximately the same payoff (because they have similar attributes). For the purposes of this explanation assume that they all receive the same payoff. Suppose that a player $i \in C$ uses the innovation heuristic, has an innovation opportunity of strategy B, and therefore chooses strategy B in period $t+1$. Ceteris paribus, the payoff of player i increases by at least ε. Provided that the population is sufficiently large (and thus the influence of player i is sufficiently small) player i will be a success example to all members of the clique C in period $t+1$. Thus, if all members of clique C use the imitation heuristic they will all choose strategy B in period $t+2$. Assume again, for simplicity, that in period $t+2$ all members of clique C receive the same payoff. Provided that the payoff of player i is higher in period $t+2$ than in period t then the payoff of every player $j \in C$ is higher in period t than in period $t+2$. If this is the case then we have illustrated how play may evolve as outlined in the previous paragraph. This will be the case if clique C is 'sufficiently small' or if the game is a coordination game with bounds L and δ for some appropriate value of δ.

An important element in the proof of our main theorems, as sketched above, is how players appear to act collectively, within their cliques, even if they are not aware of doing so. This 'collective action' stems from the imitation. By acting within cliques, and thus in groups of size L or more, players are able to realize the gains suggested by a game $\Gamma(N, \alpha)$ being a coordination game with bound L. Without imitation players may not be able to realize any such gains. We note, for example, that if $L > 1$

[19]Note, however, that cliques are defined with respect to a reference network while societies are primarily defined with respect to a strategy vector.

and players learn by innovation then, because each player acts unilaterally (both directly and indirectly), the potential gains from 'group' action my not be realized.[20] We see, therefore, that imitation has a crucial role to play in the learning dynamic. Innovation, however, plays an equally important role in enabling cliques to 'search for more efficient strategies'. This results from each individual within a clique looking for innovation opportunities. It is worth noting that it need not be enough for one person with a clique to innovate while all others imitate; players within a clique are similar, but are also sufficiently different that one player could have an innovation opportunity that another does not.

The above discussion suggests that innovation and imitation complement each other to enable learning of an approximate Nash, imitation equilibrium. This leads us to question the role of imitation and, in particular, whether imitation can 'help' players to learn to play an approximate Nash equilibrium. We look at this issue in more detail. There are broadly two viewpoints that could be taken with respect to imitation. First, we could take the view that individuals imitate because of some inherited behavior and imitation may be of no benefit to an individual. If we take this viewpoint the task is to question whether imitation 'gets in the way' of individual learning. In particular, can imitation (or conformity) be consistent with individually rational behavior. This could be seen as the viewpoint taken by Bernheim (1994) and Wooders, Cartwright and Selten (2001). A second viewpoint is to say that individuals imitate because they derive some specific benefit from doing so. This viewpoint would lead us to ask whether imitation can 'help' learning. This could be seen as the viewpoint taken by Schlag (1999). We consider each viewpoint in turn.

Theorems 2 and 3 suggest that in large coordination games imitation can be consistent with individual rationality. This is demonstrated by the fact that individuals imitate and yet approximate Nash equilibrium play emerges. The proof of Theorem 2 allows us to be a little more specific. In particular, if play evolves as set out in the proof of Theorem 2 and as outlined above then, as discussed, any player who imitates will increase her payoff. This clearly suggests that imitation can be consistent with individual rationality. We note however that this argument, for the individual rationality of imitation, relies crucially on the fact that each player only imitates those with similar attributes to himself. If this is not the case then the individual rationality

[20] If a game $\Gamma(N, \alpha) \in \mathcal{CG}(1)$ then an innovation dynamic (with inertia) $\mathcal{I}(\varepsilon; \lambda = 1; R)$ will converge to an approximate Nash equilibrium. This is a trivial result. If, however, $\Gamma(N, \alpha) \notin \mathcal{CG}(1)$ then there is no guarantee that an innovation dynamic need converge to an approximate Nash equilibrium.

of imitation is called into question.

