

ON EQUILIBRIUM IN PURE STRATEGIES IN GAMES WITH MANY PLAYERS

by

Edward Cartwright and Myrna Wooders



Working Paper No. 05-W11

April 2005

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

On Equilibrium in Pure Strategies in Games with Many Players*

Edward Cartwright

Department of Economics, Keynes College

University of Kent,

Canterbury, Kent.

CT2 7NP. UK.

E.J.Cartwright@kent.ac.uk

Myrna Wooders

Department of Economics, Vanderbilt University

Nashville, TN 37235 USA

and

Department of Economics, University of Warwick

Coventry CV4 7AL, UK

M.Wooders@warwick.ac.uk

March 20, 2005

First WP Version, University of Warwick #686, October 2003

This Version: March 2005

*We are indebted to Phil Reny and Frank Page for tutelage on Reny's 1999 results on the existence of equilibrium in discontinuous games and express our appreciation for their intellectual generosity.

Abstract: Treating games of incomplete information with countable sets of actions and types and finite but large player sets we demonstrate that for every mixed strategy profile there is a pure strategy profile that is ‘ ε -equivalent’. Our framework introduces and exploits a distinction between crowding attributes of players (their external effects on others) and their taste attributes (their payoff functions and any other attributes that are not directly relevant to other players). The main assumption is a ‘large game’ property,’ dictating that the actions of relatively small subsets of players cannot have large effects on the payoffs of others. Since it is well known that, even allowing mixed strategies, with a countable set of actions a Nash equilibrium may not exist, we provide an existence of equilibrium theorem. The proof of existence relies on a relationship between the ‘better reply security’ property of Reny (1999) and a stronger version of the large game property. Our purification theorem are based on a new mathematical result, of independent interest, applicable to countable strategy spaces.

1 Motivation for the study of purification

The concept of a Nash equilibrium is at the heart of much of economics and game theory. It is thus fundamental to question when Nash equilibrium provides a good description of human behavior. A number of challenges are posed by the evidence. Experimental evidence, for example, supports the view that individuals typically do not play mixed strategies (cf., Friedman 1996) and if they do, there may be serial correlation.¹ Challenges are

¹This has been demonstrated in a number of papers; see Walker and Wooders (2001) for a recent contribution and references therein.

also posed by the observed imitative nature of human behavior (cf., Offerman, Potters and Sonnemans 2002). The importance of equilibrium in pure strategies is evidenced by numerous papers in the literatures of game theory and economics (from, for example, Rosenthal 1973 to Cripps, Keller and Rady 2002).

In this paper we demonstrate that strategy profiles can be ‘purified’ in a wide class of games with a large but finite player set. Informally, a (mixed) strategy profile can be purified if there exists a pure strategy profile that yields approximately the same payoffs to all players. The main assumption required is a ‘large game’ property,² dictating that the actions of relatively small subsets of players cannot have large effects on the payoffs of others. As a corollary of our purification results we obtain that, for the class of games considered, ‘close’ to any Bayesian (Bayesian-Nash) equilibrium in mixed strategies is an approximate equilibrium in pure strategies. Our purification results are obtained in a setting where the existence of Bayesian equilibrium, even in mixed strategies, is not immediate from existing results; this motivates introduction of a stronger version of the large game property under which we demonstrate the existence of an exact equilibrium in mixed strategies. Our proof of existence relies on a relationship between the ‘better reply security’ property of Reny (1999) and the large game property.

Within our framework a player is characterized by his attribute, a point in a given set of attributes. An important feature incorporated into our model is a distinction between the crowding attribute of a player and his taste attribute.² A player’s crowding attribute reflects those characteristics of the player that directly affect other players – for example, whether one

²This terminology is taken from Conley and Wooders (2001) and their earlier papers.

chooses to go to a particular club may depend on the gender and composition of the membership and how attractive one finds a particular economics department may depend on the numbers of faculty engaged in various areas of research. We assume, in this paper, that the space of crowding attributes is a compact metric space but no assumptions are made on the space of taste attributes.

We treat games of imperfect information. Thus, as well as having a certain attribute, ‘nature’ randomly assigns each player a (Harsanyi) type, as in a standard game of incomplete information. We allow a countable set of pure actions and a countable number of types. A new mathematical result, allowing us to approximate a mixed strategy profile by a pure strategy profile in which each player plays a pure strategy in the support of his mixed strategy, underlies our purification results and allows the non-finiteness of strategy and type sets.

In prior research we addressed the question of whether social conformity – that is, roughly, situations where most individuals imitate similar individuals – can be consistent with approximate Nash equilibrium.³ It was assumed, throughout this research, that social conformity requires the use of pure strategies. In this paper, we treat in isolation the most basic question – the existence of an approximate equilibrium pure strategies. The framework of the current paper is, in important respects, more general than that treated in our prior research. In particular, our earlier work treated finite action and finite type sets and a compact set of attributes. As discussed in Section 5, with finite action and type sets (and at most a finite number of attributes), purification results have already been obtained. But a finite

³Wooders, Cartwright and Selten (2001).

number of types, especially of taste types, may be a strong restriction. It is therefore crucial to consider, as we do in this paper, a framework where the set of taste attributes need not be compact and the set of strategies need not be finite nor compact.

It is especially noteworthy that to the best of our knowledge the first results relating to purification of Bayesian equilibrium are provided in Kalai (2004). In particular, Kalai shows that for games with many players, with high probability the play of a Bayesian equilibrium will yield, ex-post, an approximate Nash equilibrium of the game of complete information that results after player types are revealed. We compare our purification results to Kalai's 'ex-post' results in our discussion of the literature in Section 5. We comment here on a related literature concerning purification of Bayesian equilibria in finite games with imperfect information. This literature demonstrates that if there sufficient uncertainty over the signals (or types) that players receive then any mixed strategy can be purified (e.g. Radner and Rosenthal 1982, Aumann et. al. 1983). Given that we model games of imperfect information it is important to emphasize that we do not treat this form of purification and our results also hold for games of perfect information.

We proceed as follows: Section 2 introduces definitions and notation. In Section 3 we treat purification, providing a simple example before defining the large game property and providing our main results. In Section 4 we provide a brief discussion of the literature and Section 5 concludes the paper. Additional proofs are provided in an Appendix.

2 Bayesian games and noncooperative pregames

We begin this section by defining a Bayesian game and its components. The pregame framework is then introduced and we demonstrate how Bayesian games can be induced from a pregame. Next, we consider the strategies available to players in a Bayesian game and discuss expected payoffs. We finish by defining Nash equilibrium and purification.

2.1 A Bayesian game

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

Let $N = \{1, \dots, n\}$ be a finite player set, let \mathcal{A} denote a countable set of *actions* and let \mathcal{T} denote a countable set of *types*.⁴ ‘Nature’ assigns each player a type. Informed of his own type but not the types of his opponents, each player chooses an action. Let $A \equiv \mathcal{A}^N$ be the set of *action profiles* and let $T \equiv \mathcal{T}^N$ be the set of *type profiles*. Given action profile a and type profile t we interpret a_i and t_i as respectively the action and type of player $i \in N$.

