The Robustness of Repeated Game Equilibria to Incomplete Payoff Information

V.Bhaskar
Dept. of Economics, University of Essex, UK

Current version: April 2000

Abstract

We analyse the role of mixed strategies in repeated (and other dynamic) games where players have private information about past events. History-dependent mixed strategies require a player to play distinct continuation strategies σ and σ' at information sest ω and ω' respectively, although the player the player is indifferent between these strategies at both information sets. Such equilibria are not robust to a small amount of incomplete payoff information, as in Harsanyi (1973). On the other hand, "history-independent" randomizations survive the Harsanyi critique, and may be useful in supporting cooperative outcomes in repeated games with private monitoring. Our analysis is applicable to a variety of economic situations including credence goods, double moral hazard, repeated games with private monitoring, and repeated games played between a long-run and a sequence of short-run players.

Keywords: Mixed strategies, dynamic games, imperfect information, purification.

1 Introduction

Equilibria in games and mechanisms are very often not strict. Mixed strategy equilibria are a case in point — a player is indifferent between two actions, and equilibrium requires that his opponent believes that he will behave in an appropriately random way. In a famous paper, Harsanyi [14] showed that in generic normal form games, all equilibria are robust. Specifically, if a player entertains uncertainty about the exact payoffs of the other player and this uncertainty is sufficiently small, then for any equilibrium of the complete information game, one can find an equilibrium of the incomplete information game which induces approximately the same behavior. Furthermore, the equilibrium of the incomplete information game will such that players play strict best responses for almost all realizations of their private information.[1] Harsanyi's result has meant that economists have been, with good reason, relaxed about equilibria with indifferences.

This main purpose of this paper is to argue that this relaxed attitude is unwarranted, in the context of a specific class of dynamic games and mechanisms. In the game theoretic context, our focus is on games with several "stages, where the players' payoffs are additively separable across these stages. Classic repeated games, where history does not affect future payoff functions, are a case in point. With perfect monitoring or with imperfect public monitoring, it is not difficult to construct equilibria where players have strict incentives to take their equilibrium actions. However, when monitoring is private, one typically cannot construct equilibria where players have strict incentives. Examples of private monitoring include standard repeated games with imperfect private monitoring, repeated games with overlapping generations of players with imperfect information, and games between an infinitely lived long-run player and a sequence of short run players, where the short run players also do not observe all past events. In these games, a growing literature demonstrates that one has to resort to mixed strategies in order to support cooperative outcomes. Our argument is that *some* equilibria in these games do not pass the Harsanyi test. In particular, if there is slight amount of payoff uncertainty, the behavior corresponding to these equilibria cannot be the outcome of an equilibrium of the incomplete information game. In some contexts, such as the overlapping generations model or the chain store paradox, our application of the Harsanyi test yields powerful results — one has play of the stage game Nash/perfect equilibria in every period. In other contexts, such as the repeated prisoners' dilemma with private monitoring, some approaches to sustaining cooperation pass the Harsanyi test while others do not.

[1]This is ensured if one assumes that payoffs in the incomplete information game are given by a probability measure which is absolutely continuous with respect to Lebesgue measure.

Our basic point can be made quite simply. Suppose that there is one player, say player i, who is required to behave in a history-dependent way today in order to induce another player(s), say player j, to take the appropriate action at an earlier date. Suppose also that i has some information about j's behavior, but this information is private, i.e. no other player knows what i's private information is. Specifically, suppose that there are two information sets for player i, ω and ω', and the equilibrium requires that i take (a possibly random) action α at ω and a different action α' at ω'. Suppose that i's beliefs about the continuation strategies of other players do not vary between ω and ω'. One way in which one can induce the required behavior by i is to construct the equilibrium so that i is indifferent between α and α' at ω, and these are optimal actions. This implies that he is also indifferent between α and α' at ω' (since his beliefs about others' continuation strategies do not vary across ω and ω'). Hence, we can get i to play α at ω and α' at ω'. This is a relatively simple way of constructing an equilibrium where j behaves as required, and it is not surprising that many papers utilize this.

Consider now the implications of incomplete information regarding i's payoffs. If we perturb payoffs, for (almost) any realization of his payoff, i will behave in exactly the same way at ω and ω'. To see this, note that for fixed beliefs about the continuation strategies of other players, i will have a strict best response for almost any realization of his private information. Since his beliefs at ω and ω' are identical, i will play this strict best response at both these information sets. Hence, although his best response can vary as a function of his payoff information, it cannot vary as a function of his payoff irrelevant information (i.e. his information regarding j's behavior). Hence j should not believe that i will reward/punish him, as required by the equilibrium. Hence this equilibrium of the complete information game is not robust to slight payoff uncertainty.

In other words, we propose a refinement of Nash equilibrium for dynamic games with imperfect information, based on Harsanyi's approach. We argue that game theorists should restrict attention to *robust* equilibria, which can, in principle be approximated by equilibria of games with incomplete information. Such robust equilibria always exist. The question is, do they allow cooperative behavior or the folk theorem, in repeated and other dynamic games with private information. The answer here depends on the context. In some games of imperfect private information, cooperation can be supported via robust equilibria. In others however, there is no such possibility, so that robustness implies an anti-folk theorem. We provide some examples of both cases.

The remainder of this (incomplete) paper is as follows. We set out our basic idea by analyzing an example of a two-stage game with imperfect information with time separable payoffs. This game has a sequential equilibrium where the informed player takes history-dependent random actions, which is not a sequential information when the other player is also informed. In fact our example has infinitely many sequential equilibrium outcomes, which is significant in view of Kreps and Wilsons' [18]theorem for generic extensive form games. We then show that these equilibria are not robust, by analyzing a game where players have small amount of incomplete information regarding other players' payoffs, as in Harsanyi [14]. We find that the unique robust equilibrium of this game is in history-independent strategies. We go on to analyze the same question in a mechanism design context. Finally, we will present applications to repeated and other dynamic games where information about past actions is private. These include a detailed discussion of repeated games played between a long-run player and a sequence of short-run players.

