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

October 30, 2006

Abstract

Treating games of incomplete information, we demonstrate that the existence of an ex post stable strategy vector implies the existence of an approximate Bayesian equilibrium in pure strategies that is also ex post stable. Through examples we demonstrate the 'bounds obtained on the approximation' are tight.

*This paper originally appeared as University of Warwick Department of Economics Working Paper #710, February 2004. Related research was presented at the National University of Singapore conference "Uncertainty and Information in Economics" (9 May - 3 July 2005); we thank the participants for the comments and the NUS for its support and hospitality.

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 stability' of Bayesian equilibrium, see, for example, Cremer and McLean (1985), Wilson (1987), Green and Laffont (1987), McLean and Postlewaite (2002) and Kalai (2004). If an equilibrium is ex post stable then no agent has an incentive to change his behavior *after observing the actions and types of others*. This implies many desirable characteristics of equilibrium such as the revelation principle holds (Cremer and McLean 1985) and implementation without binding contracts is possible (Green and Laffont 1987). In this paper we demonstrate a further important property; the existence of a Bayesian equilibrium that is ex post stable implies the existence of a Bayesian equilibrium in *pure strategies* that is also ex post stable. The desirability of pure strategy equilibria has long been recognized in the literature (Harsanyi 1973 and Schmeidler 1973).

We derive our result using the notion of *approximate* ex post stability introduced by Kalai (2004). For a strategy vector to be (ε, ρ) ex post Nash it must be that with high probability $(1-\rho)$ no player can gain, ex post, by more than some small amount (ε) by deviating from his given strategy. For games of incomplete information, the existence of a $(0, 0)$ ex post Nash strategy vector immediately implies the existence of a Bayesian equilibrium in pure strategies that is $(0, 0)$ ex post Nash (we shall formally demonstrate this below). For *games of complete information* the existence of an (ε, ρ) ex post Nash strategy vector implies the existence of a Nash ε equilibrium in pure strategies.[1] This leaves open the important and interesting question of whether, for games of incomplete information, the existence of an (ε, ρ) ex post Nash strategy vector, with either ε or ρ (or both) greater than zero, implies the existence of an approximate Bayesian equilibrium in pure strategies.[2]

In this paper, we demonstrate that the existence of an (ε, ρ) ex post Nash strategy vector implies the existence of a Bayesian α-equilibrium in pure strategies that is also (ε, ρ) ex post Nash, where $\alpha \leq (1-\rho)\varepsilon + \rho D$ and D is an upper bound on payoffs. The bound on α is shown to be tight and thus, in contrast to the special case of games of complete information, the existence of an (ε, ρ) ex post Nash stable strategy vector is not enough to imply the existence of a Bayesian ε-equilibrium in pure strategies. We provide an example to illustrate this point.

One corollary of our result is a purification result for Bayesian games. Recall that Kalai (2004) demonstrates that in large semi-anonymous games every Bayesian equilibrium is approximately ex-post Nash. Thus, applying our result, we see that in a large semi-anonymous game there exists an approximate

[1] If a strategy vector is (ε, ρ) ex post Nash then with positive probability it yields a profile of actions where no player can gain by more than ε by deviating and thus, there must exist a Nash ε equilibrium in pure strategies.

[2] See Section 3.2 of Kalai (2004) for further discussion of these issues.

Bayesian equilibrium in pure strategies. For games of complete information, purification results have been of interest since Schmeidler (1973); see, for examples Mas-Colell (1984), Pascoa (1993, 1998) and Khan, Rath and Sun (1997). 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.

It is worth highlighting that our result is not restricted to semi-anonymous games with many players. Thus, in general, the existence of an ex post Nash strategy vector implies the existence of an ex post stable Bayesian equilibrium in pure strategies. This serves to illustrate the usefulness of ex post stability. Existence of an ex post Nash strategy vector has been demonstrated in specific contexts such as auctions (Cremer and McLean 1985) and bargaining (Green and Laffont 1987). In situations where directly proving the existence of a pure strategy Bayesian equilibrium is more difficult than proving the existence of an ex-post Nash strategy vector our result may be of practical use.