A player’s payoff depends on the actions and types of players. Formally, in game Γ , for each player $i \in N$ there is given a *utility function* $u_i : A \times T \rightarrow \mathbb{R}$. In interpretation $u_i(a, t)$ denotes the payoff of player i if the action profile is a and the type profile t . Let u denote the set of utility functions.

⁴In fact, we could allow the sets of actions and types to each consist of a countable collection of compact sets. This, however, would increase complexity of proofs without substantial gain in understanding of the fundamental problems.

When taking an action, a player does not know the types of other players. Thus, once informed of his own type he selects an action based on his beliefs about the types of the other players. These beliefs are represented by a function p_i where $p_i(t_{-i}|t_i)$ denotes the probability that player i assigns to type profile (t_i, t_{-i}) given that he is of type t_i . Throughout we will assume *consistent beliefs*. Formally, for some probability distribution over type profiles g , we assume:

$$p_i(t_{-i}|t_i) = \frac{g(t_i, t_{-i})}{\sum_{t'_{-i} \in T_{-i}} g(t_i, t'_{-i})} \quad (1)$$

for all $i \in N$ and $t_i \in \mathcal{T}$.⁵ We denote by \mathcal{T}_i the set of types $t_i \in \mathcal{T}$ such that $\sum_{t'_{-i} \in T_{-i}} g(t_i, t'_{-i}) > 0$.

2.2 Noncooperative pregames

To treat a family of games all induced from a common strategic situation we make use of a pregame. A *pregame* is given by a tuple $\mathcal{G} = (\Omega, \mathcal{A}, \mathcal{T}, b, h)$, consisting of an *attribute space* Ω , countable sets of actions and types \mathcal{A} and \mathcal{T} , a *universal beliefs function* b and a *universal payoff function* h . We introduce and define in turn the components Ω , b and h .

A space of *player attributes* is denoted by Ω . An attribute $\omega \in \Omega$ is composed of two elements - a taste attribute and a crowding attribute. In interpretation, the crowding attribute of a player describes those characteristics that might affect other players, for example, gender, ability to do the salsa, educational level, and so on. A taste attribute describes that players preferences. Let \mathcal{P} denote a set of *taste attributes* and let \mathcal{C} denote a set of

⁵We do not require (1) to hold if $\sum_{t'_{-i} \in T_{-i}} g(t_i, t'_{-i}) = 0$; i.e. if there is no probability that player i is type t_i .

crowding attributes. We assume that $\mathcal{P} \times \mathcal{C} = \Omega$. We will assume throughout that \mathcal{C} is a compact metric space (while no assumptions are made on \mathcal{P}).

We next introduce the concepts of universal beliefs and a universal payoff function, which together induce a set of beliefs and a payoff function for each player in any game induced from the pregame. Denote by D the set of all mappings from $\mathcal{C} \times \mathcal{T}$ into \mathbb{Z}_+ (the non-negative integers). A member of D is called a *type function*. A type function $d \in D$ will be interpreted as listing the number of players of crowding type $c \in \mathcal{C}$ with type $t \in \mathcal{T}$ in some induced game. A *universal beliefs function* b maps D into $[0, 1]$ where $b(d)$ gives the probability of type profile d in some induced game.

We denote by W the set of all mappings from $\mathcal{C} \times \mathcal{A} \times \mathcal{T}$ into \mathbb{Z}_+ . A member of W is called a *weight function*. A weight function $w \in W$ will be interpreted as listing the number of players of crowding type $c \in \mathcal{C}$ and type $t \in \mathcal{T}$ who are playing action $a \in \mathcal{A}$ in some induced game. A *universal payoff function* h maps $\Omega \times \mathcal{A} \times \mathcal{T} \times W$ into \mathbb{R}_+ . In interpretation h will give the payoff to a player where his payoff depends on his attribute, his action choice, his type and the action choices, types and attributes of the complementary player set as described by a weight function.

2.3 Populations and induced games

Let N be a finite *player set*. A function α mapping from N to Ω is called an *attribute function*. The pair (N, α) is a *population*. While an attribute consists of a taste attribute/crowding attribute pair, crowding attributes play a special role and require separate notation. Thus, given an attribute function α we denote by κ the projection of α onto \mathcal{C} . Given population (N, α) the attribute of player i is therefore $\alpha(i)$ and the crowding attribute

of player i is $\kappa(i)$ satisfying $\alpha(i) = (\pi, \kappa(i))$ for some $\pi \in \mathcal{P}$.

A population (N, α) induces (through the pregame) a Bayesian game $\Gamma(N, \alpha) \equiv (N, A, T, g^\alpha, u^\alpha)$ as we now formalize. Given the population (N, α) we say that weight function $w_{\alpha, a, t} \in W$ is *relative to action profile a and type profile t* if,

$$w_{\alpha, a, t}(c, a^l, t^z) = \left| \left\{ i \in N : \kappa(i) = c, a_i = a^l \text{ and } t_i = t^z \right\} \right|$$

for each $(c, a^l, t^z) \in \mathcal{C} \times A \times \mathcal{T}$. Thus, $w(c, a^l, t^z)$ denotes the number of players with crowding attribute c and type t^z who play action a^l . The function h determines the payoff function u_i^α of each player $i \in N$; formally, given action profile $a \in A$ and type profile $t \in T$,

$$u_i^\alpha(a, t) = h(\alpha(i), a_i, t_i, w_{\alpha, a, t}).$$

We say that type function $d_{\alpha, t} \in D$ is *relative to type profile t* if,

$$d_{\alpha, t}(c, t^z) = |\{i \in N : \kappa(i) = c \text{ and } t_i = t^z\}|.$$

Thus, $d_{\alpha, t}(c, t)$ denotes the number of players with crowding attribute c and type t^z .⁶ The function b determines the beliefs of players; formally, players are assumed to have consistent beliefs with respect to function g^α where,

$$g^\alpha(t) = b(d_{\alpha, t})$$

for any type profile $t \in T$.⁷

⁶Note that $d_{\alpha, t}$ is a projection of $w_{\alpha, a, t}$ onto $\Omega \times \mathcal{T}$.

⁷Note the differences between functions g^α and b . Function g^α is defined relative to a population (N, α) and its domain is \mathcal{T}^N . Function b , however, is defined independently of any specific game and has domain D . Thus, summing g^α over its domain gives a value of one - because it describes a unique population - while the sum of b over its domain is non-finite - because it describes beliefs for any population.

2.4 Strategies and expected payoffs

Take as given a population (N, α) and induced Bayesian game $(N, A, T, g^\alpha, u^\alpha)$. Knowing his own type but not those of his opponents, a player chooses an action. A *pure strategy* details the action a player will take for each type $t^z \in \mathcal{T}$ and is given by a function $s^k : \mathcal{T} \rightarrow \mathcal{A}$ where $s^k(t^z)$ is the action played by the player if he is of type t^z . Let \mathcal{S} denote the set of strategies.