2 The Basic Example

STAGE 1 Player 1

U	a_U, b_U, c_U
D	a_D, b_D, c_D

STAGE 2

2 \ 3	L	R
T	a_{11}, b_{11}, c_{11}	a_{12}, b_{12}, c_{12}
B	a_{21}, b_{21}, c_{21}	a_{22}, b_{22}, c_{22}

Consider the following two-stage game with three players. In stage 1 player 1 chooses between U and D, and this choice generates stage payoffs for all three players, as shown in Fig. 1 above, where a, b, c denote the payoffs of $1, 2, 3$ respectively, and subscripts index actions. Assume that $a_U > a_D$ and $b_U < b_D$. In the second stage players 2 and 3 simultaneously choose from $\{T, B\}$ and $\{L, R\}$ respectively, which generate payoffs for all three players as given in Fig.1. Assume that the game in stage 2 has one (and only one) completely mixed equilibrium.[2] This mixed equilibrium can be denoted by the pair $\mathbf{p}= (p, 1-p)$,$\mathbf{q}= (q, 1-q)$, where p and q solve $p(c_{11} - c_{12}) + (1-p)(c_{21} - c_{22}) = 0$ and $q(b_{11} - b_{21}) + (1-q)(b_{12} - b_{22}) = 0$ respectively. Let the payoffs to any player in the overall game be given by the sum of payoffs in the two stages (this is why we call this a two-stage game).

It remains to specify the information of the each player at the (single) time that she chooses her action. Since player 1 moves first, she moves in ignorance of the actions of the actions of 2 and 3. Player 2 observes player 1's action, while player 3 is not informed of 1's action. Players 2 and 3 move simultaneously, without knowing each other's action. This completes the speficiation of the game, which we shall call Γ.

Assume that $a_U - a_D < (\mathbf{e} - \mathbf{p})\mathbf{A}\mathbf{q}$, where $\mathbf{e} = (0, 1)$ corresponds to 2 playing B. We now show that there is a sequential equilibrium (call this s^*) where player 1 chooses D, player 2 chooses $\mathbf{p}$ if she observes D and $\mathbf{e}$ if she observes U, and player 3 chooses $\mathbf{q}$. To verify this, observe that under this equilibrium player 3 assigns probability one to the event that player 1 has chosen U. Hence he believes that player 2 is choosing $\mathbf{p}$ with probability one, which makes the choice of $\mathbf{q}$ optimal. Given that 3 is choosing $\mathbf{q}$, player 2 is indifferent between all mixtures of T and B. Hence it is optimal for player 2 to behave as prescribed at both her information sets. Given player $2's$ strategy, and the assumed inequality regarding player $1's$ payoffs, it is optimal for player 1 to play D.[3] Hence we have an example of a history dependent equilibrium, where 1 is deterred from playing U because of the threat of player 2 to play B. Such a threat is credible precisely because 2 has private information about 1's choice.Indeed, our game has infinitely many history dependent equilibria, as is shown by the following proposition:

Proposition 1 Γ *has a continuum of sequential equilibrium outcomes.*

Proof. We show that there is a class of equilibria with the following properties. In each of these class, player 3 plays a fixed strategy, $\mathbf{q}$; 1 randomizes

[2] I.e. either $b_{11} > b_{21}, b_{22} > b_{12}, c_{12} > c_{11}, c_{21} > c_{22}$, in which case the game has a unique equilibrium, or the game has two strict pure Nash equilibria and a single completely mixed equilibrium.

[3] Kandori's[15] provides an example which is somewhat similar. He considers a twice repeated game where the stage game has a unique equilibrium in mixed strategies, and shows that if players do not observe first period actions but instead observe a private signal, there exists a sequential equilibrium where players play differently from the mixed equilibrium in the first stage.

between U and D, .with probabilities x and $(1-x)$ respectively; player 2 plays the local strategies $\mathbf{y}(U)$ and $\mathbf{y}(D)$ after observing U and D respectively with $\mathbf{y} = (\mathbf{y}(U), \mathbf{y}(D))$. Different equilibria in the class may be indexed by different values of x, and for each value of x, there corresponds a unique $\mathbf{y}$, while $\mathbf{q}$ does not vary with x. We show that one may select any x in an interval, thereby proving the theorem.

If $x, \mathbf{y}(U), \mathbf{y}(D), \mathbf{q}$ are all interior, then each of these implies an equalitiy to pure actions of the relevant player at the corresponding information set. If player 3 plays $\mathbf{q}$, 2 is indifferent between both her actions at both her information sets. Consider now player 1; if x is interior, $1's$ payoff from each of her actions must be equal, i.e:

$$\mathbf{y}(D) - \mathbf{y}(U)]\mathbf{Aq} - (a_U - a_D) = 0 \tag{1}$$

We show first that there is one $\mathbf{y}$ such that this equation (1) is satisfied. Let $\mathbf{z} = (z, 1-z)$ and consider the expression $[\mathbf{p} - \mathbf{z}]\mathbf{Aq} - (a_U - a_D)$. Our assumptions on payoffs imply that if $z = 0$, this expression is strictly positive, while if $z = p$, it is strictly negative. Hence there exists $\mathbf{z}^*$ with $p > z^* > 0$ such that this expression is zero, which implies that if we select $\mathbf{y}(D) = \mathbf{p}, \mathbf{y}(U) = \mathbf{z}^*$, player 1 will be indifferent between U and D. To see that there is a continuum of values of $\mathbf{y}$, observe that $\mathbf{y}(D), \mathbf{y}(\mathbf{U})$ also satisfy provided that

$$y(D) - y(U) = p - z^* \tag{2}$$

Hence any pair $\mathbf{y}(D), \mathbf{y}(U)$ satisfying ,where $y(D) \in [p, 1]$ will satisfy the required conditions ensuring that 1 is indifferent.