Let us now consider whether imitation can 'help' learning. At first this may seem unlikely given the proceeding discussion. In particular, if a player who imitates increases his payoff then imitation may appear to be a form of innovation. This, however, is not the case for two reasons. First, a player, through imitation, can realize individual gains of less than ε. That is, a player may be motivated to change strategy through imitation when no innovation opportunity exists. This suggests that we could just set ε to be zero. Note, however, that if $\varepsilon = 0$ there need not exist a Nash ε-equilibrium and thus play may fail to converge to a Nash ε-equilibrium. Second, imitation suggests a dynamic in which players within the same clique play the same strategy. This allows the clique to behave 'as a group' and realize the gains suggested by a game being a coordination game with bound L. If players innovate then different players may have different, and many, innovation opportunities. Thus, players within the same clique may end up playing different strategies. This 'lack of coordination' suggests that players may fail to realize the gains suggested by a game being a coordination game with bound L. The potential for imitation to 'help ' learning is discussed further in the conclusion.

We finish this section with a simple example that may help to illustrate some of the discussion. We discuss potential applications of the imitation with innovation dynamic in the conclusion.

Example 7: The strategy space is given by $S = \{1, 2, 3\}$ and the attribute space by $\Omega = [0, 3]^3$. Given attribute $\omega = (\omega_1, \omega_2, \omega_3)$ the value of ω_k could be thought of as a player's preference for strategy k. For any population (N, α) and any weight function $w \in W_\alpha$ let $y_w : S \to \mathbb{R}$ be defined by

$$y_w[k] = \sum_{\omega \in support(\alpha)} w(\omega, k)$$

for all $k \in S$. The value $y_w[k]$ is thus total number of players playing strategy k. For any population (N, α) and any player $i \in N$ the payoff function of player i is given by[21]

$$h_{\alpha(i)}(k, w) = \frac{1}{|N|} \left[\omega_k y_w[k-1] + k(y_w[k] + 1)\right]$$

for all $k \in S$ and $w \in W_{\alpha-\omega}$ where $\alpha(i) = \omega = (\omega_1, \omega_2, \omega_3)$. We note that if all players $i \in N$ play strategy k then each player receives a payoff of k.

[21] If $k = 1$ then set $y_w[0] = 0$.

Thus, for 'most games' the Nash equilibrium $(3, 3, ..., 3)$ will be the Pareto optimum. It is easily checked that this pregame satisfies the large game property.[22]

In the context of Theorem 3 it is possible to set $\delta_3(\varepsilon) = 2$ and $\eta_3(\varepsilon) = 1$ for any $\varepsilon > 0$. We will, however, consider just one specific population (N, α). Let $|N| = 300$ and suppose player 1 has attribute $\omega^1 = (0, 2, 0)$, player 2 has attribute $\omega^2 = (0, 0, 3)$ and players $3, 4, ..., 300$ have attribute $\omega^0 = (0, 0, 0)$. Further, assume the reference network is such that $R_i = N$ for all $i \in N$. Suppose that $\varepsilon = 0$.

Assume an initial state $(1, 1, ..., 1)$. Each player receives a payoff of 1. Player 1 can increase his payoff to 2 by playing strategy 2 while any other player switching to strategy 2 would see her payoff fall to $\frac{2}{300}$. Play will thus evolve to strategy vector $(2, 1, 1, ..., 1)$ whereby player 1 becomes a success example to all players. If players use the imitation heuristic then play may thus evolve to strategy vector $(2, 2, ..., 2)$. Note that the transition from strategy vector $(2, 1, 1, ..., 1)$ to $(2, 2, ..., 2)$ may come in one step or through a gradual process. Given strategy vector $(2, 2, ..., 2)$ player 2 has an innovation opportunity - he can improve his payoff by playing strategy 3. Play may thus evolve to strategy vector $(2, 3, 2, ..., 2)$ and then onto $(3, 3, ..., 3)$. In reality the evolution of play may not be 'as neat as above' in the sense that player 3 may take an innovation opportunity and play strategy 3 while some players are still playing strategy 1. That the imitation with innovation dynamic dynamic will converge to the Pareto optimum state $(3, 3, ..., 3)$ is, however, not in doubt. In this example we clearly see the process of innovation and imitation that was outlined above. We can also look at the role played by imitation by assuming that players do not imitate. We note that strategy vector $(2, 1, 1, ..., 1)$ is a Nash equilibrium. Thus, if players do not imitate and play commences with state $(1, 1, ..., 1)$ the innovation dynamic $\mathcal{I}(\varepsilon = 0; \lambda = 1; R)$ will converge to state $(2, 1, 1, ..., 1)$ and not to the Pareto optimal state.

