

Social conformity and approximate purification in games with incomplete information.

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

Abstract

Interpret a set of players with similar attributes who are all playing the same strategy 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 a framework of approximate substitutes in non-cooperative games we are able to put a bound on the ‘inefficiency’ of such social conformity for arbitrary games. This is then applied, using non-cooperative pre-games, to show that for sufficiently large games there exists an approximate Nash equilibrium in pure strategies for which the population is partitioned into a relatively small number of societies.

1 Introduction

A society or culture is a group of individuals who have commonalities of language, social and behavioral norms, and customs. Social learning consists, at least in part, in learning the norms and behavior patterns of the society into which one is born and in those other societies which one may join – our professional associations, our workplace, and our community, for example. Social learning may also include learning a set of skills from others that will enable one to fit into a society. The society in question may be broad as “Western civilization” or Canada, or as small as the Econometric Society. A fundamental question is whether this can be consistent with self-interested behavior. Such consistency requires the existence of a Nash equilibrium where individuals within the same society play the same or similar strategies.

In an earlier paper (Wooders, Cartwright and Selten (2001)) we provide conditions under which such an equilibria will exist. To understand the motivation for the present paper we briefly summarize this work. We take as given a non-cooperative pregame of which a key element is a set Ω of attributes. A component of attribute space is a complete description of the possible characteristics of a player. As such, a player set N and a function α allocating an attribute to each player, induce, through the pregame, a game $\Gamma(N, \alpha)$. We say that a pregame satisfies the large game property if, for sufficiently large games induced by the pregame, (1) payoffs are only slightly altered by a small perturbation of the attributes of players, and (2) each player’s payoff is primarily a function of their own strategy and of the numbers of players of each attribute playing each strategy, relative to the size of the population. We interpret a set of players, all with attributes in some convex subset of attribute space and all playing the same pure strategy, as a society. Our main result shows, *that given any $\varepsilon > 0$, there is an integer $J(\varepsilon)$ and real number $\eta(\varepsilon)$, such that for any game of complete information which has at least $\eta(\varepsilon)$ players and is induced by a pregame satisfying the large game property, there exists an ε -equilibrium in pure strategies that induces a partition of the set of players into at most $J(\varepsilon)$ societies.*

The purpose of this paper is to extend that result which we do in essentially three different ways. These can be summarized as (1) to generalize to games of incomplete information, (2) to look at approximate purification and social conformity as a property of a game rather than a pregame, and (3) to consider alternative notions of social conformity. The explanation for the first of these extensions will become clear as we explain the latter two, which we do in turn.

In order to look at approximate purification and social conformity as a property of a game rather than as a property of a pregame we introduce the notion of approximate substitutes of a non-cooperative game. This is a counterpart to the notion of approximate substitutes in cooperative games (Kovalenkov and Wooders (2001)). Informally, two players are approximate substitutes if they, (1) have similar payoff functions, and (2) are such that if they ‘exchanged strategies’ a third player would be relatively indifferent to this exchange. By putting a bound δ on the similarity of players, in terms of these criteria, we talk of a δ -substitute partition as a partition of the player set into approximate classifications such that any two players with the same approximate classification are approximate (δ) substitutes. A game is then defined as a (δ, Q) classification game if there exist a δ -substitute partition that partitions the player set into Q classes. It should also be noted that we restrict attention throughout to games with a finite number of strategies.

Within this framework we derive our Theorem 1, namely, *if a game Γ is a (δ, Q) classification game then there exists an ε -equilibrium in pure strategies for any $\varepsilon \geq 2\delta$* . One key point to note, is that for any game Γ and any Q there exists some δ for which the game Γ is a (δ, Q) classification game. As such, for any game, possibly with incomplete information, we can put a bound on the ε for which there exists an ε -equilibrium in pure strategies. Alternatively, the value of 2δ could be interpreted as a bound on the ‘inefficiency’ of players using pure strategies as opposed to mixed strategies. We judge efficiency according to individual rationality.

The approximate substitute framework allows us to draw conclusions about arbitrary games. It is also useful to have some general examples of (δ, Q) classification games for arbitrary values of δ . As such, in the final section, we return to the pregame framework of our earlier paper, extended to the incomplete information case. 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 Theorem 1 in deriving an ε -purification result such that, *given any real number $\varepsilon > 0$ there exists a real number $\eta(\varepsilon) > 0$ such that any game induced by a pregame satisfying the large game property with more than $\eta(\varepsilon)$ players has an ε -equilibrium in pure strategies. Moreover, for any mixed strategy equilibrium there exists an ε -purification*. This result extends the ε -purification result of Wooders, Cartwright and Selten (2001). It contrasts from existing purification results in not assuming a continuum of players (cf. Schmeidler 1973, Aumann (1983), Mas-Colell 1984, Khan 1989, Pascoa 1993,1998, Khan et al. 1997, Araujo and Pascoa 2000).

These purification results illustrate the potential applications of both

frameworks of approximate substitutes in non-cooperative games and pregames of non-cooperative games. We apply these frameworks further in addressing the efficiency of social conformity. We define a society such that any two players belonging to a society have the same approximate classification and play the same strategy. An immediate consequence of Theorem 1 is that *if a game Γ is a (δ, Q) classification game then there exists an ε -equilibrium in pure strategies for any $\varepsilon \geq 2\delta$ which partitions the player set into no more than QK societies*, where K is the number of strategies. This essentially provides a benchmark result which we look to refine.

By way of motivation suppose that a society is playing a form of n person matching pennies, and, as such, in equilibrium half of the time ‘heads’ should be played and half of the time ‘tails’ should be played. If players only use pure strategies then in this example we would get two distinct societies between those players who play ‘heads’ and those who play ‘tails’. This would have been the conclusion from our earlier paper and in some instances would probably seem an appropriate distinction. In some instances, however, it may not be appropriate. For example, the game may be driving and the strategies are to give way or to not give way at road junctions. At any one instance we might expect half of the drivers to give way and half to not give way but, if this is because players are conforming to some highway code, we would clearly not want to think of the player set as being split into two distinct societies. Instead, we would say that players are merely taking different roles within the game and make *actions conditional on their roles*. Making strategy choice conditional on roles to symmetrize a game is standard in the evolutionary game theory literature (e.g. Selten (1980) and Young (2001)).

Clearly, if social conformity is efficient when any two players in a society must use the same action it will be efficient when any two players in the same society may potentially use different actions. The converse, however, need not be true. Thus, by relaxing the definition of a society we can hope to demonstrate that it is efficient to have a higher level of social conformity. More formally, let the inverse of the number of societies be a measure of the level of social conformity. Our first result shows that it can be (2δ) efficient to have a level of social conformity of $\frac{1}{QK}$. By relaxing the criteria by which we judge a society can we put, for the same criteria of efficiency, a larger bound on the level of social conformity?

The extension to incomplete information immediately implies, that despite the fact that every player within a society plays the same strategy, players within the same society may use different actions dependent upon their type. That is, players may take up different roles in their society as

determined by the type that ‘nature’ deals them. In the context of the example, it may be that men play ‘heads’ and women play ‘tails’ for example - a persons action being determined by their type which in this case is gender. Allowing imperfect information does not imply, however, that social conformity is any more efficient than in games of perfect information. The reasons for this are clear in that the probabilities with which the types of players are drawn by nature may not be appropriate for the equilibrium of the game. In the context of our example, a highway code based on gender may not be entirely satisfactory.

It clearly seems plausible that in some games players will try to find some signalling mechanism by which players can be assigned roles within the society; if gender is not an appropriate signal then try something else. As such we suppose that players can endogenise the set of types and their probability distribution over types. We refer to the set of endogenised types as roles. Thus, instead of being given a type by nature we say that players are assigned to roles within a society. Strategies can be made conditional on a player’s role. A society is then essentially such that any two players belonging to the same society have the same approximate type, play the same strategy, and agree on the number of people that should be allocated to each role within the society. In the context of our example, drivers would agree on a strategy, say “give way on minor roads”, but, furthermore, they would also agree on which roads should be classed as minor roads and which major roads. We are able to show, under such a framework, that *if a game Γ is a (δ, Q) classification game then there exists an ε -equilibrium in pure strategies for any $\varepsilon \geq 4\delta$ which partitions the player set into no more than Q societies*. Thus, if players have some way of endogeneising types, that is of allocating players to different roles within the society, then we can put a higher bound on the level of social conformity. Indeed, given our constraint that two players in the same society must be of the same approximate class, we get the highest level of social conformity to be expected.

A second way in which we consider refining the notion of a society is to assume that players can play mixed strategies. In most games players probably are restricted to pure strategies (such as driving) but this does not mean there are not games in which mixed strategies can be played. Allowing the use of mixed strategies, we show that *if a game Γ is a (δ, Q) classification game then there exists an ε -equilibrium strategy vector for any $\varepsilon \geq 4\delta$ which partitions the player set into less than Q societies*. That is, there exists an ε -equilibrium such that any two players with the same approximate class play the same mixed strategy. In the context of our earlier example this would equate to every player playing ‘heads’ with probability one half and

‘tails’ with probability one half.