For player 3 to be indifferent between L and R, we require:

$$xy(D) + (1-x)y(U) = p \tag{3}$$

For any pair $y(D), y(U)$, we can select x to satisfy this equality as long as $y(D) \geq p, y(U) \leq p$. Since $y(D)$ has already been restricted to $[p, 1]$, the relevant constraint is the latter. Since $z^* < p$, there is a continuum of values satisfying this condition, and each of these possible values for $y(D)$ defines a value for $y(U)$ by (2) and a value of x by (3). ■

The significance of this result is its relation to Kreps and Wilsons' theorem [18], which shows that a generic extensive form game has finitely many sequential equilibrium outcomes. Furthermore, the set of Nash equilibria for any extensive form game consists of finitely many connected components (Kohlberg and Mertens, [17]). Observe that the Nash equilibria set out in the proof of the above proposition belong to a single connected component. Nevertheless, the game has infinitely many Nash equilibrium outcomes. Hence Γ must be a non-generic extensive form game, as defined by Kreps and Wilson. Observe that all our assumptions regarding payoffs in the two stages are written as inequalities. Hence in R^{18}, the set of payoffs such that our results apply is of positive Lebesque measure, and contains an open set.. However we have assumed a *multi-stage game*, with payoffs in the overall game being given by the *sum of payoffs in the two stages.* Put differently, our assumption corresponds to one of additive seperability in payoffs between the two stages. This constitutes the violation of Kreps and Wilsons' genericity assumption. We shall discuss this issue further later, noting only at this point that our example is a generic within the class of multi-stage games. Repeated games are an example of multi-stage games - not only are payoffs additively separable across stages as in our example, but also payoffs functions in each stage game are identical.

Note that all the equilibria are very robust in terms of standard refinement criteria. For example, it can be verified that all these equilibria satisfy the most stringent criterion of *strict perfectness* (Okada [23]), and hence can be considered to strategically stable [17] as singleton sets.

2.1 The Perturbed Game

Our basic purpose in this section is to show that if each player entertains a small amount of uncertainty about other players' payoffs, in the manner first set out by Harsanyi [14], all the history-dependent sequential equilibria of the game Γ disappear. The unique sequential equilibrium which is *robust* to this perturbation, is the history independent equilibrium $\mathbf{s}'$.

Let n be the cardinality of the set of action profiles A — in our example $n = 8$. Let $Z_i \subset R^n$ be an interval of positive Lebesgue measure for $i = 1, 2, 3$ and let X_1, X_2, X_3 be independent random variables distributed on $Z_i \subset R^n$ by probability measures μ_1, μ_2, μ_3 respectively, which are assumed to be absolutely continuous with respect to Lebesgue meaure. The disturbed game is as follows:

1. Nature picks a realization $x_i, i \in \{1, 2, 3\}$. Player i is informed of the realizations of ω_i only.

2. Players play the extensive form game Γ, with payoffs being given by x_i.

Let Ω_i be the information sets of a player in Γ. Note that Ω_i is a singleton set for players 1 and 3, and has two elements for player 2.A *behavior strategy* for player i is a Borel-measurable function $\sigma_i : \Omega_i \times Z_i \to \Delta A_i$. The *aggregate strategy* corresponding to this behavior strategy is given by $s_i := \int \sigma_i(x_i)\mu_i(x_i)d\mu_i$. Hence $s_i : \Omega_i \to \Delta A_i$. A behavior strategy profile $\sigma = (\sigma_1, \sigma_2)$ is an equilibrium if for $i = 1, 2, \forall(\omega_i, x_i) \in \Omega_i \times Z_i, \forall a_i \in A_i$,

$$\int_{Z_j} \sum_{\omega_j \in \Omega_j} \{x_i[(\sigma_i(\omega_i, x_i), \sigma_j(\omega_j, x_j)] - x_i[a_i, \sigma_j(\omega_j, x_j)]\}\pi(\omega_j)\mu_j(x_j)dx_j \geq 0$$

The following proposition shows that the disturbed version of the game Γ has an essentially unique equilibrium.

Proposition 2 *If $\sigma = (\sigma_1, \sigma_2, \sigma_3)$ is an equilibrium of the incomplete information game and if s_2 is the aggregate strategy of player 2, then s_2 is constant on Ω_i.*

Proof. It is sufficient to show that $\sigma_i(\omega_i, x_i) = \sigma_i(\omega_i', x_i)\forall\omega_i, \omega_i' \in \Omega_i$ for almost all $x_i \in Z_i$. Define $\alpha_j = \sum_{\omega_j \in \Omega_j} s_j(\omega_j)\pi(\omega_j)$. By the independence assumption, player i beliefs over the actions player j takes in A_j are given by α_j. Consider the set $\Phi_i(a_i, \alpha_j)$ defined by:

$\Phi_i(a_i, \alpha_j) := \{x_i \in Z_i : \sum_{a_j \in A_j}[x_i(a_i, a_j) - x_i(a_i', a_j)]\alpha_j(a_j) \geq 0 \forall a_i' \in A_i\}$

If $a_i \neq a_i', \Phi_i(a_i, \alpha_j) \cap \Phi_i(a_i', \alpha_j)$ has μ_i-measure zero. This follows since this is a subset of the set $\{x_i \in Z_i : \sum_{a_{ij} \in A_j}[x_i(a_i, a_j) - x_i(a_i', a_j)]\alpha_j(a_j) = 0\}$, which is a subset of a hyperplane of Lebesgue measure zero. ■

Call a sequential equilibrium s of a complete information game *robust* if there exists a sequence of disturbed games converging weakly to the complete information game, such that there exists a corresponding sequence $< \sigma_n >$ of sequential equilibria of the disturbed game, such that σ_n converges to σ. The preceding proposition implies the following proposition regarding robust sequential equilibria of the game Γ.

Proposition 3 Γ *has a unique robust sequential equilibrium where player* 1 *plays* U, *player* 2 *plays* $\boldsymbol{p}$ *at both information sets and player* 3 *plays* $\boldsymbol{q}$**.**

Proof. Immediate from proposition 2. ■

We note here that we considering disturbed games which are also multi-stage games, where the payoffs are additively separable. We now consider the implications of relaxing this assumption.