A (mixed) *strategy* is given by a probability distribution over the set of pure strategies. The set of strategies is thus $\Delta(\mathcal{S})$. Given a strategy x we denote by $x(k)$ the probability that a player chooses pure strategy $k \in \mathcal{S}$ and we denote by $x(a^l|t^z)$ the probability that a player chooses action a^l given that he is of type t^z . Let $\Sigma = \Delta(\mathcal{S})^N$ denote the set of *strategy profiles*. We refer to a strategy profile m as *degenerate* if for all $i \in N$ and $t^z \in \mathcal{T}$ there exists some a^l such that $m_i(a^l|t^z) = 1$.

We assume that players are motivated by expected payoffs.⁸ Given a strategy profile σ , a type $t^z \in \mathcal{T}_i$ and beliefs about the type profile p_i^α the probability that player i puts on the action profile-type profile pair $a = (a_1, \dots, a_n)$ and $t = (t_1, \dots, t_{i-1}, t^z, t_{i+1}, \dots, t_n)$ is given by:

$$\Pr(a, t_{-i}|t^z) \stackrel{\text{def}}{=} p_i^\alpha(t_{-i}|t^z)\sigma_1(a_1|t_1)\dots\sigma_i(a_i|t^z)\dots\sigma_n(a_n|t_n).$$

Thus, given any strategy profile σ , for any type $t^z \in \mathcal{T}$ and any player i of type t^z , the expected payoff of player i can be calculated. Let $U_i^\alpha(\cdot|t^z) : \Sigma \rightarrow \mathbb{R}$ denote the expected utility function of player i conditional on his

⁸We use the vNM assumption for convenience but our results do not depend on it: The large game property is sufficiently strong to obtain our results but does not imply the vNM assumption.

type being t^z where:

$$U_i^\alpha(\sigma|t^z) \stackrel{\text{def}}{=} \sum_{a \in A} \sum_{t_{-i} \in T_{-i}} \Pr(a, t_{-i}|t^z) u_i^\alpha(a, t_z, t_{-i}).$$

2.5 Purification and Bayesian equilibrium

Given a game $\Gamma(N, \alpha)$ we say that two strategy profiles σ and m are ε -equivalent if, for all $i \in N$, $x \in \Delta(\mathcal{S})$ and $t^z \in \mathcal{T}_i$:

$$|U_i^\alpha(x, m_{-i}|t^z) - U_i^\alpha(x, \sigma_{-i}|t^z)| \leq \varepsilon.$$

We say that a strategy profile σ can be ε -purified if there exists a strategy profile m that is degenerate, ε -equivalent to σ and satisfies $\text{support}(m_i) \subset \text{support}(\sigma_i)$ for all $i \in N$.⁹

The standard definition of a Bayesian equilibrium applies. A strategy profile σ is a *Bayesian ε -equilibrium* (or informally an approximate Bayesian equilibrium) if and only if:

$$U_i^\alpha(\sigma_i, \sigma_{-i}|t^z) \geq U_i^\alpha(x, \sigma_{-i}|t^z) - \varepsilon$$

for all $x \in \Delta(\mathcal{S})$, all $t^z \in \mathcal{T}_i$ and for all $i \in N$. We say that a Bayesian ε equilibrium m is a *Bayesian Nash ε -equilibrium in pure strategies* if m is degenerate.

⁹A related notion of ε -purification was introduced by Aumann et. al. (1983). There, the notion of ε -purification is relative to strategies and not strategy vectors. Thus, two strategies p and t are ε -equivalent for player i if $|U_i^\alpha(p, \sigma_{-i}) - U_i^\alpha(t, \sigma_{-i})| < \varepsilon$ for any $\sigma_{-i} \in \Sigma^{N \setminus \{i\}}$. This definition proves useful in considering games of incomplete information but is too restrictive to be of use in considering games of complete information.

3 Purification

Before providing our main results it may be useful to provide a simple example:

Example 1: There are two crowding attributes - *rich* and *poor*. Players must choose one of two pure strategies or locations A and B . A poor player prefers living with rich players and thus his payoff is equal to the proportion of rich players whose choice of location he matches. A rich player prefers to not live with poor players and thus his payoff is equal to the proportion of poor players whose choice of location he does not match.

Any game induced from this pregame has a Nash equilibrium. It is simple to see, however, that if there exists an odd number of either rich or poor players then there does not exist a Nash equilibrium in pure strategies. Also, if either the number of rich players or the number of poor players is small then there need not exist an approximate Nash equilibrium in pure strategies, no matter how large the total population.

Theorem 2 demonstrates that if a pregame satisfies a large game property then, in any induced game with sufficiently many players, any Nash equilibrium can be approximately purified. The pregame of Example 1 does not satisfy the large game property; the large game property requires that any small group of players have diminishing influence in populations with a larger player set.

3.1 Approximating mixed strategy profiles by pure strategy profiles

Our first result, Theorem 1, shows that given any strategy profile σ , there exists a degenerate strategy profile m such that (i) each player i is assigned a pure strategy k in the support of σ_i , and (ii) the number of players who play each pure strategy k is ‘close’ to the expected number who would have played k given strategy profile σ . With this result in hand our main results can be easily proved. We note now that, in the application of Theorem 1 in the proof of Theorem 2, the strategy profile σ is not (necessarily) thought of as ‘the strategy profile of the population’ but more as the strategy profile restricted to those players who have the same crowding attribute.

Theorem 1: For any strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ there exists a degenerate strategy profile $m = (m_1, \dots, m_n)$ such that:

$$\text{support}(m_i) \subset \text{support}(\sigma_i) \tag{2}$$

for all i and:

$$\left| \sum_{i=1}^n m_i(k) - \sum_{i=1}^n \sigma_i(k) \right| \leq 1 \tag{3}$$

for all $k \in S$.

Observe that if σ , in Theorem 1, were a Bayesian equilibrium, then Theorem 1 states that there is an approximating pure strategy profile m where *every* player plays a pure strategy in his best response set for σ . This is crucial in proving our subsequent theorems in that it allows us to ‘aggregate’ the strategies of players who have the same crowding attribute yet potentially different taste attributes. We highlight that the related but

distinct Shapley-Folkman Theorem will not suffice for our purposes in that it does not allow us to treat non-finite strategy sets.¹⁰

3.2 Continuity in crowding attributes

To derive our purification results we make use of a natural and mild continuity assumption on crowding attributes, introduced in Wooders, Cartwright and Selten (2001), that will be assumed throughout. Given the strategy choices of other players, it is assumed that each player is nearly indifferent to a minor perturbation of the crowding attributes of other players (provided his own crowding attribute and the strategy choices of players are unchanged). Formally:

Continuity in crowding attributes: We say that a pregame \mathcal{G} satisfies continuity in crowding attributes if: for any $\varepsilon > 0$, any two populations (N, α) and $(N, \bar{\alpha})$ and any strategy profile $\sigma \in \Sigma^N$ if:

$$\max_{j \in N} \text{dist}(\kappa(j), \bar{\kappa}(j)) < \varepsilon$$

then for any $i \in N$ where $\alpha(i) = \bar{\alpha}(i)$:

$$|U_i^\alpha(\sigma_i, \sigma_{-i}|t^z) - U_i^{\bar{\alpha}}(\sigma_i, \sigma_{-i}|t^z)| < \varepsilon$$

all $t^z \in \mathcal{T}_i$. Where ‘*dist*’ is the metric on the space of crowding attributes \mathcal{C} .