Together this provides three different ways of looking at social conformity and the notion of a society, (1) all players within a society play the same pure strategy, (2) all players within a society play the same, possibly, mixed strategy, or (3) within a society there is some agreed upon framework within which players are allocated to roles and all players given the same role play the same action. We recall that for any game Γ and any Q there exists some δ for which the game Γ is a (δ, Q) classification game. As such, the results we derive in these three contexts apply to any game and for any level of social conformity as given by $\frac{1}{Q}$ or $\frac{1}{QK}$. One implication of this is that we can put a bound on the possible inefficiency of social conformity in an arbitrary game and for an arbitrary number of societies.

We finish by applying these social conformity results to games induced by a pregame satisfying the large game property. One derived result is that, *for any $\varepsilon > 0$ there exists real numbers $\eta(\varepsilon)$ and $Q(\varepsilon)$ such that any game Γ with at least $\eta(\varepsilon)$ players induced by a pregame satisfying the large game property has (1) a Bayesian Nash ε equilibrium which partitions the player set into $C \leq Q(\varepsilon)$ societies, and (2) a Bayesian Nash ε equilibrium in pure strategies which partitions the player set into $C \leq Q(\varepsilon)K$ societies.* It is important to note that the number of societies is bounded independently of the size of the population. Thus, the size of societies become arbitrarily large as the size of the population increases.

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 social conformity in games with approximate substitutes. Section 6 looks at large games and applies earlier results before we conclude in section 7.

2 A Bayesian Game - definitions and notation

A Bayesian game Γ is given by the tuple (N, A, T, g, u) where N is the player set, A the set of action profiles, T the set of type profiles, g the probability function over type profiles and u the set of utility functions. We define these in turn.

Let $N = \{1, \dots, n\}$ 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 $A = \times_i A_i$ the set of *actions profiles*. We assume throughout, for convenience, that $T_i = T_\Gamma$ and $A_i = A_\Gamma$ for all $i \in N$ and for some T_Γ

and A_Γ . We will typically index a type as $t_h \in T_\Gamma$ and an action as $a_l \in A_\Gamma$.

A *pure strategy* of a player i is given by a vector $s_i = \{s_i(t_1), \dots, s_i(t_{|T_i|})\}$ where $s_i(t_i)$ is interpreted as the action chosen by player i when of type t_i . Denote the set of pure strategies for player i by S_i or equivalently S_Γ . For any player i we allow choice among any feasible pure strategy as allowed by the set of feasible types T_i and actions A_i . Thus, $|S_i| = |A_i|^{|T_i|}$. Let $K = |A_\Gamma|^{|T_\Gamma|}$ be the number of pure strategies.

A *strategy* of a player i is given by a vector $\sigma_i = \{\sigma_{i1}, \dots, \sigma_{iK}\}$ where σ_{ik} is interpreted as the probability player i plays pure strategy $s_k \in S_i$. A strategy σ implies a vector $\{\sigma_i(\cdot|t_1), \dots, \sigma_i(\cdot|t_{|T_i|})\}$ where $\sigma_i(\cdot|t_i)$ is interpreted as a probability distribution over the set of actions A_i 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 $\Delta(S_i)$ 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 $S = \times_{i \in N} \Delta(S_i)$ denote the set of *strategy vectors*. We refer to a strategy vector σ as *degenerate* if σ_i places unit weight on a unique pure strategy for all $i \in N$. We will typically index a strategy as $s_k \in \Delta(S_\Gamma)$.

Let $C_i = T_i \times A_i$ denote the feasible *compositions of player i* . That is, a composition is a type-action pair. Let $C = \times_{i \in N} C_i$ denote the set of *composition profiles*. For each player $i \in N$ there exists a *utility function* $u_i : C \rightarrow \mathcal{R}$. 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 $c_r \in C_\Gamma$. Let $u = \{u_1, \dots, u_n\}$ denote the set of player utility functions.

For each player $i \in N$ there exists a *prior probability distribution over types* g_i . That is, $g_i(t_i)$ denotes the probability that player i is of type $t_i \in T_i$ if the types of the remaining players $N \setminus \{i\}$ are undetermined. Let g denote a *probability function over the set of type profiles*. Thus, $g(t)$ denotes the probability of type profile $t \in T$. Each player $i \in N$ forms their own beliefs about the types of other players as given by a function p_i mapping T_i into the set of probability distributions over $T_{-i} = \times_{\substack{j \in N \\ j \neq i}} T_j$. The distribution $p_i(t_i)$ is interpreted as the probability function over the type profile of the remaining players in the population conditional on player i knowing his or her own type is t_i . 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 or she is of type t_i . We make two assumptions over probability distributions:

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 $t = (t_1, \dots, t_n)$.

2. *consistent beliefs*: for all $i \in N$ and for all $t_i \in T_i$,

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

We make both these assumptions for simplicity and intuitive appeal. In reality, none of the subsequent results in this paper seem dependent upon these assumptions. To gain some intuition for why this is the case we highlight that our main objective in this paper is to take a Nash equilibrium strategy σ and show that there exists a ‘nearby’ approximate equilibria m in pure strategies. The beliefs of players therefore have little bearing, because we take the initial equilibrium σ as given. That is, if σ was a Nash equilibrium *for some set of beliefs* then we can always find an approximate equilibrium m for *those same set of beliefs*. We say that the probability over type profiles is degenerate if $g(t) = 1$ for some $t \in T$. In this case we say game $\Gamma = (N, T, A, g, u)$ is a *game of perfect information*.

The strategy of a player i , σ_i , and their prior probability distribution over types, g_i , determine a distribution over player i ’s compositions, γ_i , where $\gamma_i(c_i) = g_i(t_i)\sigma_i(a_i|t_i)$ for composition $c_i = (a_i, t_i)$. We can then derive a *probability distribution over outcomes of the game* γ where $\gamma(c) = \prod_i \gamma_i(c_i)$ for all $c \in C$. Any strategy vector σ and probability function over the set of type profiles g induce a particular probability distribution over outcomes of the game which we index $\gamma_{\sigma, g}$. Thus, given strategy vector σ and a probability function over the set of type profiles g the probability of composition profile $c = ((a_1, t_1), \dots, (a_n, t_n))$ is given by

$$\gamma_{\sigma, g}(c) = \prod_i g_i(t_i)\sigma_i(a_i|t_i).$$

Players are assumed to act according to expected payoffs. Thus let E_γ denote the expectations operator where expectations are taken according to the probability distribution over outcomes of the game γ . For each player $i \in N$, let $U_i(\cdot|g)$ denote the *expected utility function of player i conditional on the distribution over player types g* , mapping strategy vectors into the real line, such that

$$U_i(\sigma|g) = E_{\gamma_{\sigma, g}}(u_i(c)).$$

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 $s_k \in S_i$ 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

We consider approximating games with many players, all of whom may be distinct, by games with a finite set of player classes. In particular, a *δ -substitute partition*, given game $\Gamma = (N, T, A, g, u)$, is a partition of the player set N into subsets such that any two players in the same subset are “within δ of each other”. The game Γ is referred to as a *(δ, Q) -classification game* if there is a Q -member δ -substitute partition $\{\Omega_1, \dots, \Omega_Q\}$ of N . Players in the same element of a δ -substitute partition we call δ -substitutes. The set Ω_q is referred to as an approximate class.

This, of course, begs the question of how we measure the distance between two players. In cooperative game theory this distance is typically measured by the difference in value that players can add to coalitions. Informally, a δ -substitute partition is such that, given any coalition structure, ‘swapping’ around players of the same approximate class, has relatively little effect on the value of the coalitions. The counterpart in a non-cooperative framework would appear to be one in which the distance between players is measured by the effect that players can have on each others payoffs. A δ -substitute partition would now be such that given any strategy profile, ‘swapping’ around the strategies of players with the same approximate class, has relatively little effect on the payoff to some player who keeps the same strategy. We formalize this below but first introduce the preliminary concept of a representative strategy.

3.1 Representative strategies

Take as given a partition of the player set N into subsets $\Omega_1, \dots, \Omega_Q$. We then introduce the following,

Representative strategy for class q (relative to σ): Given any strategy $\sigma \in S$ and any subset Ω_q of N let σ_{ω_q} denote the *representative*

strategy for class q defined as,

$$\sigma_{\omega_q}(a_l|t_h) = \frac{1}{\sum_{i \in \Omega_q} g_i(t_h)} \sum_{i \in \Omega_q} g_i(t_h) \sigma_i(a_l|t_h)$$

for all $a_l \in A_\Gamma$ and all $t_h \in T_\Gamma$.¹

We should check that σ_{ω_q} is indeed a strategy for any σ and any Ω_q . That is, we must show that σ_{ω_q} maps T_Γ into the set of probability distributions over A_Γ . For arbitrary $t_h \in T_\Gamma$,

$$\begin{aligned} \sum_{a_l \in A_\Gamma} \sigma_{\omega_q}(a_l|t_h) &= \frac{1}{\sum_{i \in \Omega_q} g_i(t_h)} \sum_{a_l \in A_\Gamma} \sum_{i \in \Omega_q} g_i(t_h) \sigma_i(a_l|t_h) \\ &= \frac{1}{\sum_{i \in \Omega_q} g_i(t_h)} \sum_{i \in \Omega_q} g_i(t_h) \left(\sum_{a_l \in A_\Gamma} \sigma_i(a_l|t_h) \right) \\ &= 1 \end{aligned}$$

and clearly $1 \geq \sigma_{\omega_q}(a_l|t_h) \geq 0$ for all a_l, t_h . Thus, σ_{ω_q} is indeed a strategy.

