SOCIAL CONFORMITY AND BOUNDED RATIONALITY IN ARBITRARY GAMES WITH INCOMPLETE INFORMATION: SOME FIRST RESULTS

Edward Cartwright

And

Myrna Wooders

No 672

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Social conformity and bounded rationality in arbitrary games with incomplete information; Some first results.*

Edward Cartwright
Department of Economics
University of Warwick
Coventry CV4 7AL, UK
E.J.Cartwright@warwick.ac.uk

Myrna Wooders
Department of Economics
University of Warwick
Coventry CV4 7AL, UK
M.Wooders@warwick.ac.uk
http://www.warwick.ac.uk/fac/soc/Economics/wooders/

January 2002
Draft version

Abstract

Intepret a set of players all playing the same pure strategy and all with similar attributes as a society. Is it consistent with self interested behaviour for a population to organise itself into a relatively small number of societies? By introducing the concept of approximate substitute players in non-cooperative games we are able to put a bound on the rationality of such social conformity for an arbitrary game and arbitrary number of societies.

*The main results of this paper were presented at the the 2002 General Equilibrium Conference held in Athens in May 2002 and at Northwestern University in August 2002. We thank the participants for their interest and comments.

1 Introduction

The economic literature is increasingly addressing the question of whether a population of boundedly rational individuals can learn to behave as if rational or nearly rational. A standard way to look at this issue is to study the evolution of play when players are repeatedly matched to play a stage game and where, in each period, each player use a simple heuristic or 'rule of thumb' (Young 1993, Blume 1993, Ellison and Fudenberg 1995). Underlying most such heuristics are two characteristics that seem intuitively appealing in any model of bounded rationality, namely, that players (1) only use pure strategies and (2) there is social conformity. By social conformity we mean that individuals conform to observed strategy choices of some group of similar individuals. In this paper, we ask whether, and when, such behavior can be consistent with individual rationality. We address this issue by determining, for an arbitrary game: (a) a bound on the irrationality of using pure strategies, that is, more technically, how large ε must be, as a function of the parameters describing a game, for the existence of an ε-equilibrium in pure strategies and; (b) a bound on the irrationality of using pure strategies that exhibit social conformity as a function of the parameters describing the game. In modelling social conformity we introduce the notion of *endogenous roles* and discuss how the existence of these roles may imply social conformity amongst individuals who are not necessarily playing the same actions. In characterising games we also lead to provide a definition of similarity amongst players or, in other words, to define when players can be seen as near substitutes for each other. In a companion paper (Cartwright and Wooders 2002) the results of this paper are applied to provide a general class of large games where the use of both pure strategies and social conformity can be consistent with boundedly rational behavior.

There are two fundamental observations motivating the notion that people only use pure strategies and not (non-degenerate) mixed strategies. The first observation is the fact that a person never has a strictly positive incentive to play a mixed strategy; given the strategies of the other players, the payoff from playing a mixed strategy can always be achieved by playing one of the pure strategies in the support of that mixed strategy.[1] The second

[1] This rationale will depend on the game in question. Von Neumann and Morgentsern (1944), for example, point out that in zero sum games, such as matching pennies, randomization can be a deliberate ploy to leave ones opponent guessing ones intentions. In many games, however, mixed strategy equilibria reflect the inability of players to coordinate actions, as in, for example, the battle of the sexes. Justifying the use of mixed strategy equilibria in these cases is much harder.

observation is the seeming inability of people to generate random sequences. Thus, even when players use strategies that may look like mixed strategies in terms of the variance in observed actions the actions may not actually be chosen at random (see Kagel and Roth 1995 and Walker and Wooders 2001).

There is little experimental literature in economics questioning whether social conformity plays a role in individual choice behavior. The psychology literature, however, supports the idea that people do have a tendency to conform to the actions of similar people. In an economic context there are many ways to explain such behavior. For example, if a person is boundedly rational or has imperfect information then he may imitate a person he believes is better informed (Gale and Rosenthal 1999, Schleifer 2000). Alternatively, in a coordination game with multiple equilibria, a player may be able to make a more informed strategy choice by observing the actions of others (Scharfstein and Stein 1990 Ellison and Fudenberg 1995, Young 2001). Finally, due to normative influences, a person may be motivated by desires for prestige, popularity or acceptance or more generally, to 'fit in' with a social norm (Bernheim 1994). It is worth highlighting that the psychology literature suggests that players only conform with the actions of others who they see as similar to themselves. With respect to the causes for conformity above, we can see some justification for this. For example, a person may not seek to gain prestige amongst those he views as very different to himself.

It would appear, from the above, that the use of pure strategies and conformity may appropriately be elements of models of bounded rationality, particularly, in games with many players. Wooders, Cartwright and Selten (2001), WCS, demonstrate that for a broad class of large games such behavior can be consistent with boundedly rational behavior. There are two important limitations to WCS. First, only games of complete information are treated; the limitations of this approach will become clear as we proceed. Second, the framework is restricted to noncooperative pregames. A noncooperative pregame, like a cooperative pregame or a pre-economy is a structure that enables one to treat large finite games.[2] The shortcoming of the pregame framework is that it does not treat an individual game in terms of parameters describing the game. Moreover the results have limitations

[2] By a pre-economy we mean simply a space of endowments of private commodities and preferences. This is a familiar sort of construct in economics; see, for example, Hildenbrand (1974). As in the cooperative and noncooperative pregame frameworks, given a set of agents, these are described by their attributes – in the case of a pre-economy, their endowments and preferences.

imposed by the pregame framework. One such limitation is that the social conformity result of WCS cannot treat social conformity in games where, ex-post, all players are identical.

In this paper we treat individual games and, in terms of the parameters describing a game, determine the lowest possible bound on ε for a given game to have a Nash ε-equilibrium in pure strategies exhibiting social conformity. To describe a game, we first introduce the notion of *approximate substitute players* of a non-cooperative game. This is a counterpart to the notion of approximate substitutes in cooperative games (Kovalenkov and Wooders 2001, for example). We also define the concept of a (δ, Q)-class game. A (δ, Q)-class game has the property that the player set can be partitioned into Q *classes* of players. Players in the same class are seen as approximate substitute players, where the dissimilarity of players in a class is bounded by the parameter δ. A (δ, Q)-class game also is required to have a sort of 'large game' property; the effect of a change in strategy of just one member of a class on the remaining players in the game must be small. This would be a natural property of standard models of large exchange economies games or large economies with local public goods or clubs, as in Kovalenkov and Wooders (1997), for example. We note that any finite game is a (δ, Q)-class game for any Q and some δ. An advantage, therefore, of introducing the notion of a (δ, Q)-class game is that it allows us to draw conclusions on arbitrary games for an arbitrary level of conformity, as measured by Q.

The first conclusion we draw relates to the use of pure strategies and can be, informally, stated as *if a game* Γ *is a* (δ, Q) *class game then, for any* $\varepsilon \geq 4\delta$ *the game has a Nash* ε*-equilibrium in pure strategies.* This result allows us to put a lower bound on the size of ε permitting existence of an ε-equilibrium in pure strategies for an arbitrary game (with a finite player set). Alternatively, the value of 2δ could be interpreted as a bound on the distance from full rationality of players using pure strategies as opposed to mixed strategies.

To address social conformity, we define a society as a collection of players in the same class and playing the same strategy. An immediate consequence of Theorem 1 is that *if a game* Γ *is a* (δ, Q) *class game then, for any* $\varepsilon \geq 4\delta$ *there exists a Nash* ε*-equilibrium in pure strategies that induces a partition of the player set into no more than* QK *societies*, where K is the number of strategies. This is essentially the analogue of the social conformity result due to WCS.

In the context of incomplete information there are other natural notions of social conformity. In particular, we may wish to permit the possibility that players in the same society perform different actions. We motivate this view with two examples. First, consider the example of driving automobiles on roadways with the choice of whether or not to give way at road junctions. We may think of the strategies in this instance as 'to give way' or 'not give way'. We would not want to interpret those players who give way as being part of a different society to those who do not give way. It may be that each player in the population conforms to the behavior of 'agree to a classification into major and minor roads and obey a highway code that dictates giving way at minor roads and not giving way at major roads'. What we observe in this scenario is that players who use different actions can still be thought of as belonging to the same society; their choice of action is conditional upon the role they are playing in society at any one instance, i.e. whether they are on a minor road or road.

As a second example consider the division of labour within a large firm (or university, or economy). Let the choice of job be the strategy. In equilibrium we may find a variety of strategy choices, with some people working in finance and others in engineering for example. The interpretation that there exists a society of finance workers and a distinct society of engineers is a much more intuitively justifiable one in this instance; there may be distinct finance and engineering departments, for instance. We may still, however, wish to take the view that all of the members of this firm belong to the same society, namely, the firm; the fact that different people have different jobs merely reflects the differing roles they take within the firm or society.

In both the above examples the 'aggregate strategy' of the society looks like a mixed strategy. Indeed the players within the society may also play mixed strategies. If we do not require that players choose pure strategies we demonstrate that *if a game* Γ *is a* (δ, Q) *class game then for any* $\varepsilon \geq 4\delta$ *there exists a Nash* ε*-equilibrium that induces a partition of the player set into no more than* Q *societies.* That is, there exists a Nash ε-equilibrium with the property that any two players of the same class play the same mixed strategy. The key point to note is how the number of societies no longer depends on the number of strategies K. This suggests a higher level of social conformity.

A large literature, which we review in a later section, is concerned with the purification of mixed strategy equilibria through incomplete information. In particular, in a game of incomplete information a player's action can be made conditional on her randomly determined type. In this way a player's chosen action may vary, and resemble a mixed strategy, even though she may play a pure strategy. It is conventional to assume that there is some exogenous uncertainty that makes this possible, for example uncertainty over payoffs (Harsanyi 1973). In contrast, we assume that players within a society can endogenously create incomplete information by creating a set of roles to which players in that society are randomly assigned. The players within the society choose the probability distribution with which they are assigned to roles. We now interpret a society as a set of players with the property that all players in the society (a) are in the same class, (b) have the same probability of being assigned each role, and (c) play the same pure strategy, where action choice may be conditional on role. We derive the following result *if a game* Γ *is a* (δ, Q) *class game then for any* $\varepsilon \geq 10\delta$ *there exists a Nash* ε*-equilibrium in pure strategies that induces a partition of the player set into no more than* Q *societies.*

We see that if players have some way of endogenising types, that is allocating players to different roles within the society, then each player can play a pure strategy and yet the number of societies need not be conditional on the number of strategies. This result suggests a higher level of conformity is possible than shown by WCS. We should, however, check that in allowing players to endogenise roles we still retain an interpretable and intuitive notion of a society. One key point to highlight in terms of the definition of a society is that two players within a society have the same probability distribution with which they are assigned to roles. It is non-trivial that players should desire this and it seems to be a criterion that requires players within a society to share some 'common plan' even if they perform different actions within that plan. In the driving example, for instance, this requirement could be interpreted as requiring that everybody in the population agrees on a classification on roads into major and minor roads.