Note that it is not essential to have the same bound of ε in both the above expressions, but it does simplify notation. The definition of continuity in crowding attributes takes the strategy profile as held constant. Thus, the attributes of players may change but their strategies do not. For example,

¹⁰As discussed in Section 4, Rashid (1983) does make use of the Shapley-Folkman Theorem in proving a special case of our Theorem 2 in which the strategy set is finite.

the wealths of other players may change by some small amounts while their strategic choices, such as location of residence, are held constant. Continuity in crowding attributes appears to be a mild assumption.

3.3 Large game property

To define the large game property, some additional notation and definitions are required. Denote by EW the set of functions mapping $\mathcal{C} \times \mathcal{A} \times \mathcal{T}$ into \mathbb{R}_+ , the set of non-negative reals. We refer to $ew \in EW$ as an *expected weight function*. Given a population (N, α) we say that an expected weight function $ew_{\alpha, \sigma}$ is relative to strategy profile σ if and only if:

$$ew_{\alpha, \sigma}(c, a^l, t^z) = \sum_{a \in \mathcal{A}} \sum_{t \in \mathcal{T}} w_{\alpha, a, t}(c, a^l, t^z) \Pr(a, t)$$

for all ω, a^l and t^z . Thus, $ew_{\alpha, \sigma}(\omega, a^l, t^z)$ denotes the *expected* number of players of crowding-attribute c who will have type t^z and play action a^l . Note that this expectation is taken before any player is aware of his type.

Fix a population (N, α) . Let EW_{α} denote the set of expected weight functions that may be realized given population (N, α) . We define a metric on the space EW_{α} :

$$dist(ew, eg) = \frac{1}{|N|} \sum_{a^l \in \mathcal{A}} \sum_{t^z \in \mathcal{T}} \sum_{c \in \mathcal{C}} \left| ew(c, a^l, t^z) - eg(c, a^l, t^z) \right|$$

for any $ew, eg \in EW_{\alpha}$. Thus, two expected weight functions are ‘close’ if the expected proportion of players with each crowding attribute and each type playing each action are close. We can now state our main assumption:

Large game property: We say that a pregame \mathcal{G} satisfies the large game property if: for any $\varepsilon > 0$, any population (N, α) and any two strategy

profiles $\sigma, \bar{\sigma} \in \Sigma^N$ with expected weight functions $ew_{\alpha, \sigma}, eg_{\alpha, \bar{\sigma}}$ satisfying:

$$\text{dist}(ew_{\alpha, \sigma}, eg_{\alpha, \bar{\sigma}}) < \varepsilon$$

if $\sigma_i = \bar{\sigma}_i$ then:

$$|U_i^\alpha(\sigma_i, \sigma_{-i}|t^z) - U_i^\alpha(\bar{\sigma}_i, \bar{\sigma}_{-i}|t^z)| < \varepsilon$$

for all $t^z \in \mathcal{T}_i$.

If a pregame satisfies the large game property then we can think of games induced from the pregame as satisfying two conditions on payoff functions:

1. A player is nearly indifferent to a change in the proportion of players of each attribute playing each pure strategy (provided his own strategy is unchanged); thus, any one individual has near-negligible influence over the payoffs of other players.
2. A player is ‘risk neutral’ in the sense that the expected weight function largely determines his payoff; thus two strategy profiles that induce the same expected weight function give a similar payoff.

The first condition is reflective of the type of game under consideration and is crucial to obtaining our main result; Example 1, for instance, does not satisfy the large game property in this respect. The second condition is relatively mild given that we consider games with many players; it follows, for example, from the law of large numbers that in the case of a finite strategy set, with high probability, in a game with many players the realized weight function will be close to the expected weight function.¹¹

¹¹Thus, it is not so much that players are risk neutral but rather that there is little risk.

Note that the large game property relates to changes in the strategies of players while their attributes do not change; this contrasts with the assumption of continuity in crowding attributes that relates to changes in attributes while strategies do not change. As a consequence a pregame may satisfy the large game property and yet there need not be continuity in attributes and vice-versa.

3.4 Approximate purification

Our central result demonstrates that in sufficiently large games any strategy profile can be approximately purified.

Theorem 2: Consider a pregame $\mathcal{G} = (\Omega, \mathcal{A}, \mathcal{T}, b, h)$ satisfying continuity in crowding attributes and the large game property. Given any real number $\varepsilon > 0$ there is an integer $\eta(\varepsilon)$ with the property that in any induced game $\Gamma(N, \alpha)$ satisfying $|N| > \eta(\varepsilon)$ any strategy profile can be ε -purified.

Proof: Suppose not. Then there is some $\varepsilon > 0$ such that for each integer ν there is an induced game $\Gamma(N^\nu, \alpha^\nu)$ with $|N^\nu| > \nu$ and strategy profile σ^ν that cannot be ε -purified.

Use compactness of \mathcal{C} to write \mathcal{C} as the disjoint union of a finite number of non-empty subsets $\mathcal{C}_1, \dots, \mathcal{C}_Q$, each of diameter less than $\frac{1}{3}\varepsilon$. For each $q = 1, \dots, Q$, choose and fix a point $c_q \in \mathcal{C}_q$. For each ν , without changing taste attributes of players, we define the crowding attribute function $\bar{\kappa}^\nu$ by its coordinates $\bar{\kappa}^\nu(\cdot)$ as follows:

$$\text{for each } j \in N, \bar{\kappa}^\nu(j) = c_q \text{ if and only if } \kappa(j) \in \mathcal{C}_q.$$

Define new attribute functions $\bar{\alpha}^\nu$ by $\bar{\alpha}^\nu(j) = (\pi(j), \bar{\kappa}^\nu(j))$ when $\alpha^\nu(j) = (\pi(j), \kappa^\nu(j))$ for each $j \in N^\nu$. By applying Theorem 1 to each $c \in \bar{\kappa}^\nu(N)$, i.e.

c_1, \dots, c_Q it follows that there exists a sequence $\{m^\nu\}$ of degenerate strategy profiles such that:

1. for all $c \in \mathcal{C}$, $a^l \in \mathcal{A}$ and $t^z \in \mathcal{T}$

$$\lim_{\nu \rightarrow \infty} \frac{eg_{\bar{\alpha}^\nu, m^\nu}^\nu(c, a^l, t^z)}{|N^\nu|} = \lim_{\nu \rightarrow \infty} \frac{ew_{\bar{\alpha}^\nu, \sigma^\nu}^\nu(c, a^l, t^z)}{|N^\nu|}, \text{ and} \quad (4)$$