In interpretation, suppose that $g_i(t_h) = g_j(t_h)$ for all $i, j \in \Omega_q$ and all t_h . We can see that a ‘player’ ω_q truly is a ‘representative’ of the class Ω_q in the sense that

$$\sigma_{\omega_q}(a_l|t_h) = \frac{1}{|\Omega_q|} \sum_{i \in \Omega_q} \sigma_i(a_l|t_h).$$

That is, the probability the representative ω_j has composition (t_h, a_l) equals the average probability players of approximate class Ω_q have that composition. If for some players $i, j \in \Omega_q$, $g_i(t_h) > g_j(t_h)$ then the representative strategy, given type t_h , is weighted towards the action chosen by player i .

3.1.1 δ -substitute partition

With the notion of a representative strategy introduced we can now formally define the notion of δ -substitutes. A partition $\{\Omega_1, \dots, \Omega_Q\}$ is a δ -substitute partition if there exists $\beta \in [0, 1]$, such that²:

Similarity of prior probabilities: for all $q = 1, \dots, Q$, all $t_h \in T_\Gamma$ and all $i, j \in \Omega_q$,

$$|g_i(t_h) - g_j(t_h)| \leq \beta$$

¹If $\sum_{i \in \Omega_q} g_i(t_h) = 0$ then let $\sigma_{\omega_q}(\cdot|t_h)$ be any probability distribution over A_Γ .

²There may be a range of values of β that produce a δ -substitute partition.

and,

Anonymity: for any two strategy profiles $\sigma^1, \sigma^2 \in S$ if,

$$\max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \left(\sigma_{\omega_q}^1(a_l|t_h) - \sigma_{\omega_q}^2(a_l|t_h) \right) \sum_{i \in \Omega_q} g_i(t_h) \right| < 1 + \beta |\Omega_q|$$

for all $\Omega_q, q = 1, \dots, Q$, then,

$$|U_i(s_k, \sigma_{-i}^1) - U_i(s_k, \sigma_{-i}^2)| < \delta$$

for any player $i \in N$ and any strategy $s_k \in \Delta(S_\Gamma)$.

and,

Similarity of payoffs: for all $q = 1, \dots, Q$, for any two players $i, j \in \Omega_q$ and for any strategy profile $\sigma \in S$,

$$|U_i(s_k, \sigma_{-i}) - U_j(s_k, \sigma_{-j})| < \delta$$

for any $s_k \in \Delta(S_\Gamma)$.

The similarity of prior probabilities constraint and similarity of payoffs constraint should be self explanatory. The first constraint basically says that if two players are approximate substitutes then their prior probability distribution over types should be similar, dependent on β . The latter constraint says that two players who are approximate substitutes should have similar payoff functions.

The anonymity constraint is a key requirement which we explain in more detail. This condition is essentially composed of two elements. First, the anonymity condition requires that payoffs should depend primarily on just the representative strategies for the Q approximate classes.³ Thus, there is anonymity in that ‘exchanging’ the strategies of players with the same approximate classification leaves payoffs relatively unaffected. Second, the anonymity condition requires that payoffs should be relatively invariant to bounded changes in the representative strategies of the Q classes. We illustrate by considering two extremes.

Suppose that $|\Omega_q| = 1$ for all $i \in N$. The representative of a class q would then clearly be identical to the actual player of that class. As such,

³They, of course, also depend on the player’s own strategy choice.

that payoffs should depend only on the representative strategies for the Q approximate classes, is trivial. Note, however, the definition of anonymity can now be phrased, $|\sigma_i^1(a_i|t_h) - \sigma_i^2(a_i|t_h)| g_i(t_h) \leq 1 + \beta$ for all $i \in N$. Thus, any player $i \in N$ can change their strategy any way they wish. It seems, therefore, unlikely that anonymity can hold, in this circumstance for any meaningful value of δ . Essentially, we would require that the game is such that each player is indifferent to what their opponents play (which is not much of a game).

By way of contrast, suppose that $|\Omega_1| = n$ and so there is only one class of player. Further, suppose $\beta = 0$. It is now much more restrictive to say that payoffs should depend only on the representative strategies for the Q approximate classes. It 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) but clearly fairly restrictive. Suppose, however, it was the case that payoffs did depend only on the population average. The definition of anonymity now reads $|\sigma_\omega^1(a|t) - \sigma_\omega^2(a|t)| \sum_{i \in N} g_i(t_h) \leq 1$. Thus, the definition of anonymity requires that payoffs be relatively unchanged for only small changes in the representative strategy; this seems reasonable.

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 there exists a δ -substitute partition would seem to involve a trade-off when varying the size of Q ; if Q is large then it seems more plausible that payoffs should depend only on the representative strategies for the Q approximate classes, while, if Q is small, it seems more likely that payoffs should be relatively invariant to bounded changes in the representative strategies of the Q classes.

It is perhaps important to emphasize what the anonymity constraint does not require. Anonymity does not require that players of the same approximate class be similar in all respects. In particular, they need not have similar preferences. To illustrate: suppose people may be more comfortable working in an environment in which a fair proportion of fellow workers are of the same sex. A person’s payoff is then a function of their own strategy and the representative strategy of men and that of women. As such, sex is criteria to partition the player set into approximate classes. We do not require, however, that any two men be similar in any respect other than sex. That is, players must only be similar in terms of the characteristics for which the payoffs of other players depend.

The condition of similarity of payoff functions does, of course, make this extra assumption. This later requirement is not needed, however, in our

purification results. We thus say that a partition $\{\Omega_1, \dots, \Omega_Q\}$ is a *weak δ -substitute partition* if there exists $\beta \in [0, 1]$ such that similarity of prior probabilities and anonymity are satisfied. That is, similarity of payoffs need not exist.

Let $\mathcal{H}((\delta, Q))$ denote the set of all (δ, Q) -classification games. We recall that a game Γ is (δ, Q) -classification game if there is a Q -member δ -substitute partition $\{\Omega_1, \dots, \Omega_Q\}$ of N . Note that for any q and any game Γ , $\Gamma \in \mathcal{H}((\delta, Q))$ for some δ . We Let $\mathcal{H}^W((\delta, Q))$ denote the set of all *weak* (δ, Q) -classification games where a game Γ is weak (δ, Q) -classification game if there is a Q -member weak δ -substitute partition $\{\Omega_1, \dots, \Omega_Q\}$ of N . Note that $\mathcal{H}((\delta, Q)) \subset \mathcal{H}^W((\delta, Q))$.

4 Purification of mixed strategies

We begin with two technical results and then state and prove out main ε -purification result. First, we introduce some notation. Given a vector $\sigma = (\sigma_1, \dots, \sigma_n)$ (where $\sigma_i = (\sigma_{i1}, \dots, \sigma_{iK}) \in \Delta^K$ for $i = 1, \dots, n$) let $\mathcal{M}(\sigma)$ denote the set of vectors $m = (m_1, \dots, m_n)$ such that for all $i = 1, \dots, N$,

1. $\text{support}(m_i) \subset \text{support}(\sigma_i)$ 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 $s_k \in \text{support}(\sigma_i)$ 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 vector $\sigma = (\sigma_1, \dots, \sigma_n)$ and for any vector $\bar{g} \in \mathcal{Z}_+^K$ such that $\sum_i \sigma_i \geq \bar{g}$, there exists a vector $m = (m_1, \dots, m_N) \in \mathcal{M}(\sigma)$ such that:

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

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 vector $\sigma = (\sigma_1, \dots, \sigma_n)$ there exists a vector $m = (m_1, \dots, m_n) \in \mathcal{M}(\sigma)$ such that:

$$\left\lceil \sum_i \sigma_i \right\rceil \geq \sum_i m_i \geq \left\lfloor \sum_i \sigma_i \right\rfloor.$$

Proof: Denote by $\mathcal{M}^*(\sigma)$ the set of vectors $m = (m_1, \dots, m_n) \in \mathcal{M}(\sigma)$ 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 $s_k \in \mathcal{S}$ 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 $s_{\hat{k}}$ a pure strategy such that $\sum_i m_{i\hat{k}} > \lceil \sum_i \sigma_{i\hat{k}} \rceil$.