3 Genericity Issues

We have assumed so far that Γ is a multi-stage game, where the payoffs to each player in the overall game is the sum of stage game payoffs. This can also be seen as a separability assumption - that payoffs are additively separable, so that first stage actions do not affect the *payoff functions* in the second stage. This implies that Γ is not a generic extensive form game, as defined by Kreps and Wilson [18]or Kohlberg and Mertens [17]. We now consider how our results would differ if we had gave up the separability assumption. In this case, payoff functions in the second stage would have to be defined separately after the event that player 1 chooses U and the event that she chooses D. Let $(\mathbf{A}_U, \mathbf{A}_D), (\mathbf{B}_U, \mathbf{B}_D), (\mathbf{C}_{U,}\mathbf{C}_D)$ be the payoff matrices of each of the three players after the two events. We make the same qualitative assumptions regarding payoffs as before. In particular, the bi-matrix game with payoff matrices B_U and C_U has a unique Nash equilibrium in completely mixed strategies, $(\mathbf{p}_U, \mathbf{q}_U)$; similarly, $(\mathbf{p}_D, \mathbf{q}_D)$ is the unique completely mixed equilibrium corresponding to $(\mathbf{B}_D, \mathbf{C}_D)$. Observe that for generic payoff matrices $(\mathbf{C}_U, \mathbf{C}_D)$, $\mathbf{p}_U \neq \mathbf{p}_D$; similarly, for generic $(\mathbf{B}_U, \mathbf{B}_D), \mathbf{q}_U \neq \mathbf{q}_D$. The important factor, as compared to our previous analysis, is the latter inequality, i.e. the lack of separability in player $2's$ payoffs, i.e. the player who was required take history dependent actions. We continue to assume that $a_U - a_D < \mathbf{p}_D \mathbf{A}_D \mathbf{q}_D - \mathbf{e} \mathbf{A}_U \mathbf{q}_D$, which ensures that in principle player 1 can be persuaded to play action D.

Consider first the case when $q_U > q_D$. This implies that if player 3 is playing $\mathbf{q}_D$, (since he expects player 1 to play D), and player 1 in fact plays U, then player 2's unique optimal response is to play the pure action T, i.e. to play the mixed action $(1, 0)$. Hence 1 will not play $\mathbf{e}$ on observing D. The unique sequential equilibrium of this game has player 1 choosing U and player 3 choosing $\mathbf{q}_U$, while player chooses $\mathbf{p}_U$ if she observes U, and $\mathbf{e}$ if she observes D. Hence there are no history-dependent sequential equilibria in this case.

Consider now the case when $q_U < q_D$. In this case, if player 3 is playing

$\mathbf{q}_D$, and player 1 chooses U, then player $2's$ unique optimal response is to play $\mathbf{e}$. Hence there exists a history-dependent equilibrium where player 1 plays D, player 3 plays $\mathbf{q}_D$, and player 2 plays $\mathbf{p}_D$ if she observes D and plays $\mathbf{e}$ if she observes U. Furthermore, since player 2 has strict incentives to vary her behavior in response to her information, this equilibrium appears to be robust to some incomplete information about player payoffs.

Assume now that Γ has separable payoffs, but consider now the implications of a perturbation of Γ, where we allow payoff perturbations to violate the separability assumption. Specifically, assume that for player 2, his payoffs in the incomplete information game are given by a probability measure μ with support in a small interval in R^8 which contains the original payoff in the complete information game. Hence in the incomplete information game, player 2's continuation payoffs can differ between information sets U and D. Our above arguments imply that that the history dependent equilibrium can be approximated if the perturbation places a relatively large probability mass on points such that $q_U < q_D$. Hence if for some payoff perturbations, the history dependent equilibrium can be approximated. However, if the perturbation places a relatively large probability mass on points such that $q_U > q_D$, then one cannot approximate the history dependent equilibrium.

In this context one may recall the content of Harsanyi's theorem — for generic normal form games, all Nash equilibria can be obtained as limit points of *any* sequence of incomplete information game. If one allows a larger class of perturbations, so that incomplete information payoffs are not additively separable, non-robust equilibria in dynamic games can be obtained as limit points of *some* but not all sequence of incompete information games.

4 On Robustness in Mechanism Design

Our argument also applies to mechanism design when preferences are separable. Consider the standard principal agent problem. Nature selects the realization of a state θ, from a finite set of states Θ. The agent observes nature's choice, the principal does not. The principal must choose an action a from a finite set A, and may make a transfer $t \in \mathbf{R}$. The principal's utility is given by $U(a, \theta, t)$. The agent's utility is assumed to be quasi-linear, and is given by $V(\theta, a) + t$. The principal-agent problem studies the design of a mechanism which is optimal from the principal's point of view. A direct mechanism is one where the agent makes a report, $m \in \Theta$. The principal commits to a mechanism, i.e. to a pair (ϕ, t), where $\phi : \Theta \to \Delta(A)$ and $t : \Theta \to \mathbf{R}$. A direct mechanism (ϕ, t) is *incentive-compatible* if $\forall \theta, \theta' \in \Theta$,

$$V(\phi(\theta), \theta) + t(\theta) \geq V(\phi(\theta'), \theta) + t(\theta')$$

We shall assume henceforth that preferences are *separable,* i.e. $V(a, \theta) = g(\theta) + f(a)$.

Proposition 4 *Let $\phi : \Theta \to \Delta(A)$ be an arbitrary function. If preferences are separable, then ϕ is implementable.*

Proof. Let θ_0 be an arbitrary state and set $t(\theta_0)$ to an arbitrary number t_0. For any $\theta \in \Theta$,define $t(\theta) = t_0 + f(\phi(\theta_0) - f(\phi(\theta)$. Let the realized state be θ. By announcing $m = \theta$, the agent gets utility

$$g(\theta) + f(\phi(\theta)) + t_0 + f(\phi(\theta_0) - f(\phi(\theta)) = g(\theta) + f(\phi(\theta_0) + t_0 \qquad (4)$$

Whereas, by announcing θ', she gets

$$g(\theta) + f(\phi(\theta')) + t_0 + f(\phi(\theta_0) - f(\phi(\theta')) = g(\theta) + f(\phi(\theta_0) + t_0 \qquad (5)$$

Since the payoffs to announcing θ and θ' are equal, it is optimal to announce $m = \theta$, for every θ. Hence ϕ is implementable. ■

4.1 Incomplete information

The above proposition depends critically upon the assumption that the principal has full knowledge of the agent's utility function, and therefore make her indifferent between the different actions that the principal takes under ϕ. We now assume that there is incomplete information about the agent's preferences, specifically f. Under complete information, if A has k elements, then $f \in R^k$. To allow for incomplete information, let $F \subset R^k$ and assume that the agent's preferences are selected by nature according to the probability measure μ. We make the following assumption about F and μ.