6 Conclusion

This paper has provided sufficient conditions under which a population of boundedly rational individuals will learn to play an approximate Nash equilibria. Indeed, we go further by showing that aggregate play converges towards an approximate Nash, imitation equilibrium in pure strategies. We focussed on learning in coordination games with many players and learning

[22] A suitable metric on Ω is $dist(\omega, \overline{\omega}) = \max_k |\omega_k - \overline{\omega}_k|$.

through imitation with innovation. We demonstrated that the convergence of an imitation with innovation dynamic is dependent on the reference network through which players refer to each other; if the reference network has a clustering coefficient of one and if each player refers to players similar to himself then convergence is more likely. Our main results suggest that imitation can be consistent with individually rational behavior. Through example we demonstrate that imitation may even aid learning in the sense that players learn to play a 'more efficient' strategy vector, when using both imitation and innovation, than they do when just using innovation.

Two potential applications of these results appear to be in modelling market interaction and technological or scientific evolution. 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). The imitation with innovation dynamic may also be appropriate for modelling market interaction; consumers and producers are involved in an adaptive process of choosing products to buy and sell and deciding what prices to pay or accept. Adaption in 'Cournot like' market interaction games has 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.

The notion that imitation can aid learning is another avenue we feel is worth exploring further. After all, if individuals do imitate and conform then there should be some reason for this. Intuitively, one advantage of imitation would appear to be the *speed* that it can give to learning. If we see innovation as being difficult and thus relatively rare while imitation is much easier to perform then 'learning should be quicker' if players imitate. This is surely the case with technologically and scientific evolution. The implications of imitation for the speed of learning are explored by Levine and Pesendorfer (2000, 2001). In focussing on long run convergence this paper has not addressed such short to medium run issues. We do provide some evidence that imitation can potentially 'help' learning even in the long run; we feel, however, that its ability to do so is somewhat limited. What we can say with more confidence is that imitation need not hinder learning in the sense that it can be consistent with individually rational behavior in the long run. Putting this together, we might suggest that imitation may be an aid to learning in the short run while not hindering learning in the long run. Future research hopes to consider this in more detail.

The evolution of an imitation dynamic is fundamentally dependent on the reference network that players use. Some analysis is presented in an

Appendix on the implications of players choose their reference group as play evolves. A related literature concerns network formation (see for example Jackson and Wolinsky 1996, Bala and Goyal 2000 and A. Watts 2001). 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. The question of how sensitive the convergence of the imitation with innovation dynamic is to changes in the reference network is also an open question.

As a final remark we note that any interpretation of our results must take into account the realism of our model of learning. We believe that our model of learning through imitation with innovation captures key aspects of individual learning in games with many players. 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). Experiments to test the importance of social learning in large games would be of interest.

7 Appendix 1: an evolving reference network

In this section we generalize the analysis contained in the main body of the paper by allowing players to change their reference group as play evolves. In particular, as well as choosing a strategy in each period, players are also required to choose a reference group. We provide sufficient conditions on how players choose their reference group such that Theorems 1 and 2 can be extended.

We assume that players are constrained in the reference groups that they can choose. For a player set N, let $U, L \in \mathbb{R}^N$ denote respectively *upper* and *lower limits on the size of reference groups* where $U_i > L_i$ for all i. Let $D = \{D_1, ..., D_n\}$ denote a *topological structure on reference groups* where $D_i \subset N$, $\{i\} \in D_i$ and $|D_i| \geq U_i$, for all $i \in N$. A set of *reference group constraints* is given by a triple (U, L, D) consisting of upper and lower limits on the size of reference groups and a topological structure on reference groups. In interpretation, the values U_i and L_i are interpreted respectively as the *upper* and *lower limits on the size of reference group for player* $i \in$

N. The set D_i is interpreted as the set of players to whom player i may potentially refer. Given the set of reference group constraints (U, L, D), we denote by $\Psi_{i,(U,L,D)}$ the set of *feasible reference groups of player* i where $R_i \in \Psi_{i,(U,L,D)}$ if and only if $R_i \subset D_i$ and $U_i \geq |R_i| \geq L_i$. That is, a reference group is feasible for player i when they are referring to a subset of D_i and when the number of players referred to is between the two bounds U_i and L_i. We denote by $\Psi_{(U,L,D)}$ the set of *feasible reference networks* where $R \in \Psi_{(U,L,D)}$ if and only if $R_i \in \Psi_{i,(U,L,D)}$ for all $i \in N$.