The division of labour example raises a further issue. It is possible that players are only likely to ever take one role within the society. This raises the question of whether expected utility is a relevant criteria on which to judge individual rationality. In particular, while all members of the society agree before the allocation to roles, will there be the same agreement ex-post once players are aware of the role that they have been allocated? Recent results due to Kalai (2000) suggest that this issue need not be a concern. In particular, any Nash ε-equilibrium is likely to be ex-post stable in the sense that following the allocation of players to roles no one individual will have an incentive to change her behavior.

The approximate substitute framework allows us to draw conclusions about arbitrary games. It is also useful to have some general examples of (δ, Q) class games for arbitrary values of δ. In a companion paper (Cartwright and Wooders 2002) we extend the pregame framework of WCS to allow incomplete information and to permit local interaction. We are able to connect the concept of games with approximate substitutes to that of games induced by a pregame satisfying the large game property. This allows us to apply the results of this paper in terms of arbitrarily small ε.

We proceed as follows: Section 2 introduces the notation and Section 3 defines the notion of approximate substitutes. Section 4 looks at approximate purification and Section 5 considers social conformity. Section 6 looks at social conformity in pure strategies, Section 7 looks at related literature and Section 7 concludes.

2 A Bayesian Game - definitions and notation

A Bayesian game Γ is given by the tuple (N, A, T, p, u) where N is a finite *player set*, A is a set of *action profiles*, T is a set of *type profiles*, p is a *set of player beliefs* and u the a set of *utility functions.* We define these in turn.

Let $N = \{1, ..., n\}$ be a finite player set. For all $i \in N$ there exists a finite set T_i of *feasible types of player* i and a finite set A_i of *feasible actions of player* i (independent of type). Let $T = \times_i T_i$ be the set of *type profiles* and let $A = \times_i A_i$ be the set of *action profiles.* We assume throughout, for convenience, that $T_i = \mathcal{T}$ and $A_i = \mathcal{A}$ for all $i \in N$ and for some finite sets $\mathcal{T}$ and $\mathcal{A}$. We will typically index a type as $t^z \in \mathcal{T}$ and an action as $a^l \in \mathcal{A}$.

Each player can make their action conditional on their type. Thus, a *pure strategy* of a player i is given by a vector $s^k = \{s^k(t^1), ..., s^k\left(t^{|\mathcal{T}|}\right)\}$ where $s^k(t^z)$ is interpreted as the action chosen by player i when she is of type t^z. For any player i we allow choice of any pure strategy consistent with the set of feasible types $\mathcal{T}$ and actions $\mathcal{A}$.[3] Denote the set of pure strategies by $\mathcal{S}$ where we let $K = |\mathcal{A}|^{|\mathcal{T}|} = |\mathcal{S}|$ be the number of pure strategies.

A *strategy* (possibly non-degenerate) of a player i is given by a vector $\sigma_i = \{\sigma_{i1}, ..., \sigma_{iK}\}$ where σ_{ik} is interpreted as the probability player i plays pure strategy $s^k \in \mathcal{S}$. A strategy σ implies a vector $\{\sigma_i(\cdot|t_1), ..., \sigma_i(\cdot|t_{|T_i|})\}$ where $\sigma_i(\cdot|t_i)$ is interpreted as a probability distribution over the set of actions $\mathcal{A}$ to be used by player i when of type t_i. The value $\sigma_i(a_i|t_i)$ is interpreted as the probability player i uses action a_i given he or she is of type t_i. Let $\Delta(\mathcal{S})$ denote the set of strategies for player i. Given strategy σ_i let *support*(σ_i) denote the pure strategies played with strictly positive probability. Let $S = \times_{i \in N} \Delta(\mathcal{S})$ denote the set of *strategy vectors.* We refer to a strategy vector σ as *degenerate* if σ_i places unit weight on a unique

[3] We assume that every player can be of any type and can choose any action. This greatly simplifies the analysis but could be relaxed. The complicating issue is how we can say that two players are playing the same strategy when they can potentially be of different types or choose different actions.

pure strategy for all $i \in N$.

Let $\mathcal{C} = \mathcal{T} \times \mathcal{A}$ denote the feasible *compositions of player i*. That is, a composition is a type-action pair. Let $C = \times_{i \in N} \mathcal{C}$ denote the set of *composition profiles*. For each player $i \in N$ there exists a *utility function* $u_i : C \to \mathbb{R}$. The interpretation is that $u_i(c)$ denotes the payoff of player i if the composition profile is c. We will typically index a composition as $c^r \in \mathcal{C}$. Let $u = \{u_1, ..., u_n\}$ denote the set of player utility functions.

Each player $i \in N$ forms her own beliefs about the types of other players as given by a function p_i. The function p_i is a probability distribution over the set of type profiles T. The value $p_i(t)$ is interpreted as the probability player i puts on the type profile being t. We refer to p_i as the *beliefs of player i about the type profile.* With a slight abuse of notation let $p_i(t_{-i}|t_i)$ denote the probability that player i puts on the type profile being $t = (t_{-i}, t_i) \in T$ given that he is of type t_i. We assume the marginal distribution $p_i(t_i = t^z)$ is strictly positive for all $t^z \in \mathcal{T}$ and all $i \in N$. That is, player i puts a strictly positive probability on being of each possible type. Let $p = \{p_1, ..., p_n\}$ be the set of beliefs about the type profile.

For the most part we make no assumptions about beliefs as given by set p. We occasionally, however, will make reference to two some standard assumptions on player beliefs. For each player $i \in N$ there exists some *prior probability distribution over types g_i*. That is, $g_i(t_i)$ denotes the probability that player i is of type $t_i \in \mathcal{T}$ if the types of the remaining players $N \backslash \{i\}$ are undetermined. Let g denote a *probability function over the set of type profiles.* Thus, $g(t)$ denotes the probability of type profile $t \in T$. Two standard assumptions are[4]

1. *Independent type allocation:* for all $i \in N$, g_i is independent of the type profile over the remaining players. That is, $g(t) = \prod_i g_i(t_i)$ where $t = (t_1, ..., t_n)$.

2. *Consistent beliefs:* for all $i \in N$ and for all $t_i \in \mathcal{T}$,

$$p_i(t_{-i}|t_i) = \frac{g(t_{-i}, t_i)}{\sum_{l_{-i} \in T_{-i}} g(l_{-i}, t_i)}.$$

Players are assumed to act according to expected payoffs. In particular, knowing his type (but not the type of the other players), a player is assumed

[4] In fact, because this paper treats finite Bayesian games permitting general utility functions, these two assumptions could be made without loss of generality. This follows from the fact that any finite Bayesian game Γ is equivalent to some finite Bayesian game Γ^E in which beliefs are consistent and the type allocation is independent (see, for example, Myerson 1997 pp 72-73).

to choose an action which maximizes his expected payoff. Given a strategy vector σ and beliefs about the type profile p_i a player can form expectations about the likely composition profile. For instance, the probability that player i puts on composition $c = ((a_1, t_1), ..., (a_n, t_n))$ is given by

$$prob(c) = p_i(t_1, ..., t_n)\sigma_1(a_1|t_1)...\sigma_n(a_n|t_n).$$

Let $U_i(\cdot|p_i)$ denote the *expected utility function of player i given beliefs p_i* mapping strategy vectors into the real line.

We note how the function U_i accounts for both the uncertainty over player types, and the uncertainty due to mixed strategy vectors.

A strategy vector σ is a *Bayesian Nash ε-equilibrium* if,

$$U_i(\sigma) \geq U_i(s^k, \sigma_{-i}) - \varepsilon$$

for all $s^k \in \mathcal{S}$ and for all $i \in N$. We say that a Bayesian Nash ε equilibrium m is a *Bayesian Nash ε-equilibrium in pure strategies* if m_i is degenerate for all i.

3 Approximate substitutes

Given a Bayesian game $\Gamma = (N, T, A, p, u)$ we consider partitioning the player set N into subgroups with the property that any two players in the same group can be viewed as approximate substitutes for each other. This requires us to formulate a metric by which to compare players. In cooperative game theory the distance between players is typically measured by the maximum difference in value that players can add to coalitions. Informally, in cooperative game theory a δ-substitute partition has the property that, given any coalition structure, 'swapping' players who are δ-substitutes between coalitions, has an effect of less than δ on the worth of the coalitions. In proposing the analogue in non-cooperative games we consider two different ways of measuring the distance between players. Informally, we say that two players i and j are interaction substitutes if i and j are seen as similar by those with whom they interact. In contrast, we say that players i and j are individual substitutes if they have similar payoff functions. Combining both measures together, we refer to players i and j as approximate substitutes if they are both interaction and individual substitutes. We formally introduce these terms below.

We consider two players i and j as being approximate *interaction substitutes* when the payoff to any player k is relatively invariant if player's i and

j exchange strategies. Formally, a partition $\{N_1, ..., N_Q\}$ is a δ-*interaction substitute partition* for Bayesian game Γ when for any two strategy vectors $\sigma^1, \sigma^2 \in S$, if

$$\sum_{i \in N_q} \sigma_{ik}^1 = \sum_{i \in N_q} \sigma_{ik}^2,$$

for all N_q and all $s^k \in \Delta(\mathcal{S})$, then,

$$\left| U_i(s^k, \sigma_{-i}^1) - U_i(s^k, \sigma_{-i}^2) \right| \leq \delta$$

for any player $i \in N$ and any strategy $s^k \in \Delta(\mathcal{S})$.

A second measure we put on the distance between players is the similarity of their payoff functions. In particular, a partition $\{N_1, ..., N_Q\}$ is a δ-*individual substitute partition* for Bayesian game Γ when for any N_q, for any two players $i, j \in N_q$ and for any strategy vector $\sigma \in S$ such that $\sigma_i = \sigma_j$

$$\left| U_i(s^k, \sigma_{-i}) - U_j(s^k, \sigma_{-j}) \right| \leq \delta$$

for any strategy $s^k \in \Delta(\mathcal{S})$.

We say that a partition $\mathcal{N} = \{N_1, ..., N_Q\}$ is a δ-*substitute partition* if $\mathcal{N}$ is both a δ-interaction substitute partition and a δ-individual substitute partition. In this instance we say that two players belonging to a subset N_q are δ-substitutes for each other.

