On Purification of Equilibrium in Bayesian Games and Ex-Post Nash Equilibrium

by

Edward Cartwright and Myrna Wooders

Working Paper No. 05-W12

April 2005

DEPARTMENT OF ECONOMICS
VANDERBILT UNIVERSITY
NASHVILLE, TN 37235

www.vanderbilt.edu/econ

On purification of equilibrium in Bayesian games and ex-post Nash equilibrium

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, CV3 7AL, UK
myrna.wooders@vanderbilt.edu
www.myrnawooders.com

March 20, 2005

Abstract

Kalai (2002) demonstrates that in semi anonymous Bayesian games with sufficiently many players any Bayesian equilibrium is approximately ex-post Nash. In this paper we demonstrate that the existence of an approximate expost Nash property implies a purification result of the standard sort for the original Bayesian game. We also provide an example showing that the bound we obtain on the distance of a purified approximate equilibrium from an exact equilibrium is tight.

1 Ex post Nash equilibrium and purification of Bayesian equilibrium

For games of incomplete information, a number of papers have highlighted the importance of 'ex post' properties of an equilibrium, that is, the properties of an equilibrium after actions and player types are revealed. See, for example, Cremer and McLean (1985), Wilson (1987), Green and Laffont (1987) and Postlewaite and McLean (2002). Of special interest to the current paper, Kalai (2004) introduces a notion of approximate ex post stability and demonstrates that in semi-anonymous 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. For a strategy vector to be approximately ex post stable it must be the case that no player has a strong incentive to change his action *even after* he has observed the types and actions of all other players. Ex post stability appears to be a powerful concept and Kalai highlights a number of interesting consequences of his result.

In this paper we demonstrate that the existence of an ex post stable strategy vector implies the existence of an approximate Bayesian equilibrium in *pure strategies* of the original incomplete information game. One consequence is a purification result for Bayesian equilibrium- in any semi-anonymous game there exists an approximate Bayesian equilibrium in pure strategies. For games of complete information, purification results have been of interest since Schmeidler (1975); see, for examples. Mas-Colell (1984), Pascoa (1993,1998), and Khan, Rath and Sun. Khan and Sun (2004) provide a review of this literature. In Cartwright and Wooders (2002) we obtain a purification result for Bayesian equilibrium in games with many players. The current paper serves to connect the existence of approximate ex post Nash equilibria with purification of Bayesian equilibrium of the original game of incomplete information.

That ex post stability implies the existence of a pure strategy Bayesian approximate equilibrium may appear trivial. It is not so, however, using Kalai's approximate notion of ex post stability: a strategy vector is (ε, ρ) ex post stable if with high probability $(1-\rho)$ no player can gain, ex post, by more than some small amount (ε) by deviating. A purification result can still easily be obtained, as highlighted by Kalai, for the case of *normal form games*. To explain, if a strategy vector is (ε, ρ) ex post stable then with positive probability it yields a profile of actions where no player can gain by more than ε by deviating - thus, there must exist a Nash ε equilibrium in pure strategies. In treating *Bayesian games* things are not so straightforward and more work has to be done in obtaining a purification result. This is highlighted by our results: we show that the existence of an (ε, ρ) ex post stable strategy vector implies the existence of a Nash α-equilibrium in pure strategies where $\alpha \leq (1-\rho)\varepsilon + \rho D$ and where D is an upper bound on payoffs. The bound on α is tight and thus, in contrast to the special case of normal form games, the existence of an (ε, ρ) ex post stable strategy vector is not enough to imply the existence of a Bayesian

ε-equilibrium in pure strategies. This seems an interesting property and so we provide an example of a semi-anonymous game to illustrate.

It is worth pointing out that our result applies to a general class of Bayesian games and is not restricted to semi-anonymous games with many players. That is, in any Bayesian game the existence of an ex post stable strategy vector implies the existence of an approximate Bayesian equilibrium in pure strategies. Semi-anonymous games merely provide one instance where the existence of an ex post strategy vector is guaranteed. Our result thus serves to demonstrate the strength of the ex post stability concept.

2 Model

There exists a finite set of possible player *actions*, denoted $\mathcal{A}$, and finite set of possible player *types*, denoted $\mathcal{T}$. Set $\mathcal{C} \equiv \mathcal{T} \times \mathcal{A}$ will be used to denote the possible *type-action characters* of a player.[1] A *Bayesian game* is given by a tuple $G = (N, T, p, A, u)$ where:

$N = \{1, ..., n\}$ is a finite *player set*.

$T = \times_i T_i$ is the set of *type profiles* where each $T_i \subseteq \mathcal{T}$ describes the feasible types of player i.

$p : T \to [0, 1]$ is a *prior probability function* where $p(t)$ gives the probability of type profile $t \in T$.