2. for all ν and $i \in N^\nu$,

$$\text{support}(m_i^\nu) \subset \text{support}(\sigma_i^\nu). \quad (5)$$

Pick an arbitrary ν and player $i \in N^\nu$. Consider the attribute function $\bar{\alpha}^\nu$ where $\bar{\alpha}^\nu(i) = \alpha^\nu(i)$ and $\bar{\alpha}^\nu(j) = \bar{\alpha}^\nu(j)$ for all $j \notin i$. By continuity in crowding attributes:

$$\left| U_i^{\alpha^\nu}(x, \sigma_{-i}^\nu | t^z) - U_i^{\bar{\alpha}^\nu}(x, \sigma_{-i}^\nu | t^z) \right| < \frac{\varepsilon}{3}$$

for all $t^z \in \mathcal{T}_i$ and $x \in \Delta(\mathcal{S})$, and:

$$\left| U_i^{\alpha^\nu}(x, m_{-i}^\nu | t^z) - U_i^{\bar{\alpha}^\nu}(x, m_{-i}^\nu | t^z) \right| < \frac{\varepsilon}{3}$$

for any $t^z \in \mathcal{T}_i$ and $x \in \Delta(\mathcal{S})$. In view of (4) and the large game property it is clear if ν was sufficiently large:

$$\left| U_i^{\bar{\alpha}^\nu}(x, \sigma_{-i}^\nu | t^z) - U_i^{\bar{\alpha}^\nu}(x, m_{-i}^\nu | t^z) \right| < \frac{\varepsilon}{3}$$

for any $t^z \in \mathcal{T}_i$ and $x \in \Delta(\mathcal{S})$. Thus, for ν sufficiently large and for any $i \in N^\nu$:

$$\left| U_i^{\alpha^\nu}(x, m_{-i}^\nu | t^z) - U_i^{\alpha^\nu}(x, \sigma_{-i}^\nu | t^z) \right| < \varepsilon$$

for any $t^z \in \mathcal{T}_i$ and $x \in \Delta(\mathcal{S})$. This gives the desired contradiction. ■

An immediate application of Theorem 2 is the following result showing that in sufficiently large games (approximate) Bayesian equilibrium can be approximately purified.

Corollary 1: Consider a pregame $\mathcal{G} = (\Omega, \mathcal{A}, \mathcal{T}, b, h)$ satisfying continuity in crowding attributes and the large game property. Given any real numbers ε and λ where $\varepsilon > \lambda > 0$ there is an integer $\eta(\varepsilon, \lambda)$ with the property that, for any induced game $\Gamma(N, \alpha)$ where $|N| > \eta(\varepsilon, \lambda)$ and for any Bayesian λ -equilibrium σ of game $\Gamma(N, \alpha)$, there exists a Bayesian ε -equilibrium in pure strategies m that is an ε -purification of σ .

Proof: Let $\theta = \frac{1}{2}(\varepsilon - \lambda)$. By Theorem 2 there exists integer $\eta(\theta)$ such that in any induced game $\Gamma(N, \alpha)$ where $|N| > \eta(\theta)$ any strategy profile can be θ -purified. Consider a game $\Gamma(N, \alpha)$ where $|N| > \eta(\theta)$ and let σ be a Bayesian λ -equilibrium of that game. Thus,

$$U_i^\alpha(m_i, \sigma_{-i}|t^z) \geq U_i^\alpha(x, \sigma_{-i}|t^z) - \lambda$$

for all $i \in N, x \in \Delta(\mathcal{S})$ and $m_i \in \Delta(\mathcal{S})$ where $\text{support}(m_i) \subset \text{support}(\sigma_i)$. Given that σ can be θ -purified let m be a strategy profile that is degenerate and θ -equivalent to σ . Thus,

$$|U_i^\alpha(x, \sigma_{-i}|t^z) - U_i^\alpha(x, m_{-i}|t^z)| < \theta$$

for all $i \in N$ and $x \in \Delta(\mathcal{S})$. This implies,

$$U_i^\alpha(m_i, m_{-i}|t^z) \geq U_i^\alpha(x, m_{-i}|t^z) - \lambda - 2\theta$$

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

3.5 The existence of a Bayesian equilibrium

With a countable set of strategies, a Bayesian (or simply a Nash) equilibrium, even one in mixed strategies, may not exist. This is easy to see; consider, for example, a game of choosing integers where the prize goes to the player who announces the highest integer. Notice, however, that to apply Corollary 1 all we require is the existence of an *approximate* Bayesian equilibrium (in mixed strategies). Also, we only require the existence of an equilibrium in games satisfying the large game property; observe, for instance, that the large game property is not satisfied in games where the prize goes to the player announcing the highest integer.

We will restrict attention to games of complete information. Note that this implies the sets \mathcal{A} and \mathcal{S} are equivalent. For our existence result we need some assumption to ensure that the set of mixed strategies is compact. For specificity, let us assume that if $a \in \mathcal{A}$ then $a = \frac{1}{\ell}$ for some positive integer ℓ or $a = 0$. Note that with the Euclidean topology, \mathcal{A} is a closed set. A mixed strategy can be written as a sequence (x_1, x_2, x_3, \dots) where $x_\ell \in [0, 1]$ is the probability of playing $a_\ell := \frac{1}{\ell}$ and $1 - \sum_{\ell=1}^{\infty} x_\ell$ is the probability of playing $a = 0$. Note that $1 \geq \sum_{\ell=1}^{\infty} x_\ell \geq 0$. We use the following metric on the space of mixed strategies:

$$\rho(x, y) = \sum_{\ell=1}^{\infty} \frac{1}{2^i} \frac{|x_\ell - y_\ell|}{1 + |x_\ell - y_\ell|} \text{ for any } x, y \in \Delta(\mathcal{S}).$$

Note that metric ρ makes $\Delta(\mathcal{S})$ a metric space furnished with the product topology. It follows that, with metric ρ , the set of mixed strategies $\Delta(\mathcal{S})$ is a compact metric space.

We make a form of continuity assumption on payoffs that strengthens

the large game property. We define a metric on the space EW_α :¹²

$$dist^*(ew, eg) = \frac{1}{|N|} \sum_{c \in \mathcal{C}} \sum_{l \in \mathcal{A}} \frac{1}{2^j} \frac{|ew(c, l) - eg(c, l)|}{1 + |ew(c, l) - eg(c, l)|}$$

for any $ew, eg \in EW_\alpha$. We say that a pregame \mathcal{G} satisfies the *strong large game property* if: for any $\varepsilon > 0$, any population (N, α) and any two strategy profiles $\sigma, \bar{\sigma} \in \Sigma^N$ with expected weight functions $ew_{\alpha, \sigma}, eg_{\alpha, \bar{\sigma}}$ satisfying:

$$dist^*(ew_{\alpha, \sigma}, eg_{\alpha, \bar{\sigma}}) < \varepsilon$$

if $\sigma_i = \bar{\sigma}_i$ then:

$$|U_i^\alpha(\sigma_i, \sigma_{-i}) - U_i^\alpha(\bar{\sigma}_i, \bar{\sigma}_{-i})| < \varepsilon.$$

We now make use of a result due to Reny (1999) to provide a Nash equilibrium existence result.

Theorem 3: Consider a pregame $\mathcal{G} = (\Omega, \mathcal{A}, \mathcal{T}, b, h)$ satisfying the strong large game property and where $|\mathcal{T}| = 1$ and $\mathcal{A} = \mathbb{Q}_1$. Any induced game $\Gamma(N, \alpha)$ has a Nash 0-equilibrium (in mixed strategies).