Given a stage game $\Gamma = (N, S, \{u_i\}_{i=1}^n)$ and a set of reference group constraints $Z = (U, L, D)$ we refer to an *action* as a choice of both strategy for the stage game Γ and as a choice of reference group relative to the set of constraints Z.[23] For each player $i \in N$, the *action set* of player i is thus given by the set $S \times \Psi_{i(U,L,D)}$, which we subsequently denote by $\Sigma_{i,\Gamma,Z}$. An *action profile* is given by a vector $\sigma = (\sigma_1, ..., \sigma_n)$ where $\sigma_i \in \Sigma_{i,\Gamma,Z}$ denotes the action of player i. Let $\Sigma_{\Gamma,Z} = \times_{i \in N} \Sigma_{i,\Gamma,Z}$ be the set of action profiles relative to stage game Γ and a set of reference group constraints Z.

As play evolves over periods $t = 0, 1, 2, ...$ all players simultaneously choose an action in each period. We assume that players make action choice conditional on events of the last two periods; this is a departure from the main text where only the last period is used. We model the evolution of play as a discrete time homogenous Markov chain $\{h(t)\}_{t \geq 0}$ on state space $\Sigma_{\Gamma,Z} \times \Sigma_{\Gamma,Z}$.

We assume that each player i uses a *good advice heuristic* in choosing a strategy conditional on state $a = (\overline{\sigma}, \overline{R}, \sigma, R)$. The heuristic can be summarized under a *good advice probability function* $g_i : \Sigma_{\Gamma,Z} \times \Sigma_{\Gamma,Z} \to \Delta(N)$. The value $g_i(j|a)$ is interpreted as the probability that player i would *select* player j conditional on action profile a. If player j is selected, $j \notin R_i$ and $|R_i| < U_i$ then player i will choose a reference group $R_i \cup \{j\}$. If player j is selected, $j \in R_i$ and $|R_i| > L_i$ then player i will choose reference group $R_i \backslash \{j\}$. Otherwise, player i chooses reference group R_i. Thus, reference groups evolve by the selective addition and subtraction of members to and from the group. We assume that $\sum_{j \in N} g_i(j|a) < 1$ for all a. Thus, player i may always take the option to leave the reference group unchanged. Other assumptions on g_i are as follows:

1. *achieves aspiration:* if $\sigma = \overline{\sigma}$ then $g_i(j|a) = 0$ for all $j \in N$.
2. *good advice:* if $u_i(\sigma) > u_i(\overline{\sigma})$ then $g_i(j|a) > 0$ for all $j \in R_i \backslash \{i\}$.

[23] We assume that payoffs are not directly dependent upon reference group choice.

3. *bad advice*: if $u_i(\sigma) < u_i(\overline{\sigma})$ then $g_i(j|a) > 0$ for all $j \in D_i \backslash R_i$.
4. *indifferent advice:* if $u_i(\sigma) = u_i(\overline{\sigma})$ and $\sigma \neq \overline{\sigma}$ then $g_i(j|a) > 0$ for all $j \in D_i \backslash \{i\}$.

If a player receives good advice, i.e. her payoff has increased over the previous period, then she may remove a player from her reference group. If a player receives bad advice, i.e. her payoff has declined over the previous period, then she may add an extra player to her reference group. If a player achieves her aspiration, the strategy vector remains unchanged, then she does nothing. If a player receives indifferent advice, gets the same payoff even though the strategy vector has changed, then she may add or remove a player from her reference group. As it stands the good advice heuristic does not give much leeway in reference group choice. It can easily be generalized, however, with no effect on Theorem 5, to allow more large scale revisions of the reference group.

Assume that players select strategies using the imitation and/or innovation heuristics. We refer to the resulting dynamic process as the imitation with innovation and good advice dynamic, denoted $\mathcal{I}(\varepsilon; \lambda; Z)$. The following result demonstrate that, for any feasible reference network $\overline{R}$, play must either evolve to a state with reference network $\overline{R}$ or play must converge to an absorbing state of the dynamic.