We make some observations. First, it is trivial that any Bayesian game Γ has a 0-substitute partition $\{\{1\}, ..., \{N\}\}$. That is, each player is a 0-substitute for themselves. We also note that for any Bayesian game Γ and any $Q \leq N$ there exists a δ-substitute partition for some finite $\delta \geq 0$. Further, partitions into a larger number of subsets will typically reduce the minimum value of δ for which there exists a δ-substitute partition. Formally, if game Γ has a Q member δ-substitute partition then for any $\overline{Q}$ where $N \geq \overline{Q} \geq Q$ there exists a $\overline{\delta} \leq \delta$ such that Γ has a $\overline{Q}$ member $\overline{\delta}$ substitute partition. This is an immediate consequence of the fact that players are 0-substitutes for themselves.

3.1 A (δ, Q)-class Bayesian game

We begin this section by defining a third type of partition. This can be seen as a measure of how invariant a player's payoff can be to changes in the strategy vector. Formally, a partition $\{N_1, ..., N_Q\}$ is a *δ-strategy switching partition* when for any two strategy vectors $\sigma^1, \sigma^2 \in S$ if

$$\sum_{i \in N_q} \left| \sigma_{ik}^1 - \sigma_{ik}^2 \right| \leq 1, \tag{1}$$

for all N_q and all $s^k \in S$ then

$$\left| U_i(s^k, \sigma_{-i}^1) - U_i(s^k, \sigma_{-i}^2) \right| \leq \delta$$

for any player $i \in N$ and any strategy $s^k \in \Delta(\mathcal{S})$.

The definition of a δ-strategy switching partition requires us to put a maximum bound on the change in payoffs when, essentially, Q players, one from each class, change strategy. It should be clear that, unlike the previous two types of partition, the smaller is Q then the smaller is likely to be the value of δ for which there exists a δ-strategy switching partition. It is unlikely that there will be a 0-strategy switching partition. In contrast to defining δ-substitutes there is no real intuitive justification for putting two players together into a subset of a δ-strategy partition. We note, however, that how the player set is partitioned can significantly effect the minimum δ for which the partition is a δ-strategy partition. Also, as we show in example A it is often the case that 'opposites' should be grouped together. That is, players we would see as opposite in terms of the two similarity criteria above. This would suggest that a partition of the player set into approximate substitutes need not be a 'good' partition when viewed on the criteria of a strategy switching partition. This is an issue we explore in the next sub-section after defining a (δ, Q)-class Bayesian game.

We say that a Bayesian game Γ is a *$(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game* if there exists a Q member partition $\{N_1, ..., N_Q\}$ that is a δ_I-interaction substitute partition, a δ_P-individual substitute partition and a δ_C-strategy switching partition. We say that a Bayesian game Γ is a (δ, Q)-class Bayesian game if Γ is a *$(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game* where $\delta_I, \delta_P, \delta_C \leq \delta$. We refer to each N_q as a class of player and say that two players $i, j \in N_q$ are the same class of player.

In summary, if a game is a (δ, Q)-class Bayesian game then the population can be partitioned into classes such that any two players in the same class are δ-substitutes. Furthermore, a player's payoff changes by at most δ if 'one player' from each class changes strategy.

3.2 Discussion

To illustrate the notion of a (δ, Q)-class Bayesian game consider the two extreme cases highlighted below. For this discussion we assume the game is one of complete information with player set N where $|N| = n$. Let $\{N_1, , , .N_Q\}$ denote a partition of the player set into classes. First, suppose that $|N_q| = 1$ for all $i \in N$ and so $Q = n$. As we have highlighted, this is a 0-substitute partition. Note, however, that finding a δ for which there exists a δ-strategy switching partition would require comparing strategy vectors σ^1 and σ^2 in which any player $i \in N$ can change their strategy any way they wish. Thus, unless the game is trivial, it cannot be a (δ, N)-class Bayesian game for any meaningful value of δ. Indeed, in general δ would have to be as large as the maximum possible difference between the expected payoffs of a player.

In contrast, suppose that $|N_1| = n$ and so there is only one class of player. It is now much more unlikely that there will exist a δ-substitute partition for a relatively small value of δ. This would require that payoffs depend only on the 'population average' or the number of players playing each strategy. This is plausible (such an assumption is used in Kalai 2000, for example), and not as restrictive as it may seem but in general is a strong assumption. Suppose, however, there does exists a 1-member δ-substitute partition of Γ. It seems likely that in this scenario Γ will be a $(\delta, 1)$-class Bayesian game. The reason being that if $|N|$ is 'large' then finding a δ for which there exists a δ-strategy switching partition would only require comparing strategy vectors σ^1 and σ^2 in which the strategies of players are slightly perturbed. Informally, if payoffs are only a function of the population average then we would expect players to be relatively indifferent to small changes in this population average.

Between these two extremes we clearly find a trade off between a small or large number of classes Q. In particular, for an arbitrary game Γ, finding the minimum δ for which Γ is a (δ, Q)-class Bayesian game would seem to involve a trade-off when varying the size of Q. If Q is large then it seems more plausible there should exist a δ-substitute partition for small δ while if Q is small then it seems more likely that there should exists a δ-strategy switching partition for large δ. This issue is illustrated by Example A which appears in the appendix. This example also serves to demonstrate how for any game Γ and for any Q the minimum value of δ such that Γ is a (δ, Q)-class game can be calculated.

3.3 Games with incomplete information

In defining a (δ, Q)-class Bayesian game the role of incomplete information is not explicit. It is useful to offer an illustration of the possible role incomplete information can play. We assume throughout this subsection that beliefs are both independent and consistent. We highlight the following issue: if two players i and j use the same or similar strategies then this does not necessarily imply that their expected composition is similar. This is because the prior probability distribution over types of players i and j may differ. Thus, even though players i and j play similar strategies the expectations of what will be realized, in terms of their type and action, may be dissimilar.

Consider the definition of a δ-interaction substitute partition. This definition requires comparison of the difference in payoff from two strategy vectors $\sigma^1, \sigma^2 \in S$ where,

$$\sum_{i \in N_q} \sigma^1_{ik} = \sum_{i \in N_q} \sigma^2_{ik}, \tag{2}$$

for every member of the partition N_q. Given two such strategy vectors we can compare the probability that a typical player of any class N_q will have type t^z and play action a^l. Formally, given strategy vector σ, let $pr(\sigma, N_q, t^z, a^l)$ denote the expected probability that a player of class N_q will be of type t^z and play action a^l. We note that,

$$pr(\sigma, N_q, t^z, a^l) = \frac{1}{|N_q|} \sum_{i \in N_q} \sigma_i(a^l | t^z) g_i(t^z).$$

Thus,

$$\begin{aligned}
& \left| pr(\sigma^1, N_q, t^z, a^l) - pr(\sigma^2, N_q, t^z, a^l) \right| && (3) \\
\leq \; & \frac{1}{|N_q|} \left| \sum_{i \in N_q} \sigma^1_i(a^l | t^z) g_i(t^z) - \sum_{i \in N_q} \sigma^2_i(a^l | t^z) g_i(t^z) \right| \\
\leq \; & \frac{1}{|N_q|} \max_{j,k \in N_q} \{ |g_j(t^z) - g_k(t^z)| \} \sum_{i \in N_q} (\sigma^1_i(a^l | t^z) - \sigma^2_i(a^l | t^z)) \\
\leq \; & \max_{j,k \in N_q} \{ |g_j(t^z) - g_k(t^z)| \}. && (4)
\end{aligned}$$

It is clear that this inequality may be binding. For example, suppose that there are two players 1 and 2 in a class N_q. Player 1 is always of type t^1

and player 2 is always of type t^2. That is, $g_1(t^1) = g_2(t^2) = 1$.[5] There are two actions a^1 and a^2. Let σ^1 be such that player 1 plays the strategy 'if type t^1 play action a^1 and if type t^2 play action a^2' and player 2 plays the strategy 'if type t^1 play action a^2 and if type t^2 play action a^1'. Consider a strategy σ^2 in which both players exchange strategies. It is clear that (2) holds and (3) holds with equality. Thus, even though the aggregate strategy of the class is invariant between strategy vectors σ^1 and σ^2 the expectations of what will happen, in terms of actions, change completely.

What we learn from this discussion is how the definition of a δ-interaction substitute partition, and also a δ-strategy switching partition, implicitly measures the variability in prior operability distributions over types and the importance of such variations on payoffs. In particular, if a game is to be a (δ, Q)-class Bayesian game for small δ, then we would expect that either players of the same class have similar prior probability distributions over types or payoffs are relatively invariant to the type profile. This, however, seems a reasonable assumption; an assumption of common priors, for example, makes such issues irrelevant.

4 Purification of mixed strategies

Our first result places a bound on the rationality of using pure strategies. To derive this result we require three lemmas which are stated and proved in the Appendix.

Theorem 1: Let $\Gamma = (N, A, T, p, u)$ be any Bayesian game that is a $(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game. Let ε be a positive real number where $\varepsilon \geq 2(\delta_I + \delta_C)$. The game Γ has a Bayesian Nash ε-equilibrium in pure strategies.

Proof: Using Nash's Theorem there must exist a Nash Equilibrium strategy σ^*. This implies, for all $i \in N$, that,

$$U_i(\sigma_i, \sigma^*_{-i}) \geq U_i(s^k, \sigma^*_{-i}) \tag{5}$$

for all σ_i where $support(\sigma_i) \subset support(\sigma^*_i)$ and for all $s^k \in \Delta(\mathcal{S})$.

Given that Γ is a $(\delta_I, \delta_P, \delta_C, Q)$-class game there is a partition of N into Q classes. Let $\{N_1, ..., N_Q\}$ be a such a partition of N. We apply Lemma 2 in turn to each N_q. Doing so implies that there exists a strategy vector

[5] To be in keeping with the analysis of the rest of the paper we require $g_i(t^z) > 0$ for all t^z. This could clearly be done without changing the substance of the example.

$m \in S$ where *support*(m_i) $\subset$*support*(σ_i^*), m_i is degenerate for all $i \in N$ and where,

$$\left\lceil \sum_{i \in N_q} \sigma_i^* \right\rceil \geq \sum_{i \in N_q} m_i \geq \left\lfloor \sum_{i \in N_q} \sigma_i^* \right\rfloor .$$

for all $q = 1, ..., Q$. Thus,

$$\left| \sum_{i \in N_q} m_{ik} - \sum_{i \in N_q} \sigma_i^* \right| \leq 1$$

for all $s^k \in S$ and all q. By Lemma 3 this implies that there exists a strategy vector $\overline{\sigma}$ such that $\sum_{i \in N_q} \overline{\sigma}_{ik} = \sum_{i \in N_q} \sigma_{ik}^*$ and $\sum_{i \in N_q} |\overline{\sigma}_{ik} - m_{ik}| \leq 1$ for all s^k and all N_q.