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. Finiteness of $\mathcal{C}$ is assumed for transparency and the analysis can be extended to a non-finite $\mathcal{C}$. 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 a 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 $u_i : C \to [0, D]$ where D is some real number.

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 (possibly 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 $t_i \in T_i$. 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. By use of mathematical expectation, 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 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 $\beta_1, ..., \beta_M$ satisfying

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.

Every strategy vector σ must have 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:

$$U_i(\sigma_i, \sigma_{-i}|t_i) \geq U_i(\sigma_i', \sigma_{-i}|t_i) - \varepsilon \tag{1}$$

for all σ_i', $t_i \in T_i$ and $i \in N$. Thus, if σ is a Bayesian ε-equilibrium no player i *expects* to gain 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.*

The notion of an ex post equilibrium, as discussed earlier, appears in a number of papers, (e.g. Cremer and McLean 1985). We now introduce the notion of approximate *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 σ 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 can now state our main result.

Theorem 1: Take as given a Bayesian game G and 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 (ε, ρ) ex post Nash and a Bayesian α-equilibrium where $\alpha \leq (1 - \rho)\varepsilon + \rho D$.

Note that both ε and ρ could be zero implying that α is also zero. 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 is of type H with probability ρ and type L with probability $1-\rho$. Player 1 always receives a payoff of zero. Players 2 and 3 are always of type L. When player 1 is of type L players 2 and 3 play the matrix game:

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

and if player 1 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 because with probability $1-\rho$ player 1 is of type L and no player can gain by more than 1 by changing action.

3.1 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^*. Denote by C^* the set of ε Nash type-action character profiles of game Γ. Given a strategy vector σ' let $y(c, \sigma')$ denote the probability of a type-action character profile c occurring.

From the definition of a support 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,

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

That σ^* is (ε, ρ) ex post Nash implies,

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

Using (3) and (4) and the fact that $\sum_m \beta_m = 1$ there must exist some $s \in P$ such that

$$\sum_{c \notin C^*} y(c,s) \leq \rho. \tag{5}$$

By definition strategy vector s is (ε, ρ) ex-post Nash. Also, given strategy vector s, with probability less than ρ a type-action character profile $c \notin C^*$ occurs and each player can gain by at most D by a change of action. With probability greater than $1-\rho$ an ε Nash type-action character profile arises and each player can gain by at most ε. Ex-ante, therefore, the maximum a player can expect to gain by changing his strategy is $(1-\rho)\varepsilon + \rho D$. Thus, s is a Bayesian α equilibrium.■

3.2 Semi-anonymous games

A game is semi-anonymous if the payoffs of players depend only on the proportion of players choosing each strategy. Informally, Kalai (2004), with $D = 1$, demonstrates that for any ε and any family of semi-anonymous games Γ there exist positive constants α and β such that all the equilibria of games with n or more players are $(\varepsilon, \alpha\beta^m)$ ex post Nash. Values for α and β, dependent on the primitives of family Γ, are derived by Kalai. An immediate Corollary of Theorem 1 is that any game belonging to family Γ with n or more players has a Bayesian α-equilibrium in pure strategies where $\alpha = \varepsilon + \alpha\beta^m(1-\varepsilon)$. Thus, one obtains a purification result for Bayesian games. Note, however, 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 games of complete information.] Further this holds when $\varepsilon = 0$. We provide an example, treating a family of semi-anonymous games, to illustrate this point.

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,

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

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 type-action character profile.[3]

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 α equilibrium in pure strategies for any $\alpha \leq \frac{1}{n}$.

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 type-action character 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,[4]

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

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,

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

[3] Intuitively 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 necessary for the example.

[4] 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, 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 $n \to \infty$.