$A = \times_i A_i$ is the set of *action profiles* where each $A_i \subseteq \mathcal{A}$ describes the feasible actions of player i.

$u = (u_1, ..., u_n)$ is a vector describing the players *utility functions*. Let $C_i = T_i \times A_i$ denote the feasible type action characters of player i and let $C = \times_i C_i$ denote the set of feasible *profiles of type-action characters*. Each u_i takes the form $u_i : C \to [0, D]$.

A Bayesian game G is played as follows: According to the prior probability function p each player i is assigned a type t_i. Informed of his type (but not the types of the other players) a player chooses an action (possible using some randomization). This determines the type-action character of each player and payoffs can be calculated according to the realized profile of type-action characters.

A *strategy* of player i is defined by a vector σ_i where $\sigma_i(a_i|t_i)$ gives the probability of player i choosing action a_i if of type t_i. Given a vector of strategies σ and the prior probability function p one can determine the probability of each possible profile of type-action characters. This allows utility functions to be extended to strategy vectors by assuming that $U_i(\sigma) = E[u_i(c)]$ for each i.

We say that strategy σ_i is a *pure strategy* if for each $t_i \in T_i$ there exists some a_i such that $\sigma_i(a_i|t_i) = 1$. A strategy vector σ is said to be a *pure strategy*

[1] For example, a character may be *weak type who eats quiche* etc.

vector if σ_i is a pure strategy for each i. We say that a set of pure strategy vectors $\{s^1, ..., s^M\}$ constitute a *support* for a strategy vector σ if and only if there exists real numbers $\beta_1, ..., \beta_M$ where

1. $1 \geq \beta_m > 0$ for all m,
2. $\sum_m \beta_m = 1$ and
3. $\sigma_i(a_i|t_i) = \sum_m \beta_m s_i^m(a_i|t_i)$ for all i, a_i and t_i.

   Clearly every feasible strategy vector σ has a support.

3 A purification result for Bayesian games

We begin by defining two distinct equilibrium concepts. As is standard, we say a strategy vector σ is a *Bayesian* (Bayesian Nash) *ε-equilibrium* if and only if:

$$U_i(\sigma_i, \sigma_{-i}|t_i) \geq U_i(\sigma'_i, \sigma_{-i}|t_i) - \varepsilon$$

for all σ'_i, $t_i \in T_i$ and $i \in N$.[2] Thus, if σ is a Bayesian ε-equilibrium no player i *expects* to gain by more than ε by deviating from σ_i. If s is a pure strategy vector and a Bayesian ε-equilibrium then we say that s is a *Bayesian ε-equilibrium in pure strategies.*

We now introduce the notion of *ex post Nash* as defined by Kalai (2004). A profile of type-action characters $c = (c_1, ..., c_n) = ((t_1, a_1), ..., (t_n, a_n))$ is an ε *best response* for player i if

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

for every action $a'_i \in A_i$. A profile of type-action characters is ε *Nash* if it is an ε best response for every player $i \in N$. Finally, a strategy profile is (ε, ρ) *ex post Nash* if the probability that it yields an ε Nash profile of type-action characters is at least $1 - \rho$.

We provide our main result.

Theorem 1: Take as given a Bayesian game G and small, non-negative real numbers ε and ρ (both less than 1). If a strategy vector σ is (ε, ρ) ex post Nash then in the support of σ there is a pure strategy vector s that is a Bayesian α-equilibrium where $\alpha \leq (1 - \rho)\varepsilon + \rho D$.

Before detailing the proof we provide a simple example to illustrate Theorem 1 and demonstrate that the bound provided is tight. There are three players, two types H and L and two actions B and G. Player 1, called nature, is of type H with probability ρ and type L with probability $1 - \rho$. Nature always receives a payoff of zero. Players 2 and 3 are always of type L. When nature is of type L players 2 and 3 are seen to play the matrix game:

[2] More formally we only require for $t_i \in \mathcal{T}$ where there is a positive probability that player i may be of type t_i.

	B	G
B	$1,0$	$0,1$
G	$0,1$	$1,0$

and if nature is of type H players 2 and 3 play matrix game:

	B	G
B	$D,0$	$0,D$
G	$0,D$	$D,0$

where $D > 1$. Consider the pure strategy vector $s = (B, B, B)$. Given that player 2 is playing B player 3 expects to gain by $(1-\rho) + \rho D \equiv k$ by deviating to G instead of B. From this, it can be seen that s is a Nash k equilibrium and, furthermore, there can be no pure strategy vector that is a Bayesian α-equilibrium for any $\alpha < k$. Next note that strategy vector s is $(1, \rho)$ ex post Nash. This follows in that with probability $1 - \rho$ nature is of type L and when this happens no player can gain by more than 1 by deviating.