A1. F is an interval of positive Lebesgue measure and μ is absolutely continuous with respect to Lebsegue measure.

A direct mechanism in this context is as follows. The principal commits to a mechansim, (Φ, T) where $\Phi : \Theta \times F \to \Delta(A)$ and $T : \Theta \times F \to \mathbf{R}$.

The agent reports a pair (m, u), where $m \in \Theta$ and $u \in F$. The principal takes the (random) action $\Phi(m, u)$ and makes the transfer $T(m, u)$. A direct mechanism (Φ, T) is implementable if $\forall(\theta, f) \in \Theta \times F, \forall(m, u) \in \Theta \times F$,

$$f(\Phi(\theta, f)) + T(\theta, f) \geq f(\Phi(m, u)) + T(m, u)$$

Given any mechanism (Φ, T), let $\phi(.; f, (\Phi, T)) : \Theta \to \Delta(A)$ be defined by $\phi(\theta) = \Phi(\theta, f)$.

The following theorem is the main result of this paper.

Theorem 5 *If* (Φ, T) *is implementable, then* $\phi(.; f, (\Phi, T))$ *is a constant function on* Θ *for almost all* $f \in F$.

To prove this theorem, we define the following property of a mechansim.

Definition 6 Φ *is* monotone *if* $\forall(\theta, f), (\theta', f') \in \Theta \times F$,

$$[\Phi(\theta, f) - \Phi(\theta', f')][f - f'] \geq 0$$

Intuitively, if an action $a \in A$ is assigned higher utility under f than under f', with the payoffs to all other actions being the same under both preferences, then a montone mechanism must (weakly) increase the probability with which a is chosen when the agent reports f rather than f'. This must be done *independently* of the report regarding the state θ.

The following lemma says that one may restrict attention to monotone mechanisms.

Lemma 7 *If a mechanism* (Φ, T) *is implementable, then* Φ *is monotone.*

Proof. Suppose that the realized state is (θ, f), and the agent considers announcing (θ', f'). Incentive compatibility requires

$$[\Phi(\theta, f) - \Phi(\theta', f')]f \geq T(\theta', f') - T(\theta, f) \tag{6}$$

Similarly, suppose that the realized state is (θ', f'), and the agent considers announcing (θ, f). Incentive compatibility requires

$$[\Phi(\theta', f') - \Phi(\theta, f)]f' \geq T(\theta, f) - T(\theta', f') \tag{7}$$

■

Adding these inequalities we get

$$[\Phi(\theta, f) - \Phi(\theta', f')][f - f'] \geq 0$$

which is the required condition. ■

Proof of the theorem.

Let $i \in \{1, 2., , , k\}$, and fix $h_{-i} = (h_1, h_2, ., h_{i-1}, h_{i+1}.., h_k)$. Let $L_i(h_{-i}) = \{f = (f_1, f_2, ..., f_k) \in F : f_j = h_j \forall j \neq i\}$. $L_i(h_{-i})$ is the intersection with F of the line with coordinates h_{-i} which is parallel to the i-th coordinate axis.

Let Φ_i be the i-th component of Φ, i.e. the probability assigned by Φ to the i-th action.Fix an arbitrary state $\bar{\theta}$, and consider the restriction of Φ_i to the set $\{\bar{\theta}\} \times L_i(h_{-i})$. Hence $\Phi_i : \{\bar{\theta}\} \times L_i(h_{-i}) \to [0, 1]$ is a real valued function defined on an real interval.
Since Φ is monotone, Φ_i is an increasing function, and has at most a countable number of discontinuities on $L_i(h_{-i})$.

We now show that if $\Phi_i(\theta', h) \neq \Phi_i(\theta'', h)$ for some $\theta', \theta'' \in \Theta$ and $h \in F$, then $\Phi_i : \{\bar{\theta}\} \times L_i(h_{-i}) \to [0, 1]$ must have a discontinuity at h. This follows from the fact that Φ is monotone. Let $\underline{p} = \min_\theta \Phi_i(\theta, h)$ and let $\bar{p} = \max_\theta \Phi_i(\theta, h)$. Since $\Phi_i(\theta', h) \neq \Phi_i(\theta'', h), \underline{p} < \bar{p}$. Consider $h^-, h^+ \in L_i(h_{-i})$ such that $h_i^- < h_i$ and $h_i^+ > h_i$. Since Φ is monotone, we have that $\Phi_i(\bar{\theta}, h^-) \leq \underline{p}$ and $\Phi_i(\bar{\theta}, h^+) \geq \bar{p}$. Hence $\Phi_i : \{\bar{\theta}\} \times L_i(h_{-i}) \to [0, 1]$ has a discontinuity at h.

Let $E_i(h_{-i}) = \{ h_i : \exists \theta, \theta' \in \Theta : \Phi_i(\theta, (h_i, h_{-i})) \neq \Phi_i(\theta', (h_i, h_{-i}))$. $E(h_{-i})$ is the set of points on the cross-section at h_{-i} such that Φ_i is not constant on Θ.Since this is countable, we have established that for any h_{-i}, $\lambda(E(h_{-i})) = 0$, where λ is Lebesgue measure on the line.

Let λ^z denote Lebesgue measure on $\mathbf{R}^z$. Let $E_i = \{ h \in F : \exists \theta, \theta' \in \Theta : \Phi_i(\theta, h) \neq \Phi_i(\theta', (h)\}$. $\mu(E_i) \leq \int \lambda d\lambda^{k-1} = 0$. Let $E = \{ h \in F : \exists \theta, \theta' \in \Theta : \Phi(\theta, h) \neq \Phi(\theta', (h)\} = \cup_{i=1}^k E_i$. Since E is the finite union of sets of measure zero, $\mu(E) = 0$. ■

5 General Setting

Γ is a multi-stage game

I : set of players

A_i^t : set of actions at stage t — independent of what happens at earlier dates.

$A^t = \times_{i \in I} A_i^t$ — set of action profiles at date t, with cardinality $n(t)$.

Nature may also move at any date.

Ω_i^t : set of possible signals observed by a player at date t (this includes his own action).