Proof: We introduce some definitions from by Reny (1999). A player i can *secure* a payoff of $\alpha \in \mathbb{R}$ at σ if there exists strategy x such that $U_i(x, \sigma'_{-i}) \geq \alpha$ for all σ'_{-i} in some open neighborhood of σ_{-i} . A game $\Gamma(N, \alpha)$ is *better reply secure* if whenever (σ^*, U^*) is in the closure of the graph of its vector payoff function and σ^* is not a Nash equilibrium, some player i can secure a payoff strictly above U_i^* at σ^* . From Theorem 3.1 of Reny (1999) we obtain that there exists a (mixed strategy) Nash 0-equilibrium in any game $\Gamma(N, \alpha)$ that has a compact set of mixed strategies and is better reply secure. [The Theorem also requires quasi-concavity of

¹²Given we are treating games of complete information we drop the t^z notation.

payoffs in own strategy but this trivially holds here given we are assuming mixed strategies and von-Neumann Morgenstern payoff functions.]

It remains to show that game $\Gamma(N, \alpha)$ is better reply secure. This, however, follows from the strong large game property. Let σ^* be a non Nash equilibrium strategy profile with (σ^*, U^*) in the closure of the graph of the vector payoff function. There must exist some player $i \in N$, strategy x and real number $\varepsilon > 0$ such that,

$$U_i^\alpha(x, \sigma_{-i}^*) > U_i^* + \varepsilon. \quad (6)$$

By varying the the strategy profile σ^* by no more than a sufficiently small amount we obtain a set of strategy profiles Σ' with the property that,

$$\sum_{l \in \mathcal{A}} |ew_{\alpha, \sigma^*}(c, l) - eg_{\alpha, \sigma'}(c, l)| < \frac{\varepsilon}{|N|} \quad (7)$$

for all $\sigma' \in \Sigma'$. Note that (7) implies that $dist^*(ew_{\alpha, \sigma^*}, eg_{\alpha, \sigma'}) < \varepsilon$ for all $\sigma \in \Sigma'$. We then have, from the strong large game property and (6) that

$$U_i^\alpha(x, \sigma'_{-i}) > U_i^* \quad (8)$$

for all $\sigma' \in \Sigma'$. It follows that player i can secure a payoff greater than U_i^* in game $\Gamma(N, \alpha)$. ■

4 Some relationships to the literature

Two authors that provide related results on purification with large but finite player sets are Rashid (1983) and Kalai (2004). Kalai provides sufficient conditions for the existence of an approximate ex-post Nash equilibrium. An ex-post Nash equilibrium is a strategy vector that results, with high probability in a Nash equilibrium of the induced game of complete information

that is determined by the revelation of player types. In Cartwright and Wooders (2004) we demonstrate that if any realization of a strategy vector for a Bayesian game is, with high probability (ρ) an ε -Nash equilibrium of the induced game of complete information, then there is a purification of that strategy that is an approximate (α) equilibrium of the original game, where $\alpha = (1 + \rho)\varepsilon + \rho D$ and D is an upper bound on payoffs. Thus, Kalai's result, in combination with that of Cartwright and Wooders (2004) implies approximate purification in the sense defined in the current paper. In contrast to this paper and Wooders, Cartwright and Selten (2001), Kalai requires both a finite number of actions and a finite number of crowding types.¹³ See also Blonski (2004).

With a finite set of pure strategies and games of complete information, Rashid (1983) makes use of the Shapley-Folkman Theorem to prove his result on existence of approximate equilibrium in pure strategies. By assuming a linearity of payoff functions Rashid demonstrates that 'near' to any Nash equilibrium there is an approximate Nash equilibrium in which $|N| - K$ players use pure strategies (where K is the number of strategies) and K players may play mixed strategies. (See also Carmona 2004 who argues that an additional condition, equicontinuity of payoff functions for example, is

¹³Mas-Colell (1984) remarks that strategy sets can encode for a player's attribute. For example, the payoff function may be set up in such a way that a male would never rationally choose from a particular subset of strategies while a female may only rationally choose from that subset. Similarly, in games of incomplete information (as in Kalai 2004) a player's type may encode his attribute. If, however, the set of strategies and the set of types are finite, as in Mas-Colell and in Kalai, then at most a finite number of crowding attributes can be encoded. We remark that, in contrast to our research in this paper and also in Wooders, Cartwright and Selten (2001), these authors make no further use of dependence of payoffs on crowding attributes.

required). We reiterate that the Shapley-Folkman Theorem is not sufficient for our purposes in treating a non-finite set of pure strategies.

Many authors have contributed to the literature on the existence of a pure strategy non-cooperative equilibria in games with a continuum of players (including Schmeidler 1973, Mas-Colell 1984, Khan 1989, 1998, Khan et al. 1997, Pascoa 1993a, 1998 and Khan and Sun 1999).¹⁴ This literature, given various assumptions on the strategy space, has demonstrated the existence of a non-cooperative equilibrium when payoffs depend on opponent's strategies through the induced distribution over pure strategies. Our Theorem 2 can be seen as providing a finite analogue to some of these continuum results.

Within the literature on non-atomic games, the approach of Pascoa (1993a) appears most similar to our own. Pascoa (1993a) deals with non-anonymous games as introduced by Green (1984). A player in a non-anonymous game has a type (which could be thought as an attribute in our framework) and a player's payoff depends on his opponent's strategies through the distribution over types and pure strategies. More formally, let T denote a set of types and D the set of Borel probability measures over $T \times S$.¹⁵ The payoff to a player of type t from playing strategy s when the strategies of opponents is $\mu \in D$ is given by $v(t, s, \mu)$. To obtain his results Pascoa assumes that $v(t, \cdot, \cdot)$ is jointly continuous, with respect to the weak* topology on D .¹⁶ This corresponds to our assumption of a pregame that satisfies the large game property and continuity in crowding attributes.

¹⁴Note that these authors consider games of complete information with a continuum player set.

¹⁵Where S denotes as previously the set of strategies.

¹⁶Pascoa (1993a) assumes a compact metric space of strategies.

5 Conclusions

This paper introduces a framework for studying properties of strategic games with large but finite numbers of players. Our framework extends those already in the literature in a number of respects. The major innovations of the framework itself are our mathematical result (Theorem 1), allowing countable sets of actions and types, and the formalization of the separation of crowding and taste attributes of players. This separation plays a role in other research on noncooperative games, particularly on games with many players where similar players conform (see Wooders, Cartwright and Selten 2001 and Cartwright and Wooders 2003). Our purification result is the first to demonstrate approximate purification of Bayesian equilibrium in games with many players. Our existence of equilibrium result is, to the best of our knowledge, the first result allowing a countable strategy set.

6 Appendix

We introduce some additional notation. Let $a = (a_1, \dots, a_n)$, $b = (b_1, \dots, b_n) \in \mathbb{R}^n$. We write $a \geq b$ if and only if $a_i \geq b_i$ for all $i = 1, \dots, n$. Given any strategy profile σ let $M(\sigma)$ denote the set of strategy profiles such that $m \in M(\sigma)$ if and only if (1) m is degenerate and (2) $\text{support}(m_i) \subseteq \text{support}(\sigma_i)$ for all $i \in N$. It is immediate that $M(\sigma)$ is non-empty for any σ .