Proof of Theorem 1: Let σ^* be (ε, ρ) ex post Nash and let $P \equiv \{s^1, ..., s^M\}$ be a support of σ^*. We proceed by contradiction. Thus, suppose that there exists no $s^m \in P$ such that s^m is a Bayesian α-equilibrium for $\alpha = (1-\rho)\varepsilon + \rho D$.[3]

We introduce some notation: Let C^* denote the set of ε Nash composition profiles of game Γ. Given a strategy vector σ' let $y(c, \sigma')$ denote the probability of composition profile c occurring. [2]

Take any $s^m \in P$. By our supposition, s^m is not a Bayesian α equilibrium. Given that s is not a Bayesian α-equilibrium it must be that the probability of a composition profile $c \notin C^*$ occurring is greater than ρ; that is,

$$\sum_{c \notin C^*} y(c, s^m) > \rho. \tag{1}$$

Suppose otherwise: with probability at least $1-\rho$ an ε Nash composition profile arises; if a composition profile $c \notin C^*$ arises then each player can gain at most D by changing his action; thus, ex-ante the maximum a player can gain by changing his strategy is $(1-\rho)\varepsilon + \rho D$ leading to the desired contradiction.

The set $P = \{s^1, ..., s^M\}$ is a support for strategy vector σ and thus there exists real numbers $\beta_1, ..., \beta_M$ where (1)$1 \geq \beta_m > 0$ for all m, (2) $\sum_m \beta_m = 1$ and (3) $\sigma_i^*(a_i|t_i) = \sum_m \beta_m s_i^m(a_i|t_i)$ for all i, a_i and t_i. Thus,

$$y(c, \sigma^*) = \sum_m \beta_m y(c, s^m) \tag{2}$$

for all $c \in C$. Thus,

$$\sum_{c \notin C^*} y(c, \sigma^*) = \sum_{c \notin C^*} \left[\sum_m \beta_m y(c, s^m) \right] = \sum_m \beta_m \left(\sum_{c \notin C^*} y(c, s^m) \right) \tag{3}$$

[3] Note that if s^m is not a Bayesian α-equilibrium then it cannot be a Bayesian α' equilibrium for any $\alpha' < \alpha$.

Note, however that σ^* is (ε, ρ) ex post Nash which by definition implies,

$$\sum_{c \notin C^*} y(c, \sigma^*) < \rho. \tag{4}$$

Clearly (1), (3) and (4) are incompatible if $\sum_m \beta_m = 1$. This gives the desired contradiction.■

A corollary of this result and results due to Kalai (2004) is that, given any $\varepsilon > 0$, for any semi-anonymous game with sufficiently many players and for any equilibrium σ of that game there exists a Bayesian ε-equilibrium in pure strategies in the support of σ.[4]

3.1 Example: (ε, ρ) ex post Nash does not imply Bayesian ε purification

In this section we provide an example to demonstrate that the existence of an (ε, ρ) ex post Nash strategy vector does not imply the existence of a Bayesian ε-equilibrium in pure strategies even if ρ is arbitrarily small. As pointed out in the introduction this is not the case in normal form games. In normal form games the existence of an (ε, ρ) ex post Nash strategy vector implies the existence of a Nash ε-equilibrium in pure strategies irrespective of ρ. The example treats a class of semi-anonymous games as defined by Kalai (2004).

There are, for notational simplicity, $3n$ players where n is odd. There are two actions B and G and four types $Poor$ (P), $Rich$ (R), $High$ (H) and Low (L). Players $1, 2, ..., n$ (called rich) have type R with probability 1. Players $n+1, n+2, ..., 2n$ (called poor) have type P with probability 1. Players $2n+1, ..., 3n$ (called managers) have type H with probability $\frac{1}{n}$ and type L with probability $(1 - \frac{1}{n})$. Managers are assigned types independently.

Given an action profile a, type t' and action a' let $w(t', a', a)$ be the number of players with type t' who choose action a'. Thus, for example, $w(R, B, a)$ denotes the number of players who are rich and choose action B. If player i is poor then his payoff function is given by,

$$u_i(a_i, a_{-i}, t) = \frac{w(R, a_i, a)}{n} .$$

Thus, the payoff of a poor player depends positively on the proportion of rich players who choose the same action as himself. Given a type profile t let $h(t)$ denote the proportion of managers who are type high. If player i is rich then his payoff is given by,

$$u_i(a_i, a_{-i}, t) = D - \tfrac{w(P, a_i, a)}{n} \quad \text{if } h(t) \leq \tfrac{2}{3},$$
$$u_i(a_i, a_{-i}, t) = D - \tfrac{w(P, a_i, a)}{n} - (D-1)\left(\tfrac{h(t) - \frac{2}{3}}{\frac{1}{3}}\right)\tfrac{w(P, a_i, a)}{n}$$
$$\text{otherwise}$$

[4]See Kalai (2004) for a definition of a semi-anonymous game.