$H_i^t : \Omega_i^1 \times \Omega_i^2 \times ... \times \Omega_i^{t-1}$

Player's information (history) at date t is h_i^t

H_i^t set of all possible histories at t

$u_i^t : A^t \rightarrow R$ — only depends upon current actions. Total payoff is additively separable:

$$U_i(< a^t >_{t=1}^T) = \sum_{t=1}^{T} \delta^{t-1} u_i^t(a^t) \tag{8}$$

A strategy at date t for player i is given by

$s_i^t : H_i^t \longrightarrow \Delta(A_i^t)$

$s_i =< s_i^t >_{t=1}^T$

Let Σ_i^t be the set of continuation strategies for a player at date t. I.e. $\sigma_i^t \in \Sigma_i^t$ specifies a current mixed action and a sequence of mappings $\sigma_i^{\tau t}$: $\Omega_i^t \times ... \times \Omega_i^{\tau-1} \longrightarrow \Delta(A_i^\tau)\ \forall \tau > t$.

Any strategy, in conjunction with a history, defines a continuation strategy. Write $s_i(h_i^t) = s_i'(\hat{h}_i^t)$ if the continuations strategies so defined are identical

Given any strategy profile, s_{-i},for all players except player i, and any history h_i^t, let $V^i(.|h_i^t) : \Sigma_i^t \rightarrow R$ denote the continuation payoff function for player i given that he is at history h_i^t.

Write $V^i(.|h_i^t) = V^i(.|\hat{h}_i^t)$ if for any $\sigma_i^t \in \Sigma_i^t, V^i(\sigma_i^t|h_i^t) = V^i(\sigma_i^t|\hat{h}_i^t)$

Definition 8 *An equilibrium strategy profile* $(s_i)_{i \in I}$,*is* robust *if for any player* i *and any date* t, $h_i^t, \hat{h}_i^t \in H_i^t$ *and* $V^i(.|h_i^t) = V^i(.|\hat{h}_i^t) \Rightarrow s_i(h_i^t) = s_i(\hat{h}_i^t)$.

6 Applications

We now turn to some applications of our analysis to equilibria in games and contracts.

6.1 Repeated games with private monitoring

Repeated games, where players monitor each other's actions via private signals, are currently the subject of much research. In such games, it is difficult to sustain cooperative outcomes via pure strategy equilibria. Consider for example the repeated prisoners' dilemma where each player observes a private signal which is informative about their rival's action. If the private signals are independent, the profile where each player plays the grim trigger strategy fails to be a Nash equilibrium. The reason for this is straightforward: with independent signals and pure strategies, a player's belief about his opponent's continuation strategy at date $t = 2$ is independent of the realilzation of his private signal. In consequence, if a player has a strict incentive to continue to cooperate when he receives a good signal, he will similarly have a strict incentive to continue to cooperate when he receives a bad signal. Hence it is not optimal to switch to defection on observing a bad signal, and the trigger strategy profile cannot be an equilibrium.

This difficulty demonstrates that one must rely on mixed strategies in order to support cooperative outcomes. The literature has generated two distinct lines of approach. The essence of the first approach, exemplified by Bhaskar and van Damme [5], Sekiguchi [25], Mailath and Morris [21] and Bhaskar [4] and Obara [22], is to allow the players to randomize in the initial period, between different repeated game pure strategies. Since player i does not know player j' s pure strategy, his beliefs about j's continuation strategy will now be responsive to the private signal that he observes. Hence one can construct equilibria where player's have strict incentives to play their equilibrium strategy, except in the initial period. Furthermore, the randomization in the intial period is at a single information set, and hence satisfies the criterion of robustness set out in the present paper.

A second approach, exemplified by Piccione [24] and Ely and Valimaki [9] in the context of the repeated prisoners' dilemma, is to have players randomize in almost every period, where the randomization probabilities of player i are chosen so as to make player j indifferent between cooperation (C) and defection (D) in every period. Since j is so indifferent, j will also be willing to randomize systematically across information sets in such a way as to make i indifferent. In this appoach, player i is indifferent between C (cooperate) and D (defect) at every information set; however, the probability with which he plays C must differ across information sets — it must be greater when he has observed a good signal than when he has observed a bad signal, so as to make j also indifferent. In other words, i plays different continuation strategies σ and σ' at information sets ω and ω' respectively, only because he is indifferent between both this strategies at both these information sets. Hence the equilibria so constructed are not robust to payoff perturbations.

Finally, we also note that our analysis is also relevant to the analysis of Compte [6] and Kandori and Matsushima [16]. These authors allow players to communicate their private information at the end of each period, and construct equilibria where continuation strategies depend only upon the messages communicated. Since such communication is public, they can employ the methods of analyzing perfect public equilibria and the techniques developed in Fudenberg, Levine and Maskin [13]. The critical problem in this approach is to ensure truthful communication of private signals. Whereas strict incentives for truth telling can be ensured if there are three or more players and signals are correlated, strict incentives cannot be ensured if signals are independent. For the case of two players, independent signals are required in order to prove the folk theorem in these papers. In the latter case, equilibria and the associated continuation payoffs are designed so that each player is indifferent between telling the truth and lying, no matter what the realization of his private signal. If payoffs are perturbed, the player will choose the same message irrespective of the realization of his private signal, and hence truth-telling cannot be ensured.

6.2 Contracts

We have two applications to contract theory: credence goods and double moral hazard in employment contracts.

A credence good is one where the buyer has no information on the quality of service provided by the seller — see Darby and Karni [8] and Emons [10],[11]. Doctors, mechanics and dentists are typical examples. More formally, suppose that the agent (the supplier) observes whether state θ or θ' has occured, and the principal (the customer) would like the agent to choose a at θ and a' at θ'. Emons defines a credence good as one where the state is payoff irrelevant for the agent — hence in such a case, additive separability as defined in section 4 is satisfied. The mechanism that solve the credence good problem is one where the supplier is made indifferent between both a and a' at either state, and breaks this indifference by truth-telling. Clearly, this hinges on common knowledge of the supplier preferences as section 4 makes clear. Our view is that one must rely on non-economic factors, such as the Hippocratic oath, to sustain truth telling in such contexts.