From the definition of a δ-interaction substitute partition we have that,

$$\left| U_i(s^k, \overline{\sigma}_{-i}) - U_i(s^k, \sigma_{-i}^*) \right| \leq \delta_I$$

for any player $i \in N$ and any strategy $s^k \in \Delta(\mathcal{S})$. By the definition of a δ-strategy switching partition we have that,

$$\left| U_i(s^k, \overline{\sigma}_{-i}) - U_i(s^k, m_{-i}) \right| \leq \delta_C$$

for any player $i \in N$ and any strategy $s^k \in \Delta(\mathcal{S})$. Thus, for any $s^k \in \Delta(\mathcal{S})$ and for all $i \in N$

$$\left| U_i(s^k, \sigma_{-i}^*) - U_i(s^k, m_{-i}) \right| \leq \delta_I + \delta_C.$$

Therefore, given (5),

$$\begin{aligned} U_i(m_i, m_{-i}) - U_i(s^k, m_{-i}) &\geq -\left| U_i(m_i, \sigma_{-i}^*) - U_i(m_i, m_{-i}) \right| - \left| U_i(s^k, \sigma_{-i}^*) - U_i(s^k, m_{-i}) \right| \\ &\geq -2(\delta_I + \delta_C) \geq -\varepsilon \end{aligned}$$

for all $i \in N$ and all $s^k \in \Delta(\mathcal{S})$.■

We note that the value of δ_P has no effect on the bound for which there exists a Bayesian Nash ε-equilibrium. That is, the existence of an approximate Bayesian Nash equilibrium in pure strategies does not require that players in the same class should have similar payoff functions. This will not be the case when we consider social conformity below.

As previously remarked, any game Γ is a $(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game for some δ. Theorem 1 allows us, therefore, to put a bound on the rationality of using pure strategies as opposed to mixed strategies. An important issue is, of course, how useful this bound is. Two simple examples illustrate some of the issues.

Example 1.

The number of players is given by $1,000,001$. Player 1 is labelled the *leader* and players 2 to $1,000,001$ are labelled *citizens*. The citizens play a sort of matching pennies game with the leader. There are two pure strategies H and T. The strategy of the citizens is defined as the mean average strategy of the citizens; thus, if $500,000$ citizens choose to play H and $500,000$ choose T the strategy of the citizens is to play H with probability one half and to play T with probability one half. The payoff matrix is as follows where a citizen receives a payoff according to the strategy of the citizens,

		strategy of citizens	
		H	T
leader	H	$1, -1$	$-1, 1$
	T	$-1, 1$	$1, -1$

Thus, if the leader chooses H and the strategy of the citizens is T the leader gets a payoff of -1 and each citizen gets a payoff of 1.

This game has no Bayesian Nash ε-equilibrium in pure strategies for any value of $\varepsilon < 2$. It is trivial to note that there is no strategy δ_C-switching partition of the player set for any value of δ_C significantly below 2. Thus, Theorem 1 shows there exists a Bayesian Nash 4-equilibrium.$\blacklozenge$

The first thing we observe in this example is that the use of pure strategies does not appear to be boundedly rational. The primary reason for this is that one player, namely the leader, can significantly alter the payoffs of other players. We also see that the bound provided by Theorem 1 is a significant overestimate. Indeed, given that payoffs only range between 1 and -1 we can immediately conclude there exists a Bayesian Nash 2-equilibrium.[6]

Generally, as in Example 1, the bound on ε given in Theorem 1 is unlikely to be binding. To explain, we begin by highlighting that the method of

[6] For the reader with an interest in cooperative game theory, it may be interesting to observe that here we have the failure of "small group negligibility" – a small group of players (one, in this instance) can have a significant effect on aggregate payoffs.

proof is one of 'purifying' a Bayesian Nash equilibrium in mixed strategies. With this in mind we see that while the definition of a (δ, Q)-class Bayesian game requires us to compare any two strategy vectors of a game, in reality, demonstrating the existence of a pure strategy Bayesian Nash ε-equilibrium merely requires us to look at strategy vectors 'near to' any mixed strategy Bayesian Nash equilibrium. A second reason why Theorem 1 is unlikely to be binding is motivated by an example. Suppose there exists a mixed strategy Bayesian Nash equilibrium in which a player i plays a pure strategy A with probability 0.01 and a strategy B with probability 0.99. The proof of Theorem 1 allows for the possibility that player i, after the purification, may end up playing strategy A with probability 1. Indeed, while one player having to switch in this manner is plausible, the proof of Theorem 1 permits that every player in the population may have to switch strategies to this extent. This seems unlikely to ever have to be the case.

That the bound on ε provided by Theorem 1 will typically be an overestimate does not prevent Theorem 1 of being of use in demonstrating cases where the use of pure strategies is boundedly rational. A simple example illustrates this point.

Example 2.

The situation is virtually identical to that in Example 1 except the game is now played between two different groups of citizens, that is, there is no longer a leader. Suppose that there are $1,000,002$ players with equal numbers in each group. The game has no Bayesian Nash ε-equilibrium in pure strategies for any value of ε below $1/500,001$. Theorem 1 implies that there exists a Bayesian Nash $4/500,0001$ equilibrium.♦

In Cartwright and Wooders (2002) by using a pregame framework we provide a general class of large game for which Theorem 1 shows the existence of an approximate equilibrium in pure strategies. We discuss this further in Section 7.

5 Social conformity

We begin by defining a *society*. Take as given a game $\Gamma = (N, A, T, p, u)$, a partition of N into classes $\{N_1, ..., N_Q\}$ and a strategy vector $\sigma \in S$. For any strategy $s^k \in \Delta(\mathcal{S})$ and any q, define the subset N_q^k of N such that $i \in N_q^k$ if and only if $i \in N_q$ and $\sigma_i = s^k$. If N_q^k is non-empty then we refer to the set N_q^k as a society. Thus, a society is (a maximal set) such that every player belonging to that society plays the same strategy and has the same class.

Given a partition of N into classes $\mathcal{N} = \{N_1, ..., N_Q\}$ and a strategy vector $\sigma \in S$ there exists a unique partition $\{N_1, ..., N_C\}$ of the player set N into societies. We say that $\mathcal{N}$ and σ *induce* the partition into societies $\{N_1, ..., N_C\}$.

Given a Q member partition into classes $\mathcal{N}$ we say that a Bayesian Nash ε-equilibrium m is a Bayesian Nash ε-equilibrium *with social conformity* if $\mathcal{N}$ and σ induce a partition into Q societies. That is, any two players in the same class play the same strategy.

Suppose that players can choose mixed strategies. This may appear unmotivated in view of the previous focus of the paper. The main motivation will become clear, however, in the following section. Theorem 2 is also an interesting result in its own right in focussing purely on the bounded rationality of social conformity.

Theorem 2: Let Γ be any $(\delta_I, \delta_P, \delta_C)$-class Bayesian game. Let ε be a positive real number where $\varepsilon \geq 2(\delta_I + \delta_P)$. The game Γ has a Bayesian Nash ε equilibrium m with social conformity.

Proof: By Nash's existence of equilibrium Theorem there exists a Bayesian Nash equilibrium σ^* of the game Γ. Given that Γ is a $(\delta_I, \delta_P, \delta_C)$-class Bayesian game there exists a partition $\mathcal{N}$ =$\{N_1, ..., N_Q\}$ that is both a δ_I-interaction substitute and δ_P-individual partition. For each N_q and for each $s^k \in \mathcal{S}$ let $\sigma^*(q, k)$ be defined as

$$\sigma^*(q,k) = \sum_{i \in N_q} \sigma^*_{ik}.$$

Consider a strategy vector m satisfying the property that, for all $i \in N$, if $i \in N_q$ then $m_{ik} = \sigma^*(q, k)$; the strategy vector m assigns each player some pure strategy in his best response set. Clearly $\mathcal{N}$ and m induce a partition into societies $\{N_1, ..., N_Q\}$.

It is trivial that

$$\sum_{i \in N_q} \sigma^*_{ik} = \sum_{i \in N_q} m_{ik}$$

for all q and all $s^k \in \mathcal{S}$. Given that σ^* is a Bayesian Nash equilibrium and $\mathcal{N}$ a δ_I-interaction substitute partition

$$U_i(\sigma^*_i, m_{-i}) \geq U_i(s^k, m_{-i}) - 2\delta_I$$

for all $i \in N$ and all $s^k \in \Delta(\mathcal{S})$. Given that $\mathcal{N}$ is a δ_I-individual substitute partition

$$\left| U_i(s^k, m_{-i}) - U_j(s^k, m_{-j}) \right| \leq \delta_P$$

for any players $i, j \in N_q$ for some q and for any $s^k \in \Delta(\mathcal{S})$. Given the construction of m it must be the case that

$$U_i(m) \geq U_i(s^k, m_{-i}) - 2(\delta_I + \delta_P)$$

for all $i \in N$ and all $s^k \in \Delta(\mathcal{S})$. This completes the proof.■

Theorem 2 shows that if players are allowed to use mixed strategies then each class can be taken as a society; thus, the number of societies can be bounded by the number of classes. That is, if Γ is a (δ, Q) class game then for any $\varepsilon \geq 4\delta$ there exists a Bayesian Nash ε-equilibrium such that any two players belonging to the same class play the same strategy. We note that the value of δ_C is irrelevant for the bound on which there exists a Bayesian Nash ε-equilibrium with social conformity.

We note that Theorem 2 encompasses the special case in which $Q = n$. In this case, there exists a 0-substitute partition and so there exists a Bayesian Nash equilibrium in which there are n societies. This is, of course, just an immediate application of the Nash Existence Theorem. In interpretation of Theorem 2 this observation makes clear that we need to have some notion of how large classes need to be.

Making judgements on how many people a class should have is clearly somewhat arbitrary. Indeed a class size of one may not be unreasonable; this may reflect players who choose not to conform. The model of conformity used by Bernheim (1994), for example, leads, in some instances, to a 'central' group of people who conform around a standard norm with other 'extreme' individuals choosing 'to do their own thing' by not conforming to such a norm.

We propose two ways in which the issue of class size can be, at least partially, overcome. First, by moving to a pregame framework the number of classes can be fixed independently of the size of the player set. This is pursued further in Cartwright and Wooders (2002) as discussed in Section 7. Second, the concept of ex post stability can be used to give a criterion for judging how large a class need be. We pursue this further in the next section.

Before continuing we briefly return to consider Example 1. It is trivial that for the game in example 1 there is a partition into 2 classes which is both a 0-interaction substitute partition and a 0-individual substitute partition. Thus, Theorem 2 shows that there exists a Bayesian Nash 0-equilibrium in which the population is partitioned into two societies. We find in Example 1 a game for which the use of pure strategies cannot be considered boundedly rational while social conformity can be. This appears to be a fairly general property as illustrated in Cartwright and Wooders (2002) by applying Theorems 1 and 2.