Thus, the payoff of a rich player depends negatively on the proportion of poor players playing the same action as himself. As the proportion of managers who have type H increases above $\frac{2}{3}$ then his payoff is influenced more by the actions of the poor players. Let the payoff of a manager be 1 independent of the composition profile.[5]

First, consider the existence of a Bayesian ε equilibrium in pure strategies. Given a strategy vector in which all rich players or all poor players play the same strategy there must exist at least one player who can gain by 1 or more by changing strategy. Thus, assume there to be at least one rich player and one poor player playing G and one rich player and one poor player playing B. As n is odd the number of poor players playing G is distinct to the number playing B. Given that $\Pr[h(t) > 2/3] > 0$, for any pure strategy vector s there must be at least one rich player i who can expect, ex-ante, to gain by strictly more than $\frac{1}{n}$ if he changes strategy. Thus, there does not exist a Bayesian ε equilibrium in pure strategies for any $\varepsilon \leq \frac{1}{n}$.

Let s' be the pure strategy vector whereby $\frac{n-1}{2}$ rich players choose action B and $\frac{n+1}{2}$ choose action G and similarly $\frac{n-1}{2}$ poor players choose action B and $\frac{n+1}{2}$ choose action G. With some probability $1-\rho'$ strategy vector s' will yield a composition profile c where $h(t) \leq 2/3$. When this occurs c is $\frac{1}{n}$ Nash. Thus, s' is $(\frac{1}{n}, \rho')$ ex post Nash. We shall now show that $\rho' \to 0$ as $n \to \infty$. Assuming, for simplicity that n is divisible by 3, we obtain,[6]

$$\rho' = \sum_{x=\frac{2}{3}n}^{n} \binom{n}{x} \left(\frac{1}{n}\right)^{x} \left(1-\frac{1}{n}\right)^{n-x} = \Pr\left[F_{v_1,v_2} \leq \frac{v_2 \frac{1}{n}}{v_1\left(1-\frac{1}{n}\right)}\right]$$

where F_{v_1,v_2} is the F distribution with parameters v_1 and v_2 and where $v_1 = \frac{4}{3}n$ and $v_2 = \frac{2}{3}n + 2$. Note that,

$$\frac{v_2 \frac{1}{n}}{v_1\left(1-\frac{1}{n}\right)} = \frac{\frac{2}{3}+\frac{2}{n}}{\frac{4}{3}n - \frac{4}{3}} \to 0 \text{ as } n \to \infty.$$

Also note that $v_1, v_2 \to \infty$ as $n \to \infty$. It follows that $\rho' \to 0$ as $n \to \infty$. An alternative, if less formal, way of obtaining the same result is to note that if we let p denote the probability that a manager has type H then as $n \to \infty, p \to 0$ but $np = 1$. Thus, as n becomes large the binomial distribution determining the number of managers who have type H can be approximated by a Poisson distribution with parameter 1. It follows that the $\Pr[x \geq \frac{2}{3}n] \to 0$ as $n \to \infty$.

[5] Intuitevely it may be that if managers have type H they prefer some policy or action that makes the payoff of rich players more sensitive to the actions of poor players. This is, however, not neccessary for the example.

[6] A known result (see p110 of Johnson, Kotz qnd Kempis 1993) is that,

$$\sum_{x=r}^{n} \binom{n}{x} p^x q^{n-x} = \Pr\left[F_{v_1,v_2} \leq \frac{v_2 p}{v_1 q}\right]$$

where F_{v_1,v_2} is the F distribution with parameters $v_1 = 2r$ and $v_2 = 2(n-r+1)$. See,

References

[1] Cartwright, E. and M. Wooders (2002) "On equilibrium in pure strategies in games with many players," University of Warwick Department of Economics Working Paper #686 (Revised 2005).

[2] Cremer, J. and R.P. McLean (1985) "Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent," *Econometrica* 53: 345-361.

[3] Green, E. and J-J Laffont (1987) "Posterior implementability in a two-person decision problem," *Econometrica* 55: 69-94.

[4] Johnson, N.L., S. Kotz, A. Kemp (1993), *Univariate discrete distributions* (Second Edition) - Wiley Series in Probability and Mathematical Statistics, John Wiley and Sons.

[5] Kalai, E. (2004) "Large robust games," *Econometrica* 72: 1631-1665.

[6] Khan, A. and Y. Sun (2002) "Noncooperative games with many players," *Handbook of Game Theory*, R. Aumann and S. Hart, eds. North Holland.

[7] 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.

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

[9] McLean, R.P. and A. Postlewaite (2002) "Informational size and incentive compatibility," *Econometrica* 70: 2421-2453.

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

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

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

[13] Wilson, R. (1987) "Game theoretic advances in trading processes," in *Advances in Economic Theory: Fifth World Congress*, T. Bewley (ed), Cambridge University Press, 33-70.