Theorem 4: Let $Z = (U, L, D)$ be a set of reference group constraints and $\overline{R} \in \Psi_{(U,L,D)}$ be any feasible reference network. From any state $a \in \Sigma_{\Gamma,Z} \times \Sigma_{\Gamma,Z}$ the imitation with innovation and good advice dynamic $\mathcal{I}(\varepsilon; \lambda; Z)$ either, almost surely, converges on an absorbing state of the dynamic or will pass through a state $\overline{a}$ with reference group $\overline{R}$.

Proof: Suppose not. Then there exists a reference group $\overline{R} \in \Psi_{(U,L,D)}$ and initial state a such that play does not either converge to an absorbing state or on a state with reference group $\overline{R}$. Given two sets A and B we denote by $A - B$ the set $A \backslash (A \cap B)$. Suppose that each player i chooses his reference group in the following way, where he is selecting a player, according to the good advice heuristic, and has current reference group R_i

1. if *good advice* and $R_i - \overline{R}_i \neq \phi$ then select a player $j \in R_i - \overline{R}_i$. If $R_i - \overline{R}_i = \phi$ then select no one.

2. if *bad advice* and $\overline{R}_i - R_i \neq \phi$ then select a player $j \in \overline{R}_i - R_i$. If $\overline{R}_i - R_i = \phi$ then select no one.

3. if *indifferent advice*, $\overline{R}_i - R_i \neq \phi$ and $|R_i| < U_i$ then select a player $j \in \overline{R}_i - R_i$. Else, if $R_i - \overline{R}_i \neq \phi$ select a player $j \in R_i - \overline{R}_i$. Otherwise, select no one.

If play evolves as above with transition matrix P and does not converge to an absorbing state then there must exist a non-singleton set of states Ψ, indexed $a(t) = (\sigma(t-1), R(t-1), \sigma(t), R(t))$, $t = 1, 2, 3, ..., T$, where $P_{a(t-1)a(t)} > 0$ for all $T > t > 1$ and $P_{a(T)a(1)} > 0$.

There must exist a state $a\left(\overline{t}\right) \in \Psi$ such that $\sigma\left(\overline{t}\right) \neq \sigma\left(\overline{t}-1\right)$. That is, at some point some player must change strategy. Any player $i \in N$ either receives indifferent advice, bad advice or good advice in state $a\left(\overline{t}\right)$. If a player i receives bad advice (good advice) in state $a\left(\overline{t}\right)$ then there must exist a state $a\left(\overline{\overline{t}}\right) \in \Psi$ in which player i receives good advice (bad advice). Further, according to the assumed behavior the only players that can be added to a reference group R_i are those players $j \in \overline{R}_i$ while the only players that can be taken out of reference group R_i are those players $j \notin \overline{R}_i$. This must imply that $R_i(t) = \widehat{R}_i$ for some $\widehat{R}_i \in \Psi_{i,(U,L,D)}$, all $i \in N$ and for all $a(t) \in \Psi$.

If $\overline{R}_i - \widehat{R}_i = \phi$ and $\widehat{R}_i - \overline{R}_i = \phi$ then $\widehat{R}_i = \overline{R}_i$. Thus, either $\overline{R}_i - \widehat{R}_i \neq \phi$ or $\widehat{R}_i - \overline{R}_i \neq \phi$ for some $i \in N$. Suppose that $\overline{R}_i - \widehat{R}_i \neq \phi$. Given the assumed behavior (assumptions 2 and 3) this would imply that $\left|\widehat{R}_i\right| = U_i$ which, in turn, implies (assumptions 1 and 3) that $\widehat{R}_i - \overline{R}_i = \phi$ (where we recall that $U_i > L_i$). If $\overline{R}_i - \widehat{R}_i \neq \phi$ and $\widehat{R}_i - \overline{R}_i = \phi$ this implies that $\left|\overline{R}_i\right| > U_i$ which contradicts that $\overline{R} \in \Psi_{(U,L,D)}$. Thus, $\overline{R}_i - \widehat{R}_i = \phi$ and $\widehat{R}_i - \overline{R}_i \neq \phi$. Using similar arguments to those immediately above this implies that $\left|\widehat{R}_i\right| = L_i$ and $\left|\overline{R}_i\right| < L_i$ which again contradicts that $\overline{R} \in \Psi_{(U,L,D)}$. Thus, $\widehat{R}_i = \overline{R}_i$ for all $i \in N$ and this completes the proof.■