6 Social conformity in pure strategies

Having considered the rationality of using pure strategies in Section 3 and the rationality of social conformity in Section 4, we now turn to the rationality of both social conformity and the use of pure strategies. We can begin with a result that follows immediately from Theorem 1 and should need no proof.

Corollary 1: Let Γ be a $(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game. Let ε be a positive real number where $\varepsilon \geq 2(\delta_I + \delta_C)$. The game Γ has a Bayesian Nash ε-equilibrium m in pure strategies and partition of $\mathcal{N}$ into Q classes such that $\mathcal{N}$ and m induce a partition into societies $\{N_1, ..., N_C\}$ where $C \leq QK$.

This is clearly an immediate consequence of the fact that any partition of the player set N into societies induced by a partition of N into Q classes must have no more than QK societies. This result, however, can still be of interest if the number of societies can be fixed independently of the size of the player set. This is possible using the pregame framework of Wooders et al. (2000).[7]

We note that given the type of conformity in Corollary 1 there may be players in the same class who are playing different pure strategies. We interpret such players as belonging to different societies. As highlighted in the introduction, in certain circumstances such a distinction may not be appropriate. We thus wish to suggest alternative notions of conformity in pure strategies.

A large literature (reviewed in Section 7) has shown that incomplete information may permit the purification of mixed strategies. In particular, given that action choice can be conditional on type, a player can use the randomness in his type to randomize over his actions even when playing a

[7] As we note in Wooders et al. (2001), this can also be established in the framework of Kalai (2000).

pure strategy. The possibility of such purification is dependent on sufficient randomness of the type profile (Aumann et al. 1983); for example, such purification is not possible if there is perfect information.

We now assume that players can endogenously create imperfect information if sufficient randomness does not exist. In particular, given a Bayesian game Γ we consider a *Bayesian game with endogenous roles* $\Gamma(f)$. We assume that there exists a set of roles $\mathcal{R} = \{r^1, ..., r^K\}$. The number of roles is as large as the number of actions. Let $R = \mathcal{R}^n$ be the set of *role profiles*. We assume that there exists a *probability distribution over the set of role profiles* $f : R \rightarrow [0,1]$. We also assume, throughout the remainder of the paper, that beliefs are consistent and independent. Given a Bayesian game $\Gamma = (N, A, T, g, u)$ a Bayesian game with endogenous roles $\Gamma(f) = (N, A, T(f), g(f), u(f))$ is defined to satisfy:

1. $\mathcal{T}(f) = \mathcal{T} \times \mathcal{R}$ for all $i \in N$,
2. $g(f)(t, r) = g(t)f(r)$ for all $t \in T$ and all $r \in R$,
3. $u_i(f)(a, t, r) = u_i(a, t)$ for all $a \in A$, $t \in T$, $r \in R$ and all $i \in N$.

The reader can observe that roles are basically equivalent to types. A player can, for instance, make action choice conditional on her role. Further, players are assumed to have consistent beliefs with respect to the conditional distribution over role profiles. One important distinction between a role and a type is that a player's role can have no effect on his payoff or the payoff of any other player.

Roles are signals through which players can coordinate their actions. There is an equivalence between the roles as defined here and the signals players are assumed to receive in defining correlated equilibria. We discuss this equivalence further in Section 7. We will not provide a 'story' of how these roles come to be an accepted means of coordination, or how they are assigned to players, but merely take as given the set $\mathcal{R}$ and function f. As we will discuss in section 7, this is not atypical in the literature.

We place an important restriction on the form that the probability distribution over roles f can take. To do so, we take as given a partition $\{N_1, ..., N_Q\}$ of the player set into classes. A probability distribution over roles f satisfies *within class anonymity* if the probability that a player from a class N_q will have role r^k is identical for all players belonging to the class. Formally, if $i, j \in N_q$ for some q then,

$$\sum_{r \in R: r_i = r^k} f(r) = \sum_{r \in R: r_j = r^k} f(r)$$

for all $r^k \in \mathcal{R}$. We assume that any Bayesian game with endogenous roles $\Gamma(f)$ has a probability distribution of roles f satisfying within class anonymity.

The importance of within class anonymity lies in how we interpret a society. A society, in a Bayesian game with endogenous roles, is defined, as before, as a maximal set of players such that every player in the society plays the same strategy and is of the same class. Given that actions can be conditional on role, in such a Bayesian game, we have to be careful to retain a meaningful notion of a society. In particular, does the fact that players within the same society play the same strategy imply that they can be seen as conforming to some norm or of sharing some common identity? The assumption of within class anonymity would suggest so. This is because, while there are different roles within the society and players with different roles may perform different actions, each player is equally likely to take each role within the society.

The presence of roles implies that mixed strategies can be purified by making action conditional on role. This allows us to model conformity in pure strategies more generally. Indeed an immediate corollary of Theorem 2 is the following.

Corollary 2: Let Γ be any $(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game. Let ε be a positive real number where $\varepsilon \geq 2(\delta_I + \delta_P)$. There exists a partition $\mathcal{N}$ of the player set into Q classes and a Bayesian game with endogenous roles $\Gamma(f)$ (where f is within class anonymous) such that $\Gamma(f)$ has a Bayesian Nash ε equilibrium m in pure strategies with social conformity.

This result demonstrates that if players have some endogenous system by which players can be assigned roles then we can conceive of societies in which players play different actions. Within this framework the number of societies is again equal to the number of classes. Thus, any two players of the same class play the same strategy.

Example 1 raises some questions about how to interpret Corollary 2. In particular, Corollary 2 suggests that roles can be used to imply the existence of a pure strategy Bayesian Nash equilibrium in the game of Example 1. This is a conclusion that is difficult to justify. A Bayesian Nash equilibrium requires the leader to play each strategy with probability one half and the existence of roles really seems to do nothing to purify this strategy.

To try and overcome such problems we use the concept of ex-post information proofness as introduced by Kalai (2000). A composition profile c is said to be ε *information proof* if for all $i \in N$

$$u_i(c) \geq u_i(a^k, t_i, c_{-i}) - \varepsilon$$

for all $a^k \in A$. A strategy profile σ is said to be a *Bayesian Nash ε information proof equilibrium* if it yields an ε information proof composition profile with probability one. If a strategy is a Bayesian Nash information proof equilibrium then, as discussed further by Kalai (2000), no player would wish to change their action after knowing the types (roles) and the actions of the other players. This seems to be a highly desirable property. We note that the Bayesian Nash equilibrium of Example 1 is not information proof; once the action of the leader is revealed some of the citizens would wish to change their strategy.

In seeking to show the existence of a Bayesian Nash information proof equilibrium we assume a particular form to the probability distribution over types. Given a partition into classes $\{N_1, ..., N_Q\}$ and given a role profile r let $h(r, k, q)$ be the number of players in class q who have role r^k. We say that a probability distribution over roles f is *within class determined* if for any class q and for any two role profiles r and $\overline{r}$, if $f(r), f(\overline{r}) > 0$ then $h(r, k, q) = h(\overline{r}, k, q)$ for all classes q and for all $r^k \in \mathcal{R}$. Thus, irrespective of the role profile, the number of players in each class that will have each role is determined. For example, a husband and wife may have the choice of going out to work or doing the housework. If the distribution of roles is within class determined then it follows that one of them stays at home and does the housework while the other goes out to work. Chance will decide which person takes which role.

Consider a game $\Gamma(f)$ where f is within class determined and within class anonymous and consider a Bayesian Nash ε-equilibrium with social conformity m of that game. The equilibrium m is such that every player within the same class has the same probability of being assigned each role. Further, two players of the same class assigned the same role will play the same action. Finally, the number of players in each class who will be assigned each role is known with certainty. It seems intuitive that we could interpret m as an equilibrium with social conformity; certainly the players appear to be conforming to some standard of behavior and it also appears that players could learn by imitation in such a society.

We note that a probability distribution over the set of role profiles f that is within class determined suggest a high level of coordination amongst members of a society. It would appear to suggest the presence of some central coordinating body. This is, however, not unrealistic as Example

3 below illustrates. We should also note that players can receive higher payoffs from a Bayesian Nash equilibrium of game $\Gamma(f)$, if f is within class determined, than they could from game Γ. This is similar to the way in which payoffs form a correlated equilibria can exceed those of Nash equilibrium in the same game. This suggests a reason why roles may endogenously evolve.

We now state our final result. We note that the values of δ_I, δ_p and δ_C all figure in the bound on rationality provided in Theorem 3. This contrasts with Theorems 1 and 2. We also note that we assume Γ is a game of perfect information; this could easily be relaxed subject to a relaxation on the level of information proofness.

Theorem 3: Let Γ be any $(\delta_I, \delta_P, \delta_C, Q)$-class Bayesian game with perfect information. Let ε be a positive real number where $\varepsilon \geq 2(2\delta_I + \delta_P + 2\delta_C)$. There exists a partition $\mathcal{N}$ of the player set into Q classes and a Bayesian game with endogenous roles $\Gamma(f)$ (where f is within class anonymous and within class determined) such that $\Gamma(f)$ has a Bayesian Nash ε information proof equilibrium m in pure strategies with social conformity.

Proof: From Theorem 1 the game Γ has a Bayesian Nash $2(\delta_I + \delta_C)$ equilibrium in pure strategies m. Given that game Γ is a $(\delta_I, \delta_P, \delta_C, Q)$-class game let $\mathcal{N} = \{N_1, ..., N_Q\}$ be a partition of N into classes. For any strategy profile s let $h(s, k, c)$ be the number of players i such that $i \in N_q$ and $s_i = s^k$.

Consider any strategy profile $\overline{m}$ such that $h(\overline{m}, k, q) = h(m, k, q)$. Pick an arbitrary player $i \in N_q$. Suppose that $\overline{m}_i = s^k$. There must exist some player $j \in N_q$ such that,

$$U_j(m_j, m_{-j}) \geq U_j(s^k, m_{-j}) - 2(\delta_I + \delta_C)$$

for all $s^k \in \Delta(\mathcal{S})$. Thus, given that $\mathcal{N}$ is a δ_I-interaction substitute partition,

$$U_j(m_j, \overline{m}_{-j}) \geq U_j(s^k, \overline{m}_{-j}) - 2(2\delta_I + \delta_C)$$

for all $s^k \in \Delta(\mathcal{S})$. Let $\overline{\overline{m}}$ be such that $\overline{\overline{m}}_j = \overline{m}_i$ and $\overline{\overline{m}}_l = \overline{m}_l$ for all $l \neq j$. We note that

$$U_j(m_j, \overline{\overline{m}}_{-j}) \geq U_j(s^k, \overline{\overline{m}}_{-j}) - 2(2\delta_I + \delta_C)$$

for all $s^k \in \Delta(\mathcal{S})$.