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

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

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 $\{i_0, i_1, \dots, i_t\} \in W^t$ and set of strategies $\{s_{k_1}, \dots, s_{k_t}\}$ such that,

$$\begin{aligned} m_{i_0 \hat{k}}^0 &= 1 \text{ and } \sigma_{i_0 k_1} > 0, \\ m_{i_r k_r}^0 &= 1 \text{ and } \sigma_{i_r k_{r+1}} > 0 \text{ for all } r = 1, \dots, t-1, \\ m_{i_t k_t}^0 &= 1 \text{ and } \sigma_{i_t k^*} > 0, \end{aligned}$$

where we allow the possibility that $t = 0, 1$. Suppose there exists a $k^* \in L$ such that $\sum_i m_{ik^*} \leq \sum_i \sigma_{ik^*}$. Given the chain of players $\{i_0, i_1, \dots, i_t\} \in W^t$

given above, consider the vector m^* constructed as follows,

$$\begin{aligned} m_{i_0 \hat{k}}^* &= 0 \text{ and } m_{i_0 k_1}^* = 1, \\ m_{i_r k_r}^* &= 0 \text{ and } m_{i_r k_{r+1}}^* = 1 \text{ for all } r = 1, \dots, t-1, \\ m_{i_t k_t}^* &= 0 \text{ and } m_{i_t k^*}^* = 1, \\ m_{ik}^* &= m_{ik}^0 \text{ for all other } s_k \in S \text{ and } i \in N. \end{aligned}$$

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 $\sum_{i \in W} \sum_{s_k \in L} \sigma_{ik} = |W| = \sum_{i \in N} \sum_{k \in L} m_{ik}$. Thus, if $\sum_{i \in N} m_{i \hat{k}}^* > \sum_i \sigma_{i \hat{k}}$ there must exist some $k^* \in L$ such that $\sum_{i \in W} m_{ik^*} \leq \sum_i \sigma_{ik^*}$ ■

With this we can now state and prove our first main result:

Theorem 1: For any Bayesian game $\Gamma \in \mathcal{H}^W((\delta, Q))$. Let ε be a positive real number. If, $\varepsilon \geq 2\delta$ then there exists a Bayesian Nash ε -equilibrium in pure strategies. Furthermore, for any mixed strategy Bayesian Nash Equilibrium σ there exists, a not necessarily unique, Bayesian Nash ε -equilibrium in pure strategies m such that $\text{support}(m_i) \subset \text{support}(\sigma_i)$ for all $i \in N$.

Proof: By definition there exists a δ -substitute partition of N into non-empty subsets $\Omega_1, \dots, \Omega_Q$. Furthermore, using Nash's Theorem there must exist a Nash Equilibrium strategy $\sigma^* \in \Sigma$. This implies, for all $i \in N$, that,

$$U_i(\sigma_i, \sigma_{-i}^*) \geq U_i(s_k, \sigma_{-i}^*) \quad (1)$$

for all σ_i where $\text{support}(\sigma_i) \subset \text{support}(\sigma_i^*)$ and for all $s_k \in \Delta(S)$.

We apply Lemma 2 in turn to each Ω_q . Doing so implies that there exists a strategy vector $m \in S$ where $\text{support}(m_i) \subset \text{support}(\sigma_i^*)$ and m_i is degenerate for all $i \in N$ and where,

$$\left[\sum_{i \in \Omega_q} \sigma_i^* \right] \geq \sum_{i \in \Omega_q} m_i \geq \left[\sum_{i \in \Omega_q} \sigma_i^* \right].$$

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

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

for all $s_k \in S$ and all q . This implies that,

$$\max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \sum_{i \in \Omega_q} m_i(a_l|t_h) - \sum_{i \in \Omega_q} \sigma_i^*(a_l|t_h) \right| \leq 1. \quad (2)$$

For each q pick a player $j_q \in \Omega_q$. Then for all q , using the identity $g_i(t_h) = g_{j_q}(t_h) - (g_{j_q}(t_h) - g_i(t_h))$,

$$\begin{aligned} & \max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \sum_{i \in \Omega_q} g_i(t_h) m_i(a_l|t_h) - \sum_{i \in \Omega_q} g_i(t_h) \sigma_i^*(a_l|t_h) \right| \\ & \leq \max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| g_{j_q}(t_h) \left(\sum_{i \in \Omega_q} m_i(a_l|t_h) - \sum_{i \in \Omega_q} \sigma_i^*(a_l|t_h) \right) \right| + \\ & \quad \max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \sum_{i \in \Omega_q} (g_{j_q}(t_h) - g_i(t_h)) (m_i(a_l|t_h) - \sigma_i^*(a_l|t_h)) \right|. \end{aligned}$$

Given (2), that $g_i(t_h) \leq 1$ and $g_{j_q}(t_h) - g_i(t_h) \leq \beta$ for all $j \in N$ and all $t_h \in T_\Gamma$, and finally that $m_i(a_l|t_h) - \sigma_i^*(a_l|t_h) < 1$ we get,

$$\max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \sum_{i \in \Omega_q} g_i(t_h) m_i(a_l|t_h) - \sum_{i \in \Omega_q} g_i(t_h) \sigma_i^*(a_l|t_h) \right| < 1 + \beta |\Omega_q|$$

for all q . Thus, denoting $\sigma_{\omega_q}^*$ and m_{ω_q} as the representative strategy for class $q = 1, \dots, Q$, we see that,

$$\max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \left(\sigma_{\omega_q}^*(a_l|t_h) - m_{\omega_q}(a_l|t_h) \right) \sum_{i \in \Omega_q} g_i(t_h) \right| < 1 + \beta |\Omega_q| \quad (3a)$$

for all q .

By anonymity and (3a),

$$|U_i(s_k, \sigma_{-i}^*) - U_i(s_k, m_{-i})| < \delta$$

for all $s_k \in \Delta(S_i)$ and for all $i \in N$. Given (1),

$$\begin{aligned} U_i(m_i, m_{-i}) - U_i(s_k, m_{-i}) & \geq -|U_i(m_i, \sigma_{-i}^*) - U_i(m_i, m_{-i})| - |U_i(s_k, \sigma_{-i}^*) - U_i(s_k, m_{-i})| \\ & > -2\delta \geq -\varepsilon \end{aligned}$$

for all $i \in N$ and all $s_k \in \Delta(S_i)$. This completes the proof ■

As previously remarked, any game Γ is a (δ, Q) classification game for some δ . If we measure efficiency by individual rationality, then Theorem 1 allows us to put a bound on the level of inefficiency associated with using pure strategies as opposed to mixed strategies. It would be interesting to have a class of games which are (δ, Q) substitute games for arbitrarily small δ . We consider this when looking at large games. First, we consider social conformity in more detail.

5 Social conformity

We begin by defining a *society*. Take as given a δ -substitute partition $\{\Omega_1, \dots, \Omega_Q\}$ and a strategy vector $\sigma \in S$. For any strategy $s_k \in \Delta(S_\Gamma)$ and any q , consider the subset N_q^k of N such that $i \in N_q^k$ if and only if $i \in \Omega_q$ 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 such that every player belonging to that society plays the same strategy and has the same approximate class.

Given a δ -substitute partition $\pi = \{\Omega_1, \dots, \Omega_Q\}$ and a strategy vector $\sigma \in S$ there exists a unique partition $\{N_1, \dots, N_C\}$ of the player set N into societies. We say that π and σ induce the partition into societies $\{N_1, \dots, N_C\}$.

Our first result of this section essentially summarizes the immediate implications of Theorem 1 and should need no proof.

Corollary 1: Let $\Gamma \in \mathcal{H}^W((\delta, Q))$ be any Bayesian game and let π be any δ -substitute partition with Q approximate classes. Let ε be a positive real number where $\varepsilon \geq 2\delta$. Then there exists a Bayesian Nash ε -equilibrium in pure strategies m such that π and m induce the partition into societies $\{N_1, \dots, N_C\}$ where $C \leq QK$.

This is clearly an immediate consequence of the fact that any partition into societies induced by a δ -substitute partition with Q classes must have no more than QK societies. It still, however, is an interesting result in that the number of societies is fixed independently of the size of the population. Thus, if we can envisage a ‘family of games’ in which as the population grows the number of approximate classes remains the same, then Corollary 1 is an important result. We consider this point in the latter half of the paper. The additional condition of similarity of payoffs allows us, however, to go further than corollary 1. In the following three sub-sections we consider three contrasting extensions to this initial result. In the first extension we

strengthen the definition of social conformity, in the later two sections we weaken the definition of social conformity.

5.1 Social conformity with connected societies

Assume that there is a one-to-one characteristic function y mapping N into $[0, 1]$. We interpret y as ordering the player set in the sense that we can put some significance to the fact that $y(i) > y(j)$. In particular, the characteristics of a player as given by y may be a measure of similarity. For example, we will assume for any δ -substitute partition $\{\Omega_1, \dots, \Omega_Q\}$ and three players $i, j, k \in N$ that if $i, k \in \Omega_q$, for some q , and $y(i) > y(j) > y(k)$ then $j \in \Omega_q$. We say that a set $Y \subset N$ is *connected with respect to y* when, for any $i, k \in N$, if there exists a player $j \in N$ such that $y(i) > y(j) > y(k)$ then $j \in Y$. Thus, the set of players with the same approximate classification is connected.