Consider now the double moral hazard problem, as in Malcomson [19], [20]. Suppose that a firm employ's n workers, who must choose between high and low effort. The firm does not observe the effort choice, but observes a non-verifiable vector of signals, such as individual outputs. In order to induce high effort, the individual worker's wage must be made contingent on the signal vector. However, since the signal is non-verifiable, the firm must be indifferent between all vectors of wage payments that are feasible under the contract. Malcomson proposes a tournament — the total wage bill is fixed under the contract, but the firm can vary individual wage payments in line with the observed signals. Since the firm is assumed to be concerned only its total payments, it is indifferent between all vectors of wage payments that are feasible under the contract, and can therefore be relied upon to honour the contract. Suppose however that workers differ along some dimension (eg. sex) and suppose that the firm may have discriminatory preferences. This in itself creates no problem — given the firm's utility function, one can simply make the firm indifferent, in utility space, between all contractually feasible wage vectors. However, this relies upon the firm's utility function being common knowledge, and as section 4 demonstrates, the slightest uncertainty about firm preferences suffices to rule out all possibilities of supporting high effort. In this context, the literature on repeated games with private monitoring shows that if the worker and the firm repeatedly interact, one can sustain high effort equilibria even if the signal of worker effort is privately observed by the firm.

7 Chain store paradox

We now consider a repeated game between a long run player and a sequence of short-run players, as in Fudenberg and Levine [12]. We show that if the short run player has finite memory, i.e. she only observes the outcomes in the last m periods, then the long run player loses all committment power. In any robust sequential equilibrium, the backward induction outcome of the stage game must be played in every period.

Our analysis is perfectly general, but consider the chain store game for concreteness. The long run player is the incumbent firm, infinitely lived. In each period, she plays the same game with a short run player. The short run player at date t can either choose IN or OUT. If she chooses OUT, the stage game ends, and if she chooses IN, the long run player must choose between accomodate (A) and fight (F). Payoffs are as given below, and the *outcome* is the realized element from the set {OUT,A,F}, i.e the long run player's choice is not observed if OUT is chosen.

	A	F
IN	1,0	-1,-c
OUT	0,1	0,1

Consider first the case where the entrant has one-period memory, i.e.only observes the outcome in one period — the argument generalizes for arbitrary m period memory on the part of entrant. The incumbent however has infinite memory, and observes all past event. Hence at any date $t > 1$, the entrant's history $\in$ {OUT,A,F}. Consider the following pure strategies for the entrants, ρ :

OUT at $t = 1$

At $t > 1$: IN if a_{t-1} =A; OUT if $a_{t-1} \in$ {OUT, F}

For the incumbent, consider the strategy σ :

F as long as he as never played A; A otherwise.

Althoug this strategy profile is a Nash equilibrium if δ is sufficiently large ($\delta > \frac{c}{1+c}$), this is not a sequential equilibrium. At information set where $a_{t-1} = A$, it is not optimal to play A today if the entrant plays IN. By playing A, the incumbent's total discounted payoff is 0, whereas by deviating and fighting, he gets:

$$V(F, \rho) = -(1-\delta)c + \delta > 0 = V(A, \rho) \tag{9}$$

However, if incumbent plays F even after he has played A,the entrant should not enter even when he observes accomodation.

But if the entrant never enters, even if he observes $a_{t-1} = A$, the incumbent always plays A, and hence the entrant always plays IN

This argument is an example of the point made in Ahn [1] about pure strategy equilibria — generically, there exists no pure strategy sequential equilibrium where entry is deterred, with finite memory.

7.1 Mixed Equilibrium

We now consider mixed equilibria.

Entrant's strategy ρ :
At $t = 1$, play OUT
At $t > 1$, OUT if $a_{t-1} \in$\{OUT,F\}; IN with probability q if $a_{t-1} =$A
Incumbent's strategy σ : F with prob $p = \frac{1}{2}$ if entrant enters.

Note that the entrant is indifferent between IN and OUT at any information set, given incumbent's strategy. If the entrant enters

$$V(F, \rho) = -(1 - \delta)c + \delta \tag{10}$$

$$V(A, \rho) = \delta\{qV(F) + (1 - q)1\} \tag{11}$$

$$q = \frac{\delta - (1 - \delta)c}{\delta + \delta^2 - \delta(1 - \delta)c} \tag{12}$$

Hence σ is optimal for this value of q. This is an example of an equilibrium where no entry takes place on the path of play.

Note that the incumbent plays the same strategy p at every information set. So his strategy is robust. However, the entrant's payoff satisfies

$$U(IN|A) = U(OUT|A) = U(IN|F) = U(OUT|A) \tag{13}$$

However, the incumbent's continuation strategy differs across the information sets where he has played A in the previous period and where he has played F in the previous period — his continuation strategy is OUT at F, q at A. Hence entrant's equilibrium strategy is not robust.

7.2 Mixed Equilibrium 2

Alternatively, one may construct a different mixed equilibrium by modifying the incumbent's strategy so that it also depends upon (one period history).
If $a_{t-1} =$OUT or F, F with prob. 1
If $a_{t-1} =$A, F with prob $\frac{1}{2}$

Now the entrant's beliefs about the continuation strategy of the incumbent vary — differ between A and \{OUT,F\}. Hence entrant's strategy is robust.

For the incumbent,

$$V(F|IN, F) = V(A|IN, F) = V(F|IN, A) = V(A|IN, A) \tag{14}$$

But his continuation strategy differs between the information sets (IN,F) and IN,A).

Hence this equilibrium is not robust.

Suppose that in each period, both the long-run and short run players' payoffs are random. These are independent across players and across time. Both players are informed privately of their realizations in that period before they move.

Proposition 9 *Let G be a generic two-player extensive form game of perfect information played where the short run player moves once and the long run player moves once. Let $G^{\infty}(\delta, m)$ be the game between the long run player and an infinite sequence of short run players, with finite memory m. In any robust equilibrium of $G^{\infty}(\delta, m)$, the backward induction outcome of G is played every period.*

Proof. h_t and h'_t are $k-$equivalent ($h_t \sim_k h'_t$) if the last m action profiles are identical.