It follows, given $\mathcal{N}$ is a δ_P-individual substitute partition, that,

$$\left| U_j(s^k, \overline{\overline{m}}_{-j}) - U_i(s^k, \overline{\overline{m}}_{-i}) \right| \leq \delta_P$$

for all $s^k \in \Delta(\mathcal{S})$. Further, given that $\mathcal{N}$ is a δ_C-strategy switching partition,

$$\left| U_i(s^k, \overline{\overline{m}}_{-i}) - U_i(s^k, \overline{m}_{-i}) \right| \leq \delta_C$$

for all $s^k \in \Delta(\mathcal{S})$. Thus,

$$U_i(\overline{m}_i, \overline{m}_{-i}) \geq U_i(s^k, \overline{m}_{-i}) - 2(2\delta_I + \delta_P + 2\delta_C)$$

for all $s^k \in \Delta(\mathcal{S})$.

The statement of the theorem is now more or less immediate. Let f be such that for any role profile $r \in R$, if $f(r) > 0$ then $h(r,k,q) = h(m,k,q)$. Then let everybody in the population have the pure strategy (for Bayesian game with endogenous roles $\Gamma(f)$) which says 'play action a^l if role r^l'.■

We finish this section with an example.

Example 3:

This example concerns driving on roadways. The players choose between two actions - 'give way' (G) or 'not give way' (N). Players from a large population are randomly matched to play the stage game with payoff matrix

		player 2	
		G	N
player 1	G	$0,0$	$0,2$
	N	$2,0$	$-2,-2$

Players are not able to distinguish amongst each other. Thus they must play the same strategy against any opponent they meet. A player's payoff is the sum of payoffs received from playing each stage game.

Without any roles the only outcome we would expect is the unique Bayesian Nash equilibrium in which each player randomizes in giving way half of the time and not giving way otherwise. The result is a lack of coordination with 'crashes' happening on one in four occasions.

Suppose there exists two roles - 'to be on a major road' and 'to be on a minor road'. On each meeting one player is selected to be on a major road and one to be on a minor road. Each player can play the pure strategy 'if on a minor road give way and if on a major road do not give way'. Using these roles we see conformity in pure strategies amongst the whole population.♦

This example highlights many issues. First, players are able to realize average payoffs of 1 which would not be possible without roles. This in turn suggests a reason why a set of roles, or a highway code, would evolve. We note, however, that for such a set of roles to evolve we would ultimately appear to need some central coordinating body to classify roads as major and minor. Finally, we note that the probability distribution over role profiles in this example is within class determined; there is always one person on a minor road and one on a major road.

7 Relationships to the literature

A large related literature addresses the possible motivations for players to use mixed strategies (a recent paper on this topic is Govindan, Reny and Robson 2002). The central issue is whether a mixed strategy equilibrium can be seen as approximately equivalent to a pure strategy equilibrium. This is plausible because imperfect information, and the resultant exogenous uncertainty, make explicit randomization unnecessary. Aumann at el. (1983) provide sufficient conditions, on the exogenous uncertainty, such that any mixed strategy vector can be approximately purified. Harsanyi (1973) argues that a game with perfect information should be considered as an idealization of nearby games in which there is a small amount of payoff uncertainty. Harsanyi (1973) shows that such uncertainty implies any mixed strategy equilibrium can be approximately purified.

This literature would suggest that for a wide class of games and for any mixed strategy equilibrium of such games there exists an approximately equivalent pure strategy equilibrium. This may appear to generalize the results in this paper as our results suggest there are games for which an approximate equilibrium in pure strategies does not exist. Further, this remains the case even when we allow players to introduce some exogenous uncertainty. The literature, however, is concerned with questioning the motivations of rational players while the focus of this paper is one of questioning the bounded rationality of using pure strategies. The implications of this different approach are felt in both the assumptions of social conformity and that the number of roles is no larger than the number of actions. For example, it is typical of the literature, as in Aumann et al (1983) and Harsanyi (1973), to assume a continuum of types. It is difficult to envisage a boundedly rational player being able to condition actions on a continuum of types.

The model of this paper is also related to the literature on correlated equilibria (for an introduction see Myerson 1997). Correlated equilibrium are motivated as reflecting the incentives on players to coordinate their actions. For example, in a Battle of the sexes game the two players can coordinate their actions by utilizing exogenous signals. Such signals may reflect pre-play communication or readily observable 'sunspots'. Theorem 3 could be interpreted as showing the existence of an approximate correlated equilibrium with conformity. As above, however, we can point to a slight difference in emphasis between this paper and the literature. In particular, the literature on correlated equilibria is motivated by considerations of how rational players can coordinate their actions. Roles, as introduced in this paper, are motivated by considerations of how boundedly rational players may be able to approximate rational behavior through conformity.

A further related literature concerns the evolution of institutions (see, for example, Durlauf and Young 2001, Young 2001 and references therein). This literature addresses the question of how conventions or institutions can evolve, through individual interactions, to create coordination on a large scale. Such a literature helps in understanding how roles could become endogeneised in the way we assume in this paper.

We conclude this section by relating the results of this paper to those of Wooders, Cartwright and Selten (2001) and Cartwright and Wooders (2002). Wooders, et al (2001) take as given a non-cooperative pregame, consisting of three elements – a space of attributes, , a set of pure strategies S, and a payoff function h. A component of attribute space is a complete description of the possible characteristics of a player. The set of pure strategies S is assumed to be finite. The payoff function h determines a payoff function for any player in any game derived from the pregame. Given any finite player set N and and an attribute function α ascribing an attribute to each player, the pregame induces a game $\Gamma(N, \alpha)$ on the population (N, α) determined by N and α. A set of players, all with attributes in some convex subset of attribute space and all playing the same pure strategy is interpreted as a society. Roughly, the main result of WCS demonstrates that for any $\varepsilon > 0$, for any sufficiently large game induced by the pregame satisfying a certain 'large game property,' there exists an ε-Nash equilibrium in pure strategies and this equilibrium induces a partition of the player set into at most $J(\varepsilon)K$ societies, where $J(\varepsilon)$ is fixed independently of the size of the player set. This result establishes properties of a set of games where both the use of pure strategies and social conformity can be consistent with rational behavior.

The approach of Wooders et al. (2001) has the virtue of providing a general class of games where the use of pure strategies and social conformity is boundedly rational. This contrasts with the results of this paper which, while applying to any game, only provide a bound on the rationality of such behavior given certain properties of the game. In Cartwright and Wooders (2002) we connect the concepts of a pregame induced by a large game property and that of a (δ, Q)-class game. Doing so allows us to apply the results of this paper. In particular, it allows us provide a general class of game that are (δ, Q)-class games for arbitrarily small δ and fixed number of societies Q. By applying the results of this paper we are also able to significantly generalise the results of WCS. For example, we demonstrate conformity in pure strategies even if all players are identical. Further, we relax the requirement of a global interaction assumption to one of local interaction.

8 Conclusion

This paper introduces the concept of approximate substitutes in non-cooperative games. Doing so allows us to put a bound on the rationality, or irrationality, of using pure strategies and of social conformity. We us a definition of a society which allows players within the same society to perform different actions. Thus, players who are conforming to some norm may perform different actions. This is possible through imperfect information and the existence of roles. In particular, players can make action choice conditional on their role and roles are assigned to players randomly. Thus, players in the same society can play the same strategy and yet perform different actions. To retain a meaningful notion of society we impose two restrictions on how roles are allocated; first, while roles are assigned randomly the number of players in a society who will have each role is not random; second, any player in the same society must have the same probability of being assigned each role. We argue that players within the same society can be seen to conform to some norm or convention. In research in progress, we relate our results to the experimental research of Friedman (1996) and Van Huyck, Battalio and Rankin (1997).

9 Appendix

9.1 Example A

The player set is given by $N = \{1, ..., 10\}$. The *attribute* of player i, denoted $\alpha(i)$, is given by the pair $(i, 11 - i) = (\alpha_1(i), \alpha_2(i))$ for all $i \in N$. Each player has a unit of time to devote to production. A player chooses between two actions - 'to produce good A' or 'to produce good B'. The attribute of a player is a measure of how much of the good she is capable of producing in a unit of time. In particular, given any strategy vector σ the amount of the good A that will be produced is given by,

$$a(\sigma) = \sum_{i \in N} \sigma_{iA} \alpha_1(i).$$

The amount of good B produced is given by,

$$b(\sigma) = \sum_{i \in N} \sigma_{iB} \alpha_2(i).$$

Each player $i \in N$ has the same payoff function. The payoff function for any player $i \in N$ is given by,

$$U_i(s^k, \sigma_{-i}) = \frac{a(\sigma_{-i}) + b(\sigma_{-i}) + s^{kA}\alpha_1(i) + s^{kB}\alpha_2(i)}{|N|}$$

for any $s^k \in \Delta(\mathcal{S})$ and any strategy vector σ. That is, a player's payoff is simply the per capita amount of both goods produced. This completes the definition of the game.

Given any partition $\mathcal{N}$ of the player set into Q classes we can calculate the minimum δ for which the game above is a (δ, Q)-class game. We illustrate with three partitions.

First, suppose that $Q = 1$; that is, there exists a unique class. Consider the two strategy vectors $\sigma^1, \sigma^2 \in \Sigma$ such that $\sigma^1 = (B, B, B, B, B, A, A, A, A, A)$ and $\sigma^2 = (A, A, A, A, A, B, B, B, B, B)$. For example, given strategy σ^1 player 1 produces good B with probability 1. Given that $a(\sigma^1) = b(\sigma^1) = 40$ and $a(\sigma^2) = b(\sigma^2) = 15$ we have that,

$$\left|U_5(\sigma^1) - U_5(\sigma^1_5, \sigma^2_{-5})\right| = 4.9.$$

This shows that for $Q = 1$ there exists no δ-interaction substitute partition for any $\delta \leq 4.9$. It is clear that there does exist a δ-interaction substitute partition for any $\delta > 4.9$. It is easily checked that this also implies Γ is a $(\delta, 1)$-class game for any $\delta > 4.9$.

Moving to the other extreme suppose that $Q = 10$; that is, each player is there own class. We have already noted there will exist a 0-substitute partition. By re-using the two strategy vectors σ^1 and σ^2 above we see, however, that Γ is not a $(\delta, 10)$-class Bayesian game for any $\delta \leq 4.9$. It is a $(\delta, 10)$-class for any $\delta > 4.9$.