In summary: There does not exist a Bayesian α-equilibrium in pure strategies for any $\alpha \leq \frac{1}{n}$. There does exist strategy vector s' that is $(\frac{1}{n}, \rho')$ ex post Nash where ρ' can be made arbitrarily small. A minor change in the example allows us to illustrate the case where $\varepsilon = 0$. Let the payoff function of a rich player i for $h(t) \leq \frac{2}{3}$ be

$$u_i(a_i, a_{-i}, t) = D \text{ if } h(t) \leq \tfrac{6}{10} \text{ and}$$
$$u_i(a_i, a_{-i}, t) = D - \tfrac{w(P, a_i, a)}{n}(15h(t) - 9) \text{ if } h(t) \leq \tfrac{2}{3}.$$

Consider the pure strategy vector s'' where all players choose G. If $h(t) \leq \frac{6}{10}$ then the type-action character profile is 0 Nash. If ρ'' is the probability that $h(t) \leq \frac{6}{10}$ then we can repeat the above argument to show that $\rho'' \to 0$ as $n \to \infty$. Thus, strategy profile s'' is $(0, \rho'')$ ex post Nash for arbitrarily small ρ''. There does not exist, however, a Bayesian 0 equilibrium in pure strategies for the same arguments given above.

It is clear that the existence of an (ε, ρ) ex-post Nash strategy vector need not guarantee the existence of a Bayesian ε-equilibrium in general and not just in specific examples. Indeed, given the nature of ex-post stability when applied to Bayesian games, it is almost to be expected. Ex-post stability requires that with probability $1 - \rho$ a type-action character profile occurs where any gains from a change of action must be less than ε. This means, however, that with probability ρ a type-action character profile occurs where some player could gain well in excess of ε from a change of action. If a player will be able to gain by ε from changing action with probability $1 - \rho$ and will be able to gain by more than ε from changing action with probability ρ then this clearly need not be consistent with a Bayesian ε equilibrium.

3.3 Does existence of a pure strategy Bayesian equilibrium imply ex-post Nash?

One may ask whether there is a converse to Theorem 1: Is it the case that a pure strategy Bayesian α equilibrium is (ε, ρ) ex post Nash for some ε and ρ where $(1 - \rho)\varepsilon + \rho D \leq \alpha$ or, more generally, ε and ρ are small when α is small? In general, the answer is no as two simple examples illustrate.

First, even if every player has a very small probability of being able to gain by ε ex post, the probability that *some player* may have the possibility to gain by more than ε ex post can be very high. Suppose that there are n players and two types H and L. A player is type H with probability α. There are two strategies B and G. A player of type L gets a payoff of 1 irrespective of his action. A player of type H gets a payoff of 1 if he plays B and 0 if he plays G. Note that $D = 1$ and the pure strategy vector s where all players play G irrespective of type is a Bayesian α equilibrium. Suppose that $\alpha = \frac{1}{n}$ and the prior probability function over types is such that one (and only one) player is of type G. With probability 1 a type-action character profile occurs where one player can gain $D = 1$ from a change of strategy. Thus, s is $(0, 1)$ ex post Nash. Clearly, $\rho D > \alpha$. [Alternatively one could put $\varepsilon = 1$ and $\rho = 0$ to say that s is $(1, 0)$ ex post Nash but now $(1 - \rho)\varepsilon > \alpha$.] As an alternative suppose that types are distributed independently to players but each player still has probability α of being type H. The number of players with type H is given by a binomial distribution. The probability that all players have type L is given by $(1 - \alpha)^n$. Thus, with probability $1 - (1 - \alpha)^n$ a type-action character profile occurs where at least one agent can gain $D = 1$ from a change of strategy. So, s is $(0, 1 - (1 - \alpha)^n)$ ex post Nash. Clearly, for $\alpha < 1$ and n large enough $\rho D > \alpha$.

Second, there may be strategic value in knowing another player's type. For example consider a game with two players, two types H and L and two strategies G and B. Suppose that player 1 maximizes his payoff by playing G if he is of type H and by playing B if he is of type L. The probability that player 1 is type H is 0.6. Player 2 is always of type H and gets 1 if he matches the choice of player 1 and 0 otherwise. The pure strategy vector where player 1 maximizes his payoff and player 2 plays G is a Bayesian 0 equilibrium. Clearly, however, there is a 0.4 probability that once player 2 knows the type of player 1 he can gain 1 by changing action. Thus, this strategy vector is, at best, $(0, 0.6)$-ex post Nash.

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] Harsanyi, J.C. (1973) "Games with randomly distributed payoffs: a new rationale for mixed strategy equilibrium points:" *International Journal of Game Theory* 2: 1-23.

[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.