It may be intuitive for societies to be connected with respect to some characteristic function y . An example may be useful. Suppose that the characteristics of a player can be summarized by a number from the unit interval. Then, if societies are connected this would seem to imply that the majority of players belong to the same society as those players with the characteristics most similar to their own. This example is taken further in the final part of this paper leading to corollary 5.

We now state our first social conformity result,

Corollary 2: Let $\Gamma \in \mathcal{H}((\delta, Q))$ be any Bayesian game and let π be any δ -substitute partition with Q approximate classes. Let ε be a positive real number where $\varepsilon \geq 6\delta$. Let y denote any characteristic function. Then there exists a Bayesian Nash ε -equilibrium in pure strategies m such that π and m induce the partition into societies $\{N_1, \dots, N_C\}$ where N_c is connected with respect to y for all $c = 1, \dots, C$ and where $C \leq JK$.

Proof: By Theorem 1 there exists a Bayesian Nash 2δ equilibrium in pure strategies σ . Thus,

$$U_i(\sigma) \geq U_i(s_k, \sigma_{-i}) - 2\delta$$

for all $s_k \in \Delta(S)$ and for all $i \in N$. For all $s_k \in S_\Gamma$ and all $q = 1, \dots, Q$, let n^{kq} denote the number of players $i \in \Omega_q$ such that $\sigma_{ik} = 1$. Assume, subject to a possible reordering of the player set that, for all $i, j \in N$, if $i > j$ then $y(i) > y(j)$. Suppose, further, that player 1 belongs to Ω_1 . Let m be informally defined as the strategy vector where for players $i \in \{1, \dots, n^{11}\}$,

$m_i = s_1$, for players $i \in \{n^{11} + 1, \dots, n^{11} + n^{21}\}$, $m_i = s_2$ and so on. Strategy vector m is such that π and δ induce a partition into societies $\{N_1, \dots, N_C\}$ where N_c is connected with respect to y for all $c = 1, \dots, C$. We also note that,

$$\sum_{i \in \Omega_q} g_i(t_h) m_i(a_l | t_h) - \sum_{i \in \Omega_q} g_i(t_h) \sigma_i(a_l | t_h) \leq \beta |\Omega_q|$$

for all q and all a_l, t_h . Thus, by anonymity for any $i \in N$,

$$|U_i(s_k, m_{-i}) - U_i(s_k, \sigma_{-i})| < \delta$$

for any strategy $s_k \in \Delta(S)$. Furthermore, for any two players $i, j \in \Omega_q$,

$$|U_i(s_k, m_{-i}) - U_j(s_k, m_{-j})| < \delta$$

for any $s_k \in \Delta(S_\Gamma)$.

For every player $j \in \Omega_q$ and for all q , if $m_{jr} = 1$ then there exists a player $i^{rq} \in \Omega_q$ such that $\sigma_{i^{rq}} = 1$. This implies that $U_{i^{rq}}(\sigma) \geq U_{i^{rq}}(s_k, \sigma_{-i^{rq}}) - 2\delta$. Note it is possible that $i = j$. Thus,

$$\begin{aligned} U_j(m) &\geq U_j(s_k, m_{-i^{rq}}) - 2\delta \\ &\quad - |U_{i^{rq}}(\sigma) - U_{i^{rq}}(m)| - |U_{i^{rq}}(m) - U_j(m)| \\ &\quad - |U_{i^{rq}}(s_k, \sigma_{-i^{rq}}) - U_{i^{rq}}(s_k, m_{-i^{rq}})| - |U_{i^{rq}}(s_k, m_{-i^{rq}}) - U_j(s_k, m_{-j})| \\ &> U_j(s_k, m_{-i^{rq}}) - 6\delta \end{aligned}$$

for all $s_k \in \Delta(S_\Gamma)$. This completes the proof \blacksquare

For the statement of this result, the definition of a society is essentially strengthened in that societies are connected. As such, the bound that we can put on the inefficiency of social conformity is higher at 6δ .

5.2 Social conformity in mixed strategies

For this sub-section we suppose that players can choose mixed strategies. We retain the notion of a society as being such that two players belonging to the same society must have the same approximate class and same strategy. We can derive the following result

Theorem 2: Let $\Gamma \in \mathcal{H}((\delta, Q))$ be any Bayesian game and let π be any δ -substitute partition with Q approximate classes. Let ε be a positive real number where $\varepsilon \geq 4\delta$. Then, there exists a Bayesian Nash ε equilibrium m

such that π and m induce the partition into societies $\{N_1, \dots, N_C\}$ where $C = Q$.

Proof: By Nash's Theorem there exists a Bayesian Nash equilibrium σ of the game Γ . That is, for any $i \in N$ and any strategy $s_k \in \Delta(S_\Gamma)$,

$$U_i(\sigma) \geq U_i(s_k, \sigma_{-i}).$$

There also exists a δ substitute partition $\pi = \{\Omega_1, \dots, \Omega_Q\}$. For each q let σ_{ω_q} denote the representative strategy of class q . Consider now the strategy vector m such that, for all $i \in N$,

$$\text{if } i \in \Omega_q \text{ then } m_i = \sigma_{\omega_q}.$$

Clearly π and m induce a partition into societies $\{N_1, \dots, N_C\}$ where $C = Q$.

We note that the representative strategy for each class q is now,

$$\begin{aligned} m_{\omega_q}(a_l|t_h) &= \frac{1}{\sum_{i \in \Omega_q} g_i(t_h)} \sum_{i \in \Omega_q} g_i(t_h) \sigma_{\omega_q}(a_l|t_h) \\ &= \sigma_{\omega_q}(a_l|t_h). \end{aligned}$$

Thus, by the assumption of anonymity,

$$|U_i(s_k, \sigma_{-i}) - U_i(s_k, m_{-i})| < \delta$$

for all $i \in N$ and all $s_k \in \Delta(S_\Gamma)$.

For each q let $S^q = \{s_k \in S_\Gamma : \text{there exists a player } i \in \Omega_q \text{ such that } \sigma_{ik} > 0\}$. We note that for all q and for each degenerate strategy s_r where $\text{support}(s_r) \subset S^q$ there exists a player $j^{rq} \in \Omega_q$ such that

$$U_{j^{rq}}(s_r, m_{-j^{rq}}) \geq U_{j^{rq}}(s_k, m_{-j^{rq}}) - 2\delta$$

for all $s_k \in \Delta(S_\Gamma)$. We also note for all $i \in N$ that $\text{support}(m_i) \subset S_q$. For any q , for any two players $i, j \in \Omega_q$ and for any strategy s_r where $\text{support}(s_r) \subset S^q$,

$$|U_i(s_r, m_{-i}) - U_j(s_r, m_{-j})| < \delta$$

by similarity of payoffs. Thus, for any $i \in N$ and for any $s_r \in S^q$,

$$U_i(m) \geq U_i(s_k, m_{-i}) - 4\delta.$$

for any $s_k \in \Delta(S_\Gamma)$. This completes the proof ■

This result shows that if players can use mixed strategies then the number of societies can be bounded by the number of approximate classes. That is, if Γ is a (δ, Q) class game then for any $\varepsilon \geq 4\delta$ there exists an ε -equilibrium such that any two players belonging to the same approximate class play the same strategy.

It is interesting to note that despite Theorem 1 being a purification result it can be used to prove a similar result to Theorem 2 although the bound becomes ‘for all $\varepsilon \geq 6\delta$ ’.

5.3 Social conformity with roles

As discussed in the introduction it is sometimes appropriate that players belonging to the same society should play different actions. This is possible through games of incomplete information by making actions conditional on a player’s type - a player’s type being determined by nature. In this final sub-section on social conformity we suppose, by way of extension, that players could endogenously choose a set of roles and a probability distribution over those roles.

Formally, suppose that there exists a set of *roles* $R = \{r_1, \dots, r_K\}$. Each player can then choose their own *probability distribution over roles* $f_i : R_i \rightarrow [0, 1]$. Given a Bayesian game $\Gamma = (N, A, T, g, u)$ we then consider a *two stage Bayesian game with endogenous types* Γ^R defined such that

1. In stage 1 each player independently chooses a probability distribution over roles f_i , the choices of players as given by $f = (f_1, \dots, f_n)$ are freely communicated,
2. In stage 2 the Bayesian game $\Gamma^f = (N, A, T^f, g^f, u^f)$ is played where,

- (a) $T_i^f = T_i \times R_i$ for all $i \in N$,
- (b) $g_i^f(t_h, r_k) = g_i(t_h)f_i(r_k)$ for all $i \in N$, all $t_h \in T_\Gamma$ and all $r_k \in R$
- (c) $u_i^f((a_1, t_1, r_1), \dots, (a_n, t_n, r_n)) = u_i((a_1, t_1), \dots, (a_n, t_n))$ for any composition profile .