Suppose, finally that we set $Q = 3$ where $N_1 = \{1, 2, 3, 4\}$ and $N_2 = \{5, 6\}$. To find a δ-interaction substitute partition we compare a strategy vector σ^1 with σ^2 where, $\sigma^1 = (B, B, A, A, B, A, B, B, A, A)$ and $\sigma^2 =$

$(A, A, B, B, A, B, A, A, B, B)$. We note that $a(\sigma^1) = b(\sigma^1) = 32$ and $a(\sigma^2) = b(\sigma^2) = 23$. Thus,

$$\left|U_5(\sigma^1) - U_5(\sigma_i^1, \sigma_{-i}^2)\right| = 1.7.$$

It follows that there is a δ-substitute partition for any value of $\delta > 1.7$. Compare two strategy vectors $\sigma^1 = (B, B, B, B, B, A, A, A, A, A)$ and $\sigma^2 = (A, B, B, B, B, B, A, A, A, B)$. We can calculate that,

$$\left|U_5(\sigma^1) - U_5(\sigma_i^1, \sigma_{-i}^2)\right| = 1.9.$$

This leads us to claim that Γ is a $(\delta, 3)$-class Bayesian game for any $\delta > 1.9$.

For each value of Q, by looking at all partitions of the player set into Q subsets we can calculate the minimum value of δ such that the game is a (δ, Q)-class game. This value of δ, as a function of Q, is plotted in Figure 1 below with more details given in Table 1. The minimum value of δ, for which there exists a δ-substitute partition or δ-strategy switching partition can also be calculated as a function of Q. These are also plotted in figure 1 with more details given in Tables 2 and 3.

In Figure 1 we see the trade-off between a large and small value of Q. If Q is large then there exists δ-substitute partition for small δ. These partitions, however, do not imply the game is a (δ, Q)-class game for small δ. Each class has relatively few players and so payoffs are not near-invariant when one player from each class changes strategy. By contrast, if Q is small then the δ for which there exists a δ-substitute partition is larger. This simply reflects the fact that we are grouping together players with more diverse attributes. Given, however, that the number of players in a class is now relatively large any δ-substitute partition is likely to be a (δ, Q)-class game.

It is interesting to highlight that the Q member partition which implies the minimum δ for which there is a δ-substitute partition may differ from the Q member partition consistent with the minimum δ for which the game is a (δ, Q)-class game. For example, set $Q = 9$. The partition in which players 5 and 6 are put in the same subset is a 0.2-substitute partition. By using this same partition we can say that the game is a $(4.9, 9)$-class game but cannot put a lower bound on δ than 4.9. Consider the partition in which players 1 and 10 are put together. This partition is only a 1.8-substitute partition, however, this partition demonstrates that the game is a $(4.1, 9)$-class game.

Table 1:

Q	partition	values of δ such that a (δ, Q)-class game
1	$\{1,2,3,4,5,6,7,8,9,10\}$	> 4.9
2	$\{1,2,3,4,5\}$ $\{6,7,8,9,10\}$	> 2.4
3	$\{1,2,3,4\},\{5,6\},$ $\{7,8,9,10\}$	≥ 1.9
4	$\{1,2,3,4\},\{5\},$ $\{6\},\{7,8,9,10\}$	≥ 2.0
5	$\{1,2,3\},\{4,7\},\{5\},$ $\{6\},\{8,9,10\}$	≥ 2.3
6	$\{1,2,3\},\{4\},\{5\},$ $\{6\},\{7\},\{8,9,10\}$	≥ 2.6
7	$\{1,10\},\{2,3\},\{4\},$ $\{5\},\{6\},\{7\},\{8,9\}$	≥ 3.1
8	$\{1,10\},\{2,9\},\{3\},$ $\{4\},\{5\},\{6\},\{7\},\{8\}$	≥ 3.4
9	$\{1,10\},\{2\},\{3\},\{4\},$ $\{5\},\{6\},\{7\},\{8\},\{9\}$	≥ 4.1
10	$\{1\},\{2\},\{3\},\{4\},\{5\},$ $\{6\},\{7\},\{8\},\{9\},\{10\}$	≥ 4.9

Table 2:

Q	partition	values of δ such that a δ-substitute partition
1	$\{1,2,3,4,5,6,7,8,9,10\}$	> 4.9
2	$\{1,2,3,4,5\}$ $\{6,7,8,9,10\}$	> 2.4
3	$\{1,2,3\},\{4,5,6,7\},$ $\{8,9,10\}$	> 1.6
4	$\{1,2\},\{3,4,5\},$ $\{6,7,8\},\{9,10\}$	> 1.2
5	$\{1,2\},\{3,4\},\{5,6\},$ $\{7,8\},\{9,10\}$	> 1.0
6	$\{1,2\},\{3\},\{4,5,6\},$ $\{7,8\},\{9\},\{10\}$	> 0.8
7	$\{1\},\{2\},\{3\},\{4,5,6\},$ $\{7,8\},\{9\},\{10\}$	> 0.6
8	$\{1\},\{2\},\{3\},\{4,5\},$ $\{6,7\},\{8\},\{9\},\{10\}$	> 0.4
9	$\{1\},\{2\},\{3\},\{4\},\{10\}$ $\{5,6\},\{7\},\{8\},\{9\}$	> 0.2
10	$\{1\},\{2\},\{3\},\{4\},\{5\},$ $\{6\},\{7\},\{8\},\{9\},\{10\}$	≥ 0

Table 3:

Q	partition	values of δ such that a δ-strategy switching partition
1	$\{1,2,3,4,5,6,7,8,9,10\}$	≥ 0.9
2	$\{1,2,3,4,5,7,8,9,10\},\{6\}$	≥ 1
3	$\{1,2,3,4,7,8,9,10\}$, $\{5\},\{6\}$	≥ 1.1
4	$\{1,2,3,8,9,10\}$, $\{4,7\},\{5\},\{6\}$	≥ 1.4
5	$\{12,3,8,9,10\}$, $\{4\},\{5\},\{6\},\{7\}$	≥ 1.7
6	$\{12,9,10\},\{3,8\}$, $\{4\},\{5\},\{6\},\{7\}$	≥ 2.2
7	$\{1,2,9,10\},\{3\}$, $\{4\},\{5\},\{6\},\{7\},\{8\}$	≥ 2.7
8	$\{1,10\},\{2,9\},\{3\}$, $\{4\},\{5\},\{6\},\{7\},\{8\}$	≥ 3.4
9	$\{1,10\},\{2\},\{3\},\{4\}$, $\{5\},\{6\},\{7\},\{8\},\{9\}$	≥ 4.1
10	$\{1\},\{2\},\{3\},\{4\},\{5\}$, $\{6\},\{7\},\{8\},\{9\},\{10\}$	≥ 4.9

9.2 Proof of Theorem 1

First, we introduce some notation. Given a strategy vector $\sigma = (\sigma_1, .., \sigma_n)$ (where $\sigma_i = (\sigma_{i1}, ..., \sigma_{iK}) \in \Delta^K$ for $i = 1, ..., n$) let $\mathcal{M}(\sigma)$ denote the set of strategy vectors $m = (m_1, .., m_n)$ such that for all $i = 1, ..., N$,

1. $support(m_i) \subset support(\sigma_i)$ for all $i \in N$ and,
2. m_i is degenerate.

Informally, given a strategy vector σ the strategy vector $m \in \mathcal{M}(\sigma)$ if, for all i, strategy m_i is such that player i plays some pure strategy $s^k \in support(\sigma_i)$ with probability one.

Our main result makes use of the following Lemma from Wooders, Cartwright and Selten (2001),

Lemma 1 (Wooders, Cartwright, Selten): For any strategy vector $\sigma = (\sigma_1, ..., \sigma_n)$ and for any vector $\overline{g} \in \mathbb{Z}_+^K$ such that $\sum_i \sigma_i \geq \overline{g}$, there exists a vector $m = (m_1, ..., m_n) \in \mathcal{M}(\sigma)$ such that:

$$\sum_i m_i \geq \overline{g}.$$

We extend Lemma 1. First, we introduce further notation. Given real number h let $\lfloor h \rfloor$ denote the nearest integer less than or equal to h and $\lceil h \rceil$ the nearest integer greater than or equal to h (i.e. $\lfloor 9.5 \rfloor = 9$ and $\lceil 9.5 \rceil = 10$ etc.). Given vector h denote by $\lfloor h \rfloor$ the vector such that $\lfloor h \rfloor_k = \lfloor h_k \rfloor$ for all k with a similar definition for $\lceil h \rceil$.

Lemma 2: For any strategy vector $\sigma = (\sigma_1, ..., \sigma_n)$ there exists a strategy vector $m = (m_1, ..., m_n) \in \mathcal{M}(\sigma)$ such that:

$$\left\lceil \sum_{i=1}^{n} \sigma_i \right\rceil \geq \sum_{i=1}^{n} m_i \geq \left\lfloor \sum_{i=1}^{n} \sigma_i \right\rfloor.$$

Proof: Denote by $\mathcal{M}^*(\sigma)$ the set of vectors $m = (m_1, ..., m_n) \in \mathcal{M}(\sigma)$ such that $\sum_i m_i \geq \lfloor \sum_i \sigma_i \rfloor$. By Lemma 1 this set is non-empty. Proving the Lemma thus amounts to showing that there exists a vector $m \in \mathcal{M}^*(\sigma)$ such that $\lceil \sum_i \sigma_i \rceil \geq \sum_i m_i$. Suppose not. Then, for every vector $m \in \mathcal{M}^*(\sigma)$

there exists some strategy $s^k \in S$ such that $\sum_i m_{ik} > \lceil \sum_i \sigma_{ik} \rceil$. Choose a vector $m^0 \in \mathcal{M}^*(\sigma)$ such that

$$C \equiv \sum_{s^k : \sum_i m_{ik} > \lceil \sum_i \sigma_{ik} \rceil} \left(\sum_i m_{ik} - \left\lceil \sum_i \sigma_{ik} \right\rceil \right)$$

is minimized. That is, m^0 comes as close as any vector to satisfying the statement of the Lemma. Denote by $s^{\widehat{k}}$ a pure strategy such that

$$\sum_{i=1}^n m_{i\widehat{k}} > \left\lceil \sum_{i=1}^n \sigma_{i\widehat{k}} \right\rceil .$$

We then introduce the following sets W^t and L^t, $t = 0, 1, 2, ...,$

$$\begin{aligned} W^0 &= \{i : m_{i\widehat{k}} = 1\} \\ L^t &= \{s^k : \sigma_{ik} > 0 \text{ for some } i \in W^t\} \text{ for } t \geq 0 \\ W^t &= \{i : m_{ik} = 1 \text{ for some } s^k \in L^t\} \text{ for } t > 0. \end{aligned}$$