Lemma 2: Let $N = \{1, \dots, n\}$ be a finite set. For any strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and for any function $\bar{g} : S \rightarrow \mathbb{Z}_+$ such that $\sum_i \sigma_i \geq \bar{g}$, there exists $m \in M(\sigma)$ such that

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

Proof: Suppose the statement of the lemma is false. Then there exists a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and a function \bar{g} where $\sum_{i \in N} \sigma_i \geq \bar{g}$, such that, for any vector $m = (m_1, \dots, m_n) \in M(\sigma)$ there must exist at least one \hat{k} where $\hat{k} \in S$ and $\sum_i m_i(\hat{k}) < \bar{g}(\hat{k})$. For each vector $m \in M(\sigma)$ let L be defined as follows:

$$L(m) = \sum_{k \in S: \sum_i m_i(k) < \bar{g}_k} \left(\bar{g}(k) - \sum_i m_i(k) \right)$$

We note that $L(m)$ must be finite and positive for all m .¹⁷ Select $m^0 \in M(\sigma)$ for which $L(m)$ attains its minimum value over all $m \in M(\sigma)$. Intuitively the vector m^0 is ‘as close’ as we can get to satisfying the lemma. We remark that the method of proof will be one of ‘shuffling’ the pure strategies that players use so as to demonstrate the existence of a strategy profile m^* where $L(m^*) = L(m^0) - 1$. Providing the desired contradiction.

Pick a strategy \hat{k} such that $\bar{g}(\hat{k}) - \sum_i m_i^0(\hat{k}) > 0$. For any subset I of N let the set $S(I) \subset S$ be such that:

$$S(I) = \{ \hat{k} \} \cup \{ k \in S : m_i^0(k) = 1 \text{ for some } i \in I \}$$

We can now define sets N^t for $t = 0, 1, \dots$ as follows:

$$\begin{aligned} N^0 &= \{ i \in N : m_i^0(\hat{k}) = 1 \} \text{ and for all } t > 0 \\ N^t &= N^{t-1} \cup \left\{ \begin{array}{l} j \in N : \sigma_j(k) > 0 \text{ and } m_j^0(k) = 0 \\ \text{for some } k \in S(N^{t-1}) \end{array} \right\} \end{aligned}$$

Ultimately, for some $t^* \geq 1$ we must have that $N^{t^*+1} = N^{t^*} \equiv \bar{N}$. This is an immediate consequence of the finiteness of the player set. Let $S(\bar{N}) \equiv \bar{S}$.

¹⁷Note that the set of k such that $\sum_i m_i(k) < \bar{g}_k$ need not be finite. Given, however, that $\sum_k \sum_i \sigma_i(k) = |N|$ it must be that $\sum_k \bar{g}(k) \leq |N|$ and thus $L(m)$ is finite.

Consider any pure strategy $k^* \in \bar{S}$. The construction of \bar{N} and \bar{S} imply that there must exist a chain of players $\{i_1, \dots, i_{\bar{t}}\} \subset \bar{N}$ where (1) $m_{i_t}^0(k_t) = 1$ for $t = 1, \dots, \bar{t} - 1$, (2) $m_{i_{\bar{t}}}(k^*) = 1$, (3) $\sigma_{i_t}(k_{t-1}) > 0$ for $t = 2, \dots, \bar{t}$ and (4) $\sigma_{i_1}(\hat{k}) > 0$. Thus, there exists a vector $m^* \in M(\sigma)$ such that:

$$\begin{aligned} m_{i_1}^*(k_1) &= 0 \text{ and } m_{i_1}^*(\hat{k}) = 1, \\ m_{i_{\bar{t}}}^*(k^*) &= 0 \text{ and } m_{i_{\bar{t}}}^*(k_{\bar{t}-1}) = 1 \\ m_{i_t}^*(k_t) &= 0 \text{ and } m_{i_t}^*(k_{t-1}) = 1, \text{ for all } t = 2, \dots, \bar{t} - 1, \text{ and} \\ m_i^*(k) &= m_i^0(k) \text{ for all other } i \text{ and } k. \end{aligned}$$

Suppose that:

$$\sum_{i \in \bar{N}} m_i^0(k^*) > \bar{g}(k^*).$$

This implies that:

$$\sum_{i \in \bar{N}} m_i^0(k^*) \geq \bar{g}(k^*) + 1$$

and thus $L(m^*) = L(m^0) - 1$.

To avoid a contradiction we need:

$$\sum_{i \in \bar{N}} m_i^0(k) \leq \bar{g}(k). \quad (9)$$

for all $k \in \bar{S}$. Using the definition of \bar{S} there can exist no player $j \in N \setminus \bar{N}$ such that $\sigma_j(k) > 0$ for some $k \in \bar{S}$ unless $m_j^0(k) = 1$. This implies that:

$$\sum_{i \in N \setminus \bar{N}} m_i^0(k) \geq \sum_{i \in N \setminus \bar{N}} \sigma_i(k) \quad (10)$$

for all $k \in \bar{S}$. Using the definition of \bar{S} we have that:

$$\sum_{k \in \bar{S}} \sum_{i \in \bar{N}} m_i^0(k) \geq \sum_{k \in \bar{S}} \sum_{i \in \bar{N}} \sigma_i(k). \quad (11)$$

Combining (10) and (11) and using the statement of the lemma, we see that:

$$\sum_{k \in \bar{S}} \sum_{i \in N} m_i^0(k) \geq \sum_{k \in \bar{S}} \sum_{i \in N} \sigma_i(k) \geq \sum_{k \in \bar{S}} \bar{g}(k)$$

However, by assumption:

$$\bar{g}(\hat{k}) > \sum_{i \in N} m_i^0(\hat{k})$$

and also by assumption, $\hat{k} \in \bar{S}$. Thus, there must exist at least one $k \in \bar{S}$ such that:

$$\bar{g}(k) < \sum_{i \in N} m_i^0(k).$$

This contradicts (9) and completes the proof. ■

We introduce some additional 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 h (i.e. $\lfloor 9.5 \rfloor = 9$ and $\lceil 9.5 \rceil = 10$. Also note that $\lfloor 9 \rfloor = 9$ and $\lceil 9 \rceil = 10$).

Theorem 1: For any strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ there exists a degenerate strategy profile $m = (m_1, \dots, m_n)$ such that:

$$\text{support}(m_i) \subset \text{support}(\sigma_i) \tag{12}$$

for all i and:

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

for all $k \in S$.

Proof: Denote by $M^*(\sigma)$ the set of vectors $m = (m_1, \dots, m_n) \in M(\sigma)$ such that $\sum_i m_i(k) \geq \lfloor \sum_i \sigma_i(k) \rfloor$ for all k . By Lemma 2 this set is non-empty. Proving the Lemma thus amounts to showing that there exists a vector $m \in M^*(\sigma)$ such that $\lceil \sum_i \sigma_i(k) \rceil \geq \sum_i m_i(k)$ for all $s_k \in S$. Suppose not. Then, for every vector $m \in M^*(\sigma)$ there exists some strategy $k \in S$ such

that $\sum_i m_i(k) > \lceil \sum_i \sigma_i(k) \rceil$. For any strategy profile $m \in M^*(\sigma)$ define $L(m)$ by:

$$L(m) \equiv \sum_{k: \sum_i m_i(k) > \lceil \sum_i \sigma_i(k) \rceil} \left(\sum_{i=1}^n m_i(k) - \left\lceil \sum_{i=1}^n \sigma_i(k) \right\rceil \right).$$

We note that $L(m)$ is always positive and finite. Pick strategy profile $m^0 \in M^*(\sigma)$ where the value of $L(m)$ is minimized. We note that m^0 comes as close as any profile to satisfying the statement of the Lemma.