We highlight that in the Bayesian game Γ^f players are allowed to make their action choice conditional on their role in the same as action choice can be made conditional on type. The standard definition of an approximate Nash equilibrium still applies. We refer to a strategy of a two stage Bayesian game with endogenous types Γ^R as a pure strategy if each player $i \in N$ chooses a pure strategy for game Γ^f given any choice of f in stage 1.

The definition of a society remains the same - namely that two players in the same society have the same attribute and the same strategy. This means they have the same choice of probability distributions over roles and the set strategy for the resulting Bayesian game.

We can now state our final result of this section.

Corollary 3: Let $\Gamma \in \mathcal{H}((\delta, Q))$ be any Bayesian game and let π be any δ -substitute partition with Q approximate classes. Let ε be a positive real number where $\varepsilon \geq 4\delta$. Then, in the two stage Bayesian game with endogenous types Γ^R there exists a Bayesian Nash ε equilibrium in pure strategies m such that π and m induce the partition into societies $\{N_1, \dots, N_C\}$ where $C = Q$.

Proof: Consider an arbitrary player $i \in N$. We note that any mixed strategy σ_i of the game Γ can be decomposed as a choice of strategy m_i in a game Γ^f and choice of probability distribution over roles f_i . In particular, suppose that player i plays pure strategy s_k (from game Γ) if given role r_k , for all k . Then, put $f_i(r_k) = \sigma_i(s_k)$ for all k . Having noted this the result is immediate give Theorem 2 ■

This result demonstrates that if players have some way of endogenous system by which players can be assigned roles then we can conceive of societies in which players play different actions. There actions are determined by the role that they are playing within the society.

6 Large Games

We begin by reaffirming that for any game Γ and any Q there exists some δ such that game Γ is a (δ, Q) classification game. Thus, the results of above apply to any game. Clearly, in interpretation, however, we would want that δ is small. So, what characteristics of a game imply that there will exist a δ -substitute for small δ ? We would expect such a game to have the characteristics that, (1) a player's payoff is not largely dependent upon the actions of any small subset of the population, and (2) there is a natural way of grouping players with similar characteristics. Games induced from a pregame with the large game property satisfy both these requirements. We thus use this section to demonstrate the later claim and provide an example of how earlier results can be applied in practice.

6.1 Definitions

6.1.1 Pregames

We assume a compact metric space Ω of *player attributes*. There is a finite set A_Γ of actions and a finite set T_Γ of types. The set of pure strategies is given by S_Γ and the set of strategies by $\Delta(S_\Gamma)$. A function from $\Omega \times A_\Gamma \times T_\Gamma$ into \mathcal{R}_+ is said to be a *weight function* if it satisfies $\sum_{a_l \in A_\Gamma} \sum_{t_h \in T_\Gamma} w(\omega, a_l, t_h) \in \mathbb{Z}_+$ for all $\omega \in \Omega$. Let W denote the set of weight functions.

A *pregame* is given by a tuple $\mathcal{G} = (\Omega, A_\Gamma, T_\Gamma, g, U)$, consisting of a compact metric space Ω , finite sets A_Γ and T_Γ , a function $g : \Omega \times T_\Gamma \rightarrow [0, 1]$ and a function $U : \Omega \times \Delta(S_\Gamma) \times W \rightarrow \mathcal{R}_+$.

6.1.2 Populations and games

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*.

Given a population (N, α) and a strategy vector σ for the population (N, α) we say that weight function $w_{\alpha, \sigma}$ is *relative to strategy vector σ and attribute function α* if,

$$w_{\alpha, \sigma}(\omega, a_l, t_h) = \sum_{i \in N: \alpha(i) = \omega} \sigma_i(a_l | t_h) g_i(t_h)$$

for all $a_l \in A_\Gamma$, $t_h \in T_\Gamma$ and all $\omega \in \Omega$. In interpretation, given the population (N, α) and the strategy vector σ ,

$$\frac{w_{\alpha, \sigma}(\omega, a_l, t_h)}{|\alpha^{-1}(\omega)|}$$

is the expected proportion of times composition $c_r = (a_l, t_h)$ will be seen by a player of attribute ω . We note that,

$$\begin{aligned} \sum_{a_l \in A_\Gamma} \sum_{t_h \in T_\Gamma} w_{\alpha, \sigma}(\omega, a_l, t_h) &= \sum_{a_l \in A_\Gamma} \sum_{t_h \in T_\Gamma} \sum_{i \in N: \alpha(i) = \omega} \sigma_i(a_l | t_h) g_i(t_h) \\ &= |\alpha^{-1}(\omega)|. \end{aligned}$$

Given a population (N, α) and player $i \in N$, define α_{-i} as the restriction of α to $N \setminus \{i\}$. Thus, given an attribute function α and strategy vector σ , for all $\omega \in \Omega$, all $s_k \in S$ and for all $i \in N$,

$$w_{\alpha_{-i}, \sigma}(\omega, a_l, t_h) = \begin{cases} w_{\alpha, \sigma}(\omega, a_l, t_h) - \sigma_i(a_l | t_h) g_i(t_h) & \text{if } \alpha(i) = \omega \\ w_{\alpha, \sigma}(\omega, a_l, t_h) & \text{otherwise.} \end{cases}$$

Weight functions modified by the property that one player of some particular attribute is not included will play a role in the definition of games. We will use $W_{\alpha-\omega}$ to denote the set of weight functions corresponding to α_{-i} where $\omega = \alpha(i)$.

6.1.3 Induced games

Given a population (N, α) , a *game*

$$\Gamma(N, \alpha) = \left((N, \alpha), A_\Gamma, T_\Gamma, \{g_\omega : T_\Gamma \longrightarrow [0, 1] \mid \omega \in \alpha(N)\} \right. \\ \left. \{U_\omega : \Delta(S_\Gamma) \times W_{\alpha-\omega} \longrightarrow \mathcal{R}_+ \mid \omega \in \alpha(N)\} \right)$$

is induced from the pregame $\mathcal{G} = (\Omega, A_\Gamma, T_\Gamma, g, U)$ by defining

$$g_\omega(t_h) \stackrel{\text{def}}{=} g(\omega, t_h)$$

and

$$U_\omega(s_k, w) \stackrel{\text{def}}{=} U(\omega, s_k, w)$$

for all $s_k \in \Delta(S_\Gamma)$, $w \in W_{\alpha-\omega}$ and $\omega \in \alpha(N)$. In interpretation, $U_{\alpha(i)}(s_k, w)$ is the payoff received by a player $i \in N$ of attribute $\alpha(i)$ from playing the strategy s_k when the strategies of other players are summarized by w . Thus, players of the same attribute have the same payoff function, inherited from the pregame. Similarly, $g_\omega(t_h)$ is the probability that a player $i \in N$ of attribute $\alpha(i)$ is of type t_h .

We impose the standard Von-Neumann, Morgenstern assumptions on the linearity of payoffs with respect to mixed strategies. To explain further it is useful to relate the utility function induced from the pregame $U_\omega : \Delta(S_\Gamma) \times W_{\alpha-\omega} \longrightarrow \mathcal{R}_+$ to the expected utility function $U_i : S \longrightarrow \mathcal{R}_+$ as used in the first half of this paper. Consider a game $\Gamma(N, \alpha)$ induced from the pregame \mathcal{G} . We assume that this game $\Gamma(N, \alpha)$ is equivalent to the Bayesian game $\bar{\Gamma}(N, A, T, \bar{g}, \bar{U})$ where $A = \times_{i \in N} A_\Gamma$, $T = \times_{i \in N} T_\Gamma$, $\bar{g}(t) = \prod_{i \in N} g_\omega(t_i)$ for all $t \in T$ and $\bar{U}_i(\sigma) = U_{\alpha(i)}(\sigma_i, w_{\alpha-i, \sigma})$ for all σ and all $i \in N$.

6.2 Large game property

We continue by defining the concepts of global interaction and continuity in attributes. These two concepts allow us to introduce the *large game property*. This property constitutes an assumption on a pregame $\mathcal{G} = (\Omega, A_\Gamma, T_\Gamma, g, U)$. In particular, it places restrictions on the payoff function U and distribution over types g of the pregame. As a preliminary step, let $F(\mathcal{G}, n)$ denote the

set of games induced by the pregame \mathcal{G} by societies of size n . That is game $\Gamma(N, \alpha) \in F(\mathcal{G}, n)$ if and only if $|N| = n$.

Global Interaction: Given positive real numbers $\delta > 0$ and $\tau > 0$ the game $\Gamma(N, \alpha)$ is said to satisfy δ, τ -global interaction when for any two weight functions w_α and g_α , both relative to attribute function α , if,

$$\frac{1}{|N|} \sum_{a_l \in A_\Gamma} \sum_{t_h \in T_\Gamma} \sum_{\omega \in \alpha(N)} |w_\alpha(\omega, a_l, t_h) - g_\alpha(\omega, a_l, t_h)| < \tau$$

then,

$$|U_{\alpha(i)}(s_k, w_{\alpha-i}) - U_{\alpha(i)}(s_k, g_{\alpha-i})| < \delta$$

for all $i \in N$ and all $s_k \in \Delta(S_\Gamma)$.