For some t^*, $W^{t^*} = W^{t^*+1} \equiv W$ and $L^{t^*} = L^{t^*+1} \equiv L$. The construction of W^t and L^t imply that for any $s^{k^*} \in L^{t^*}$ there must exist a set of players $\{i_0, i_1, ..., i_{t^*}\} \in W^t$ and set of strategies $\{s^{k_1}, ..., s^{k_t}\}$ such that,

$$\begin{aligned} m^0_{i_0\widehat{k}} &= 1 \text{ and } \sigma_{i_0k_1} > 0, \\ m^0_{i_rk_r} &= 1 \text{ and } \sigma_{i_rk_{r+1}} > 0 \text{ for all } r = 1, .., t-1, \\ m^0_{i_tk_t} &= 1 \text{ and } \sigma_{i_tk^*} > 0, \end{aligned}$$

where we allow the possibility that $t = 0, 1$. Suppose there exists a $k^* \in L$ such that

$$\sum_{i=1}^n m_{ik^*} \leq \sum_{i=1}^n \sigma_{ik^*}.$$

Given the chain of players $\{i_0, i_1, ..., i_{t^*}\} \in W$ given above, consider the vector m^* constructed as follows,

$$\begin{aligned} m^*_{i_0\widehat{k}} &= 0 \text{ and } m^*_{i_0k_1} = 1, \\ m^*_{i_rk_r} &= 0 \text{ and } m^*_{i_rk_{r+1}} = 1 \text{ for all } r = 1, ..., t^*-1, \\ m^*_{i_{t^*}k_t} &= 0 \text{ and } m^*_{i_{t^*}k^*} = 1, \\ m^*_{ik} &= m^0_{ik} \text{ for all other } s^k \in S \text{ and } i \in N. \end{aligned}$$

It is easily checked that the vector $m^* \in \mathcal{M}(\sigma)$ leads to the desired contradiction by reducing by one the value C. We note, however, that

$$\sum_{i=1}^{n} \sum_{s^k \in L} m_{ik} = |W| = \sum_{i \in W} \sum_{s^k \in L} \sigma_{ik}.$$

Thus, if

$$\sum_{i=1}^{n} m_{i\widehat{k}} > \sum_{i=1}^{n} \sigma_{i\widehat{k}} \geq \sum_{i \in W} \sigma_{i\widehat{k}}$$

there must exist some $s^{k^*} \in L$ such that

$$\sum_{i=1}^{n} m_{ik^*} \leq \sum_{i \in W} \sigma_{ik^*} \leq \sum_{i=1}^{n} \sigma_{ik^*}$$

giving the desired contradiction.■

We require one final preliminary result.

Lemma 3: Given any two strategy vectors m and σ where m is degenerate and where

$$\left| \sum_{i=1}^{n} m_{ik} - \sum_{i=1}^{n} \sigma_{ik} \right| \leq 1$$

for all k, there exists a strategy vector $\overline{\sigma}$ such that,

$$\sum_{i=1}^{n} \overline{\sigma}_{ik} = \sum_{i=1}^{n} \sigma_{ik}$$

and

$$\sum_{i=1}^{n} |\overline{\sigma}_{ik} - m_{ik}| \leq 1$$

for all k.
Proof: Given such a σ and m we proceed by constructing an appropriate $\overline{\sigma}$.

Let K^- denote the set of pure strategies for which $\sum_{i=1}^{n} m_{ik} = \lfloor \sum_{i=1}^{n} \sigma_{ik} \rfloor$

and let K^+ denote the set of strategies for which $\sum_{i=1}^{n} m_{ik} = \lceil \sum_{i=1}^{n} \sigma_{ik} \rceil$. We note that for every s^k either $s^k \in K^-$ or $s^k \in K^+$. For each s^k let,

$$A(k) = \sum_{i=1}^{n} \sigma_{ik} - \left\lfloor \sum_{i=1}^{n} \sigma_{ik} \right\rfloor.$$

and let,

$$A^+ = \sum_{k:s^k \in K^+} A(k).$$

Provisionally, set $\overline{\sigma}_i = m_i$ for all i. Then, for each $s^{\overline{k}} \in K^+$ identify a player $i_{\overline{k}}$ such that $m_{i_{\overline{k}}} = 1$. For each $s^{\overline{k}} \in K^+$ re-allocate player $i_{\overline{k}}$ the strategy defined, for each k, as follows,[8]

$$\overline{\sigma}_{i_{\overline{k}}k} = \begin{cases} A(k) \text{ if } k = \overline{k} \\ 0 \text{ if } s^k \in K^+ \text{ and } k \neq \overline{k} \\ A(k)\frac{(1-A(\overline{k}))}{|K^+|-A^+} \text{ otherwise.} \end{cases}$$

We conjecture that this strategy vector $\overline{\sigma}$ satisfies the required conditions. First, we should check that $\overline{\sigma}$ as defined above is indeed a strategy vector. For those players i for whom $\overline{\sigma}_i = m_i$ there is nothing to check. Consider a player $i_{\overline{k}}$ for some $s^{\overline{k}} \in K^+$. We begin by noting that $1 \geq \overline{\sigma}_{i_{\overline{k}}k} \geq 0$ for all k as $|K^+| - A^+ \geq (1 - A(\overline{k}))$. We also note that

$$\begin{aligned} \sum_k A(k) &= \sum_k \left(\sum_{i=1}^{n} \sigma_{ik} - \left\lfloor \sum_{i=1}^{n} \sigma_{ik} \right\rfloor \right) \\ &= n - \sum_k \left\lfloor \sum_{i=1}^{n} \sigma_{ik} \right\rfloor \\ &= n - (n - |K^+|) \\ &= |K^+|. \end{aligned}$$

Thus,

$$\sum_{k:s^k \in S^-} A(k) = |K^+| - A^+.$$

[8] If $|K^+| = A^+$ then set $\overline{\sigma}_i = m_i$ for all i.

Therefore,

$$\begin{aligned}\sum_k \overline{\sigma}_{i_{\overline{k}}k} &= A(\overline{k}) + \left(\left|K^+\right| - A^+\right)\frac{\left(1 - A(\overline{k})\right)}{|K^+| - A^+} \\ &= 1.\end{aligned}$$

Thus, $\overline{\sigma}$ is a strategy vector. We note that for all $s^k \in K^-$,

$$\begin{aligned}\sum_{i=1}^n \overline{\sigma}_{ik} &= \left\lfloor \sum_{i=1}^n \sigma_{ik} \right\rfloor + \frac{A(k)}{|K^+| - A^+} \sum_{s^{\overline{k}} \in K^+} \left(1 - A(\overline{k})\right) \\ &= \sum_{i=1}^n \sigma_{ik}.\end{aligned}$$

Clearly, for all $s^k \in K^+$ we have that $\sum_{i=1}^n \overline{\sigma}_{ik} = \sum_{i=1}^n \sigma_{ik}$.

Finally, for all $s^k \in K^-$

$$\sum_{i=1}^n |\overline{\sigma}_{ik} - m_{ik}| = \frac{A(k)}{|K^+| - A^+} \sum_{s_{\overline{k}} \in K^+} \left(1 - A(\overline{k})\right) \leq 1,$$

and for all $s^k \in K^+$

$$\sum_{i=1}^n |\overline{\sigma}_{ik} - m_{ik}| = (1 - A(k)) \leq 1.$$

This completes the proof.■

References

[1] Aumann, R.J., Y. Katznelson, R. Radner, R.W. Rosenthal,and B. Weiss (1983) “Approximate purification of mixed strategies,” *Mathematics of Operations Research* 8, 327-341.

[2] Blume, L.E. (1993) “The statistical mechanics of strategic interaction,” *Games and Economic Behaviour* 5:387-424.

[3] Cartwright. E.J. and M.H. Wooders (2002) “Bounded rationality in games with many players” (typescript).

[4] Durlauf, S. and H.P. Young (2001) *Social Dynamics*, MIT Press.

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

[6] Follmer, H. (1974) "Random economies with many interacting agents," *Journal of Mathematical Economics* 1, 51-62.

[7] Friedman, D., (1996) "Equilibrium in evolutionary games: Some experimental results," *Economic Journal* 106: 1-25.

[8] Fudenberg, D. and D.K. Levine (1998) *The Theory of Learning in Games*, Cambridge: MIT press.

[9] Gale, D. and R. Rosenthal, 1999, "Experimentation, Imitation, and Strategic Stability," *Journal of Economic Theory* 84, 1-40.

[10] Hargreaves-Heap, S. and Y. Varoufakis (2002) "Some experimental evidence on the evolution of discrimination, co-operation and perceptions of fairness", *The Economic Journal*, 112, 679-703.

[11] Hildenbrand, W. (1971) "Random preferences and equilibrium analysis," *Journal of Economic Theory* 3, 414-429.

[12] Kalai, E. (2000) "Private information in large games," Northwestern University Discussion Paper 1398.

[13] Kagel, J.H. and A.E. Roth, 1995, The handbook of experimental economics, Princeton University Press.

[14] Khan, A., K.P. Rath and Y.N. Sun (1997) "On the existence of pure strategy equilibria with a continuum of players," *Journal of Economic Theory* 76:13-46.

[15] Kovalenkov, A. and M. Wooders (1996) "Epsilon cores of games with limited side payments; Nonemptiness and equal treatment," *Games and Economic Behavior* (to appear).

[16] Kovalenkov, A. and M. Wooders (1997) "An exact bound on epsilon for non-emptiness of the epsilon-core of an arbitrary game with side payments," Autonoma University of Barcelona Working Paper 393.97 revised, *Mathematics of Operations Research* (to appear).

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

[18] Myerson, R.B. (1997) *Game Theory: Analysis of Conflict*, Harvard University Press.

[19] Pascoa, M. (1998) "Nash equilibrium and the law of large numbers," *International Journal of Game Theory* 27: 83-92.

[20] Pascoa, M. (1993) "Approximate equlibrium in pure strategies for nonatomic games," *Journal of Mathematical Economics* 22: 223-241.

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

[22] Reny, P.J. (1999) "On the existence of pure and mixed strategy Nash equilibria in discontinuous games," *Econometrica* 67: 1029-1056.

[23] Scharfstein, D.S. and J.C. Stein (1990) "Herd behaviour and investment," *The American Economic Review*, vol 80, no. 3, 465-479.

[24] Schleifer, A., (2000) *Inefficient Markets, an introduction to behavioural finance*, Oxford University Press.

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

[26] Sugden, R., (1995) "A theory of focal points," *Economic Journal* 105: 533-550.

[27] Van Huyck, J., R.C. Battalio and F.W. Rankin (1997) "On the origin of convention: evidence from coordination games", *Economic Journal*, 107, 576-597.

[28] Walker, M. and J. Wooders (2001) "Minimax play at Wimbledon," *American Economic Review*, 91: 1521-1538.

[29] Wooders, M.H. (1979) "Asymptotic cores and asymptotic balancedness of large replica games" (Stony Brook Working Paper No. 215, Revised July 1980).

[30] Wooders, M. (1983) "The epsilon core of a large replica game," *Journal of Mathematical Economics* 11, 277-300.

[31] Wooders, M. (1994) "Equivalence of games and markets," *Econometrica* 62, 1141-1160.

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

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