1) Assume that $h_t \sim_k h'_t \Longrightarrow \rho_{t+1}(h_t) = \rho_{t+1}(h'_t)$. Then $h_{t-1} \sim_{k-1} h'_{t-1} \Longrightarrow \rho_t(h_{t-1}) = \rho_t(h'_{t-1})$.

Let a_s be any action of the short run player.

a) If σ is robust, and $h_{t-1} \sim_{k-1} h'_{t-1}, \sigma(\ h_{t-1}a_s) = \sigma(h'_{t-1}a_s)$

If ρ_t is robust, then (a)$\Rightarrow \rho_t(h_{t-1}) = \rho_t(h'_{t-1})$

2) $\forall t > m, h_t \sim_m h'_t \forall h_t, h'_t$.

1 & 2$\Longrightarrow \rho_t$ is history independent for $t > m$.

Hence we must have the backwards induction outcome in each period ■

In other words, with imperfect information, the long run player loses her committment power entirely. This arises since the long run player has private information about past events, since the short run player does not observe the entire history. A similar phenomenon arises in the overlapping generations model with imperfect information — see Bhaskar [3] for details.

7.3 Reputation

We now show that if there are reputation effects, this may restore the power of the long run player (with perfect information of histories, reputation is not required, since there is always an equilibrium when entry is deterred).

Prob θ that the long run player is crazy type, will always fight entry.

Prob $1 - \theta$ that payoffs are of rational type.

1 period memory
Strategy for entrant ρ
at date $t = 1$, OUT
at $t > 1$, OUT if $a_{t-1} \in \{$OUT,F$\}$
at $t > 1$, IN with prob q if a_{t-1} =A.
Rational incumbent follows Markov strategy — play q if in
Now, entrant's strategy differs between the information sets A and $\{$OUT,F$\}$
But beliefs about the incumbent also differ.

$$\mu(F|A) = q \tag{15}$$

$$\mu(F|ForOUT) \geq \theta + (1-\theta)q > q \tag{16}$$

Since beliefs vary across these informatioins sets, this equilibrium is robust. Hence with a positive probability of a crazy type, there exists a robust equilibrium with entry-deterrence.

8 Conclusion

Harsanyi's famous theorem on mixed strategies has, in the main, been a permissive result. In generic normal form games, every Nash equilibrium is robust to payoff perturbations and can be purified. The examples of equilibria which are not firm in normal form games are relatively uninteresting, since they usually hinge upon players have the same payoffs from different actions, and there is no specific story why these equalities may hold. In extensive form games, the incomplete information approach may provide an alternative way to refine imperfect Nash equilibria, as van Damme [26] notes. Nevertheless, if one is content to use a more refined equilibriuum concept such as sequential equilibrium, there did not seem to be any reason to take the incomplete information route. This paper has argued that Harsanyi's critique has powerful implications in the context of games where payoffs are additively separable across stages — standard repeated games are a leading example. This critique is of special relevance in games and contracts where players have private information about past events, as our applications demonstrate.

References

[1] Ahn, I., 1997, Repeated games with a long run and short-run players and bounded memory, mimeo, University of Pennsylvania.

[2] Aumann, R., 1987, Correlated Equilibrium as an expression of Bayesian rationality, *Econometrica*, **55**, 1-18.

[3] Bhaskar, V., 1998, Informational constraints and the overlapping generations model: Folk and Anti-Folk theorems, *Review of Economic Studies* 65, 135-149.

[4] Bhaskar, V., 1999, Sequential equilibria in the repeated prisoners' dilemma with private monitoring, mimeo, University of Essex.

[5] Bhaskar, V. , and E. van Damme, 1997, Moral hazard and private monitoring, mimeo, University of Essex.

[6] Compte, O., 1999, Communication in repeated games with imperfect private monitoring,*Econometrica*

[7] Darby, M.R., and E. Karni, 1973, Free Competition and the Optimal Amount of Fraud, *Journal of Law and Economics* 16, 67-88.

[8] Dubey, P., and M.Kaneko, 1984, Information Patterns and Nash equilibrium in extensive games, *Mathematical Social Sciences.*

[9] Ely, J., and J. Valimaki, A robust folk theorem for the repeated prisoners' dilemma, mimeo, University of Southampton.

[10] Emons, W., 1997, Credence Goods and Fraudulent Experts, *Rand Journal of Economics* 28, 107-119.

[11] Emons, W., 1997, Credence Goods Monopolists, mimeo, University of Bern.

[12] D. Fudenberg and D. Levine, 1989, Repeated games with a short run and long run player, *Econometrica.*

[13] D. Fudenberg, D. Levine and E. Maskin, 1994, The Folk Theorem with Imperfect Public Information, *Econometrica,* 62 , 997-1040.

[14] Harsanyi, J., 1973, Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points, *International Journal of Game Theory*, **2**, 1-23.

[15] Kandori, M., 1991, Cooperation in finitely repeated games with imperfect private information, mimeo.

[16] Kandori, M and H. Matsushima, 1999, Private observation, communication and collusion,*Econometrica*

[17] Kohlberg, E., and J.F.Mertens, 1986, On the strategic stability of equilibria, *Econometrica*, **54**, 1003-1037.

[18] Kreps, D., and R. Wilson, 1982, Sequential equilibrium, *Econometrica*, **50**, 863-894.

[19] Malcomson, J., 1984, Work incentives, hierarchy, and internal labormarkets, *Journal of Political Economy,* 92, 486-507.

[20] Malcomson, J., 1986, Rank-order contracts for a principal with many agents, *Review of Economic Studies*, 53, 807-817.

[21] Mailath, G., and S. Morris, 1999, Repeated games with almost public monitoring, mimeo, University of Pennsylvania.

[22] Obara, I., 1999, The repeated prisoners' dilemma: The N-player case, mimeo, University of Pennsylvania.

[23] Okada, A, 1984, On strictly perfect equilibrium points, *International Journal of Game Theory*,

[24] Piccione, M., 1998, The repeated prisoners' dilemma with imperfect private monitoring, mimeo, LSE.

[25] Sekiguchi, T., 1997, Efficiency in the repeated prisoners' dilemma with private monitoring, *Journal of Economic Theory*,

[26] van Damme, E., 1987, *Stability and Perfection of Nash Equilibrium*, Springer-Verlag, Berlin.