Continuity in attributes: Given positive real numbers $\delta > 0$ and $\tau > 0$, the set of games $\gamma(\mathcal{G}, n)$ is said to satisfy ε, τ -continuity in attributes when for any two games $\Gamma(N, \alpha)$ and $\Gamma(N, \bar{\alpha})$ in $F(\mathcal{G}, n)$, if, for all $i \in N$,

$$dist(\alpha(i), \bar{\alpha}(i)) < \tau$$

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

$$|U_{\alpha(j)}(s_k, w_{\alpha-j, \sigma}) - U_{\bar{\alpha}(j)}(s_k, w_{\bar{\alpha}-j, \sigma})| < \delta$$

for all $s_k \in \Delta(S_\Gamma)$. where $w_{\alpha, \sigma}$ and $w_{\bar{\alpha}, \sigma}$ are the weight functions relative to strategy vector σ and, respectively, attribute functions α and $\bar{\alpha}$.

We can now introduce the main assumption,

Large game property: The pregame $\mathcal{G} = (\Omega, A_\Gamma, T_\Gamma, g, U)$ satisfies the large game property if for any $\delta > 0$ and there exists real numbers $\eta_l(\delta), \tau_g(\delta) > 0$ and $\tau_c(\delta) > 0$ such that for any $n > \eta(\delta)$ the set of games $F(\mathcal{G}, n)$ satisfy $\delta, \tau_c(\varepsilon)$ -continuity of payoff functions and any game $\Gamma(N, \alpha) \in F(N, \mathcal{G})$ satisfies $\varepsilon, \tau_g(\varepsilon)$ -global interaction.

Thus, the pregame \mathcal{G} satisfies the large game property if both global interaction and continuity of payoff functions are satisfied by large games. The large game property implies a form of continuity of $U : \Omega \times \Delta(S_\Gamma) \times W \rightarrow \mathcal{R}_+$ with respect to changes in the weight function w and attribute ω while the strategy s_k remains constant. It also, implies a form of continuity on $g : \Omega \times T_\Gamma \rightarrow [0, 1]$.

A detailed motivation and explanation of the above assumptions, for games of complete information, is provided by Wooders, Cartwright and Selten (2001) so we give here only a brief discussion. Global interaction says that players payoffs are largely a function of their own strategy and on the numbers of players we each attribute of each type playing each action. As such a player's payoff is not largely dependent on the actions of any small group of individuals. This clearly has a close relationship with the notion of anonymity in the definition of δ -substitutes.

The assumption of continuity in attributes is a more wide ranging assumption. Essentially, given a strategy vector σ , it says that the attributes of all players can be slightly perturbed and the payoff to each player remains largely unaffected. The first thing we should highlight is how two distinct games $\Gamma(N, \alpha)$ and $\Gamma(N, \bar{\alpha})$ are compared. This is possible through the use of the pregame. The assumption thus formalizes the intuition that if a population changes only slightly from (N, α) to $(N, \bar{\alpha})$ then the games induced by these societies should be largely the same.

One element of continuity in attributes that should be emphasized. Namely, even though the strategies of the players remain the same, in both societies, their attributes have changed and thus their prior probability over types may have changed. As such, the same strategy can imply a different probability distribution over compositions. This would imply that implicit in the continuity in attributes assumption is the idea that players with similar attributes should have similar probability distributions over types. The continuity in attributes assumption is thus inter-related to both the similarity in prior probabilities and similarity in payoff functions of the definition of approximate substitutes.

6.3 Preliminary result

Having defined the large game property we can now go on to apply the results from the first half of this paper. To do this we need to find a connection between the games induced by the a pregame \mathcal{G} satisfying the large game property and the set $\mathcal{H}(\delta, Q)$ for some Q . This relationship is not straightforward.

We can show, however, that for any pregame \mathcal{G} and any sufficiently large game $\Gamma(N, \alpha)$ induced by that pregame there is ‘nearby’, a game $\Gamma(N, \bar{\alpha})$ which is also induced by the pregame and is a δ, Q -classification game for some Q . Formally, we have,

Lemma 3: If the pregame \mathcal{G} satisfies the large game property then given any real numbers $\delta > 0$ and $\tau > 0$ there exists real numbers $\eta(\delta, \tau)$ and $Q(\delta, \tau)$ such that for any population (N, α) , where $|N| > \eta(\varepsilon)$, there exists a population $(N, \bar{\alpha})$ such that $\max_{i \in N} \{dist(\alpha(i), \bar{\alpha}(j))\} < \tau$ and the induced game $\Gamma(N, \bar{\alpha})$ belongs to the set $\mathcal{H}(\delta, Q(\delta, \tau))$.

Proof: Suppose that the statement of the lemma is false. Then there is some $\bar{\delta} > 0$ and $\bar{\tau} > 0$, such that for any real number \bar{Q} and for each integer ν there is a population (N^ν, α^ν) where $|N^\nu| > \nu$ and such that for no population $(N^\nu, \bar{\alpha}^\nu)$ where $\max_{i \in N} \{dist(\alpha^\nu(i), \bar{\alpha}^\nu(j))\} < \bar{\tau}$ does the induced game $\Gamma(N^\nu, \bar{\alpha}^\nu)$ belong to the set $\mathcal{H}(\bar{\delta}, \bar{Q})$.

The pregame \mathcal{G} satisfies the large game property and so let $\eta_l(\delta)$, $\tau_g(\delta)$ and $\tau_u(\delta)$ be the appropriate numbers for a payoff bound of δ . Let $\tau = \min\{\tau_u(\delta), \bar{\tau}\}$. Partition Ω into subsets $\Omega_1, \dots, \Omega_Q$ each of diameter less than τ , that is, for any $\omega, \omega' \in \Omega_q$ and for any q , $dist(\omega, \omega') < \tau$. To each subset Ω_q choose and fix an attribute ω_q . Define the attribute function $\bar{\alpha}^\nu$ as follows, for all $i \in N^\nu$,

$$\bar{\alpha}^\nu(i) = \omega_q \text{ if and only if } \alpha(i) \in \Omega_q.$$

For each ν consider the game $(N^\nu, \bar{\alpha}^\nu)$. Let $\pi^\nu = \{N_1^\nu, \dots, N_Q^\nu\}$ denote, for each ν , the partition of the player set such that $i \in N_q^\nu$ if and only if $\bar{\alpha}^\nu(i) = \omega_q$. We propose π^ν as a candidate for a δ -substitute partition in the game $(N^\nu, \bar{\alpha}^\nu)$ for each ν .

Similarity of prior probabilities is satisfied for $\beta = 0$. For an arbitrary ν , take any two strategy profiles $\sigma^{\nu 1}$ and $\sigma^{\nu 2}$, of the game $\Gamma(N^\nu, \alpha^\nu)$. Suppose that,

$$\max_{a_l \in A_\Gamma} \max_{t_h \in T_\Gamma} \left| \left(\sigma_{\omega_q}^{\nu 1}(a_l | t_h) - \sigma_{\omega_q}^{\nu 2}(a_l | t_h) \right) \sum_{i \in \Omega_q} g_i(t_h) \right| < 1 \quad (4)$$

for all q , for all $a_l \in A_\Gamma$ and $t_h \in T_\Gamma$. Let $w_{\bar{\alpha}^\nu, \sigma^{\nu 1}}$ denote the weight function relative to attribute function $\bar{\alpha}^\nu$ and strategy vector $\sigma^{\nu 1}$. Let $w_{\bar{\alpha}^\nu, \sigma^{\nu 2}}$ denote the weight function relative to attribute function $\bar{\alpha}^\nu$ and strategy vector $\sigma^{\nu 2}$.

By (4) we have that,

$$\frac{1}{|N^\nu|} \sum_{a_l \in A_\Gamma} \sum_{t_h \in T_\Gamma} \sum_{\omega \in \alpha(N)} |w_{\bar{\alpha}^\nu, \sigma^{\nu 1}}(\omega, a_l, t_h) - w_{\bar{\alpha}^\nu, \sigma^{\nu 2}}(\omega, a_l, t_h)| < \frac{Q |A_\Gamma| |T_\Gamma|}{|N^\nu|}. \quad (5)$$

By global interaction and (5) there exists a ν^* such that if $\nu > \nu^*$,

$$\left| U_{\alpha(j)}(s_k, w_{\bar{\alpha}_{-j}, \sigma^{\nu 1}}) - U_{\alpha(j)}(s_k, w_{\bar{\alpha}_{-j}, \sigma^{\nu 2}}) \right| < \delta$$

for all $s_k \in \Delta(S_\Gamma)$.