Theorem 4 allows us to extend Theorems 1, 2 and 3 in allowing players to choose their reference group. For example, in applying Theorem 1, we have that: for any stage game Γ and any set of reference group constraints Z, for which there is a feasible reference network R that has a clustering coefficient of one, the imitation with innovation and good advice dynamic $\mathcal{I}(\varepsilon; \lambda = 0; Z)$ almost surely converges on an absorbing state. At this absorbing state the strategy vector chosen is an imitation equilibrium. This is immediate from

Theorem 5 above by setting $\overline{R}$ to be a reference network with a clustering coefficient of one.

8 Appendix 2: The imitation heuristic with inertia

We provide some analysis of the imitation with innovation dynamic in which players use the imitation heuristic with inertia as opposed to the imitation heuristic. We recall that the distinction between these two heuristics (as discussed in Section 2.2) lies in whether a player i will imitate a success example who is not a costrategist and is earning the same payoff as a costrategist. The following example may help to illustrate the importance of this distinction. This example demonstrates that Theorem 1 does not hold if players use the imitation heuristic with inertia. Throughout the rest of this section we assume players use the imitation heuristic with inertia when selecting a strategy through imitation.

Example A1: There are 5 players and two strategies labelled A and B. The reference network is given by $R_1 = \{1, 2, 4, 5\}$, $R_2 = \{1, 2, 4, 5\}$, $R_3 = \{2, 3, 4\}$, $R_4 = \{1, 2, 4, 5\}$ and $R_5 = \{1, 2, 4, 5\}$. This network has a clustering coefficient of one. We highlight the following payoffs,

strategy vector	payoff vector
A, A, B, B, B	$100, 10, 0, 0, 100$
A, A, A, B, B	$100, 0, 0, 10, 100$

Assume $\lambda = 0$. There exists a cycle of strategy vectors $(A, A, B, B, B) \rightarrow$

$(A, A, A, B, B) \rightarrow (A, A, B, B, B)$ in which player 3 changes strategy motivated by observing players earning a payoff of 10. Note that because players $1, 2, 4$ and 5 are using the imitation heuristic with inertia they have no desire to change strategy; if using the imitation heuristic they would have such an incentive⧫

Given a network R we say that there is a *directed path between player i and player j* if there exists a chain of players $i_1, .., i_M$ such that $i_1 \in R_i$, $i_{m+1} \in i_m$ and $j \in i_M$. We say that a network R has a *characteristic path length of one* when for any two players $i, j \in N$ if there exists a directed path between i and j then $j \in R_i$.[24] We note that the reference network

[24]See D. Watts (1999) for a definition of and discussion on the characteristic path length of a network.

in Example A1 does not have a characteristic path length of one; player 3 refers to player 2 who in turn refers to player 1; player 3, however, does not refer to player 1. The following result complements Theorem 1. Before stating Theorem A1 we modify the definition of an imitation equilibrium in the obvious way. The strategy vector σ is an *imitation equilibrium with inertia* of stage game Γ relative to reference network R if

$$\max_{l \in R_i} u_l(\sigma) \leq \max_{l \in C_i(\sigma)} u_l(\sigma)$$

An imitation equilibrium with inertia is an absorbing state of an imitation with inertia dynamic.

Theorem A1: For any stage game Γ and any reference network R that has a clustering coefficient of one and characteristic path length of one the imitation with innovation dynamic $\mathcal{I}(\varepsilon; \lambda = 0; R)$ almost surely converges on an imitation equilibrium with inertia.

Proof: The proof closely follows that of Theorem 1 and so only the differences will be explained in detail. Thus, given an arbitrary state σ we demonstrate that there exists states $\sigma(2), .., \sigma(T)$ where $P_{\sigma\sigma(2)} > 0$, $P_{\sigma(t)\sigma(t+1)} > 0$ for all $T - 1 \geq t \geq 2$ and $\sigma(T)$ is an imitation equilibrium with inertia. We assume that every player $i \in N$ always chooses the same strategy as a success example and we 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 imitates the success example playing the smallest strategy. This behavior occurs with positive probability under the imitation with innovation dynamic.