Denote by \widehat{k} a pure strategy such that:

$$\sum_{i=1}^n m_i^0(\widehat{k}) > \left\lceil \sum_{i=1}^n \sigma_i(\widehat{k}) \right\rceil.$$

We introduce sets S^t and N^t , $t = 0, 1, 2, \dots$, where:

$$N^0 = \{i : m_i^0(\widehat{k}) = 1\} \text{ and for } t > 0$$

and for $t > 0$,

$$S^t = \{k : \sigma_i(k) > 0 \text{ for some } i \in N^{t-1}\}$$

$$N^t = \{i : m_i^0(k) = 1 \text{ for some } k \in S^t\}.$$

For some t^* , $N^{t^*} = N^{t^*+1} \equiv \overline{N}$ and $S^{t^*} = S^{t^*+1} \equiv \overline{S}$. The construction of S^t and N^t imply that for any $k^* \in \overline{S}$ there must exist a set of players $\{i_0, i_1, \dots, i_{\overline{t}}\} \in \overline{N}$ such that:

$$m_{i_0}^0(\widehat{k}) = 1 \text{ and } \sigma_{i_0}(k_1) > 0,$$

$$m_{i_r}^0(k_r) = 1 \text{ and } \sigma_{i_r}(k_{r+1}) > 0 \text{ for all } r = 1, \dots, \overline{t} - 1,$$

$$m_{i_{\overline{t}}}^0(k_{\overline{t}}) = 1 \text{ and } \sigma_{i_{\overline{t}}}(k^*) > 0,$$

Suppose there exists $k^* \in \overline{S}$ such that:

$$\sum_{i=1}^n m_i^0(k^*) \leq \sum_{i=1}^n \sigma_i(k^*).$$

Given the chain of players $\{i_0, i_1, \dots, i_{\bar{t}}\} \in \overline{N}$ as introduced above, consider the vector m^* constructed as follows:

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

It is easily checked that the vector $m^* \in M(\sigma)$ leads to the desired contradiction given that $L(M^*) = L(m^0) - 1$. We note, however, that:

$$\sum_{i=1}^n \sum_{k \in \overline{S}} m_i^0(k) = |\overline{N}| = \sum_{i \in \overline{N}} \sum_{k \in \overline{S}} \sigma_i(k).$$

Thus, if:

$$\sum_{i=1}^n m_i^0(\hat{k}) > \sum_{i=1}^n \sigma_i(\hat{k}) \geq \sum_{i \in \overline{N}} \sigma_i(\hat{k})$$

there must exist some $k^* \in \overline{S}$ such that:

$$\sum_{i=1}^n m_i(k^*) \leq \sum_{i \in \overline{N}} \sigma_i(k^*) \leq \sum_{i=1}^n \sigma_i(k^*)$$

giving the desired contradiction. ■

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] Blonski, M. (2004) "Games in aggregated form," typescript.

- [3] Carmona, G. (2004) “On the purification of Nash equilibria of large games,” *Economics Letters* 85: 215-219.
- [4] Cartwright, E. and M. Wooders (2004) “On purification of equilibrium in Bayesian games and ex-post Nash equilibrium,” University of Warwick Department of Economics Working Paper no. 701.
- [5] Cartwright, E. and M. Wooders (2002) “Social conformity in arbitrary games with incomplete information,” University of Warwick Working Paper no. 672.
- [6] Conley, J. and M.H. Wooders (2001) “Tiebout economics with differential genetic types and endogenously chosen crowding characteristics,” *Journal of Economic Theory* 98: 261-294.
- [7] Cripps, M.W, G. Keller, and S. Rady (2002) “Strategic experimentation: The case of Poisson bandits,” May 2002 CES Working Paper #737.X
- [8] Friedman, D. (1996) “Equilibrium in evolutionary games: Some experimental results,” *The Economic Journal* 106: 1-25.
- [9] Green, E.J. (1984) “Continuum and finite-player noncooperative models of competition,” *Econometrica* vol.52, no. 4: 975-993.
- [10] Gross, R. (1996) *Psychology. The Science of Mind and Behaviour*, Hodder and Stoughton.
- [11] Harsanyi, J.C. (1973) “Games with randomly disturbed payoffs: a new rationale for mixed strategy equilibrium points,” *International Journal of Game Theory* 2: 1-23.

- [12] Kalai, E. (2004) "Large robust games," *Econometrica* 72, 1631-1665.
- [13] 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.
- [14] Khan, A. and Y. Sun (2002) "Noncooperative games with many players," *Handbook of Game Theory*, R. Auman and S. Hart, eds. North Holland.
- [15] Khan, A. and Y. Sun (1999) "Non-cooperative games on hyperfinite Loeb spaces," *Journal of Mathematical Economics* 31, 455-492.
- [16] 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.
- [17] Mas-Colell, A. (1984) "On a theorem of Schmeidler," *Journal of Mathematical Economics* 13: 206-210.
- [18] Offerman, T., J. Potters and J. Sonnemans (2002) "Imitation and belief learning in an oligopoly experiment," *Review of Economic Studies*, 69: 973-997.
- [19] Pascoa, M. (1998) "Nash equilibrium and the law of large numbers," *International Journal of Game Theory* 27: 83-92.
- [20] Pascoa, M. (1993a) "Approximate equilibrium in pure strategies for nonatomic games," *Journal of Mathematical Economics* 22: 223-241.
- [21] Pascoa, M. (1993b) "Noncooperative equilibrium and Chamberlinian monopolistic competition," *Journal of Economic Theory*, 69: 335-353.

- [22] Radner, R and R.W. Rosenthal (1982) "Private information and pure-strategy equilibria," *Mathematics of Operations Research* 7: 401-409
- [23] Rashid, S. (1983) "Equilibrium points of nonatomic games; Asymptotic results," *Economics Letters* 12: 7-10.
- [24] Rath, K.P., Y. Sun, S. Yamashige (1995) "The nonexistence of symmetric equilibria in anonymous games with compact action spaces," *Journal of Mathematical Economics* 24: 331-346.
- [25] Reny, P.J. (1999) "On the existence of pure and mixed strategy Nash equilibria in discontinuous games," *Econometrica* 67: 1029-1056.
- [26] Rosenthal, R.W (1973) "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, 2:65-67, 1973.
- [27] Schmeidler, D. (1973) "Equilibrium points of nonatomic games," *Journal of Statistical Physics* 7: 295-300.
- [28] Wooders, M., E. Cartwright and R. Selten (2001) "Social conformity and equilibrium in pure strategies in games with many players," University of Warwick Department of Economics Working Paper 589 (revised as 636).