We now consider similarity of payoff functions. For an arbitrary ν consider any strategy vector σ^ν of the game $\Gamma(N^\nu, \bar{\alpha}^\nu)$. Let $i, j \in N$ be any two players such that $i, j \in N_q$ for some q . Let $\bar{\sigma}^\nu$ be the strategy vector such that $\bar{\sigma}_i^\nu = \sigma_j^\nu$ and for all $l \neq i$, $\bar{\sigma}_l^\nu = \sigma_l^\nu$. By global interaction for any $\nu > \nu^*$,

$$\left| U_{\bar{\alpha}(j)}(s_k, w_{\alpha_{-j}, \sigma^\nu}) - U_{\bar{\alpha}(j)}(s_k, w_{\alpha_{-j}, \bar{\sigma}^\nu}) \right| < \delta$$

for all $s_k \in \Delta(S_\Gamma)$ ■

6.4 Approximate purification in large games

Lemma 3 allows us to apply all of the results obtained for (δ, Q) -classification games to sufficiently large games induced by a pregame satisfying the large game property. We demonstrate with three additional results. The first result is an ε -purification result,

Corollary 4: Given a real number $\varepsilon > 0$ there exists a real number $\eta(\varepsilon) > 0$ such that if the pregame \mathcal{G} satisfies the large game property, then for any population (N, α) where $|N| > \eta(\varepsilon)$, the induced game $\Gamma(N, \alpha)$ has an ε -equilibrium in pure strategies. Moreover, for any mixed strategy equilibrium there exists an ε -purification.

Proof: Let $\delta = \frac{\varepsilon}{4}$. Given that the pregame \mathcal{G} satisfies the large game property let $\tau = \tau_c(\delta)$. By Lemma 3 there exists a real number $\eta(\delta, \tau)$ such that for any population (N, α) , where $|N| > \eta(\delta, \tau)$, there exists a population $(N, \bar{\alpha})$ such that $\max_{i \in N} \{dist(\alpha(i), \bar{\alpha}(j))\} < \tau$ and the induced game $\Gamma(N, \bar{\alpha})$ belongs to the set $\mathcal{H}(\delta, Q)$ for some finite real number Q . Thus, by Theorem 1, for any population (N, α) , where $|N| > \eta(\delta, \tau)$, there

exists a population $(N, \bar{\alpha})$ such that $\max_{i \in N} \{dist(\alpha(i), \bar{\alpha}(j))\} < \tau$ and the induced game $\Gamma(N, \bar{\alpha})$ has a Bayesian Nash $\frac{\varepsilon}{2}$ -equilibrium in pure strategies m . Thus, for all $i \in N$,

$$U_{\bar{\alpha}(i)}(m_i, w_{\bar{\alpha}_{-i}, m}) \geq U_{\bar{\alpha}(i)}(s_k, w_{\bar{\alpha}_{-i}, m}) - \frac{\varepsilon}{2}$$

for all $s_k \in \Delta(S_\Gamma)$. By choice of τ and continuity in attributes,

$$|U_{\alpha(i)}(s_k, w_{\alpha_{-i}, m}) - U_{\bar{\alpha}(i)}(s_k, w_{\bar{\alpha}_{-i}, m})| < \frac{1}{4}\varepsilon$$

for all $s_k \in \Delta(S_\Gamma)$. Thus, m is an Bayesian Nash ε -equilibrium of the game $\Gamma(N, \alpha)$ ■

6.5 Social conformity in large games

Corollary 4 and its proof show how the results from the framework of approximate substitutes in non-cooperative games can be applied in a straightforward way within the framework of non-cooperative pregames satisfying the large game property. As such, in this section, we state without proof, two social conformity results.

Before, stating our first result we define a further term. Given any population (N, α) and any player $i \in N$ we say that player j is player i 's *closest neighbor* if $dist(\alpha(i), \alpha(j)) \leq \min_{k \in N} \{dist(\alpha(i), \alpha(k))\}$. A person may have more than one closest neighbor. The following result, applying corollary 2, demonstrates how we can apply the notion of a characteristic function on the set of players.

Corollary 5: Let $\mathcal{G} = (\Omega, A_\Gamma, T_\Gamma, g, U)$ be any pregame satisfying the large game property and assume $\Omega = [0, 1]$. Then given any $\varepsilon > 0$ there exists a real number $\eta(\varepsilon)$ such that for any population (N, α) where $|N| > \eta(\varepsilon)$ there is an ε -equilibrium in pure strategies of the induced game $\Gamma(N, \alpha)$ with the property that at least $|N|(1 - \varepsilon)$ players play the same strategy as a closest neighbor.

In contexts where players attributes can be ordered along the unit interval (Greenberg and Weber (1986) for example) this result clearly demonstrates that for large games it can be efficient for similar players to play similar strategies. This is demonstrated further by the following result which applies Theorem 2 and corollary 1,

Corollary 6: Let \mathcal{G} be any pregame satisfying the large game property. Then, for any $\varepsilon > 0$ there exists real numbers $\eta(\varepsilon)$ and $Q(\varepsilon)$ such that for any population (N, α) where $|N| > \eta(\varepsilon)$ the induced game $\Gamma(N, \alpha)$ has (1) a Bayesian Nash ε equilibrium σ such that π and σ induce the partition into societies $\{N_1, \dots, N_C\}$ where $C \leq Q(\varepsilon)$ for some partition π , and (2) a Bayesian Nash ε equilibrium in pure strategies m such that π and m induce the partition into societies $\{N_1, \dots, N_C\}$ where $C \leq Q(\varepsilon)K$.

7 Conclusion

Can it be efficient for individuals to form a relatively small number of societies such that all players within a society play the same strategy and have similar attributes? By using the framework of games with approximate classes we are able to put bounds on the inefficiency of such social conformity for arbitrary games and for arbitrary numbers of societies. We then apply this to large games induced by a pregame satisfying the large game property. Such games are characterized primarily by the fact that no player's payoff is overly dependent on any small subset of the population. For such games we demonstrate that there does exist an ε Nash equilibrium, for arbitrarily small ε , which partitions the player set into a relatively small number of societies. It would appear, therefore, that in such games social conformity can be efficient.

References

- [1] Araujo, A., M. Pascoa and J. Orrillo (2000) "Equilibrium with default and exogenous collateral," *Mathematical Finance*.
- [2] Aumann, R.J. (1985) "An axiomatization of the non-transferable utility value," *Econometrica* 53: 599-612.
- [3] 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.
- [4] 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.
- [5] Follmer, H. (1974) "Random economies with many interacting agents," *Journal of Mathematical Economics* 1, 51-62.

- [6] Fudenberg, D. and D.K. Levine (1998) *The Theory of Learning in Games*, Cambridge: MIT press.
- [7] Greenberg, J. and S. Weber (1986) ‘Strong Tiebout equilibrium with restricted preference domain’, *Journal of Economic Theory* 38:101-117.
- [8] Hildenbrand, W. (1971) “Random preferences and equilibrium analysis,” *Journal of Economic Theory* 3, 414-429.
- [9] Kalai, E. (2000) “Private information in large games,” Northwestern University Discussion Paper 1312.
- [10] Khan, A. (1989) “On Cournot-Nash equilibrium distributions for games with a nonmetrizable action space and upper semi continuous payoffs,” *Transactions of the American Mathematical Society* 293: 737-749.
- [11] 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.
- [12] Kovalenkov, A. and M. Wooders (1996) “Epsilon cores of games with limited side payments; Nonemptiness and equal treatment,” *Games and Economic Behavior* (to appear).
- [13] 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).
- [14] Mas-Colell, A. (1984) “On a theorem of Schmeidler,” *Journal of Mathematical Economics* 13: 206-210.
- [15] Pascoa, M. (1998) “Nash equilibrium and the law of large numbers,” *International Journal of Game Theory* 27: 83-92.
- [16] Pascoa, M. (1993) “Approximate equilibrium in pure strategies for nonatomic games,” *Journal of Mathematical Economics* 22: 223-241.
- [17] Rashid, S. (1983) “Equilibrium points of nonatomic games; Asymptotic results,” *Economics Letters* 12: 7-10.
- [18] Reny, P.J. (1999) “On the existence of pure and mixed strategy Nash equilibria in discontinuous games,” *Econometrica* 67: 1029-1056.

- [19] Schmeidler, D. (1973) "Equilibrium points of nonatomic games," *Journal of Statistical Physics* 7: 295-300.
- [20] Wooders, M.H. (1979) "Asymptotic cores and asymptotic balancedness of large replica games" (Stony Brook Working Paper No. 215, Revised July 1980).
- [21] Wooders, M. (1983) "The epsilon core of a large replica game," *Journal of Mathematical Economics* 11, 277-300.
- [22] Wooders, M. (1994) "Equivalence of games and markets," *Econometrica* 62, 1141-1160.
- [23] Wooders, M. and Zame, W.R. (1984) "Approximate cores of large games," *Econometrica* 52, 1327-1350.
- [24] Young, H.P. (2001) *Individual Strategy and Social Structure*, Princeton University Press.