Consider an arbitrary player $i \in N$ for whom there exists a player $j \in R_i$ such that $i \in R_j$. As demonstrated in Theorem 1 player i, and j, belong to a clique R_i. That is, $R_j = R_i$ for all $j \in R_i$. As play evolves, given the assumed behavior of agents, the number of distinct strategies played by members of R_i can only diminish. For example, if a player i is plays strategy s^k in period $t-1$ and then imitates, in period t, a success example who is not a costrategist, there can be no player $j \in R_i$ who plays strategy s^k in period t or in any subsequent period. Given that there are only a finite number of players there must exist some t_i such that for every state $\sigma(t)$, $t \geq t_i$, no player $j \in R_i$ can have a success example who is not a costrategist.

Consider an arbitrary player $i \in N$ for whom there does not exist a player $j \in R_i$ such that $i \in R_j$. As shown in Theorem 1 player i must refer to a subset of a clique R_k. Indeed, given that the network has a characteristic path length of one, it must be the case that $R_i = R_k \cup \{i\}$; that is, player i refers to everybody in the clique R_k. Restrict attention to those states $\sigma(t)$ such that $t \geq t_k$. That is, those states for which no player in clique R_k can have a success example who is not a costrategist. If there is a success example of player i who is not a costrategist in some state $\sigma(t)$ then any success example of player i in a subsequent state must be a costrategist of i. Given the player set is finite, there must exist, therefore, some t_i such that for every state $\sigma(t)$, $t \geq t_i$, player i does not have a success example who is not a costrategist. This completes the proof.■

The analogs of Theorem 2 and Theorem 3 hold without further qualification.[25] The analysis, however, is somewhat more involved. In particular, a complicating factor is the possibility that players in the same clique may play different strategies. As, argued in the proof of Theorem 1A, however, a dynamic can be assumed in which the number of strategies used by players in a clique can only ever diminish even if it does not fall to one.

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] Bala, V. and S. Goyal, 2000, A noncooperative model of network formation, *Econometrica*, 68: 1181-1229.

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

[5] Blume, L. 1993, The Statistical Mechanics of Strategic Interaction, *Games and Economic Behavior*, 4: 387-424.

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

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

[25] The only minor qualification is that play converges to an approximate Nash, imitation equilibrium with inertia rather than a Nash, imitation equilibrium.

[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] Ellison, G., Learning, 1993, Local Interaction and Coordination, *Econometrica*, 61: 1047-1071.

[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] Gigerenzer, G. and R. Selten, 2001, *Bounded Rationality: The Adaptive Toolbox*, The MIT Press.

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

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

[18] Jackson, M. and A. Watts, forthcoming, The Evolution of Social and Economic Networks, *Journal of Economic Theory.*

[19] Jackson, M. and A. Wolinsky, 1996, A strategic model of social and economic networks, *Journal of Economic Theory*, 71: 44-74.

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

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

[22] Kalai, E. and E. Lehrer, 1993, Rational learning leads to Nash equilibrium, *Econometrica*, 61, No 5: 1019-1045.

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

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

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

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

[27] Mas-Colell, A., 1984, On a theorem of Schmeidler, *Journal of Mathematical Economics* 13: 206-210.

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

[29] Offerman, T., J. Potters and J. Sonnemans, forthcoming, Imitation and belief learning in an oligopoly experiment, *Review of Economic Studies.*

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

[31] Robson, A.J. and F. Vega-Redondo, 1996, Efficient equilibrium selection in evolutionary games with random matching, *Journal of Economic Theory*, 70: 65-92.

[32] 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.

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

[34] Schmeidler, D., 1973, Equilibrium points of nonatomic games, *Journal of Statistical Physics* 7: 295-300.

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

[36] Selten, R. and J. Apesteguia, 2002, Experimentally observed imitation and cooperation in price competition on the circle, mimeo.

[37] Selten, R. and A. Ostmann, 2000, Imitation Equilibrium, Mimeo.

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

[39] Watts, A., 2001, A dynamic model of network formation, *Games and Economic Behavior*, 34: 331-341.

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

[41] Wooders, M., E.J. Cartwright and R. Selten, 2001, Social Conformity and Equilibrium in Pure Strategies in Games with Many Players, University of Warwick working paper no. 589.

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

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

[44] Young, H.P., 2001b, The Dynamics of Conformity, in *Social Dynamics*, edited by S.N. Durlauf and H.P. Young, The MIT Press.

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